

UNITED KINGDOM · CHINA · MALAYSIA

The measurement of particle dispersions in turbulent,

four-way coupled flows

Matthew Yates, MSci. BA

University of Nottingham

Department of Chemical and Environmental Engineering

Supervised by: Dr David Hann, Dr. Buddhika Hewakandamby, Prof. Barry Azzopardi

Sponsored by Chevron as part of the Transient Multiphase Flows Consortium

Thesis submitted to the University of Nottingham in fulfilment of the Degree of Doctor of Philosophy

This thesis is dedicated to the memory of Gordon Yates, who shall be forever missed.

Acknowledgements

This work was conducted as part of the Transient Multiphase Flows Consortium, a collaboration of universities and industrial sponsors. The author wishes to thank all parties involved, especially the contribution of Chevron and Dr. Lee Rhyne, without whom this work would never have happened.

As part of this collaboration, I had the pleasure of working with the CFD team at Cranfield. I would like to thank Chris Thompson, Robert Sawko and Marcin Traczyk. Our discussions were always deeply appreciated, and did a great deal to further my knowledge.

Of vast importance to this project were the Engineering department technicians, who went above and beyond to get experimental rigs up and running after failures. Therefore, special thanks go to Mel, Mick, Fred, and anyone else who lent a helping hand during this project.

I would also like to express my gratitude towards David Hann. Your support and guidance, particularly during times of stress, were greatly appreciated.

I would also like to thank my family for their patience, particularly my Mother. Completing a PhD can be a life-consuming activity, and I promise to spend more time with them by way of celebrating an accomplishment which, without them, I would never have achieved.

Finally, and most importantly, I would like to thank my partner Shamal. Your support and encouragement kept me driving on even when it seemed like there was a mountain to climb, and your company on late nights working kept me sane. For these things, and many many more, I will be eternally grateful.

iii

Abstract

This work contained in this thesis is the result of an industrial and academic collaboration, designed to investigate and further the present knowledge of dense turbulent dispersions. Experiments were conducted to provide support and experimental validation to a CFD code being simultaneously developed, which was able to give insight into these types of flows. Additional to this support, the aim of this thesis was to also further knowledge of key topics in this field.

The experimental methodology chosen was to use a mixture of Particle Image Velocimetry and Particle Tracking Velocimetry. To discriminate between particle and liquid phases, two approaches were adopted, depending upon the experiment. In one approach, fluorescent dyes were used to tag one phase, whilst optical filters were applied to the camera lenses. In the second approach, a size-based binary mask was applied to a single image, in order to remove phase information and produce two sets of images.

A number of different analysis techniques were researched and developed as part of this thesis. The performance of particle tracking algorithms was assessed to ascertain their most suitable usage. A number of different algorithms, designed to characterise particle positions, were validated against known test cases. These included the Box Counting Method, a Voronoi analysis, and Radial Distribution Functions. A further technique, known as the Particle Potential method, was also developed to characterise local clustering.

Two experiments were undertaken throughout this project, both of which were developed from scratch so that full control was assured over all experimental parameters. A vertical channel experiment was designed to assess the injections of particles into a rectangular channel. These experiments allowed for an ideal test case of highly concentrated particles, without the need to achieve optical visibility through a dense solution. The experiments also provided an early test of a Refractive Index Matching candidate pair; hydrogel particles and water.

The second experiment was known as the Circulating Dispersion Rig, which was designed to pump a slurry in a continuous loop in a cylindrical pipe. These experiments, due to the geometry used and dense nature of the slurry, were reliant upon trying to achieve optimum

iv

optical visibility, and so hydrogel/water mixtures were tested in advance against other, more well-utilised pairings.

The experiments conducted have provided some insight into the nature of particles in turbulent flows, in particular their clustering properties. Clustering was assessed under various concentrations. Key results included analysis of these clusters using a Voronoi diagram technique, which identified four key types of cluster structure, and the parameters under which these form. Collision probabilities of particle pairs were also assessed, using Particle Tracking data and computation of relative velocities. Such information is of importance for experimental validation of CFD codes relating to dispersed two-phase flows, where particle-particle coupling must be assessed in order to provide accurate solutions.

The key drive towards the future, should further experiments be desirable, would be to investigate the improvement of optically matching liquids and solids, which was felt to be the limiting factor towards achieving measurements at even higher concentrations. However, these experiments show some progress can be made in making measurements of four-way coupled turbulent flows.

Contents

Chapter	r 1. Introduction	1
1.1.	Background	1
1.2.	Multiphase flow terminology	2
1.3.	CFD methods and validation	4
1.4.	Project context	6
1.5.	Aims and objectives	7
1.6.	Thesis outline	9
Chapter	r 2. Literature review of turbulent dispersions phenomena	11
2.1.	Preferential accumulation	11
2.1	1.1. Inertial particle behaviour	11
2.1	1.2. Studies of preferential accumulation	13
2.2.	Interphase coupling	
2.3.	Turbulence modulation	23
2.4.	Conclusion	25
Chapter	r 3. Experimental methods	27
3.1.	Optical techniques	27
3.2.	Particle Image Velocimetry – principles and practice	
3.2	2.1. Tracer particles	31
3.3.	PIV analysis	32
3.4.	PIV reliability and error estimation	
3.5.	Phase discrimination in two phase PIV	
3.5	5.1. Application of wavelength filtering	43
3.6.	Conclusion	43
Chapter	r 4. Particle tracking and phase discrimination	45
4.1.	Information flow	45

4.2	2. Pa	article identification and tracking	46
4	4.2.1.	Particle identification	47
4	4.2.2.	Watershed algorithm and improving particle estimation	49
4.3	8. Pa	article tracking code comparisons	51
4	4.3.1.	Algorithms	52
4	4.3.2.	Comparison of performance	54
4	4.3.3.	Minimum travel constraint	57
4	4.3.4.	Algorithm verification using simulated particle movement	59
4	4.3.5.	Extension to elongated particle tracks	66
4.4	I. Pł	nase discrimination	66
2	4.4.1.	Methods from literature	66
2	4.4.2.	Size-based binary mask	68
4.5	5. G	eometric transformations between two camera viewpoints	76
2	4.5.1.	Transformation types	76
4	4.5.2.	Implementation of geometric transformations	77
4.6	6. Co	onclusion	81
Chap	ter 5.	Vertical channel rig - design	84
5.1	L. Te	est section design	85
5.2	2. Ex	perimental rig design	89
ļ	5.2.1.	Pump and flow meter choice and flow rate characteristics	91
!	5.2.2.	Particle-injection system	92
!	5.2.3.	Particle filtration	93
5.3	8. Ex	perimental operation and parameters	95
ļ	5.3.1.	Calibration grid	95
ļ	5.3.2.	Data capture – LabVIEW	96
ļ	5.3.3.	Pump stability	97

5.3	3.4.	Experimental operation	98
5.4.	Ехр	erimental parameters	100
5.5.	Inje	ection rate experiments – single phase characterisation	101
5.6.	Cor	nclusion	104
Chapte	r 6.	Analysis techniques	106
6.1.	Sim	ulated particle distributions	106
6.2	1.1.	Distribution descriptions	106
6.2	1.2.	Distribution properties	107
6.2	1.3.	Particle concentration	108
6.2.	Вох	counting method	110
6.2	2.1.	Overview of box counting method	110
6.2	2.2.	Box counting methodology	112
6.2	2.3.	Further particle distributions	114
6.2	2.4.	Process of calculation	118
6.3.	Vor	onoi cell analysis method	119
6.3	3.1.	Voronoi diagram formulation and statistics	119
6.3	3.2.	Standard deviations of cell-size distributions for simulated particles .	121
6.3	3.3.	PDF shape	122
6.3	3.4.	Voronoi code structure	126
6.4.	Par	ticle-pair potential method	129
6.5.	Rad	lial distribution functions	132
6.5	5.1.	Monte Carlo methods	133
6.5	5.2.	Radial Distribution Functions of simulated particle distributions	136
6.6.	Cor	clusion	141
Chapte	r 7.	Vertical channel results	143
7.1.	Dis	persed phase distribution	143

7.1	.1.	Box counting method	143
7.1	.2.	Voronoi analysis method	147
7.1	.3.	Further Voronoi analysis	149
7.1	.4.	Local cluster method	152
7.2.	Par	ticle pair statistics	165
7.3.	Cor	nclusions	167
Chapter	8.	Circulating dispersions rig - design	171
8.1.	Opt	tical transparency of hydrogel-water mixtures	171
8.2.	Opt	tically matched liquid-solid pairs	179
8.3.	Rig	design considerations	182
8.3	.1.	Equipment	185
8.3	.2.	Instrumentation and rig control	191
8.3	.3.	Parameter space and data acquisition	196
8.4.	Sin	gle phase characterisation	198
8.4. 8.5.	Sing Cor	gle phase characterisation	198 203
8.4. 8.5. Chapter	Sing Cor 9.	gle phase characterisation nclusion Circulating dispersions rig - results	198 203 205
8.4. 8.5. Chapter 9.1.	Sin Cor 9. Box	gle phase characterisation nclusion Circulating dispersions rig - results	198 203 205 205
8.4. 8.5. Chapter 9.1. 9.2.	Sin Cor 9. Box Par	gle phase characterisation nclusion Circulating dispersions rig - results counting method ticle-particle closing times	198 203 205 205 207
8.4. 8.5. Chapter 9.1. 9.2. 9.3.	Sin Cor 9. Box Par Par	gle phase characterisation nclusion Circulating dispersions rig - results counting method ticle-particle closing times ticle-particle Radial Distribution Functions	198 203 205 205 207 208
8.4. 8.5. Chapter 9.1. 9.2. 9.3. 9.4.	Sing Cor 9. Box Par Par Loc	gle phase characterisation nclusion Circulating dispersions rig - results counting method ticle-particle closing times ticle-particle Radial Distribution Functions al cluster method	198 203 205 205 207 208 209
8.4. 8.5. Chapter 9.1. 9.2. 9.3. 9.4. 9.4	Sing Cor 9. Box Par Par Loc	gle phase characterisation nclusion Circulating dispersions rig - results counting method ticle-particle closing times ticle-particle Radial Distribution Functions cal cluster method Local cluster method discrepancies	198 203 205 205 207 208 209 214
8.4. 8.5. Chapter 9.1. 9.2. 9.3. 9.4. 9.4	Sing Cor 9. Box Par Loc .1.	gle phase characterisation nclusion Circulating dispersions rig - results counting method ticle-particle closing times ticle-particle Radial Distribution Functions al cluster method Local cluster method discrepancies Local cluster connectivity	198 203 205 205 207 208 209 214 215
8.4. 8.5. Chapter 9.1. 9.2. 9.3. 9.4. 9.4 9.4	Sing Cor 9. Box Par Loc .1. .2. Cor	gle phase characterisation nclusion Circulating dispersions rig - results c counting method ticle-particle closing times ticle-particle Radial Distribution Functions al cluster method Local cluster method discrepancies Local cluster connectivity	198 203 205 205 207 207 209 214 215 217
8.4. 8.5. Chapter 9.1. 9.2. 9.3. 9.4. 9.4 9.4 9.5. Chapter	Sing Cor 9. Box Par Par Loc .1. .2. Cor 10.	gle phase characterisation clusion Circulating dispersions rig - results counting method ticle-particle closing times ticle-particle Radial Distribution Functions al cluster method Local cluster method discrepancies Local cluster connectivity nclusions Conclusions and future work	198 203 205 205 207 207 208 209 214 215 217 219
8.4. 8.5. Chapter 9.1. 9.2. 9.3. 9.4. 9.4 9.4 9.4 9.5. Chapter 10.1.	Sing Cor 9. Box Par Loc .1. .2. Cor 10.	gle phase characterisation clusion Circulating dispersions rig - results counting method ticle-particle closing times ticle-particle Radial Distribution Functions ticle-particle Radial Distribution Functions ticle cluster method Local cluster method discrepancies Local cluster connectivity conclusions and future work thesis conclusions	198 203 205 205 207 207 207 209 219 219 219 219

10.1.2	2. Experimental methodology	220
10.1.3	 Particle tracking methods 	220
10.1.4	 Masking and geometric transforms 	221
10.1.5	5. Vertical channel rig	222
10.1.6	5. Investigation of analysis techniques	222
10.1.7	7. Key experimental results - two dimensional channel rig	224
10.1.8	3. Hydrogels and refractive index matching	225
10.1.9	 Circulating dispersions rig 	226
10.1.1	10. Key experimental results - circulating dispersions rig	226
10.2.	Future work	227
10.2.1	L. Achieving higher particle phase fractions	227
10.2.2	2. Further data analysis	229
10.3.	Overall conclusions	230
References	5	232
Appendix A	۹	239
Conferer	nce Presentations	239
Conferer	nce paper	239
Appendix B	3. Box counting method	240
Appendix C	C. Monte Carlo performance	244
Monte	e Carlo predictions	244
Comp	utational expense of Monte Carlo routines	247
Implic	cations with respect to data	249

Table of figures

Figure 1.1: Classification of coupling regimes for (1) one-way coupling, (2) two-way coupling
where particles enhance turbulence production, (3) two-way coupling where particles
enhance turbulence dissipation and (4) four-way coupling, when considering dispersed
phase fraction and timescales for both phases4
Figure 2.1: Deviation from a Poisson distribution, compared to the length scale of the cells
used. Peaks vary depending upon the particle's Stokes number, shown in the key14
Figure 2.2: Schematic representation of trajectory bias of particles within a turbulent flow
field20
Figure 2.3: TKE contour plots for particles in forced isotropic turbulence, for various
experimental conditions (Tanaka and Eaton 2010). TKE values are normalised with respect
to an unladen condition. Particles are considered as physical barriers in the flow field, which
act as dampers24
Figure 3.1: Schematic of experimental configuration of PIV experiments
Figure 3.2: Scattering cross section as a function of particle size, for a refractive index of 1.6
Figure 3.3: Representation of cross-correlation map for a 64x64 pixel interrogation region 33
Figure 3.4: Top-down representation of interrogation, showing cross-correlation in plane of
tracer particle movement
Figure 3.5: Comparison of first peak height (red dotted line) to pre-defined threshold (black
dotted line)35
Figure 3.6: Comparison between the two highest correlation peaks under consideration.
Peak heights are shown by the red and black dotted lines
Figure 3.7: Comparison between noise level (black dotted line) and the peak height (red
dotted line)36
Figure 3.8: Absorption and emission spectra for Rhodamine 6G, demonstrating a Stokes shift
of around 25nm between the peak wavelengths. Obtained from
https://en.wikipedia.org/wiki/Stokes_shift, accessed on 12th October 201740
Figure 3.9: Illustration of the experiment set up required to achieve phase discrimination
using fluorescent particles, using a beam splitter, focusing optics, and two CCD cameras41
Figure 4.1: Flow of information, from multiphase images to advanced information

Figure 4.2: Comparison of pixels which are considered contenders for local maxima in two
neighbouring particles, with the left hand particle at a saturated scattering intensity, and the
right hand particle exhibiting a clear local maximum47
Figure 4.3: Initial position, circled in red, is altered to the position circled in blue, as a first
estimate of position of the particle48
Figure 4.4: Example of particle under-estimation and particle over-estimation in clusters49
Figure 4.5: Distance transforms of binary images of particle clusters discussed in under and
over estimation examples respectively50
Figure 4.6: Watershed transforms of under and over-estimated particle clusters respectively51
Figure 4.7: Comparison of time of computation against N, number of particles in each image,
for Hungarian/Munkres and Nearest Neighbour Tracking algorithms
Figure 4.8: Relative performance ratio, R, between the Hungarian and Nearest Neighbour
algorithms, as the number of particles, N, is varied55
Figure 4.9: Ratio of particle links found, R, between the two algorithms under investigation,
as a function of particle number, N56
as a function of particle number, N56 Figure 4.10: Comparison of (left) crosswise velocity component PDFs and (right) streamwise
as a function of particle number, N56 Figure 4.10: Comparison of (left) crosswise velocity component PDFs and (right) streamwise velocity component PDFs obtained via differing tracking algorithms, known as Hungarian
as a function of particle number, N
as a function of particle number, N
as a function of particle number, N
as a function of particle number, N
as a function of particle number, N
as a function of particle number, N
as a function of particle number, N
as a function of particle number, N
as a function of particle number, N
as a function of particle number, N
as a function of particle number, N

Figure 4.15: Comparison of Hungarian tracking methods for the streamwise velocity
component, as the minimum tracking distance is varied, with reference to the input velocity
distribution. This reference is shown as a black dashed line62
Figure 4.16: Comparison of Nearest Neighbour algorithm performance for low and high
concentrations, against the anticipated distribution63
Figure 4.17: Comparison of Hungarian tracking algorithm for low and high concentrations,
compared to the anticipated distribution64
Figure 4.18: Comparison of predicted particle velocity distributions for the Nearest
Neighbour tracking algorithm for different levels of information loss in the virtual images, as
compared to the input velocity distribution65
Figure 4.19: Comparison of predicted particle velocity distributions for the Nearest
Neighbour tracking algorithm for different levels of information loss in the virtual images, as
compared to the input velocity distribution65
Figure 4.20: Example of an out of focus tracer particle70
Figure 4.21: Example of a typical dispersed phase particle70
Figure 4.22: Dispersed phase particle after erosion using structure element 1
Figure 4.23: Dispersed particle mask image after using erosion with structure element 1 and
dilation with structure element 272
Figure 4.24: Example of dispersed phase particle image found in these experiments, for
which the rounded edges are not obvious and the particle shape is generally square73
Figure 4.25: Square-shaped dispersed phase particle after erosion using structure element 1
Figure 4.26: Particle mask image for a square-shaped particle, using erosion and dilation
which both operate using structure element 174
Figure 4.27: Schematic of camera configuration in a two-camera, simultaneous recording
experiment utilising optical filters77
Figure 4.28: Identification of points on a calibration grid compared from two different
camera aspects78
Figure 4.29: Comparison of two simultaneously captured images of a particulate flow,
showing displacement between particles on screen due to the different camera aspects
which must be utilised79

Figure 4.30: Comparison of raw "fixed image" (green) with "moving image" (purple) after a
projective transformation has been performed. Areas of white represent good agreement
between the two images
Figure 4.31: Transformation of raw "fixed image" (green) with the "moving image" (purple)
after an affine transformation is performed80
Figure 5.1: Schematic representation of the test section used in the vertical channel rig. The
dashed lines show the outline of a measurement section observed via the camera
Figure 5.2: Schematic of experimental design for the vertical channel rig
Figure 5.3: Schematic of entrance conditions for water and hydrogel particles into the
vertical channel91
Figure 5.4: Photographs of a) Grundfos pump used in vertical channel rig b) Kobold MIK
magnetic flow meter used in vertical channel rig92
Figure 5.5: CAD drawing of frame construction for tank outlet filter
Figure 5.6: Schematic of fluid flow through the particle filter attached to the tank outlet. The
direction of fluid flow is indicated using a cross-section of the designed filter. The black
dashed lines indicate the attachment of the insert, which was permeable to water95
Figure 5.7: Example calibration image, captured from a single camera investigation. PIV
tracer particles in the flow can be observed in front of the calibration grid
Figure 5.8: Characterisation of flow rate delivery according to digital inverter settings. Data
points represent flow rate at each digital inverter setting. Flow rates recorded are accurate
to 1%, and therefore error bars are not shown97
Figure 5.9: Fluctuation of flow rate delivery to test section over time. Mean and standard
deviations for this time series are overlaid as solid and dashed lines, respectively
Figure 5.10: Comparison of mean streamwise velocity profiles for differing injection
parameters, under single flow conditions, at a distance of 80 channel heights downstream
of the injection point and at a fixed Re = 20800102
Figure 5.11: Comparison of mean streamwise velocity profiles at a distance of 40 channel
heights downstream of the channel entrance, for injections at 0.4 bar and without any
injected material. These are reproduced for a) Re = 20800 b) Re = 27300 c) Re = 31600103
Figure 6.1: (Left) Plot of probability distribution in the X-position coordinate for the Gaussian
profiles (Right) Plot of probability distribution in the Y-position for the corresponding
particle distribution, with the mean probability superimposed (dashed line)108

Figure 6.2: Relationship between standard deviation of simulated particle position in x-
direction, and number of particles, N, in each image, shown for both non volume excluded
and volume excluded cases
Figure 6.3: Centralised gridding method, shown for a randomly-generated particle set. In
this example, the grid size is 55 pixels, with the generated grid being 10 x 14 in size. This
method leaves gaps of 25 either side in the x-direction, and 15 either side in the y-direction
Figure 6.4: Comparison of D against cell size, for multiple values of N, particles per image
frame. Shown for reference is the expected value of a Poisson process for this type of
analysis (black dotted line)114
Figure 6.5: Comparison of D against cell size, for three values of N, the number of particles
per image frame, for a uniformly distributed set of particles with an excluded volume.
Shown for reference is the zero line, the Poisson reference for a completely random set of
particles115
Figure 6.6: (Left) Plots of D against cell size for varying simulated particle area fractions, as
calculated using the excluded area (Right) Normalisation of curves using the square root of
volume fraction, α , showing collapse of data116
Figure 6.7: Comparison of D against cell size, for three values of N, the number of particles
per image frame, for a set of particles with Gaussian distribution. Shown for reference is the
zero line, the Poisson reference for a completely random set of particles
Figure 6.8: Comparison of D against cell size, for three values of N, the number of particles
per image frame, for a set of particles with Gaussian distribution and an excluded volume.
Shown for reference is the zero line, the Poisson reference for a completely random set of
particles118
Figure 6.9: Schematic for operation of the box counting routine119
Figure 6.10: Example of a Voronoi diagram, constructed using a single frame of data from
the uniform random distribution120
Figure 6.11: : Standard deviations for normalised Voronoi cell PDFs created from four
different simulated particle distributions, varying with simulated particle number N122
Figure 6.12: Comparison of PDF of normalised Voronoi cell areas taken from uniform
random particle distributions with the fitted Gamma equation123

Figure 6.13: Comparison of Voronoi cell area PDFs for uniform and Gaussian distributions with no volume exclusion, over a range of concentration124 Figure 6.14: Comparison between Voronoi cell area PDFs for uniformly and Gaussian distributed sets of particles at N = 2000. Also shown is the Gamma fit from the Equation 6.5 Figure 6.15: Comparison between Voronoi cell area PDFs for uniform and Gaussian particle distributions with volume exclusion at varying concentration. Also shown is the Gamma fit. Figure 6.17: Comparison of experimental (vertical channel, Chapter 7) and simulated data, with definitions of Vc and Vv determined by crossing over points for the two distributions. This splits the Voronoi cells found into three regions, depending upon size and relation to Figure 6.18: Section of Voronoi diagram taken from sample experimental data, where cells have been assigned void (blue), neutral (white) or cluster (red)129 Figure 6.19: Identification of nearest neighbour particles (blue) around each central particle (red) is performed, for each angular region. Distances, r_i, are found between each of these pairs. The formulation thereby uses an assumption of shadowing, so that particles further Figure 6.20: Comparison of PDFs of particle potential for an experimental distribution and a random simulated particle distribution. In this case, a 12-particle nearest neighbour Figure 6.21: Example of labelled particles, where red particles represent those considered to be clustered. Particles cluster along strands in this image, which bears similarities to particles clustering along lines of strain and vacating vortices, as described in Chapter 2..132 Figure 6.23: A particle located at (0,0), surrounded by an annular element with Rmax = 2 and Rmin = 1. The measurement region is specified as being bounded by x = 0 and y = 0, which means the particle is located at the corner of the measurement region. In this case, the annular element contained within the measurement region is only 25% of the total area,

Figure 6.24: Monte Carlo scheme applied to the situation illustrated in fig. 1. The annular
element is populated randomly, and then the selection criteria for the measurement area is
applied to leave only the points shown in red135
Figure 6.25: Comparison of Radial Distribution Function, g(r _i), at varying radial distances, R _i ,
for generated sets of random particle images137
Figure 6.26: (top) Change in linear fitted gradient with particles per simulated image.
(bottom) Change in linear fitted intercept with particles per simulated image138
Figure 6.27: Comparison for three different simulated particle concentrations of Radial
Distribution Function with radial separation, for an excluded volume set of simulated
particles139
Figure 6.28: (top) Radial distribution functions for three particle concentrations with
Gaussian x-dimension profiles. (bottom) corresponding data for inclusion of an excluded
volume140
Figure 7.1: Deviations from Poisson histogram statistics of particle positions, as computed
via the box counting method, at a distance of 40H from the injection point for a variety of
particulate volume fractions144
particulate volume fractions144 Figure 7.2: (left) Relationship between length scale of peak clustering, taken from the x co-
particulate volume fractions

Figure 7.6: PDFs of individual normalised particle potential, using a 12-neighbour
formulation, for varying particulate concentrations, using experimental data from the
vertical channel rig, at Re = 20800153
Figure 7.7: Illustration of angular separation in the vertical channel experiment. Angles
within A are 0° to 90°, within B are 90° to 180°, within C are -90° to 0° and within D are -180°
to 90°154
Figure 7.8: Comparison of variation in nearest neighbour at different angular separations, at
9% particulate concentration. These plots are shown for the full distribution, saturated
particles (those with neighbours in each angular region, and those which are classified as
clusters
Figure 7.9: Mean nearest neighbour particle separation as a function of angular separation,
for varying particulate concentration, for cluster particles only156
Figure 7.10: Mean nearest neighbour particle separation as a function of angular separation,
for varying particulate concentration, for saturated particles only157
Figure 7.11: Mean separation variability for varying angular separation and particulate
concentration, for cluster particles158
Figure 7.12: Nearest neighbour separation variability for varying angular separation and
particulate concentration, for saturated particles158
Figure 7.13: Mean separation profiles for saturated particles, <r>, multiplied by the square</r>
root of the bulk particle concentration, $\alpha^{0.5}$ 160
Figure 7.14: Mean separation profiles for cluster particles, <r>, multiplied by the square root</r>
of the bulk particle concentration, $\alpha^{0.5}$ 160
Figure 7.15: Plots of probability for a central cluster particle to have N local cluster
neighbours, for varying particulate concentration. Fits achieved using a smoothing spline to
indicate general trend161
Figure 7.16: Percentage of cluster particles with 12 local neighbours with changes to bulk
particulate concentration163
Figure 7.17: Percentage of cluster particles with N local neighbours, for varying bulk particle
concentration and values of N164
Figure 7.18: RDFs for a range of separation distances, for three distances downstream of the
injection point. Fits achieved using smoothing spline fit, show a peak at low radial
separation, with collapse to zero for infinitesimal radial separation165

Figure 7.19: Sample PDF of separation timescales, t*, with closing and opening particle pairs
labelled. Closing particles are defined as those which have the potential to collide, whereas
opening particles are those which are considering to be moving away from each other166
Figure 7.20: (Left) Calculated probabilities of collision for varying concentrations (Right)
Attempted collapse of data with concentration167
Figure 8.1: The temperature and transmittance over time of a 5% PMMA and NaI
suspension at the optical path length of 49mm173
Figure 8.2: Plots of peak transmittance values for Nal/PMMA suspensions, with variations in
optical path length and the particle concentration174
Figure 8.3: The temperature and transmittance over time of a 5% Hydrogel and water
suspension at the optical path length of 49mm175
Figure 8.4: 100x magnification image of PMMA particles used in the refractive index
matching experiments
Figure 8.5: 20x magnification of agarose hydrogel particles used in the refractive index
matching experiments
Figure 8.6: Schematic of design for the circulating dispersions rig
Figure 8.7: (A) Laser interlock warning light (B) Interior of laser safety cabin during
experimental work (C) Laser interlock box, including connection to door switches
Figure 8.8: Progressive cavity pump used in the circulating dispersions rig. This was fitted
with the additional safety mechanism provided, which was an overpressure switch187
Figure 8.9: Illustration of typical radial and axial flow patterns in an agitated vessel. For
suspension of solids, creating an axial flow pattern is preferable
Figure 8.10: (Left) Photograph of storage tank used in circulating dispersions rig; A is the
stirrer motor, and B is the heat jacket for use with the heat exchanger (Right) Photograph of
the interior of the tank; C is the impeller used to agitate the particulate-liquid mixture189
Figure 8.11: (A) Central measurement section (B) measurement section including
rectangular box to reduce distortions at boundaries190
Figure 8.12: Demonstration of functionality of the surrounding box in the measurement
section, achieved by filling the box to half way191
Figure 8.13: Photograph of the electromagnetic flowmeter used on the circulating
dispersions rig

Figure 8.14: T-type thermocouple, held in coupling to maintain position within the fluid flow
Figure 8.15: Photograph of differential pressure cell used across the measurement section
Figure 8.16: Temperature measured at pump outlet over standard running time of the
experimental rig without any temperature control of the pipe contents, using water as the
liquid phase for a test case
Figure 8.17: Pipe and measurement box cross section, with distance between central and
measurement plane shown, along with the pipe radius R
Figure 8.18: Axial velocity profiles for varying Reynolds numbers in single phase flow,
normalised using average axial velocity calculated from measured flow rate
Figure 8.19: Crosswise velocity profiles for varying Reynolds numbers in single phase flow,
with varying normalised Radial Distance200
Figure 8.20: Turbulent intensity profiles, u _{rms} / <u>, in streamwise direction201</u>
Figure 8.21: Turbulent velocity fluctuations in the crosswise direction, v', for varying
normalised Radial Distance, at different Reynolds numbers
Figure 8.22: (left) Raster scan presentation of mean streamwise velocities, for Re = 3.86 x
10^4 (top) and Re = 9.63 x 10^4 (bottom) (right) Deviation in streamwise velocity profiles at
subsequent distances along the pipe, for Re = 3.86 x 10 ⁴ (top) and Re = 9.63 x 10 ⁴ (bottom)
Figure 9.1: Deviations, D, from Poisson histogram statistics of particle position distributions,
as computed via a Box Counting Method for a variety of particulate volume fractions, at
as computed via a Box Counting Method for a variety of particulate volume fractions, at varying length scales of observation206
as computed via a Box Counting Method for a variety of particulate volume fractions, at varying length scales of observation206 Figure 9.2: Closing times of particle-particle pairs, as computed by considering radial
as computed via a Box Counting Method for a variety of particulate volume fractions, at varying length scales of observation206 Figure 9.2: Closing times of particle-particle pairs, as computed by considering radial separation and resolved relative velocity. This is shown in the form of probabilities of
as computed via a Box Counting Method for a variety of particulate volume fractions, at varying length scales of observation206 Figure 9.2: Closing times of particle-particle pairs, as computed by considering radial separation and resolved relative velocity. This is shown in the form of probabilities of collision, over a cumulative time frame
as computed via a Box Counting Method for a variety of particulate volume fractions, at varying length scales of observation
as computed via a Box Counting Method for a variety of particulate volume fractions, at varying length scales of observation
as computed via a Box Counting Method for a variety of particulate volume fractions, at varying length scales of observation

Figure 9.5: Definition of angular separation for particles in the circulating dispersion rig.
Angles within A correspond to 0° to 90°, within B are 90° to 180°, within C are 0° to -90°, and
within D are -180° to -90°211
Figure 9.6: Comparison of variation in nearest neighbour at different angular separations, at
3% particulate concentration. These plots are shown for the full distribution, saturated
particles (those with neighbours in each angular region), and those which are classified as
clusters211
Figure 9.7: Mean nearest neighbour particle separation as a function of angular separation,
for varying particulate concentration, for cluster particles only
Figure 9.8: Mean nearest neighbour particle separation as a function of angular separation,
for varying particulate concentration, for saturated particles only
Figure 9.9: Plots of probability for a central cluster particle to have N local cluster
neighbours, for varying particulate concentration215
Figure 9.10: Percentage of cluster particles with N local neighbours, for varying bulk particle
concentration and values of N216
Figure B.1: Corner-anchored gridding method, shown for a randomly-generated particle set.
In this example, the grid size is 55 pixels, with the generated grid being 10 x 14 in size. This
method leaves gaps of 50 on one side in the x-direction, and 30 on one side in the y-
direction240
Figure B.2: Comparison of deviation factor, D against cell size, for multiple values of N,
particles per image frame. Shown for reference is the expected value of a Poisson process
for this type of analysis (black dotted line)241
Figure B.3: Comparison between two gridding methods using identical data, for N = 2000
particles per image frame241
Figure B.4: Comparison of averaging methods to obtain D, for two different values of N,
using a centralised gridding method242
Figure B.5: Comparison of averaging methods to obtain D, for two different values of N,
using a corner-anchored gridding method243
Figure C.1: Progression of predicted fractional area for increases in the number of simulated
points, N, used per annular element, for R _{max} = 2245
Figure C.2: Variation of δ , the percentage difference between actual and predicted values of
included area, with the number of points N used per annular element. The equation for

linear line of best fit in each case is displayed on the graphs, with comparison between $R_{\rm m}$	ах
= 2 and R _{max} = 102	46
Figure C.3: Standard deviations from Monte Carlo simulations using varying number of	
points, N, per annular element, for R _{max} = 22	47
Figure C.4: Plot of average time per Monte Carlo computation for variations in number of	
simulated points per annular element, for R _{min} = 1 A.U. and R _{max} = 2 A.U2	.48

Table of tables

Table 4.1: Summary of erosion-dilation processes for different dispersed particle shapes74
Table 5.1: Comparison of development lengths using the channel height and hydraulic
diameter as the characteristic length scale, for a range of Reynolds numbers in the turbulent
flow regime
Table 5.2: Matrix of Reynolds numbers tested within the vertical channel experiment100
Table 5.3: Stokes numbers for experiments conducted within the vertical channel
experiment101
Table 6.1: Description of the four types of simulated particle sets used to test the particle
distribution codes107
Table 6.2: Area particle fractions associated with different particle numbers per image
frame in excluded volume cases109
Table 8.1: Physical parameters of storage tank used in circulating dispersions rig
Table 8.2: Liquid flow rate characteristics relevant to the experiments conducted within the
circulating dispersions rig197
Table 8.3: Particle Stokes number for each of the experiments conducted within the
circulating dispersions rig197

Table of Equations

Equation 1.1: First Navier-Stokes equation4
Equation 1.2: Second Navier-Stokes equations4
Equation 1.3: First Reynolds-averaged Navier-Stokes equation5
Equation 1.4: Second Reynolds-averaged Navier-Stokes equation5
Equation 1.5: Decomposition of momentum transfer between phases of flow5
Equation 2.1: Definition of Stokes number12
Equation 2.2: Equation of dispersed phase motion13
Equation 2.3: Relationship fitted using empirical data by Sumbekova et al16
Equation 2.4: Effective drag coefficient equation for particles falling through isotropic, zero-
mean field19
Equation 2.5: Standard equation for particle drag co-efficient19
Equation 3.1: Fringe interference pattern caused by Laser Doppler Anemometry
measurements
Equation 3.2: Relationship to determine substitution using a Universal Outlier Detection
procedure
Equation 4.1: Matrix representation of translation between co-ordinate sets 1 and 277
Equation 4.2: Matrix representation of tilt between co-ordinate sets 1 and 277
Equation 4.3: Matrix representation of affine transformation found using MatLab for the
sample images in Fig. 4.2881
Equation 4.4: Matrix transformation of projective transformation found using MatLab for
the sample images in Fig. 4.2881
Equation 5.1: Development length, l _e as function of Reynolds number, normalised by
characteristic lengthscale D
Equation 5.2: Hydraulic diameter relationship for rectangular channel with dimensions W
and H86
Equation 5.3: Definition of channel aspect ratio86
Equation 5.4: Hydraulic diameter relationship for rectangular channel using aspect ratio86
Equation 5.5: Reynolds number, for flow rate Q and channel aspect ratio n
Equation 5.6: Calibration relationship between flow rate and output current

Equation 5.7: Derivation of time delay between PIV frames	99
Equation 6.1: Area fraction calculation	109
Equation 6.2: Equation of Poisson distribution	111
Equation 6.3: Definition of Poisson mean for box counting method	111
Equation 6.4: Definition of Poisson deviation as defined for box counting metho	od111
Equation 6.5: Analytical fit for PDF of normalised Voronoi cell area distribution	for a random
process	122
Equation 6.6: Calculation of local particle potential	130
Equation 6.7: Calculation of normalisation distance	130
Equation 6.8: Definition of Radial Distribution Function	133
Equation 6.9: Calculation of fractional area within the RDF calculation	135
Equation 6.10: Equation for gradient of best fit line with particle variation	138
Equation 6.11: Equation for y-intercept with particle variation	138
Equation 7.1: Relationship between cluster structure perimeter and area	150
Equation 8.1: Definition of light transmission	172
Equation 8.2: Beer-Lambert law	173
Equation 8.3: Refractive index variation of Nal, reported by Bai and Katz	181
Equation 8.4: Refractive index variation of Nal, reported by Narrow et al	181
Equation 8.5: Refractive index variation of KSCN, reported by Agrawal et al	181
Equation 8.6: Refractive index variation of Nal, accounting only for concentration	on181
Equation 8.7: Refractive index variation of KSCN, accounting only for concentra	tion182
Equation 8.8: Temperature dependency, T, of refractive index of KSCN at optim	ium
conditions	182
Equation 8.9: Equation for heating of a liquid using a coolant	196
Equation C.1: Percentage deviation for RDF test case	245
Equation C.2: Linear line fit of percentage deviations	245
Equation C.3: Linear line fit to assess computational expense	248
Equation C.4: Estimation of computation time for complete data set	249

Nomenclature

Symbol	Subscript/Modifier	Meaning
α	P (optional)	Particle volume fraction (%)
δ		Deviation in prediction (%)
Σ		Local particle potential (m ⁻¹)
η		Kolmogorov length scale (m)
θ		Angular separation (degrees)
λ		Laser wavelength (nm)
ρ		Density (kg m ⁻³)
σ		Standard deviation (case-
		dependent)
τ		Time constant (s):
	Р	Particle
	К	Kolmogorov
	F	Fluid
μ		Dynamic viscosity (kg m ⁻¹ s ⁻¹)
Ω		Vorticity tensor (s⁻¹)
A		Area (m²)
	С	Cluster area (m ²)
С	D	Drag co-efficient
	D,eff	Effective drag co-efficient
С		Concentration by mass (%)
D		Distance (m)
	h	Hydraulic diameter (m)
d		Particle diameter (µm)
F		Force vector (kg m s ⁻²)
	d	Drag component
	I	Lift component
	vm	Virtual mass component
	b	Basset force component
	р	Inter-facial pressure

		component
f	D	Fringe frequency (s ⁻¹)
g(r)		Radial distribution function
h		Channel height (m)
I		Current (mA), Light intensity
		(Candela)
	0	Initial intensity (Candela)
L		Optical path length (mm)
	е	Entrance length (m)
Μ	j	Momentum component (kg
		m s ⁻¹)
n		Refractive index
	Subscript	RI of specific substance
р		Pressure (Pa)
	Tilde accent	Instantaneous pressure (Pa)
Р	С	Cluster perimeter (mm)
Q		Flow rate ($m^3 h^{-1}$)
		Heat transferred (J)
R		Radius (mm)
Re		Reynolds number
S		Strain rate tensor (s ⁻¹)
St		Stokes number
t		Time (s)
	0	Starting time (s)
	*	Collision time (s)
	х	x-translation (mm)
	У	y-translation (mm)
Δt		Elapsed time (s)
Т	Tilde accent + i,j	Mean strain-rate tensor (s ⁻¹)
	λ	Transmission
u	Tilde accent + i	Fluctuating velocity (ms ⁻¹)

u		Velocity vector (ms ⁻¹)
v		Velocity (ms ⁻¹)
V		Volume (m ³)
	С	Voronoi cluster threshold
		(m ²)
	V	Voronoi void threshold (m ²)
x	i	Spatial displacement (m)

Abbreviations

Abbreviation	Meaning
A.U.	Arbitrary Units
CAD	Computer Aided Design software
cDAQ	Compact Data Acquisition Platform
CFD	Computational Fluid Dynamics
KSCN	Potassium Thiocyanate
LDA	Laser Doppler Anemometry
LIF	Laser Induced Fluorescence
Nal	Sodium Iodide
PIV	Particle Image Velocimetry
PMMA	Poly(methyl methacrylate)
PTV	Particle Tracking Velocimetry
RDF	Radial Distribution Function
RI	Refractive Index
ROI	Region of Interest
Rpm	Revolutions per minute
RPP	Random Poisson Process
s.f.	Scaling factor

Chapter 1. Introduction

1.1.Background

Multiphase flows are encountered in many industrial, scientific and environmental settings. In the petrochemical industry, flows containing two and sometimes multiple different phases of matter are commonplace in pipelines, gravitational separators and combustors. These types of flows are also prevalent in environmental applications, such as distribution of pollutants in water and air flows.

In general, for these types of real-world scenarios, it is incredibly unlikely that modelling and analysis can be conducted using only a single-phase approach. Pipeline flows may contain up to four distinct phases of matter, such as mixtures of oil, water, air and sand. Therefore, a thorough comprehension of the flow mechanisms, and phase coupling and interactions, is of vital importance to our ability to understand and predict how a real-world flow will develop. Only with this understanding will it be possible to optimise the design of flow facilities, increasing their reliability and improving their function.

However, the ability to mathematically model even single-phase flows is considered very difficult, especially where there exists a complex geometry or a great range of length scales or time scales. This difficultly arises from the non-linear interactions between different structures, which results in chaotic behaviour.

These calculations are then further complicated by the addition of a separate phase. It becomes necessary to consider not only each phase and the governing Navier-Stokes equations, but also interactions which may occur between these phases. For example, the exchange of momentum must be considered at each phase boundary. Depending upon the flow constituents, it may also be necessary to consider the evolution of phase boundaries with time.

In some cases, Direct Numerical Simulation (DNS) of the governing equations may be used to produce a solution for a particular problem. However, when the range of scales which must be resolved is large, this is a method which scales poorly, and can be computationally expensive. As a result of these complications, there exist a variety of Computational Fluid Dynamics (CFD) models, which attempt to simplify the governing equations. Many of these

models have in common a requirement for validation data, since the Reynolds Averaged Navier-Stokes equations, which are commonly used for description of fluid motion, do not result in a closed system, and therefore require experimental input. Additional terms in multiphase-specific models describing phase coupling also require further experimental data to aid these descriptions.

Developing these CFD models is an incredibly competitive area of research in multiphase flows. These models are widely used as a predictive tool in a large array of applications, and therefore there is a great demand for increased accuracy and speed of computation. In order to facilitate this, relevant experimental data is required for the development of this field. In addition, experimental phenomena are of great interest themselves in determining the underlying mechanism for these types of fluid flows, which may then inform the rationale behind the models which are developed in the future.

1.2.Multiphase flow terminology

Multiphase flows describe a broad range of flow types and regimes. Initial characterisation often falls under physical description of the flow constituents; for example, a flow consisting of oil and water is called a "liquid-liquid" flow, with many other combinations possible.

Further sub-divisions of multiphase flows mostly focus upon the distribution of the contents within the contained area. For example, a "liquid-gas" flow may have its constituent parts oriented in wildly-different fashions, depending upon the relative flow rates and properties of the two phases. Some examples of these are stratified flows, "bubbly" or dispersed flows, annular flows, and "slug" or intermittent flows. Upon change of flow conditions, transitions between these can be observed, and these are often illustrated in flow regime maps, which can predict what type of flow will be experienced for the input conditions given (usually the superficial phase velocities).

In this work, the focus is upon dispersed flows. These flows are characterised by one phase, which is continuous, and another phase which is dispersed or entrained within this. For example, we might consider bubbles caught up in a liquid flow, or grains of sand carried along in a pipeline flow. Typically, for liquid-liquid and liquid-gas flows, this type of flow regime is most commonly associated with a high superficial velocity of the continuous phase.

In these types of flows, there is no clear boundary between the two phases under consideration; rather, many smaller boundaries must be considered. This is especially complicated for CFD codes to capture since there are a great many phase boundaries to consider and simulate. Further complication is provided by the fact that, in the flow domain, there are multiple areas where small length scales must be resolved, rather than just performing calculations to this accuracy across one single boundary.

Further delineation of dispersed multiphase flows is possible by considering the phase fraction of the dispersed phase (Crowe 2006). For very low dispersed phase fractions, the system is considered "one-way coupled". This means that, whilst the continuous phase has a significant effect upon dispersed phase motion, the effect is not reciprocal; the continuous phase largely behaves as it would without the presence of the particles or droplets. As the phase fraction is increased, the "two-way coupled" regime is entered. At this point, the "back-effect" of the particles upon the fluid motion must also be considered. In a relevant CFD model, phase momentum coupling must therefore be calculated in both directions.

Further increases in the dispersed phase fraction result in a regime called "four-way coupling". In these types of flows, particle-particle interactions such as collisions must now be considered, in addition to interphase coupling mechanics. For the purpose of developing CFD models, this often requires development of collision models. If the dispersed phase has flexible boundaries, such as liquid or gas droplets, aggregation and break-up models may also be required.

A phase diagram of these coupling regimes is shown in Fig. 1.1 (Elghobashi 1994). By considering only the dispersed phase fraction, we can see that in fact only moderate phase fractions of around 0.1% are considered sufficient for four-way coupling to become evident. In this diagram, τ_{K} represents the Kolmogorov fluid time scale, τ_{P} represents the Particle response time scale, and τ_{e} represents the turnover time of the largest eddies in the flow.



Figure 1.1: Classification of coupling regimes for (1) one-way coupling, (2) two-way coupling where particles enhance turbulence production, (3) two-way coupling where particles enhance turbulence dissipation and (4) four-way coupling, when considering dispersed phase fraction and timescales for both phases

1.3.CFD methods and validation

The starting point for solving incompressible fluid flows computationally are the Navier-Stokes equations, shown as follows:

$$\rho\left(\frac{\partial \boldsymbol{u}}{\partial t} + [\boldsymbol{u}.\nabla]\boldsymbol{u}\right) + \nabla p - \nabla.\boldsymbol{S} = 0$$

Equation 1.1: First Navier-Stokes equation

 $\nabla \mathbf{u} = 0$

Equation 1.2: Second Navier-Stokes equations

S represents the viscous stress tensor, p the hydrodynamic pressure, and **u** is the vector fluid velocity. It is possible to solve these equations using Direct Numerical Simulation (DNS) methods. However, the requirement for these methods to be used is that the lowest length scale of the flow, the Kolmogorov length scale, must be resolved by the cells used. For most industrial and scientific applications, this is simply too computationally expensive to be considered, as discussed previously.

In order to reduce time of computation, simplifications can be made. These vary greatly across a variety of methods. The common method is to decompose the vector fluid velocity, **u**, into a mean and fluctuating component. This is known as the Reynolds decomposition.

When this averaging procedure is applied to the Navier-Stokes equations, the Reynolds averaged Navier-Stokes equations can be derived, as shown below. Equation 1.3 describes the various momentum components of a fluid element, and its conservation along a line of motion. Equation 1.4 is called the continuity equation, and describes the conservation of mass.

$$\rho\left(\frac{\partial \tilde{u}_i}{\partial t} + u_i \frac{\partial \tilde{u}_j}{\partial x_i}\right) + \frac{\partial \tilde{p}}{\partial x_i} - \frac{\partial}{\partial x_i} \left(\tilde{T}_{ij} - \rho u_i \tilde{u}_j'\right) = 0$$

Equation 1.3: First Reynolds-averaged Navier-Stokes equation

$$\frac{\partial \tilde{u}_i}{\partial x_i} = 0$$

Equation 1.4: Second Reynolds-averaged Navier-Stokes equation

This averaging process adds an extra component to the equations. This acts as a stress term, and is therefore known as the Reynolds stress. It is a function of the fluctuating components of the fluid velocities. The equation is now in an unclosed form, which means that additional modelling is required to gain the extra information which will close the system. For singlephase flows, experimental data, along with additional computational models, are required to complete this closure.

The addition of other phases of matter to the computational problem further complicates matters. There are practical problems, such as resolving cells in which both phases are present. Also to be considered is the transfer of momentum between the two phases. Global momentum must be conserved, and so this constraint is placed upon the model.

Momentum is broken down into a series of different forces which act upon each phase, as shown:

$$\boldsymbol{M}_{j} = \frac{\alpha_{i}}{V} \left(\boldsymbol{F}_{d} + \boldsymbol{F}_{l} + \boldsymbol{F}_{vm} + \boldsymbol{F}_{b} + \boldsymbol{F}_{p} \right)$$

Equation 1.5: Decomposition of momentum transfer between phases of flow

These components represent the drag force, lift force, virtual mass, the Basset force, and the inter-facial pressure force. Depending upon the particular flow being simulated, these components may or may not be of importance. Experimental data is useful for the

development of the models which describe each of these components, since they can validate whether the behaviour being simulated is realistic or not.

1.4.Project context

This research project was conducted as part of the Consortium of Transient and Complex Multiphase Flows and Flow Assurance (TMF). This collaboration consists of researchers from a number of UK universities, and also industrial partners. The research presented in this thesis forms a companion piece to another research project (Traczyk 2016), which aimed to develop a new CFD model to accurately predict high concentration liquid-liquid dispersed flows.

The companion project was motivated by a specific shortcoming of CFD solvers with regards to predicting behaviour within industrial gravitational separators. Existing solvers were found to be unable to provide a stable dispersion in cases which actually exhibit this behaviour in real life. In particular, it is most likely that a stratified flow with no entrainment will be predicted. The theory posited within the collaboration is that a mathematical formulation is required to counteract the buoyancy force, and therefore prevent the inevitable formation of a stratified flow. The work in this counterpart thesis therefore consisted of developing a model, and using a horizontal pipe flow as a test case to see if a stable dispersed flow would be predicted.

This work is of great importance to any project in which deep-sea flow separators are required. If the design of these flow separators can be optimised, this allows for their use in oil fields under a wider variety of conditions, allowing for extraction over a greater surface area. Additionally, these flow separators will increase the efficiency of the extraction process if they are well designed for the conditions under which they operate. Therefore, there is a clear economical and environmental benefit to making these improvements. The aim of this project was therefore to provide validation data for CFD codes being developed to assist with this optimisation project.

In order to provide validation data for these codes, experiments investigating dispersions at suitable test conditions were therefore required. The work contained within this thesis focuses on solid-liquid dispersed flows, with a high phase fraction of solid particles. This was felt to be a reasonable comparison, since many of the behaviours between the two types of

flows are common. Nevertheless, it is worth addressing some of the differences which must be considered when using experimental data from this thesis to validate the CFD model created within the TMF consortium.

Firstly, it should be noted that, for solid-liquid flows in a closed system, the dispersed phase has a fixed size distribution. The dispersed phase in these systems cannot be described as monodisperse, due to irregularities and randomness in the manufacture of particles. However, the size distribution can be modelled using a narrow Gaussian distribution, due to the selection processes which the manufacturers impose on each particle batch. Therefore, the mean and standard deviations of particle diameter are well-defined. In contrast, the size distribution of the dispersed phase in a liquid-liquid flow must be considered, since this has a non-regular distribution which depends upon the flow conditions, and does not have predetermined values. The derivation of this size distribution is also examined by the counterpart thesis in terms of solving population balance models.

Secondly, the nature of the four-way coupling in these types of flows differs. In solid-liquid flows, dispersed phase particles are fixed entities, and therefore will collide and scatter. In contrast, for liquid-liquid flows, dispersed phase components may combine, and therefore the mechanisms of this aggregation must be resolved.

Despite the differences between these flow types, it was considered that experimental data from solid-liquid dispersions could prove very useful for validation of CFD codes with liquidliquid dispersions, provided that the differences between the two were kept in mind. The methodology adopted is described in the following section.

1.5. Aims and objectives

The objectives of this project were the design and construction of two large-scale rigs to make measurements of solid-liquid dispersions under a variety of different experimental conditions. The aim was to consider dispersions with high solid-particle content, thus investigating cases in which the particle-particle interactions were a dominant factor. The rigs chosen used a channel geometry in one case, and a cylindrical pipe geometry in the other.
Introduction

The data expected included detailed descriptions of the dispersed phase (solid particle) distribution for a variety of flow conditions. This would prove useful for CFD validation, since the most pressing problem was creating a model which would accurately represent this distribution. Furthermore, data to support collision models was required for representation of four-way coupled (highly dense) particle flows.

It was therefore desirable to provide detailed descriptions of the clustering of particles, and the effects of particle concentration upon this behaviour. This includes the assessment of global clustering parameters, as well as local clustering behaviour. Further information relating to particle-particle pairs was sought, such as the Radial Distribution Functions, and Relative Velocities, which provide valuable insight towards modelling particle-particle interactions, which become a dominant mechanism in dispersed flows where the dispersed phase fraction is greater than 0.1%.

To provide such data, an experimental methodology was required which would capture instantaneous images of the particle distribution which could then be analysed. These methods are discussed further in Chapter 3, and focused around the use of high speed imaging, and appropriate algorithms such as Particle Image Velocimetry (PIV) and Particle Tracking Velocimetry (PTV). PIV was predominantly used to describe the continuous phase and characterise rig performance, whilst PTV was used to create particle velocity vectors, which could be analysed to determine relative velocities between particle pairs.

In order to make relevant comparisons, the experiments were designed with some specific flow parameters in mind. As discussed in Section 1.2, dispersions are most likely to form in liquid-liquid flows when the carrier phase has a high superficial velocity. Clearly, dispersion is a fixed quantity within solid-liquid flows, since no other orientation is possible. In order to provide reasonable comparisons, experiments were conducted within a highly turbulent regime, so that the carrier fluids also had a high superficial velocity. This defined a range of Reynolds numbers for experiments between 10^4 - 10^5 .

In order to investigate four-way coupled flows, particle concentrations were required to be above a threshold at which particle collisions can be considered a dominant effect. Referring to Fig. 1.1, it can be seen that values of particle concentration exceeding 0.1% would meet this criteria. In order to guarantee high levels of particle collision within the flows being

Introduction

investigated, attempts were made to achieve particle concentrations of over 1%. Experimental design, including the use of refractive index matching, was considered which would allow optical measurements to be made in dense dispersions.

There was also a necessary requirement to use dispersed phase particles with a suitable size. In general, particle diameters were chosen which would lead to reasonable interactions between the two phases. These types of interactions are discussed in greater detail in Chapter 2.

1.6.Thesis outline

In the following chapter is a presentation of the relevant literature for the work contained within this thesis. This includes topics of interest, such as preferential accumulation of the dispersed phase and phase interactions.

In Chapter 3, an overview of the experimental methodology is provided, including justification of its use with respect to the data types which are required. Also discussed are the extensions to the technique which are relevant to this work, the relevant computations to analyse the data captured, and any sources of error which are characteristic of these types of measurement.

In Chapter 4, additional analysis techniques used to extract data are outlined. This includes algorithms used to track the dispersed phase particles, and the filtering techniques used to separate phase information in the raw data images.

The design and building of the "Vertical channel rig" is discussed in Chapter 5. This experimental rig allowed measurements of highly concentrated particle dispersions to be made, without encountering "depth-of-field" problems. The measurements made in these experiments were of particular particle injection events, rather than a continuously circulated dispersion. In this chapter, the rig is also characterised using only the single phase flow, for comparison with the multiphase experiments.

Chapter 6 covers the various types of analysis which were required in order to make sense of the particulate data. This includes coverage of various methods used to characterise clustering, such as the box counting method, Voronoi cell method, and a nearest-neighbour

Introduction

method. Radial distribution functions are also considered. For each of these methods, algorithms used are detailed, and tested throughout using test cases of simulated particles

The results of the multiphase experiments of the channel experiments are outlined in Chapter 7. This includes information about the particle distributions and collision probabilities, and information relevant to modelling dispersed phase collisions.

In Chapter 8, the design and building of the "Circulating Dispersions Rig" is presented. This rig was designed to provide a measurement of a turbulent dispersion within a cylindrical pipe geometry, which corresponds most closely to the geometry of the results presented in the counterpart thesis of this collaboration. As with Chapter 5, the rig performance is characterised using single phase results, obtained by PIV. Additionally, this chapter contains information about the refractive index matching considerations for these experiments, and the choice of flow constituents.

In Chapter 9, the multiphase results from the "Circulating Dispersions Rig" are shown. The topics presented are similar to those discussed in Chapter 7, and comparisons are made where applicable.

Finally, the key contributions of this thesis to the field are summarised in Chapter 10, along with recommendations for future work which could be conducted in these flow facilities and elsewhere.

Chapter 2. Literature review of turbulent dispersions phenomena

Studies of turbulent dispersions have identified several key areas of interest. The first key feature is the non-uniform distribution of dispersed particles, caused by deterministic eddy motions within the continuous phase. This is commonly known as 'preferential accumulation'. The second notable phenomenon is the effect of the continuous phase turbulence upon standard laminar heat and momentum-transfer correlations. This effect is called 'interphase coupling'. The final aspect to review is that of 'turbulence modulation' by particles dispersed in a turbulent flow, whereby the addition of a dispersed phase may attenuate or augment turbulence as compared to an unladen flow, depending upon the specific mechanism which dominates. The literature concerning each of these effects will be summarised in this chapter, in order to determine the most relevant phenomena for the experiments conducted within this thesis.

Within the realm of turbulent dispersed flows, there are many different types of flow constituents. It is possible to have dispersed solid particles in liquid (as in this thesis) and gaseous flows, and it is also possible to have dispersed liquid droplets within each of these types of flows. Since many of the phenomena are shared between these types of flows, multiple types of flow constituent will be considered within this review of the literature.

2.1. Preferential accumulation

2.1.1. Inertial particle behaviour

Initial theories describing turbulent particle dispersion assumed that the continuous phase applies a random forcing on the dispersed phase. This would suggest that the particle distribution should converge towards statistical uniformity, in the absence of any other external forces. For instance, in the presence of buoyancy effects caused by gravity, there would be a directional influence leading to non-uniformity.

Instead, we observe that particles are drawn into dense clusters by the deterministic eddy motions present in the continuous phase. Experiments measuring water droplet distribution in inhomogeneous, anisotropic turbulent air flows (Lázaro and Lasheras 1989) noted a separation by particle size. It was determined that this separation was caused by large two-dimensional eddies. These experiments were conducted in a dilute regime, where $\alpha = 10^{-5}$.

The underlying cause of this clustering is the inability of particles with a high inertia to follow curved streamlines. The parameter used to describe this behaviour is the ratio of the inertia of the two phases, which is called the Stokes number. This is defined by Equation 2.1 (Wood, Hwang et al. 2005), where St is the particle Stokes number, τ_p is the aerodynamic time constant of the particle, and τ_f is an appropriate fluid time scale. In their paper (Wang and Maxey 1993), Wang and Maxey surmised that the fluid time scale used should be the Kolmogorov timescale, under the assumption that preferential accumulation is controlled primarily by the intense vortical structures characteristic of small-scale turbulence. Subsequent simulations and experiments have determined this to be the most appropriate scaling, and therefore it is widely used in the definition of particle Stokes number.

$$St = \frac{\tau_p}{\tau_f}$$

Equation 2.1: Definition of Stokes number

To illustrate the effect of St, we will suppose that the fluid time scale is equivalent to the Kolmogorov time scale. For values of St approaching zero, the particles resemble fluid elements, and follow fluid streamlines faithfully. This results in a uniform distribution of the dispersed phase, because particles will sample space within the continuous phase with no preference. For values of St which are much greater than one, the particles follow ballistic trajectories, which can only be influenced by the most energetic eddies, or via inter-particle collisions. Since no coherent motion can be induced by the fluid, this scenario also results in a uniform distribution of particles.

For the case of Stokes numbers of order unity, inertia prevents particles following streamlines faithfully, and they feel a centrifugal force. Numerical simulations of particle dispersion in an axisymmetric jet provided evidence for the existence of an intermediate ratio of timescales at which particle dispersion would be optimal (Chung and Troutt 1988). Direct numerical simulations of isotropic turbulence, conducted at Stokes numbers of 0.15 (based on the integral length scale), showed preferential concentrations of particles with peak concentrations of thirty times the mean value (Squires and Eaton 1990). These simulations also showed that there exists an optimum Stokes number at which preferential accumulation occurs.

The motion of the dispersed phase can be described using Equation 2.2.

$$\nabla \boldsymbol{v} = -St(1-\beta)(\|\boldsymbol{S}\|^2 - \|\boldsymbol{\Omega}\|^2)$$

Equation 2.2: Equation of dispersed phase motion

In this equation, $\|S\|$ and $\|\Omega\|$ are the local strain-rate and rotation-rate tensors, **v** is the dispersed phase velocity vector and $\beta = \frac{3}{2\rho+1}$ where ρ is the particle-to-fluid density ratio.

Heavier-than-fluid particles, for which $0 < \beta < 1$, accumulate in regions where the strain rate dominates over vorticity, according to our understanding of divergence in flow fields. In contrast, lighter-than-fluid particles accumulate in regions of intense vorticity. The linear dependence of this expression upon the particle Stokes number is only valid for small values of St; as mentioned earlier, large St particles cannot respond to continuous phase motion, and so turbulence motions of these particles are more usually modelled by application of small random impulses.

2.1.2. Studies of preferential accumulation

The behaviour of particles in a jet dominated by vortex structures has been studied (Longmire and Eaton 1992). It was found that, for glass particles in air, concentration field is strongly affected by vortex ring structures. In this case, particles are dragged into highly-strained regions due to local fluid velocity variations, as shown in Equation 2.2. Preferential concentration between vortex rings in coaxial jets has been observed elsewhere (Wicker and Eaton 2001). Direct numerical simulations of isotropic turbulence also showed that particles would collect preferentially in regions of low vorticity and high strain rate (Squires and Eaton 1991).

The role of coherent motions in forming particle accumulations has been explored by investigating the length scale at which concentration deviates most from a random distribution, in the case of a turbulent channel flow (Fessler, Kulick et al. 1994). This was achieved by implementing the box counting method, which was the common method initially used to quantify preferential accumulation. This method is utilised within this thesis, and therefore is properly outlined in Chapter 6. Using the box counting method, a maximum deviation from random particle placement was found for a particular value of D, the length scale of interest. This length scale for maximum preferential accumulation varied with

particle Stokes number, suggesting a dependency upon the particle's time constant. Further experiments were conducted in a spherical turbulence chamber, generating isotropic and homogeneous turbulence (Wood, Hwang et al. 2005). For Stokes numbers near unity, a substantial peak was observed for a cell size D corresponding to 8-20 Kolmogorov length scales. This data is shown in figure 2.1.



Figure 2.1: Deviation from a Poisson distribution, compared to the length scale of the cells used. Peaks vary depending upon the particle's Stokes number, shown in the key

Another method used initially to quantify preferential accumulation is the radial distribution function. This is implemented in this thesis, and further described in Chapter 6. The radial distribution function (RDF) is primarily a measure of particle-pair separation. The number of particles within a radial separation of a central particle are calculated, and compared to the average separation of particle pairs throughout the whole measurement volume. This comparison indicates which separations are more or less likely than those seen throughout the bulk volume.

Initially, this was attempted using data obtained from Lagrangian/Eulerian DNS simulations. Sundaram and Collins (Sundaram and Collins 1997) used these simulations to study particle collisions in homogenous turbulence. This work showed that, for Stokes numbers near unity, the peak of the radial distribution function is higher, with particles collecting in regions of low vorticity. This preferential accumulation effect was shown to have an influence of increasing particle-particle velocity decorrelation, thereby increasing the probability of

particle-pair collisions. The preferential accumulation of particles also contributes by bringing particles into closer proximity on average. Therefore, it was shown that providing a measure of preferential accumulation could be key to describing particle collision rates.

Radial distribution functions can also be applied to experimental data, as shown in Wood et al 2005 (Wood, Hwang et al., 2005). Due to the nature of most measurement systems, these are not true radial distribution functions. In the case of LDA and PIV measurements, a 1D and 2D analogue must be used respectively, since the information which can be obtained is restricted. In their paper, Wood et al achieved qualitative agreement with data from simulations, but the peak measured values of the RDFs were lower in the experiments. This may be due to an attenuation effect known to occur for RDFs at lower dimensions, as described by Holtzer and Collins (Holtzer and Collins 2002).

RDFs have been applied to three-dimensional experimental data (de Jong, Salazar et al. 2010), captured from 3D-Particle Image Velocimetry. These measurements were made in a turbulence chamber, with a dispersion of glass spheres in air. Both RDFs and relative velocities were calculated between particle pairs. As stated by Sundaram and Collins, the combination of these quantities could allow for calculation of the particle-particle collision frequency, which could provide useful validation for CFD codes. This work also identified a method for using the Radial Distribution Function across the boundaries of the measurement volume, by imposing periodic boundary conditions which reflected the particle distribution across each boundary.

Another extension of this technique is to use the Angular Distribution Function (ADF) (Gualtieri, Picano et al. 2012), instead of the RDF. These were applied to results from modelling of a homogeneous shear flow, in which cluster formation was typically anisotropic. This method enabled the authors to identify the directionality of collisions, by again combining with relative velocities between particle pairs.

Currently, it is commonplace to use the radial distribution function to quantify the preferential accumulation of particles in numerical simulations. It is felt to be the most reliable technique for this type of data, and is also calculated in the process of determining particle collision probability. However, the limited dimensions available for most

experimental techniques create issues with making direction comparisons, as shown above. Therefore, experimental data is typically characterised using a variety of methods.

A recently popularised method of quantifying preferential accumulation is the use of Voronoi diagrams. These are utilised within this thesis, and are described in detail in Chapter 6. The principle is to create a Voronoi diagram, based upon the particle positions observed, which divides the image up into sections which are closer to one particle than another. The size distribution of these cells can then be used to infer detail about the nature of the particle distribution. For example, regions of the flow which contain very small cells represent clustering of particles, because these shapes are caused by very small separations between local particles.

The advantage of using this method is that, in addition to providing a global measure of clustering, assignments can be made of individual particles. This allows for the assignment of clusters within the captured image, the properties of which can then be analysed. This might include the area and shape of the clusters, and the concentration of particles within it.

This method was used in the context of particles in turbulent flows by Monchaux et al (Monchaux, Bourgoin et al. 2010). In this paper, the authors developed the method of comparing experimental Voronoi diagrams with a standard distribution for randomly distributed particles, and applied it to PIV images of particles in suspension in a horizontal channel flow. Further work has since been completed using the same facility. (Obligado, Teitelbaum et al. 2014, Obligado, Cartellier et al. 2015). Such studies have primarily concentrated on the effects of Particle Stokes number, St, and Particle Reynolds number, R_λ, upon the levels of preferential accumulation. Once again, it was found that preferential accumulation is enhanced for Stokes numbers of around unity. However, cluster sizes and geometry were found to be invariant upon this measure. A recent paper by this collaboration (Sumbekova, Cartellier et al. 2017) attempted to fit the level of preferential accumulation to the key variables in the system, and reported the following result:

 $\sigma_{REL} \approx 0.68 St^{0.02} Re_{\lambda}^{0.97} \phi_{\nu}^{0.47}$

Equation 2.3: Relationship fitted using empirical data by Sumbekova et al.

These results suggested an increasing dependence on clustering from Stokes number, to particulate volume fraction, ϕ_v , to particle Reynolds number. The particle volume fractions explored in these investigations were 10⁻⁵, suggesting a low incidence of particle-particle coupling.

This Voronoi cell diagram characterisation has also been adopted by other researchers. Tagawa et al. (Tagawa, Mercado et al. 2012) used both computational results and 3D-PTV images to carry out a fully 3-dimensional analysis of the Voronoi cell constructions to investigate a number of factors include Stokes number and Reynolds number. In these investigations, lighter particles showed a stronger preference to cluster, and clustering generally increases for higher particle Reynolds number.

The Voronoi method was applied to beads suspended in a turbulent square channel flow (Rabencov and van Hout 2015). The system used was a closed loop circulating a dispersion of water and polystyrene beads. This work determined the fractal behaviour in the shape of larger clusters via analysis of the power law governing the properties of the cluster shape.

Voronoi techniques have also been extended to results from Direct Numerical Simulations (Liu, Ji et al. 2017). Clustering was identified in a turbulent boundary layer, with standard deviations of the Voronoi cell size distribution measured for a variety of Particle Stokes numbers. Analysis was also performed on the shapes of the clusters found. It was determined that anisotropy was present, due to the 'streaky' nature of the structures found.

Similar work was performed for Direct Numerical Simulations of a vertical plane channel (García-Villalba, Kidanemariam et al. 2012). Again, the Voronoi technique was used to assess the aspect ratio of the structures found, and determined that these structures were aligned in the streamwise direction, an effect which was attributed to wake sheltering.

Further techniques have been developed which provide information about concentration effects within dense flows, but sacrifice detailed knowledge of particle structure. One such method was applied to a vertical wind tunnel containing polymer particles, with a Particle Stokes number of 1.4 (Lau and Nathan 2017). The method employed was to skeletonise the image by using the scattered intensity of the particles present. This provides a good

measure of particle concentration, without providing detailed knowledge on individual particle properties (position, velocity, etc).

A similar technique was used to measure glass particles injected from a nozzle into a wind tunnel (Birzer, Kalt et al. 2012). These experiments were two-way coupled flows, in which the mass flow rate of the glass particles was altered. Again, only mean quantities were determined by using intensity from the images captured. These experiments found that as the mass loading was increased, the mean properties of the plume distribution were not significantly altered. However, this technique would not be able to resolve microscopic structure formation.

In conclusion, preferential accumulation of particles in turbulent flows is a common factor, and may influence collision rates due to the clustering which occurs, along with the decorrelation of particle velocities. The phenomenon is dependent upon the Stokes number of the particles, which in turn dictates the length scales at which the clustering is most important. There is general consensus that the Stokes number should be defined using the Kolmogorov length scale, and that the effects of preferential accumulation are highest for Stokes numbers on the order of 1. The length scale of maximal clustering is dictated by the coherent structures within the flow, and their ability to influence the particles within it.

Multiple methods have been used to characterise this clustering, with developments also extending into three dimensions for experimental data where available. Typically, the majority of the focus has been put on particle Stokes number, which was initially identified as an important variable. Less research has focused upon the effect of particle concentration; however this may have an impact upon the levels of clustering observed. (Sumbekova, Cartellier et al. 2017) Therefore, measurements at varying concentrations, within the four-way coupled regime, of the levels of preferential accumulation may be revealing. Such measurements of particle dispersion would also provide useful validation for CFD analysis, and provide some insight into the likelihood of particle-particle collisions in a very dense dispersion.

2.2. Interphase coupling

A key component of any multiphase CFD model are the coupling terms, which describe how mass, momentum and energy are transferred between the phases. Of particular interest are

the forces that act between phases, which may then be used to model evolving phase boundaries. The focus of this section will be on the momentum coupling expressions which have been explored by other researchers.

In conventional treatments, the momentum coupling term is accounted for via the presence of drag and lift forces. However, experimental measurements of the drag force of particles in turbulence flows have been inconclusive, making it difficult to validate data in CFD models. Some experiments have observed an increase in drag co-efficient in turbulent flow (Brucato, Grisafi et al. 1998); others have measured the opposite effect, with reduced drag force due to turbulence (Rudoff and Bachalo 1988).

The effective drag co-efficient for particles falling through an isotropic turbulent, zero-mean field can be obtained by ignoring unsteady forces in a complete formulation (Balachandar and Eaton 2010):

$$C_{D,eff} = \frac{24}{\langle Re_v \rangle} \frac{1}{|\langle v \rangle|} \left| \langle u_r \rangle + 0.15 \left(\frac{d}{v} \right)^{0.687} \langle u_r^{1.687} \rangle \right|$$

Equation 2.4: Effective drag coefficient equation for particles falling through isotropic, zero-mean field In this equation, v represents the particle velocity, and u_r represents the fluid velocity, d represents the particle diameter, and Re_v is the particle Reynolds number. This can be compared to the standard drag coefficient:

$$C_D = \frac{24}{\langle Re_v \rangle} (1 + 0.15 \langle Re_v \rangle^{0.687})$$

Equation 2.5: Standard equation for particle drag co-efficient

For a zero-mean field, the average velocity experienced by the particle might not be zero. This means that $|\langle \mathbf{u}_{r} \rangle|$ may be larger or smaller than $|\langle \mathbf{v} \rangle|$ in Equation 2.4. As a result, the effective drag co-efficient can therefore be larger or smaller than this standard value computed in Equation 2.5, which explains the inconsistent experimental results seen in literature.

Several mechanisms have been devised to explain how particles might travel through a turbulent velocity field and either have their settling velocity enhanced or suppressed. The first of these is trapping, in which particles move towards the centre of vortices, from which

they can no longer move. This systematically reduces the overall settling velocity of a distribution of particles. This mechanism is most applicable to buoyant particles, and is therefore likely to arise in liquid-gas bubbly flows. It is not likely to be a dominating factor for solid-liquid flows.

The second mechanism of particle transport in turbulent flows is known by the terms trajectory bias or fast tracking, which is illustrated in Figure 2.2.



Figure 2.2: Schematic representation of trajectory bias of particles within a turbulent flow field

In this scenario, vortices are arranged in an alternating fashion, which is typical of arrangement when vortices are shed from a particular point. A particle seeking to travel straight down the middle of these vortices would encounter, in an alternate manner, resisting and enhancing velocities from the underlying fluid. However, a particle following the trajectory of the dashed line will experience, at all points along its path, an enhancing velocity. Therefore, differences in the particle trajectory taken can influence the settling velocity of the particle. The third mechanism which can affect particle settling velocities is loitering. In this scenario, rather than moving through the vortices in a sweeping motion, particles instead sample at random. The effect of this is that particles spend more time in regions of flow which reduce their velocity, and therefore the time taken to travel through the flow field is larger than it would be otherwise.

The fourth mechanism which affects particle settling velocities is known as a non-linear drag effect. These effects have been most commonly shown in DNS calculations which take account of non-linear drag corrections. (Yang and Shy 2003, Good, Ireland et al. 2014)

The final mechanism affecting particle transport is that of two-way coupling. Particles in turbulent flows typically cluster in regions of high strain. For higher particle loadings, this results in much larger inhomogeneities in the spatial distribution of the flow. Experimentally, there has been shown to be a dependence of the settling velocity when conditioned upon the local particulate volume fraction (Aliseda, Cartellier et al. 2002), with a monotonic increase being reported for higher particle loadings. Therefore, these two-way coupling effects play an important role in determining the settling velocity. This is therefore another reason why measurements of preferential accumulation are so valuable when trying to assess particle behaviour.

This mixture of influences therefore explains the variation of the effective drag co-efficient, and the previous inconsistent experimental results (Wang and Maxey 1993, Bosse, Kleiser et al. 2006), since varying experiments may exhibit some of these effects to a greater degree.

This approach is based upon point-particles, and it is necessary to consider complications caused by the more realistic scenario of a particle which takes up a finite volume within the fluid. This presents problems such as anticipating what flow would exist within the volume the particle occupies in the unladen case. Even if the undisturbed flow is known, a Taylor expansion must be used to assess the flow at the phase boundary using information from the particle centre, which may only be accurate for small particles. For large particles more than the leading order information of the expansion may be required, which will reduce the accuracy of the estimation achieved.

Experimental results (Bagchi and Balachandar 2003) have indicated that, for small particles $(d \ll \eta)$ the approach described is able to accurately predict the time evolution of the drag force. For such particles, the dominant force contribution is from the quasi-steady term. This simplicity is a result of length scale separation between the two phases. For larger particles $(d \ge \eta)$, the equation is able to account for slow variations in the force from length scales larger than d. However, length scales smaller than this can induce rapid variations, which are not adequately represented by the theory. Essentially, the length scale separation described above is no longer possible. The standard way of dealing with this scenario is to model the force as a superposition of a deterministic and stochastic component (Zeng, Balachandar et al. 2007). The behaviour of this stochastic component, and the small scale eddies which cause it, are areas which require further research (Sawford 2001).

Another factor to consider is the effect of vortex shedding. Self-induced vortex shedding becomes important above a critical particle Reynolds number (Zeng, Balachandar et al. 2007), the value of which decreases with an increasing level of ambient turbulence. The presence of this vortex shedding introduces extra force fluctuations, which show up predominantly in the lift component (Merle, Legendre et al. 2005) and are not correlated to the ambient turbulent flow.

A Direct Numerical Simulation of a homogeneous flow was used to show how vortex shedding and preferential accumulation can become coupled (Kajishima 2002). For high Reynolds numbers, it was found that clusters of the dispersed phase would form due to wakes in the continuous phase, and vortex shedding would enhance these clusters. However, the presence of these clusters would then also cause additional turbulence in the continuous phase fluid, which would then result in cluster break-up. The periodicity of this formation and break-up was found to be linked to the particle response time.

In conclusion, the velocity of particles in turbulent fluids can be challenging to predict, due to the many contributing and conflicting factors which may influence it. In particular, trajectory bias and two-way coupling effects have been shown to enhance the settling velocity, whilst loitering, trapping and non-linear drag contributions have been shown to suppress it. Further complications arise when considering larger particles; these drag

contributions are typically assessed by decomposing the acting forces into a deterministic and stochastic component. The interaction between small-scale eddies and large-scale particles, which have an overlap in length scale, is an area of much interest. Finally, particles may shed their own vortices under certain conditions, which will affect the lift component in a manner which cannot be tied to the ambient level of turbulent flow. It can also be seen that these effects are inextricably tied to the preferential accumulation effects described in the previous section.

2.3. Turbulence modulation

For dispersed multiphase flows with high volume fractions, extra mechanisms of turbulence production, distortion and dissipation become important. This modulation of turbulence levels from an unladen case can be noticeable, potentially changing the behaviour of engineering systems quite drastically. If single-phase turbulence production is small, then the effect of adding particles will be particularly noticeable. Factors which influence the modulation of turbulence include particle size (Kussin and Sommerfeld 2002), mass loading and Stokes number (Kulick, Fessler et al. 1994) (Squires and Eaton 1990).

The primary mechanism causing turbulence attenuation is attributed to extra dissipation due to unresponsive particles. Energy is transferred from turbulent eddies to small-scale flows around the particle boundary, from where it is dissipated rapidly (Rogers and Eaton 1991, Burton and Eaton 2005). This is amplified by the distortion of eddies at the particle surface. Therefore, particle size plays a role by affecting the surface boundary conditions. These local distortions are shown in Figure 2.3, taken from high-resolution PIV measurements in forced isotropic turbulence.

Literature review of turbulent dispersions phenomena



Figure 2.3: TKE contour plots for particles in forced isotropic turbulence, for various experimental conditions (Tanaka and Eaton 2010). TKE values are normalised with respect to an unladen condition. Particles are considered as physical barriers in the flow field, which act as dampers

The main cause of turbulence augmentation has been identified as an enhanced velocity fluctuation due to wake dynamics and self-induced vortex shedding (Zeng, Balachandar et al. 2010). Large particles create unsteady wakes above a critical Re, as described in the previous section, which augments turbulence (Bagchi and Balachandar 2004). A further mechanism to account for turbulence augmentation arises from preferential particle concentrations. These can cause buoyancy-induced instabilities, as a result of the density variation in the flow field, and as described in the previous section. Evidently, this is particularly noticeable for high particle-to-fluid density ratios.

All of these mechanisms act at different length scales, and simultaneously. Therefore, we can conceive of situations where a particle may both augment and attenuate turbulence, depending upon the length scale and direction (Geiss, Dreizler et al. 2004) under consideration. The overall effect will of course depend upon the weighting of each mechanism.

There has been some disagreement over the correlations used to determine levels of turbulence augmentation or attenuation. Experiments conducted in a high Reynolds

number flow of water, containing <1% glass beads of various sizes, found that previously developed correlations were not applicable to such a high Reynolds number (circa 300,000) (Shokri, Ghaemi et al. 2017).

In conclusion, the presence of particles in a turbulent flow can either significantly reduce or increase the turbulent kinetic energy levels, depending upon the dominant mechanism. Turbulence attenuation is typically caused by distortion of turbulent eddies at the phase boundary, thus accelerating eddy breakdown. Turbulent augmentation is typically caused by particle induced instabilities, either due to wake dynamics and vortex shedding, or as a result of inhomogeneous distributions brought about by particulate clustering.

2.4. Conclusion

This chapter has presented an overview of the main phenomena associated with turbulent dispersed flows. These are common across all types of dispersed flow, regardless of the flow constituent types (i.e. solid-liquid, liquid-liquid). Typically, particles which are denser than water will accumulate in high strain regions of the flow field, creating formations of clustered particles (Scott, Karnik et al. 2009). These clusters have been analysed in a variety of manners, in order to determine the present levels of clustering, typically as a function of particle Stokes number. A recent study has identified otherwise important contributions of particulate volume fraction and particle Reynolds number, which may have a greater influence than the Stokes number.

Clustering of particles may then have an influence upon the underlying fluid; this is known as two-way coupling. For instance, two-way coupling may affect the settling velocity and drag coefficients which are measured, and may also have an effect upon the turbulent kinetic energy levels observed.

Settling velocity is influenced by a variety of other mechanisms; most of these relate to the transport of particles through a field of rotating vortices, and also the possibility of nonlinear drag corrections coming into play. As a result, experimentally researchers have observed both enhancement and reduction of the settling velocity, depending upon the most dominant effect.

Turbulent kinetic energy levels are similarly likely to be enhanced and reduced based on a variety of mechanisms. Attenuation of turbulence is attributed to complex phase boundary problems, which can only be experimentally observed using high-resolution PIV. Turbulence enhancement is also affected by vortex shedding, and so is intrinsically linked to the problem of settling velocity changes.

As can be seen from this literature review, each of the three main topics in this field is linked, with one often being the cause of another, or providing at least some level of feedback. Therefore, during this thesis, the results obtained should have some relevance to each of the topics discussed in this chapter. This thesis will focus primarily upon results pertaining to preferential accumulation. It would not be possible to produce measurements of turbulent kinetic energy around particles using the experimental set up that has been utilised, due to lack of resolution. However, preferential accumulation can be seen to affect both phase coupling and turbulence modification, and so its measurement is of great importance to producing reliable CFD codes, and towards understanding solid-liquid flows as a whole.

Chapter 3. Experimental methods

There are many different methods of obtaining information from multiphase flows. Information required during these investigations ranged from quantities averaged over the full spatial domain, right down to detailed mechanics of particles. The best choice of experimental technique to attain information about the flows used was Particle Image Velocimetry (PIV). These measurements, combined with algorithms described in Chapter 4, may also be used to create Particle Tracking Velocimetry (PTV) information of the particles within these flows. This chapter will present the methodology used for these experiments, with a justification for the choice of PIV. Following on from this will be an explanation of the key features of this technique, which are crucial to understanding the results gained. There will also be a discussion of the errors we may expect from PIV measurements in the context of the data taken, and an extension to the technique which utilises fluorescent dyes to enable phase separation.

3.1.Optical techniques

Laser based techniques have notable advantages over other recording methodologies such as wire-mesh sensors and electrical capacitance tomography. Other techniques record only whole-field or averaged quantities, such as phase fraction profiles, whereas laser based techniques enable researchers to build fuller descriptions of fluid flows based either on scattering intensity or observed particle counts. Furthermore, laser based visualisation techniques allow for direct determination of particle velocities, whereas electrically based techniques only allow for inference of velocities between subsequent measurement planes via use of correlations.

The other advantage of laser based measurements is that they are regarded as nonintrusive. Whereas instrumentation such as a wire-mesh sensor provides a barrier to the flow, and therefore may change flow behaviour, observations using lasers are more likely to provide a true representation of the flow physics.

The main disadvantage of these techniques is that the experiments must be built to be optically transparent, and also that the transmission of light through the flow constituents must be maximised so that measurements can be made. In order to allow for optical access, clear materials are required for construction of the test sections. These must then fit with other parts of the experimental apparatus made from more conventional materials, which is a practical concern. Additionally, the processing materials chosen to conduct the experiments must be chosen to prevent significant scattering and distortions of light, which creates an additional constraint.

The earliest laser technique used to probe flows was called Laser Doppler Anemometry (LDA). The underlying principle is light-scattering interferometry. The measurement point is defined as the intersection of two laser beams, and as particles pass through this point, they scatter light from both laser beams, creating an interference pattern.

As the waves scattered from these particles interfere, they create a beat frequency. This beat frequency can be captured by photo detectors, and interpreted as a particle velocity using Equation 3.1.

$$f_D = \frac{2.U.\sin\left(\frac{\theta}{2}\right)}{\lambda}$$

Equation 3.1: Fringe interference pattern caused by Laser Doppler Anemometry measurements

In this case, U is the velocity of the particle normal to the fringes, θ is the angle created between the two beams, λ is the laser wavelength, and f_D is the beat frequency created.

These measurements are performed on single particles. For a single phase flow, tracer particles are added, which follow the flow of the single phase in a representative fashion. It is also possible to measure Doppler bursts from dispersed phase particles. However, there must be some method of discriminating between signals to determine which information belongs to which flow field.

Larger particles scatter the laser light more strongly. This provides the method of phase discrimination in LDA. To study flow modification by the dispersed phase (i.e. make simultaneous measurements), the photo detector output is sent to two processors. One of these has a low gain, and only detects dispersed phase particles. The other is set with a high gain and an amplitude threshold, so that high amplitude signals are rejected.

In general, the preferred laser measurement method is now Particle Image Velocimetry (PIV). The chief reason for this is that, by using a laser light sheet as opposed to probing the intersection of two beams, it is possible to make measurements across the whole of a two-dimensional plane. Firstly, this provides a whole-field visualisation that cannot be achieved using LDA; it is possible to visualise coherent structures in the continuous phase by moving through images taken at successive times. The further consequences of this are that it is easier to determine correlations between velocities of fluid elements which have a spatial separation. Such correlations cannot be determined by a single LDA measurement due to the one-point nature of the technique, and so previously experimentalists using LDA had to invoke Taylor's frozen turbulence hypothesis in order to determine any spatial correlation. Finally, the experimental work involved with using LDA is much more pain-staking; in order to create a flow map, an LDA system must be traversed numerous times with great accuracy in order to reproduce the same results as a PIV map.

This is not to say that LDA is totally without merits. It is generally regarded that LDA measurements are more appropriate for determination of flow statistics. This is due to the requirement for a large number of samples to be taken over a long duration of time. Whilst PIV can achieve this, the data storage required for comparable results would be massive, and even the fastest PIV systems are only now comparable with the temporal frequency which can be achieved with an LDA system. LDA also has the advantage that, since its probe volume is smaller than PIV, it can resolve smaller length scales. Additionally, with all the beam intensity concentrated at one point in space, it should be expected that the signal-to-noise ratio is preferable for LDA measurements, leading to a more accurate assessment of individual measurements

Nevertheless, PIV has superseded LDA in literature in the last twenty years to become the most dominant technique. Many researchers have noted that the techniques have become comparable in their accuracy as advances in processing algorithms and laser sheet optics have been made. Therefore, the remainder of this chapter will focus on PIV, which with added PTV for the particles entrained in the flows, was chosen as the adopted methodology for these investigations.

3.2.Particle Image Velocimetry – principles and practice

The optical arrangement in a PIV system consists of a laser source, a light guide arm, and a series of optics designed to produce a planar light sheet. The laser light sheet is aligned so that it illuminates the measurement area under investigation. In single-phase PIV applications, the liquid or air to be measured is typically doped with tracer particles, which can then be observed by a high speed camera.

In order to maximise the available laser intensity, the PIV system is operated in a pulsed mode. This increases the light available, and therefore the intensity of light scattered to the camera. For PIV measurements, two separate lasers are used, and these are activated in a double-pulsed mode. The cycle is defined by the acquisition rate and the time between the pulses from these two lasers, Δt . The cycle starts with the triggering of laser 1 at t₀, which coincides with the triggering of a camera capture. Following this, there is time lapse before the next trigger signals are sent, at t₀ + Δt , to laser 2 and the camera. The cycle ends after the period associated with the acquisition rate, and then starts again.

The result of this is that double-images are recorded sequentially to the CCD pixel array of the camera, making it possible to fully reconstruct the flow of particles and therefore the underlying fluid properties over the duration of the experiment. Parameters recorded along with the images include the time between pulses, Δt , and the acquisition rate, which allow each image to be given an appropriate time stamp.



Figure 3.1: Schematic of experimental configuration of PIV experiments

The general principle of this technique is then to analyse the particle motion in each set of double images. In the earliest days of this technique, this was done by averaging the particle displacements within certain regions, and then using the known time elapsed in order to assign a fluid velocity at that point. The method for acquiring fluid velocities is now more complex, and will be discussed in the PIV analysis section.

3.2.1. Tracer particles

Tracer particles must be chosen with respect to the fluid being used, so that they faithfully follow the flow dynamics of the underlying fluid. Ideally, the particle should have a characteristic response time which allows it to accurately reflect the fluid motion, without significant slip velocities. The particles used must also have a high scattering cross section, allowing them to be observed using the available laser power. Figure 3.2 (Melling 1997) shows the typical scattering cross-section as a function of particle size; it can be seen that the larger the particle, the stronger the signal is. Typically, these two factors represent a trade-off for the ideal particle diameter, since very large particles might scatter a lot of light, but will have very low response times to changes in fluid velocity.



Figure 3.2: Scattering cross section as a function of particle size, for a refractive index of 1.6

These factors also place a restriction on the types of materials which can be used. For waterbased flows, there are numerous different particles which can be used, including titanium oxide, aluminium oxide, various types of polymer particles, and also reflective silver-coated particles. In the experiments in this thesis, a polyamide seeding particle was used. Various sizes were available, and their usage depended principally upon the imaging length scales. Typically, 20µm diameters were used, although 50µm and 5µm were also available.

3.3.PIV analysis

Computationally, particle displacements in captured PIV images are measured using a crosscorrelation based method. These algorithms firstly split the image into interrogation regions, so that regions of the image are analysed on an individual basis. The intensity pattern within each region is then cross-correlated across the two frames, in order to determine the most common particle pixel displacement (Δx , Δy). This is shown in Fig 3.3 and Fig 3.4. Once this procedure has been repeated across all of the interrogation regions, a velocity map is then created, using the recorded time which elapsed between the capture of the two frames. In order to improve the accuracy of this measurement, sub-pixel interpolation is used on the signal peaks.



Figure 3.3: Representation of cross-correlation map for a 64x64 pixel interrogation region



Figure 3.4: Top-down representation of interrogation, showing cross-correlation in plane of tracer particle movement

An extension of this technique is the adaptive correlation, with uses N iterative steps in order to perform this calculation. In this case, the initial interrogation area has a size of N times the size of the final interrogation area. Through each step, a smaller area size is used, with the intermediary results used as input information for the next stage of the procedure.

As an example of this, if the final interrogation area size was selected as a 16x16 pixel area, and the algorithm was performed using N = 3 steps, then the initial interrogation area size would be 128x128 pixels. Additionally, it is possible to define an overlap between interrogation areas in order to increase the vector field resolution. This is typically set to 25% when using this technique, but can be set to any value independently along the vertical and horizontal directions.

There are numerous different validation methods associated with this technique, which can be used to remove spurious vectors from the measurements. The first step in this procedure is a peak validation, which places constraints upon what can be considered a reliable cross-

correlation peak to give a measurement. This can either be done by setting minimum and maximum values for the peak width, or for the peak height ratio between the first and second highest peaks. Once this step identifies an invalid peak, a local neighbourhood validation is used to replace its value with a more appropriate estimate. This is done by observing vectors in the neighbourhood, in an MxM region of interrogation areas. Typically, this is set to be a 3x3 region. If a spurious vector has been detected by the peak validation procedure, it is replaced by using a local interpolation of the vectors within this pre-defined area. The interpolation is performed using either a local median or moving average methodology, with a specified number of iterations.

Spurious vectors are also identified via another pre-set value, called the 'Acceptance Factor'. This factor defines the degree of freedom on the velocity vector gradients within the MxM region. If the calculated gradient is larger than the 'Acceptance Factor', then the central vector is removed. If this factor is therefore set too low, the vector map will be smoothed at whole-field level due to the removal of velocity gradients.

An alternative method for vector map calculation is the Adaptive PIV method. This method iteratively adjusts the size and shapes of the individual interrogation areas, in order to adapt to local seeding densities and flow gradients. Therefore, this is a useful analysis technique in circumstances where these quantities may be particularly uneven or unpredictable.

In this case, the number of interrogation areas and their centre separation is determined by the choice of the 'grid step size', which can be specified separately for the horizontal and vertical directions. A minimum and maximum interrogation area size is then defined in each direction in order to bound the algorithm to avoid small or large interrogation areas being generated.

The first iteration of the algorithm will always use the largest interrogation area allowed within this range. Subsequent iterations will then allow for a smaller size, if the local seeding density is high enough to justify it. This has the advantage of using only the very local seeding data to contribute to the displacement statistics, reducing the measurement as closely as possible to a point measurement.

In this case, validation is again used to prevent outlying vectors. Peak validation can be applied in three different forms. The first of these is to use a peak height validation, which requires the first peak height in the interrogation area, h, to be above a specified threshold value, denoted as 'th', in order to be considered a valid measure. This is illustrated in the Fig. 3.5, for a one-dimensional case.







Figure 3.6: Comparison between the two highest correlation peaks under consideration. Peak heights are shown by the red and black dotted lines

The final measure is to use the signal-to-noise ratio so that signal peaks below a certain threshold are rejected. The noise level, n, is evaluated via Root Mean Square of the negative correlation values. A one-dimensional illustration of this is shown in Fig. 3.7.



Figure 3.7: Comparison between noise level (black dotted line) and the peak height (red dotted line)

Vector values which do not meet these criteria may then either be rejected, or substituted using the Universal Outlier Detection procedure. A neighbourhood of vectors is selected,

and the rejected vector is then replaced with the median of the neighbouring vectors. The relationship below is calculated:

$$r_0 = \frac{|U_0 - U_M|}{r_M + \epsilon}$$

Equation 3.2: Relationship to determine substitution using a Universal Outlier Detection procedure where U_0 is the vector under consideration, U_M is the median vector from the neighbourhood, r_M is the medium residual of the form $r_i = |U_i - U_M|$ and ε is a normalisation constant, chosen to be 0.1 pixels. If the value of r_0 exceeds a detection threshold of 2.0, then U_0 is substituted with U_M . Otherwise, the vector remains the same. This procedure is detailed by Westerweel and Scarano (Westerweel and Scarano 2005).

Once either of these initial procedures have been performed, a vector map will have been produced for each of the double frames captured during an experimental run. It is then possible to perform a 'Moving Average Validation'. This compares each vector with the average of the other vectors in a defined neighbourhood, which can be chosen with a different width in each direction. Vectors which deviate significantly can be replaced by the average as an estimate of true velocity.

Some caution must be applied when selecting the size of the area to use for this validation. If the dimensions are set too high, then the vector field will become smoothed out due to over-sampling of the surrounding area. Only regions with similar behaviours should be used to compute these substitutions. In most cases, this means that the sizes of the neighbourhoods are limited to 7x7.

The 'acceptance factor' is used as the threshold to determine an anomalous vector defining percentage deviation from local neighbours. The typical values chosen for this are 0.12-0.15. The new vector is calculated by using a local interpolation with n iterations, with n typically being 2 or 3.

The PIV analysis in this thesis was performed using a commercial package called Dynamic Studio v4.10, which implements the algorithms above as recipes using the controllable variables described.

Experimental methods

3.4.PIV reliability and error estimation

For each step within the PIV process, it is possible to anticipate some form of information loss which may contribute to errors in the measured velocities. One example of this may be found in how accurately the tracer particles reflect the underlying flow. If there exists a slip velocity between tracer particle and fluid, this will introduce a bias in the results. This will become a particular issue in fluid flows with large accelerations. Other problems include out-of-plane motion of tracer particles, incorrect seeding levels, limited image resolution, blurring of particle images due to image noise, and variations in image intensity.

Some of these are possible to eliminate using simple checks and rules of thumb. For example, seeding particles can be added to achieve an optimal value, and the correct image resolution can be chosen using good experimental design. Some problems are more inherent to the experimental design, and must therefore be accounted for in error analysis. For example, it is hard to control out-of-plane motion, other than by orienting the laser light sheet to minimise this. Even so, out-of-plane motion will still occur, effectively lowering the seeding density which can contribute towards the cross-correlation in each interrogation region. Non-uniform image intensity may also be an issue; even with correct alignment of the two laser planes, there may still be differences in illumination between successive image frames, due to the Gaussian profile of the lasers used. As with the previous issue discussed, this problem arises due to tracer particle motion which is perpendicular to the laser light sheet, resulting in changes in particle image intensity between frames. This can cause an unavoidable error in the cross-correlation, depending upon the severity of this discrepancy. Such issues can also arise from overlapping tracer particles in images.

Generally, once a PIV system has been set up to its optimal state, it is expected that velocity vectors can be recorded with an accuracy of approximately 1-2% of the full-scale displacement range (Westerweel, Elsinga et al. 2013). Another rule of thumb which is often quoted for experimental measurements is an accuracy of 0.1 pixels. Considering the standard displacement to aim for is eight pixels for optimal cross-correlation assessments, these two estimates are roughly co-incident.

Advances in computational processing power and development of PIV algorithms have undoubtedly reduced the errors in the processing chain. By using iterative window shifting

and image deformation techniques, authors have reported accuracies on the order of 0.01 pixels when using synthetic images (Lecordier and Trinite 2003) (Nobach and Honkanen 2005). However experimental usages, as noted above, mostly report uncertainties which are an order of magnitude higher. Work has been completed to demonstrate the part played by intensity variations in consecutive particle images, which limits accuracy even in otherwise ideal conditions, and reproduced the often-cited 0.1 pixel limitation on accuracy (Nobach and Bodenschatz 2009). This work verified the existence of a lower bound on accuracy regardless of algorithm used, and also demonstrated the chief cause as the intensity variations of the PIV light sheet.

In reality, even errors of 0.1 pixel are hard to come by in an experimental set-up. Many errors (losing tracer particles out-of-plane, uneven tracer particle concentrations, image noise) compound on top of the error caused by intensity variations in the light sheet. Therefore, typical error values for single PIV vectors are 5-10%, depending upon the value being measured. For instance, in areas with a lower flow rate, the error will be a larger percentage of the overall measurement, so errors of 10% are particularly likely.

3.5.Phase discrimination in two phase PIV

The implementation of PIV in two-phase flows is more complicated than in the single phase flows which have been considered so far. In these situations, we capture images of seeding particles which track the movement of the fluid, and also of the particles which are suspended in the flow. In order to provide separate descriptions of each flow constituent, there is a requirement for a method of phase discrimination. One method of doing this is by using various image post-processing methods to create two images from the original, and this will be covered in detail in Chapter 4. The second method is an extension of the standard PIV technique, which uses optical filtering in order to provide the discrimination.

Both of these methods are equally relevant to applications in solid-liquid flows considered within this thesis, and to liquid-liquid flows which are investigated by other researchers. Since the principles are the same in either case, this review will cover any application seen in literature, regardless of the type of two-phase flow. Nevertheless, it should be pointed out that, for optical filtering, the challenges of creating a suitable experimental set-up

depend heavily upon the flow constituents, and therefore suitability may vary wildly depending upon the precise experimental configuration required.

It is possible to tag one phase in the flow with a fluorescent dye. There are examples of fluorescent tracer particles, and it is also possible to encapsulate fluorescent dye into particles which might be used in such experiments. The fluorescent dye should be characterised by a reasonable Stokes shift from the absorbed to emitted wavelength, which will then allow for reasonable separation of the optical data in wavelength-space. The dye used should also have an absorption peak close to the frequency of light of the Nd-YAG PIV laser used in these investigations, which is 527nm.





Due to the prevalence of dyes being used in these types of experiments, suitable candidates for this are well known, and typically absorb in the green and emit in the orange spectrum, resulting in a long-shifting of the emitted wavelength. Examples of these dyes include fluorescein and rhodamine B. The experiments can then be observed simultaneously through two different cameras, each fitted with an optical filter. Generally it is appropriate to use a green notch filter to observe scattered laser light, since these are readily manufactured to laser specifications. It may then be possible to use a longpass filter to observe the emitted orange light, although a notch filter may also be used for this purpose if the wavelengths are well matched.

The two cameras can be synchronised readily by using the same trigger signals to capture frames. With this set up in place, it is therefore possible to record two sets of images simultaneously from the experiments; one set of images with light scattered from the particles, and one with the emitted light.

So that the two sets of images can then be compared, there is a need to align to two camera positions as closely as possible. In an ideal set-up, the light from the experiment would be sent through a beam splitter to the two different cameras, as shown in Fig. 3.9.





Without a beam splitter, it is still possible to adjust the viewpoints of the cameras so that there is minimal difference between their vantage points in the experiment. This is achieved by ensuring that their angle to the horizontal is zero, positioning them at the same height and also by focusing at the same central point. With this achieved, it is then possible to map images onto each other, especially using the knowledge of what angles and distances do exist between the two cameras. However, such geometrical transformations in postprocessing are only possible if the transformation is minimised in the actual experiment, and so care must be taken over the camera configuration. This image transformation is described in more detail in Chapter 4.

This technique can be seen for various two-phase investigations in literature. There are plenty of examples of PIV being combined with Laser-Induced-Fluorescence (LIF) techniques. In these cases, rather than specifically doping a particle, one phase of the flow is doped wholly with dye. The optical filters are then used to identify phase regions, whilst tracers are used to produce the velocity vector maps. By combining these two sets of images, it is possible to see which velocity vector results from which phase. An example of this is seen in the analysis of the impact of an oil droplet onto an oil-water interface (Miessner, Coyajee et al. 2006). Whilst this is a viable technique for phases with large contiguous areas, it is less useful for flows containing particles or fine droplets, since the spatial resolution required to resolve phase boundaries is too high for most practical experimental set-ups.

Particle-specific doping has been used to investigate settling behaviour of glass particles in water, where the liquid phase was seeded with fluorescent (Rhodamine B) particles (Montante, Magelli et al. 2012). Another successful application of this technique was seen in the analysis of the decay of grid-generated turbulence, using varying dispersed phase parameters. The two camera viewpoints were compared using a calibration image, and then using cross-correlation techniques a disparity field between the two cameras was calculated. This map was then used to combine the two sets of phase information (Poelma, Westerweel et al. 2007).

In the latter investigation, it should be noted that only one filter was used. In general, if the fluorescence is applied to the dispersed phase, there will still be sufficient scattering from the laser light so that dispersed phase particles appear in both sets of images. Depending upon the scattering efficiency of the particles used, the experimental configuration may only be able to separate the fluorescent data out using an orange filter, whilst the green filter will then be rendered redundant. Further post-processing would then be required on this 'mixed-phase' image, which will be discussed in Chapter 4.

3.5.1. Application of wavelength filtering

During these investigations, a mixture of phase discrimination techniques were used. In the vertical channel rig, discussed in Chapters 5 and 7, the dispersed phase system chosen was a continuous phase of water and dispersed phase of hydrogels, with 500µm diameter.

Hydrogel particles are created by injecting a network of protein chains with water, which then makes up 99% of its volume. The specific structure of these particles is discussed in more detail in Chapter 8, in the context of creating a refractive index matched mixture. The mechanism of creating these particles means that fluorescent dye can be added to the aqueous solution which the protein structures absorb during manufacture. Tests were performed using hydrogels containing different levels of fluorescein, from 0.01-1%. The lowest level of dye was felt to be the most appropriate; re-emission was of the same magnitude as light scattering in the flow, albeit with enough of a difference to apply threshold filters in post-processing of images. The higher dye concentrations produced too much light re-emission, and resulted in prominent shadows behind particles, which acted to reduce the seeding density. This can result in less reliable PIV vectors.

Once these particles were selected for use in the experiments, a two-fold strategy was adopted. Some experiments were conducted using a dual-camera set up and optical filters. No beam splitter equipment was available to align the camera inputs, and so a geometric transformation was required to map images from one camera onto the other. Other experiments were conducted using only one camera, with phase separation conducted using size and intensity based masking techniques, which are discussed in Chapter 4.

The circulating dispersion rig, discussed in Chapters 8 and 9, used a dispersed phase system of 250 µm pyrex particles and potassium thiocyanate. The disadvantage was a lack of ability to incorporate fluorescent dye into the solid particles. Therefore, the strategy adopted for these experiments was to use only a single camera set-up, with all phase discrimination performed via separation algorithms.

3.6.Conclusion

In this chapter, the suitability of Particle Image Velocimetry for these investigations was discussed, in context of the results required. In order to gain insight into whole-field behaviour of dispersed phase flows, which spatially resolved information, this was the only
suitable technique. The main features of the experimental method were then discussed, along with the processing techniques required to obtain useable results. The key features of two PIV algorithms which were used are outlined, including the use of overlapping interrogation regions, iterative interrogation regions which pass information down a chain, sub-pixel interpolation of the correlation signals, and validation techniques which use information from neighbouring calculated vectors. This experimental technique produces images which can be combined with PTV algorithms, allowing for particle velocity information to be deduced.

Following on from this, the anticipated error for each PIV calculated vectors was discussed. Typically, this is found to be 1-2%, or a 0.1 pixel in terms of the cross-correlation displacement. Investigations have determined that standard commercial PIV algorithms achieve an error which is an order of magnitude smaller using synthetic images, and the most likely reason for the additional error is the uneven profile of the laser light sheet, coupled with perpendicular movement of tracer particles. However, with the addition of other experimental errors, typical PIV velocity errors seen range from 5-10%.

Finally an extension of the technique, which uses fluorescent dyes and camera filters, is outlined in the context of discrimination between two phases within a flow. This was used during some of the experiments contained within Chapter 7 of this thesis, due to the encapsulation of dye within the hydrogel particles used. The remaining experiments relied upon image processing algorithms to achieve this phase separation, which is outlined in the following chapter.

Chapter 4. Particle tracking and phase discrimination

Aside from the Particle Image Velocimetry algorithm used to deduce local fluid velocities, a variety of other analysis techniques were used to obtain information from the images taken during these investigations. In this chapter, the basic parts of these analysis routines will be described. Firstly, the process of obtaining the dispersed phase velocity statistics will be covered, including comparison between different Particle Tracking Velocimetry algorithms in terms of performance efficiency and accuracy. Following on from this, there will be a description of the masking techniques used on mixed phase images, which provides an alternative to the optical phase discrimination techniques described in Chapter 3. Finally, the geometric transformations used in a two camera configuration will be covered to describe how local points in both images can be matched spatially.

4.1.Information flow

Figure 4.1 shows the flow of information during the processing of data, with respect to the constituents outlined above. This was automated to allow for ease of processing of large quantities of data, with inputs for routines set in advance depending upon the relevant experimental conditions. The corresponding section of this chapter is noted next to each method where relevant.



Figure 4.1: Flow of information, from multiphase images to advanced information

4.2.Particle identification and tracking

An essential data component in these investigations is the dispersed phase statistics, including both the positions and velocities of the dispersed phase particles. In order to obtain these statistics, the particle positions must be identified in each image of the double pulsed capture. Following on from this, particles may then be tracked between both frames in the PIV capture. Given the known time separation between the frames and a calibration of a scale factor, s.f., in pixels per mm, it is then possible to assign a velocity to each observed particle, if tracking between frames is performed accurately.

4.2.1. Particle identification

In order to find particles in the image, a threshold was applied to the image above a known value. By doing this, an image was therefore obtained containing only pixels which are likely to make up a particle. Following on from this, pixels in this subset were then checked against each of their nearest neighbours. In cases where particle images exhibit a local maxima, this results in identification of the particle centre to the nearest integer pixel number in both spatial directions. In the less ideal case where there is saturation of the image intensity, then the algorithm is unable to distinguish a centre between neighbouring pixels, and so all pixels are retained as part of a plateau in intensity. Once all of these pixels have been checked, a list of all possible local maxima in the image is obtained.





An example of such an image is shown in Fig. 4.2 for two particles. This image has been magnified. The particle on the left hand side is saturated, and therefore we see a plateau in image intensity, and many possible particle peak locations. In contrast, the particle on the right hand side shows that there is an easily definable peak, since this saturation did not occur.

At this point, a priori knowledge of the particle size in the images is used, building on the intensity information derived above. A particle mask is applied to the pixel which is being checked, which uses the radius of a typical particle in the flow. This mask is used to pick out the intensity map around the pixel in question. The intensities in this region of interest are

then scaled by the local co-ordinates in both directions, to find the likeliest particle centre position.



Figure 4.3: Initial position, circled in red, is altered to the position circled in blue, as a first estimate of position of the particle

In the example shown in Fig. 4.3, the first pixel to be checked is located on the left hand side of the particle in the image. The circular particle mask is applied to this region of interest (ROI), with the pixel under investigation located at the centre of the ROI co-ordinates. Once these co-ordinates are scaled by the local intensity, we can see that the lower x co-ordinate numbers are less influential in the scaling than those at higher x-values. In the y-direction, the scaling either side of the pixel of interest is more even. This results in an estimation of particle centre which is shifted significantly in the direction of positive x, and marginally downwards in the y direction. The result of scaling by intensity is to move the original estimate of particle centre closer to the centre of regions with high-intensity value, which are more likely to correspond to particle centres.

Once this estimation is obtained, another mask is applied around the new estimated value, and pixels which fall within this mask are then set equal to zero, since they are determined to belong to the same particle. These are then removed from contention when considering further particles within the image, and cannot influence other particle position estimates.

Following the attainment of these particle centres at an integer pixel level, it is then required to determine position at sub-pixel level accuracy. This determination is vital in order to accurately determine the centroids of the particles; if this is done accurately, then the subsequent velocity calculations can also be considered accurate.

The sub-pixel level estimation is achieved via the same procedure as above. A mask is applied centred at the particle position estimation from the previous step, and then pixel intensities in this mask are used to weigh the co-ordinates in the direction of the particle centre. The advantage of performing the first step beforehand is that the particle position is more likely to be near to the centre of the particle than it would be by selecting a random pixel; thus, the mask takes in more of the particle of interest than it would do if the first pixel (in red, Fig 4.3) was used. Thus, we expect that this estimation is more accurate as a result of using more of the particle information.

4.2.2. Watershed algorithm and improving particle estimation

On the whole, the above methods of determining particle centroids work well. Using preknowledge of typical particle sizes in the images, it has been proved possible to apply a mask of the same size in order to determine which pixels belong to the same particles. However, in flows containing highly concentrated dispersions of particles, the overlaps may cause problems in particle identification.





Figure 4.4: Example of particle under-estimation and particle over-estimation in clusters

The left hand image of Fig 4.4 shows four particles which have clustered together in the flow. After applying the initial and sub-pixel estimates of the particle positions, we find that three positions have been identified. In this case, the bottom two particles in the cluster have been identified correctly and have also been assigned a position in the centre of the particle. However, the top-most particle has been poorly identified, with the position

allocated towards the bottom edge of the particle. As a result, the remaining particle is not accounted for, since the zone of exclusion from the top most particle has removed it from consideration.

The right hand image of Fig 4.4 shows three particles which have clustered together within the flow. However, there are some size discrepancies due to the particle positioning within the laser light sheet; this has resulted in a much larger central particle. Due to this size difference, there is a double assignment for this particle, due to it taking up a larger size in the image than would be anticipated.

In order to improve on particle estimation and eliminate these two types of errors, a watershed algorithm can be used in conjunction with the algorithms described previously. The watershed algorithm attempts to segment images by using a topological approach. It does this by transforming a gray-scale image to find "catchment basins" and "ridge lines", corresponding in this instance to particle centres and lines of bisection between particles.

To use this algorithm, the image must first be pre-processed in order to match up particle centres with "catchment basins". This is achieved by first creating a binary image using a known threshold value, to only include pixels of interest. Then, the distance transform of this image is computed. This gives the nearest distance between pixels of the same type within the binary image, and therefore contains information about the image topology. Examples are shown in Fig 4.5, following on from the situations displayed previously.







Figure 4.5: Distance transforms of binary images of particle clusters discussed in under and over estimation examples respectively

The complement of this distance transform is found, and subsequently pixels outside of the areas of interest are set of infinite value. This image may then be used in the Watershed

algorithm, which uses the Fernand Meyer algorithm (Meyer 1993). The watershed transforms appear as shown in Fig.4.6.





Figure 4.6: Watershed transforms of under and over-estimated particle clusters respectively

The watershed algorithm determines areas belonging to the same particles, and labels them accordingly. As can be seen from these examples, the algorithm has correctly identified the missing particle in the left hand image, shown in yellow, and has correctly identified the large central particle, again in yellow, in the right hand image. Using this information can then lead to more accurate calculation of the particle centres using the original algorithms, because the information given to these algorithms as inputs is determined by the topology of the images.

4.3.Particle tracking code comparisons

Two different tracking algorithms were proposed for use in analysis of particle movement in the experimental data. These are known as the Hungarian tracking algorithm (Kuhn 2009), and a Nearest Neighbour Linking algorithm (Banitalebi 2008) (Shindler, Moroni et al. 2010). Whilst the former algorithm is more computationally expensive, it may provide better results since it provides a global optimisation, rather than looking only at results on a local basis.

In this section, the mechanisms behind both algorithms will be detailed, along with their theoretical scalings. Following on from this, trial results from a sample set of data will be presented, to both confirm the scaling relationships, and in order to assess the accuracy of each algorithm. Using this information, an argument will be presented for the use of one, or indeed a combination of both, in the analysis of the data presented in this thesis.

4.3.1. Algorithms

The Hungarian algorithm is applied as follows (Lutteke 2012) (Sahbani 2016).

- 1. Particle positions found in two successive frames are listed; those in the first frame are designated as source particles, and those in the second frame are designated as target particles.
- 2. A matrix of values is created between the source and targets, with weightings of likelihood of a link between the two particles. In this case, the weighting is simply the physical distance between the two particles (for an extremely fast-moving set of particles with large times between frames, this would clearly be a poor measure of likelihood).
- 3. This matrix of values can be altered in accordance to certain conditions. For example, if it is known that there is a maximum possible distance of travel, those distances which are above this value can be changed to infinity. This prevents allocation of a link between these two particles, and also reduces the number of possible permutations, reducing time of computation.
- 4. The rows of the cost matrix are then reduced by subtracting the minimum value of each row from that row.
- 5. If there are any columns in the cost matrix that do not have a zero after step 4, these are then reduced similarly by subtracting the minimum column value.
- 6. The zero elements of the matrix are then covered with the minimum number of lines that is possible for them to be covered with. If this number of lines equals the number of possible particles which can be matched, then the matching is optimised.
- 7. If this is not the case, the minimum element in the matrix which is uncovered must then be found. This is subsequently added to every covered element in the matrix. If an element is covered at an intersection of two lines, the minimum element is added twice.
- 8. The minimum element of the matrix is found, and subtracted from every element again.
- 9. The zero elements are covered again, as in step 6. If the number of lines covering the zero elements does not equal the possible particle matches, then the algorithm returns to step 7 and starts this process again.
- 10. Once this has been achieved, a matching between source and targets is created by choosing a set of zeros in the matrix so that each row or column is selected only once i.e. a unique match.
- 11. This matching can then be applied to the original cost matrix, providing the total cost of the global matching.

This algorithm is also known as the Munkres assignment algorithm, and it is generally believed to have a time complexity of $O(n^3)$.

The Nearest Neighbour algorithm is applied as follows (Banitalebi 2008).

- 1. As with the Hungarian/Munkres algorithm, a cost assignment matrix is created between all source and target particles, and is conditioned depending upon additional requirements such as maximal linking distance.
- 2. The target with the lowest cost is found for each of the source points.
- 3. Assignments are made by sorting through the sources in order to find a minima in distances to target.
- 4. If an assignment is made, this is done, and then the source and targets are set to infinity in the cost matrix to avoid further assignments. If an assignment cannot be made, then a further pass is made to find another candidate at a later date.
- 5. Once all of the sources and targets have been assigned, or cannot be assigned due to mismatching, the algorithm completes with a list of associated particles.

This algorithm is generally known to operate with time complexity of O(n²). It is therefore much faster than the Munkres algorithm when dealing with high particle densities. However, it may not be as accurate, particularly in cases where there may be several suitable targets to pair with each source. This is increasingly likely in high particle concentrations, and therefore using this to simply save time as N increases could give some spurious results.

It should also be noted that the arbitrary ordering of source particles may adversely affect results; if there exists a situation where two source particles are in close proximity to one target, this algorithm may assign the wrong source if the false source is higher in the list provided to it. It is then more difficult to rectify this, since unique matching means that this pair would then have to be decoupled. In contrast, the Munkres algorithm uses a global optimisation approach which considers all particle pairs, and so is less likely to suffer from these problems.

Nevertheless, in situations where the majority of particles are moving with a high bulk velocity in a uniaxial direction, the Nearest Neighbour algorithm may actually provide a decent estimate (Mazzaferri, Roy et al. 2015). Since this would also provide a significant time efficiency advantage, the next section will seek to evaluate the performance of the two algorithms, including the computational efficiency, the number of links found in identical situations, and the velocity assignments created from each.

4.3.2. Comparison of performance

These algorithms were compared by using a sample sets of data in which the concentration was varied. Six different concentration ranges were used in order to provide a relationship between average time of computation, <t>, and the number of particles in the image, N. The time of computation was computed over 11 frames of data, with the time recorded over this full range. These trends are shown in Fig. 4.7.



Hungarian tracker performance

Figure 4.7: Comparison of time of computation against N, number of particles in each image, for Hungarian/Munkres and Nearest Neighbour Tracking algorithms

It was noted that each algorithm performed as expected with increases in N, the Hungarian algorithm scaling with N^3 and the Nearest Neighbour algorithm scaling with N^2 . The relative performance, R, is shown in Fig. 4.8.



Figure 4.8: Relative performance ratio, R, between the Hungarian and Nearest Neighbour algorithms, as the number of particles, N, is varied

For values of N which are close to 200, the performance of the two algorithms is roughly matched. However, for particle numbers of 1200, there is found to be a factor of 3 difference in the performance of the two algorithms.

It is also possible to compare the relative success of the two algorithms in creating a match using these sets of sample data. For each algorithm, there will exist a set of sources and targets which are not matched. This represents a data loss. Some of these might occur as particles are lost from images (either through edges or illumination effects), but some are also lost as a result of the algorithm performance. It is preferable to use an algorithm which minimises this data loss. Fig. 4.9 shows the ratio of the average particles found between the two algorithms.



Figure 4.9: Ratio of particle links found, R, between the two algorithms under investigation, as a function of particle number, N

The ratio N_{NN}/N_{H} is close to 99% for the lowest particle concentrations. However, as particle concentrations of N = 1200 are investigated, the losses due to the use of the nearest neighbour algorithm appear to be around 5% on top of other losses which might reduce the chance of successful tracks being found.

Fig 4.10 shows a comparison of the velocity statistics generated by the two different algorithms. In each case, the same pre-processing and inputs are applied, with only the tracking algorithm varied. The experimental data used for this comparison is taken from the vertical channel rig experiments.



Figure 4.10: Comparison of (left) crosswise velocity component PDFs and (right) streamwise velocity component PDFs obtained via differing tracking algorithms, known as Hungarian and Nearest Neighbour methods.

In both of the PDFs shown, the nearest neighbour method enforces the same characteristics on the PDFs. For each velocity component, use of this algorithm results in a larger spread of data, illustrated by lower peaks and higher probabilities at extreme values. Due to the less precise matching, this is likely a result of the algorithm finding a particle which is adjacent to the one required as a match. This is particularly noticeable in the lower tail of the streamwise velocity distribution, where the nearest neighbour over predicts the lower velocity occurrence and indicates an upwards trend at this point, whereas intuitively for a directional flow we would expect a drop off to zero probability.

4.3.3. Minimum travel constraint

It is also possible to vary the minimum pixel value which a particle must travel in order for a track to be eligible for consideration. In most cases, since the experiments observed in this investigation are highly directional, this is a reasonable constraint, which can be checked by manually comparing frame 1 and frame 2 in an image set. An illustration of this is shown in

Fig. 4.11. The red dashed line shows the maximum radius of travel from the starting point, represented by the blue dot at the origin. The blue dashed line represents the minimum radial distance travelled. The solid black lines show the enforced constraint of movement with the bulk flow; it is highly unlikely that particles will move against the net fluid movement, and therefore the radial distances are restricted to a semi-circle.



Figure 4.11: Illustration of the constraints placed upon particle travel by the tracking codes used. The blue dot represents the origin particle, the blue and red dashed lines the minimum and maximum travel constraints, and the solid black lines an enforced direction of travel

Fig. 4.12 shows the effect upon the PDF as this value is increased, putting a tighter constraint upon the movement possible for particles. It should be noted that the movement in each case is also capped at 12 pixels, referring to the straight line distance. These restrictions refer only to the particle movement in line with the mean direction of flow.



Figure 4.12: Comparison of PDFs for Hungarian tracking algorithm when varying constraints upon minimum distance of travel are applied

As the range is decreased, velocity PDFs with a higher kurtosis are observed. It is also noted that, despite the decrease in tracks found due to the tighter restrictions, this track does not simply relate to the 'chopping off' of the lower tail. Indeed, it appears that adding these restrictions can improve the performance of the algorithm, assigning additional tracks to the acceptable velocity range where they would not have been assigned previously.

4.3.4. Algorithm verification using simulated particle movement

In order to verify the performance of these algorithms, a set of virtual particles were created. The particle positions were generated randomly on a 600 x 800 co-ordinate grid, with uniform position in both directions. To each of these virtual particles, velocity components were assigned using a random number generator with a normal distribution. The velocities distributions assigned are as shown in Fig. 4.13.



Figure 4.13: Plots of expected PDFs for both velocity components, u being crosswise and v being streamwise components

The velocities simulated therefore represent a highly directional flow, with the v component having a mean of 8 Arbitrary Units (A.U.) and the u component having a zero mean value. The widths of both distributions were the same in both cases, with standard deviations of 1.5 A.U.

These velocities are then used to progress the positions of the particles from the previous frame. As would be expected, those particles close to the lower boundary of the virtual image are often moved to an area which is outside of the spatial domain under consideration. These particles are therefore removed from the second frame virtual image, so that realism of experimental recording techniques is simulated.

Similarly, under experimental conditions there is also some loss of behaviour, which might be incurred by movement of particles out of the laser light sheet plane, or possibly uneven illumination. In order to simulate this, some particles are removed from both virtual images in the pairs, at random, with varying levels of information loss used.

In order to provide faster analysis, it is possible to impose a minimum value for particle travel on the particles between consecutive frames. Therefore, these values are varied in order to determine the effect of varying this parameter upon the PDFs obtained in this analysis. Fig. 4.14 shows comparisons of the Nearest Neighbour tracking method for varying values of the minimum distance, in comparison to the expected values of v-component velocity shown above.



Figure 4.14: Comparison of Nearest Neighbour tracking methods for the streamwise velocity component, as the minimum tracking distance is varied, with reference to the input velocity distribution. This reference is shown as a black dashed line

As can be seen from Fig. 4.14, when the minimum tracking distance in the streamwise distance is limited only to 2 A.U., the tracking algorithm over predicts the number of particles with a low streamwise velocity. This over prediction is accompanied by a lower peak value than expected, and therefore this represents an under prediction of particles which have the median streamwise velocity. As the minimum tracking distance is increased the tracking prediction matches the input distribution more closely; however, at a distance of 5 A.U., whilst the rest of the distribution is an accurate fit, information is missing from the lower edge of the distribution.

In this case, the best fit via a least squares method is the minimum distance equal to 5 A.U. Nevertheless, it is clear from the input variables that this will miss some information, so it wouldn't be recommended to use a minimum tracking value if it were known that data fell outside of this range. As the minimum tracking distance increases further, the least square value drastically increases, as would be expected since more and more particles are lost from the lower edge of the distribution.



The same information is shown in Fig. 4.15 for the Hungarian tracking method.



For the Hungarian tracking algorithm, the tracking prediction is a good fit with the input distribution for all cases; the only exception to this is for minimum tracking distance of 5 A.U., where information is again missing from the lower tail of the distribution, and there is an over prediction of particles at the median velocity value.

The lowest least squares value in this case is for the minimum tracking distance of 2 A.U. However, the increases only become drastic at 5 A.U., where information starts to be lost from the tracking results. Therefore, the minimum tracking distance can be set to 4 A.U. in this method, with little effect upon the distribution obtained, but a faster processing time can be achieved.

The effect of particle concentration upon algorithm performance was also investigated. In a densely-packed set of particles, the Nearest Neighbour algorithm may struggle, since each particle may have more neighbours within a certain region which it could be paired with. However, due to the relative scaling of the two algorithms' computational time, the Hungarian method may take a prohibitively long time for a set of highly concentrated particles.

Low and high concentrations were investigated by varying the input number of particles in each virtual image set. For the low concentration cases, values were randomly chosen to give between 300 and 900 particles per virtual image. For the high concentration cases, values were chosen similarly to give between 1500 and 2000 particles per virtual image.



Figure 4.16: Comparison of Nearest Neighbour algorithm performance for low and high concentrations, against the anticipated distribution

Fig. 4.16 shows a comparison of the computed particle velocity PDFs for high and low concentration cases, with the input velocity distribution also shown as a reference. In each case, a minimum tracking value in the streamwise direction of 4 A.U. was used, as discussed above. In the high concentration case, all factors discussed above are further exaggerated; there is an over prediction of particles with low streamwise velocity values, which is compensated for by a lower peak value at the median particle velocity.



Figure 4.17: Comparison of Hungarian tracking algorithm for low and high concentrations, compared to the anticipated distribution

In the case of the Hungarian algorithm, there is little difference between the predicted velocity distributions as concentration is changed. Therefore, the Hungarian algorithm deals with the higher concentrations of particles much more effectively.

Algorithm performance was also characterised for different levels of information loss. As mentioned earlier, particles were deleted randomly from frame 1 and frame 2, at a certain percentage level. In all of the cases above, this value was set to 5%. For completely random deletion of particles, this will correspond to track loss of almost 10%, since the overlap will be fairly minimal. Further investigations were made at 1%, which will lead to track loss of just under 2%; the comparisons between these cases are shown in Figs. 4.18 and 4.19.



Figure 4.18: Comparison of predicted particle velocity distributions for the Nearest Neighbour tracking algorithm for different levels of information loss in the virtual images, as compared to the input velocity distribution



Figure 4.19: Comparison of predicted particle velocity distributions for the Nearest Neighbour tracking algorithm for different levels of information loss in the virtual images, as compared to the input velocity distribution

In both cases, the variation in information loss from the virtual images makes very little difference to the predicted velocity distributions obtained from the algorithms, at these low concentrations.

4.3.5. Extension to elongated particle tracks

Given the correct recording positions, it may subsequently be possible to link particles in frame 2 of capture n to particles in frame 1 of capture n+1. Typically, the conditions for this analysis being feasible rely upon limiting the movement of particles between these two images. Naturally, at low recording frequencies, it is typically impossible to link particles within a reasonable time, since the search radius is much larger. Therefore, recording frequencies must be chosen so that particle movement between these two frames is limited to around 8 pixels or less on average.

Therefore, if the time Δt between frames 1 and 2 is 300µs, then the same time should be set between frame 2 and the subsequent frame 1. This means that the recording frequency will be 1666Hz. It may be possible to reduce the time between frame 2 and the subsequent frame 1 even further, but there are limitations, such as the finite exposure time of the camera for each capture.

Practically, the individual tracks were recorded between frame 1 and frame 2 in each data capture. Following on from this, tracks were then recorded between frame 2 and frame 1 of the subsequent capture. Tracks were then built by matching "in-capture" tracks with "out-of capture" tracks. Typically, the maximum observable track under the experimental conditions listed is 50 individual velocity calculations, since most particles will traverse the experimental viewing area in this time.

4.4.Phase discrimination

4.4.1. Methods from literature

During the experimental methods section, the possibility of using optical phase discrimination was discussed. Using this procedure, we can obtain information with only one phase on the camera which is fitted with an orange notch filter, due to the inclusion of fluorescent dye in one flow constituent. This means that less image processing is required for these orange-filtered images, which is advantageous both in terms of reducing processing time and accuracy of results. However, since green light will be reflected from both phases regardless of the inclusion of a dye, the other camera in this experimental configuration will contain information from both phases, and thus some form of phase discrimination algorithm will always be required in order to perform a separation of information.

Such discrimination techniques require some a priori information, usually regarding the properties of the particles to be distinguished in terms of their scattering intensity and size. In terms of these experiments, using solely scattering intensity to discriminate between PIV tracer particles and the dispersed phase would be unreliable; whilst there is generally a larger scatter from the dispersed phase, some tracer particles do exhibit strong scattering which would confuse an algorithm based on this premise. Furthermore, there is too much dependency upon particle position within the laser light sheet to make this a concrete rule.

Another possibility is to use a method which utilises both scattering intensity and size differences between the two particle types. This technique is called a pattern correlation technique. A template image is created, either of a simulated or a pre-selected particle from the experimental images, which provides an accurate representation of a tracer particle in both size and intensity distribution. Candidate particles are then identified in the experimental images, which may belong to either phase, and are then cross-correlated with the template. For instances where the correlation is strong, it can be assumed that a tracer particle is found, and so the original image is left unaltered. If the correlation is found to be below a certain threshold, it may be assumed that the particle is part of the dispersed phase, and therefore the area corresponding to this is changed to the local background level.

The image which is obtained via following this procedure is the PIV tracer image. The dispersed phase image may be obtained similarly, or by subtracted the two images from one another. This technique has been demonstrated in literature, and is a viable method for automated phase discrimination (Kiger and Pan 1999). However, for thousands of particles, it was deemed to be too computationally expensive, since a two-dimensional correlation must be performed for each candidate in the image.

4.4.2. Size-based binary mask

The chosen method for these investigations was the creation of a binary mask, based solely on the size difference between the particles. Since the typical diameter of the particles used is at least an order of magnitude greater than the tracer particles, there is a strong difference in the images of the two which can be used to provide automatic differentiation.

The binary mask is created by thresholding the image by intensity, creating a mask which contains all the particles within the image. The mask is then processed in order to remove all tracer particle elements.

This can be achieved by combining erosion and dilation morphological operations. These operations operate upon pixels in the input image to produce a different pixel value in the output image, depending upon the values of the neighbouring pixels in the input image. The number and shape of the neighbouring pixels considered is determined by the structuring element which is applied when performing these operations.

The rules for pixel value changes are as follows:

- In a dilation, the value of the output pixel is set to the maximum value of all the surrounding pixels within the input pixel's neighbourhood.
- In an erosion, the value of the output pixel is set to the minimum value of all the surrounding pixels within the input pixel's neighbourhood.

In order to operate on the mask, it is necessary to first perform an erosion operation. When the correct structuring element is used, this erosion will remove tracer particles from the image, whilst leaving the larger dispersed phase particles. Nevertheless, the erosion will remove the edges of these particles. Therefore, a dilation is then applied to restore this region. Since the tracer particles will have been removed, there is no pixel intensity in these regions for the dilation to act upon, and so they do not reappear.

The structuring elements used for each of these processes are shown below:

0	0	1	0	0
0	1	1	1	0
1	1	1	1	1
0	1	1	1	0
0	0	1	0	0

Structure element 1: Two-dimensional disk with radius 2.

0	0	0	1	0	0	0
0	0	1	1	1	0	0
0	1	1	1	1	1	0
1	1	1	1	1	1	1
0	1	1	1	1	1	0
0	0	1	1	1	0	0
0	0	0	1	0	0	0

Structure element 2: Two-dimensional disk with radius 3

Most tracer particles in images appear as single pixels. However, there are some tracer particles which, due to being slightly out of focus, can appear as a couple of pixels. Indeed, there are some advantages to leaving tracer particles slightly out of focus when PIV autocorrelations are applied, since it reduces the likelihood of peak locking due to the added ease of interpolation to find the tracer particle centre.

An example of an out-of-focus tracer particle is shown in Fig 4.20, with the image blown-up for clarity.



Figure 4.20: Example of an out of focus tracer particle

With the application of the lower radius structure element, the particle is entirely removed from the mask. This should be clear since, for each of the four mask elements equal to one, there is a directly neighbouring pixel which equals zero, enabling the transition to be made.

A typical dispersed phase particle from these experiments is shown in Fig 4.21.



Figure 4.21: Example of a typical dispersed phase particle

The particle shown in Fig. 4.21 is well rounded, with a diameter of 7. Upon application of an erosion operation with structure element 1, the particle mask appears as shown in Fig. 4.22.



Figure 4.22: Dispersed phase particle after erosion using structure element 1

The particle now appears as a diamond; therefore the diameter has been reduced to three, since two pixels have been eroded from each side. Application of the dilation operation, with the same structure element, restores the image to its original form by adding two pixels to either side. However, use of the second structure element in this dilation produces the particle mask image in Fig. 4.23.

Figure 4.23: Dispersed particle mask image after using erosion with structure element 1 and dilation with structure element 2.

The particle mask generated in this instance has the same properties in terms of diameter and particle centre. However, the corners of this mask are filled in where they are empty in the original image.

In reality, particles do not appear as 'spherical' forms in these images very often. In general, the dispersed phase particles observed are less rounded in shape, and more likely to appear as square. The rounded definition of these particles are only obvious at magnifications which are higher than was typical for these experiments. The typical form of such a particle is shown in Fig 4.24. The diameter of seven pixels has been retained from the previous example.



Figure 4.24: Example of dispersed phase particle image found in these experiments, for which the rounded edges are not obvious and the particle shape is generally square

As with the previous example, structure element 1 is applied in an erosion operation. This operation will remove all tracer particles from the image mask, whilst leaving the dispersed phase particles as shown in the Fig. 4.25. Again, it can be noted that two pixels are lost from each side, as expected.



Figure 4.25: Square-shaped dispersed phase particle after erosion using structure element 1

As with the 'spherical-shaped' particle, a dilation can be subsequently applied with either of the structure elements used. When structure element 1 is used for this operation, the particle mask image appears as shown in Fig 4.26.



Figure 4.26: Particle mask image for a square-shaped particle, using erosion and dilation which both operate using structure element 1

For the 'spherical-shaped' particles, this combination of operations managed to perfectly recreate the input mask. However, in this case, we find that the mask has not fully developed back to its original form, leaving zero elements at the particle corners. This is in contrast to when we use structure element 2 to perform the dilation, which results in full recovery of the original mask image. These differences in methods are summarised in Table 4.1 for both types of particle shape.

Particle type	Dilation structure element radius	Mask recovery
Rounded	2	Complete
Rounded	3	Oversized mask
Square	2	Undersize mask
Square	3	Complete

Table 4.1: Summary of erosion-dilation processes for different dispersed particle shapes

For a spherical shape particle, if structure element 2 is used for the dilation process, this could result in a particle mask which is oversized. The consequence of this is that continuous phase information in the form of tracer particles may instead be conveyed in the dispersed phase images, and be lost from the PIV analysis.

For a square shaped particle, if structure element 1 is used in the dilation process, this can result in a particle mask which is undersized. The consequence of this is that dispersed phase information in the form of dispersed particle edges may instead by conveyed in the continuous phase images, and may not be considered in analysis of particle centroids and tracking.

Neither of those scenarios is ideal. However, it is preferable to have a mask which is oversized. In these cases, some information is lost which can be used in PIV, which can result in more noisy cross-correlations. However, this is preferable to dispersed particle edges influencing PIV cross-correlations with unreliable information. Furthermore, it is relatively unlikely that small tracer particles will be picked up in dispersed phase analysis due to intensity differences which are discriminated in a threshold at a later point. Thus, it was chosen to proceed with structural element 2 in the dilation process, since this will at worst result in an oversized mask when particles are more rounded, and will at best result in complete recovery of the mask. In general, particle shapes were observed to be largely square during these investigations, so complete mask recovery was considered the more likely outcome.

Once the binary mask has been obtained, it is used to process the image in two-different manners, creating the separate images containing different phase information. To create the dispersed phase images, the mask is simply inverted, and then used as a condition to set all non-dispersed phase pixels equal to zero. These images can then be readily analysed using the tracking algorithm, without the possibility of picking up tracer particles.

To create continuous phase images, the original mask is used to make replacements on all dispersed phase pixels. However, it is not advisable to simply set these pixels equal to zero in this scheme. If this is done, then the PIV cross-correlation algorithm will still pick up these areas as regions of zero intensity, moving with the dispersed phase velocity. This will therefore affect the continuous phase statistics. To counter this, a random replacement

scheme is used. This ensures that the 'dispersed phase gap' in frame 1 will not readily be matched to the same corresponding gap in frame 2. To additionally provide consistency with the images themselves, the random numbers used are weighted according to the local levels of background noise, so that the gaps are filled in accordance with the surrounding area.

4.5.Geometric transformations between two camera viewpoints

4.5.1. Transformation types

When using optically-based phase discrimination, the physical space seen by the two cameras used must be matched up if results are to be compared between them. Typically, transformations between images may be broken down into distinct types, including translations, shearing, scaling, rotations, and tilting.

A schematic of the typical camera set up for these experiments is shown in Fig. 4.27. Typically, there would be no anticipated issues with scaling, shearing or rotations, if the cameras are both set up in the same alignment (i.e. parallel to the plane of observation, at an identical distance, with the same focal point). Therefore, the transformation required between the two camera images can be broken down into a translation, and a tilt caused by the requirement of using the same focal point. In general, the aim is still to minimise the translation between the cameras, and therefore also the required tilt, but there will nevertheless be differences between the images.



Figure 4.27: Schematic of camera configuration in a two-camera, simultaneous recording experiment utilising optical filters

The transformation matrix given for a translation in a 2D space is given by:

$$[x_2 \ y_2 \ 1] = [x_1 \ y_1 \ 1] \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ t_x & t_y & 1 \end{bmatrix}$$

Equation 4.1: Matrix representation of translation between co-ordinate sets 1 and 2

where t_x and t_y represent displacement along the x and y axis respectively. A tilt transformation matrix is given the following form:

$$[x_2 \ y_2 \ 1] = [x_1 \ y_1 \ 1] \begin{bmatrix} 1 & 0 & E \\ 0 & 1 & F \\ 0 & 0 & 1 \end{bmatrix}$$

Equation 4.2: Matrix representation of tilt between co-ordinate sets 1 and 2

In this type of transformation, parallel lines in the plane being viewed may converge towards a vanishing point. As the values of E and F grow in size, this vanishing point moves closer to the origin (i.e. the parallel lines appear to converge faster).

4.5.2. Implementation of geometric transformations

Implementation of the geometric transformation was performed by MatLab during these investigations. Initially, a set of control points pertaining to equal points in space are identified in both sets of images. Typically, this can be achieved using a calibration grid placed in the laser light sheet, as shown in Fig. 4.28. Alternatively, sample images can be used of particles in the flow, if it is practical to identify the same particle between two different images.

Once this has been achieved, one set of points is designated a 'moving set', and one is designated as a 'fixed set'. Following on from this, a transformation may be determined to map the moving set of points onto the fixed set of points, matching the camera viewpoints. Naturally, the suitability of the fit depends upon the number of control points identified, plus the spread of those points across the space being mapped. Generally, around 25 points were used for each mapping process, encompassing the majority of the viewed areas.



Figure 4.28: Identification of points on a calibration grid compared from two different camera aspects

An example is shown in Fig 4.29 for a sample set of images, where the alignment is such that there are displacements between particles on the images. In this instance, there is a clear set of particles which can be easily identified in both images, making this a useful test case. In this image, the pixels in green represent areas of intensity in the 'fixed image', whereas the pixels in purple represent areas of intensity in the 'moving image'. It was therefore required to transform the 'moving image', so that the particles spatially matched with the particles in the 'fixed image'.



Figure 4.29: Comparison of two simultaneously captured images of a particulate flow, showing displacement between particles on screen due to the different camera aspects which must be utilised

In this instance, a projective transformation was applied using control point inputs. This type of transformation includes a capability of accounting for tilt in the images. As discussed, this may be required to account for difference in camera aspect. Fig. 4.30 shows the resulting transformation. The colours represent the quantities listed earlier, with white representing a good match between the two images. This small cluster of particles has been well-matched by this transformation, with pixel centres differing by less than one pixel.



Figure 4.30: Comparison of raw "fixed image" (green) with "moving image" (purple) after a projective transformation has been performed. Areas of white represent good agreement between the two images.
Initially, an assumption could be made from the original images shown in Fig. 4.29 that a simple translation would suffice in order to carry out this transformation. This would therefore be covered by an affine transformation, which omits any effects of tilt and assumes transformations are only required in the imaging plane. The last column of the matrices representing these types of transformations contain the column vector [0 0 1] out of necessity. This is because the out-of-plane spatial component must remain the same under these types of transformations, and only a matrix of this form can allow for this to happen.

The effects of an affine transformation are shown in Fig. 4.31. It can be seen from this comparison that the match is poorer. Typically, particles are well matched in the x-direction, but less so in the y-direction, where there is a clear displacement of 2-3 pixels between particle centroids. This is a consequence of the tilt which must be accounted for; whilst it is possible to use just a translation based on the mean distribution of control points, local deviations are introduced by the image tilt which will result in poor matching in less central areas.



Figure 4.31: Transformation of raw "fixed image" (green) with the "moving image" (purple) after an affine transformation is performed

Following on from this example, it is clear that projective transformation should be used for this process. At least 50% of the particles contained within these images are matched to within a distance of 1 pixel by this method. Even so, there are typically larger mismatches in

any image, which can have errors of up to 4 pixels. These mismatches most commonly occur at the edges, where the transform is typically less accurate. This therefore should be considered when considering any results achieved via these means.

Finally, it is worth considering the matrix forms of these types of transforms. For the affine transform of this camera configuration, the following was attained:

/ 0.960	-0.013	0\
0.022	0.961	0)
\-30.1	49.9	1/

Equation 4.3: Matrix representation of affine transformation found using MatLab for the sample images in Fig. 4.28

From the image co-ordinates in Fig. 4.28, and deviations from the identity matrix one would anticipate in the upper left corner, it is clear that some other types of transformations were included than a translation. These may have been in the form of a scaling, a shear, or a rotation. Therefore, even with great care of setting up camera configurations, these types of transformations may need to be considered.

The projective transformation of the camera configuration returned the following transformation matrix:

/ 0.881	-0.040	-6.02×10^{-5}
0.030	0.929	1.28×10^{-5}
\−6.27	66.9	1 /

Equation 4.4: Matrix transformation of projective transformation found using MatLab for the sample images in Fig. 4.28

This matrix now contains non trivial elements in the third column, due to the presence of some tilt. As can be seen from attempts to match particles, these elements are important for reducing errors which can occur in this process. There are also some effects upon the other matrix elements, because the tilt transformation compensates for some effects.

4.6.Conclusion

This chapter has provided some details of the analysis procedures required in order to process particle data from images captured using the PIV equipment in these experiments. Initially, methods of locating particles and finding their centroids were discussed. Initial information can be discovered by performing a watershed transformation on the original image. This algorithm attempts to find "catchment basins" in the pixel intensity, thus giving an initial approximation of particle positions. Following on from this, particle centroids can be found using this initial information; firstly by attempting to find local maxima, and then by using a local mask to weight an initial estimate followed by determination of the particle centre with sub-pixel level accuracy. The application of the watershed algorithm prior to this process was found to improve instances of over and under prediction of particles.

Following on from this, two different candidates for tracking algorithms were examined in order to quantify their performance. These were the Hungarian algorithm, which has a typical performance time of $O(N^3)$ and a nearest-neighbour algorithm, which has a typical performance time of $O(N^2)$. These relationships were then verified using real experimental images.

Experimental PDFs of particle velocity created by both algorithms were created. These showed that the nearest-neighbour algorithm typically over-predicted extreme velocities due to poor matching. However, this can also be rectified by limiting the area of search, which can be achieved using a priori knowledge of the particle images. This approach can be applied to both algorithms, thus reducing the number of partners to be searched over (and therefore time of computation) and resulting in less poor matches.

Both algorithms were also tested using virtual particle images, where particles were assigned velocities randomly based on Gaussian distributions. The Hungarian algorithm performed very well at replicating the input distribution in all situations. The nearestneighbour algorithm struggled unless suitable constraints of searching distance were assigned initially.

During these investigations, the Hungarian algorithm will be used wherever possible, with a search limit for particle partners to reduce time of computation and increase performance. However, in instances where time of computation would become prohibitive, the nearest neighbour algorithm may be temporarily used due to its preferable scaling.

Following on from this, the masking procedure was outlined, which is used to separate phase information in mixed-phase images. Discrimination was performed primarily by particle size, after a threshold has been applied to the image. Structure elements were used

to perform an erosion of this mask, followed by a dilation to recover the elements lost from the dispersed phase mask. The effects of different structure elements were discussed with respect to different types of particle image. Typically, a structure element of radius 2 was used to perform the erosion, whilst a structure element of radius 3 was used as the basis for the dilation.

Finally, the procedure of matching images taken from two cameras simultaneously was discussed. This is used as another method of particle discrimination, with optical filters attached to the camera lenses as outlined in Chapter 3. Geometric transformations were performed in MatLab using control points, determined using the calibration grid. Projective transformations were found to provide lower errors in matching particle centroids, due to their inclusion of tilt transformations.

Chapter 5. Vertical channel rig - design

For the first set of experiments to be conducted in this project, a simple experiment was to be designed which would serve as a preliminary investigation. The design of the experiment needed to fulfil three key criteria. Firstly, the experimental test section should be designed to give optical access, since Particle Image Velocimetry was the planned methodology for the experiments. Secondly, experiments conducted were to cover relevant phenomena from literature in dispersed multiphase flows, which can therefore give viable comparisons between the data produced and existing literature. Finally, it was desirable for the experiments conducted to be simulated easily in CFD models, so that comparisons can be made readily between the two. These simulations were to be performed in a separate, counterpart project.

It was determined that the experiment to be designed would be measurement of an injection of particles into an established single phase flow, conducted in a vertical channel. This experiment therefore satisfies all the above criteria. By constructing the channel from Perspex, optical access was achievable at all points along the test section. The experiment is related to both literature from turbulent dispersed jets, and also literature based on flows of particles in channel flows. Finally, this experiment was designed to allow for comparisons to CFD simulations performed by academic partners at Cranfield University. This was enabled by choice of the test section geometry and the well-defined entry conditions of the particles.

The particle-liquid system to be used for these investigations was the water-hydrogel system. The hydrogel particles used were 500µm diameter, with a tolerance of 50µm either side. A fluorescent dye (fluorescein) was encapsulated within each particle; this dye absorbs the 532nm laser light and re-emits with a longer wavelength, thus allowing for optical discrimination using camera filters within this set of experiments.

Whilst it proved to be difficult to refractive index match these particles in bulk, this was not an issue for these experiments, where depth of particles was much less than in a concentrated, evenly-dispersed pipe experiment. In these experiments, particles were injected in plumes, resulting in minimal overlap. Furthermore, there was an interest to investigate the properties of these particles in a flow, particularly for a potentially densitymatched system.

The injection mechanism was also anticipated to have some effect upon the flow profile of the channel. This is further investigated in Section 5.5 for varying injection conditions.

The experimental equipment for this investigation was built from the ground up, meaning that every facet of the experimental design could be tailored towards these experiments, without having to deal with compromises due to previous design choices.

5.1.Test section design

Some initial choices to be made concerned the test section design. The experiment required establishment of a flow in a channel that might be considered 'two-dimensional' in geometry, which therefore provided a restriction of the aspect ratio for the channel crosssection. This was for the purposes of comparison with CFD; three-dimensional effects are anticipated in all laboratory experiments, but it is useful to create a geometry with which easy comparisons can be made when CFD simulations are considering a two-dimensional condition.

Whilst there is no concrete definition for a 'two-dimensional threshold', there are a few guidelines in the published literature. One paper in particular recommends that the aspect ratio should be larger than 7:1 (Dean 1978) in order to achieve this condition. A further experiment used a ratio of 10:1 for their 'two-dimensional channel' Suzuki, Ikenoya et al. (2000). With these guidelines in mind, it was decided to use an aspect ratio of 10:1 for this test section, since this was felt to be a good guarantee of giving a two-dimensional condition within the CFD comparison.

A further design choice to consider for the test section was the development length required. Again, there are various reported relationships and evaluations of when a flow is fully-developed. For turbulent flows, a readily available analytical expression is given by Equation 5.1 (White 2003).

$$\frac{\ell_e}{D} = 4.4(Re)^{\frac{1}{6}}$$

Equation 5.1: Development length, l_e as function of Reynolds number, normalised by characteristic lengthscale D.

In this expression, the entrance length has been non-dimensionalised by a characteristic length scale of the system, D. In pipe flows, this is typically taken to be the pipe diameter, but in a channel system there is more ambiguity due to the presence of two different length scales. Typically, this is accounted for in an experimental parameter known as the hydraulic diameter, D_h, which is defined by the cross-sectional area divided by the wetted perimeter, multiplied by 4. In this case, for a channel with dimensions W and H:

$$D_h = \frac{2WH}{(W+H)}$$

Equation 5.2: Hydraulic diameter relationship for rectangular channel with dimensions W and H

Given the earlier restriction on the aspect ratio, we can define the channel width in relation to the channel height as:

W = nH

Equation 5.3: Definition of channel aspect ratio

n is the aspect ratio we have defined as 10. This being the case, the definition of the hydraulic diameter becomes:

$$D_h = \frac{2nH}{(n+1)}$$

Equation 5.4: Hydraulic diameter relationship for rectangular channel using aspect ratio

It can be observed that, for an 'infinite' two-dimensional channel (large aspect ratio) this simplifies to the 2H. In some cases, particularly in CFD simulations, this definition is ignored in favour of simply using the channel height as the characteristic length scale. During calculations relevant for these experiments, the hydraulic diameter will be used consistently, but it is important to note that some definitions may differ when comparing work.

For the aspect ratio defined, it is expected that the characteristic length scale for this experiment should be 20H/11. This therefore provides a more stringent restriction upon the

development length than the corresponding situation of simply using the channel height, as can be seen in Table 5.1 for a range of Reynolds numbers in the turbulent flow regime.

Reynolds number	l _e /D (D = H)	I_e/D (D = D_h)
1x10 ⁴	20.4	37.1
1.5x10 ⁴	21.9	39.7
2x10 ⁴	22.9	41.7
2.5x10 ⁴	23.8	43.3
3x10 ⁴	24.5	44.6
3.5x10 ⁴	25.2	45.8

 Table 5.1: Comparison of development lengths using the channel height and hydraulic diameter as the characteristic length scale, for a range of Reynolds numbers in the turbulent flow regime.

Using the more stringent condition provided by the hydraulic diameter, it appears that no more than 50 channel heights would be required in order to achieve fully developed flow conditions for this experimental configuration. Due to the lack of other analytical expressions for turbulent flows, some experimental investigations sought to try and validate this assertion. Typically, all experimental results are quoted with a normalisation by channel height.

From the available experimental results, there is not a strong consensus upon the conditions required for developed flow. Fully developed turbulent pipe flow has been claimed at 30 diameters, based on the mean velocity distribution (Laufer 1953). However, more recent studies (Lien, Monty et al. 2004) have imposed restrictions of 130 channel heights to the condition of fully developed flow. Another study, requiring that mean velocity and turbulence intensity profiles were invariant, found a development length of 71.9 diameter lengths (Perry and Abell 1978). One reason for this lack of consistency is the inability of most experimentalists to gain access to the channel or duct at enough points to give good streamwise resolution.

In order to combat these restrictions, a further study of turbulent pipe flow used a probe capable to fine streamwise movement (Doherty, Ngan et al. 2007). Invariance of mean velocity profiles was found at $I_e/D \approx 60$. Higher order statistics, such as kurtosis, skewness and variance, were considered to be invariant at $I_e/D \approx 80$, as suggested by Perry and Abell.

Questions remain about the effect of coherent flow structures, and their influence upon the definition of fully-developed turbulent flow, even after defining these restrictions.

The claim of $I_e/D \approx 60$ seen previously was identified by another source as being the point of fully developed flow (Zanoun, Kito et al. 2009). Extreme values have also been reported for normalised flow development length of 200 channel heights (Vinuesa, Bartrons et al. 2014). However these researchers claimed that, for high Reynolds numbers, 120 channel heights may be adequate.

A further attempt has been attempted at an analytical solution (Anselmet, Ternat et al. 2009), similar to that seen in Equation 5.1. This research determined that the normalised entrance length scales with Reynolds numbers with an exponent of 1/4, rather than 1/6 as seen earlier. They stated that the use of the previous relationship may not have made much difference for Reynolds numbers typically seen in pipe flows, since the development length is not greatly affected by this change in exponent. Nevertheless, this relationship may become more widely used than Equation 5.1 if further validation is found.

The results shown in Doherty et al. were determined to be sufficient to guarantee fullydeveloped flows. These provide a more stringent condition than is given by the analytical expression using the hydraulic diameter, which predicts that only 50 channel heights would be required for the highest Reynolds numbers used.

As a result of this restriction, it was decided to make the length of the test section at least 110 times as large as the channel height. This would allow for measurements in the 80-90 channel heights range with the camera, whilst also allowing some distance between the measurement point and the outlet of the channel.

The final aspect of the test section design was the inclusion of a flow straightener, in order to encourage the development of the flow profile. These devices act to create reliable flow profiles and turbulent behaviour by eliminating swirl and flow distortions caused by inlet conditions. For simplicity, a perforated plate was chosen to perform this function. This plate was 6mm thick, and perforated with 1.6mm holes, in a 5x50 array, giving a porosity of 50%. In the centre of the flow straightener, this pattern deviates to accommodate a 5mm hole, with acts as the entry point for injected particles. The mechanism for this is described in

Section 5.2.2. The flow straightener was placed at a distance of 5 channel heights downstream of the entrance to the channel.



Figure 5.1: Schematic representation of the test section used in the vertical channel rig. The dashed lines show the outline of a measurement section observed via the camera

5.2.Experimental rig design

The remainder of the test rig was then designed around this test section. The two main options were to establish a flow under gravity, or a pumped flow under pressure. It was preferred to proceed with the later technique for two reasons. Firstly, it was felt that entrainment of air bubbles may be more likely with a gravity-driven flow, due to the more open conditions. This would be undesirable because it would add a third phase to the flows which was not desirable. The presence of these bubbles would also increase optical distortions and reduce visibility, making measurements more difficult. Secondly, by using a closed-system controlled by pressure, we were able to achieve a greater range of flow rates, with better reliability.

The schematic of the overall system is shown in Fig. 5.2; the key features will be described in the following section. Generally, contents were pumped from the storage tank, through a magnetic flow meter, into a top-level tank. There was a bypass line included here, which

returned straight to the storage tank. Water from the top-level tank passed into the measurement section, where it was joined by particles, which were driven under pressure. Once the particles reached the measurement section, they were free to disperse, and the contents were then returned to the storage tank after observation of the flow.



Figure 5.2: Schematic of experimental design for the vertical channel rig.

Water passing through the top-level tank was forced through a series of baffles inside the tank. This was designed to cause flow which was directed downwards before entrance to the channel, and to eliminate the effects of the tank inlet. Water was able to enter the vertical channel across all points of its span.

The tubing carrying the hydrogel particles was 0.5mm outside diameter. This passed through the top of the top-level tank, before descending into the entrance of the vertical channel. A hole was made in the flow straightener; the end of this tube was then aligned with the bottom edge of the flow straightener, to minimise protusion. Hydrogel particles were only able to enter the channel at the centre, in a 0.3mm diameter circular area. This arrangement is demonstrated in Fig. 5.3.



Figure 5.3: Schematic of entrance conditions for water and hydrogel particles into the vertical channel 5.2.1. Pump and flow meter choice and flow rate characteristics

In order to choose a pump for these investigations, it was necessary to establish the flow rate range required for these experiments. Minimum and maximum Reynolds numbers of 1.5×10^4 and 3.5×10^4 were chosen to span a range of turbulent Reynolds numbers. Given the earlier definition of hydraulic diameter as the characteristic length scale for these experiments, the Reynolds number can be defined as:

$$Re = \frac{\rho v D}{\mu} = \frac{\rho}{\mu} \cdot \frac{2Q}{h(n+1)}$$

Equation 5.5: Reynolds number, for flow rate Q and channel aspect ratio n

where the flow rate, Q, is given in m³s⁻¹, and n is the aspect ratio of the channel. Given the water properties and the already established geometries, the corresponding minimum and maximum flow rates were then found as 2.98 m³h⁻¹ and 6.96 m³h⁻¹. Along with the head height required of 2m, this was then used to specify a pump for the experiment, which was a Grundfos CME 5-2-G ARGE AQQE Horizontal Multi-Stage Booster Pump 240V. This is a simple centrifugal pump, which could be used in these experiments since no particles were being pumped. The pump also came with a built in inverter, which allowed for a direct

control of flow rate; therefore, the bypass valve V2 could be used for small adjustments once the inverter had been set.

A magnetic flow meter was chosen to monitor flow rates through the measurement section. This was a convenient choice due to the use of water in these experiments. This was placed after the branch off to the bypass, so that all water passing through this flow meter was subsequently being displaced through the measurement section (it should be noted that valve V3 in the experimental schematic was closed for the duration of the measurements, and was only used during filling/priming of the system).

The flow meter chosen was MIK-5NA65AL model supplied by Kobold Instruments, chosen to fit the flow rate range specified above. Signal output was 4-20mA, which was captured using a National Instruments cDAQ 9219 module. The capture from the flow meter was triggered using the signal which causes the lasers in the PIV system to pulse, and therefore was synchronised with the optical measurements. Once the signal had been received, data was taken from the flow meter for the next two seconds of the flow, which therefore provides all the relevant flow data over the course of the optical measurements.



Figure 5.4: Photographs of a) Grundfos pump used in vertical channel rig b) Kobold MIK magnetic flow meter used in vertical channel rig

5.2.2. Particle-injection system

These experiments required that, once the continuous phase flow was established, particles are then injected into the bulk flow so that they can freely dispersed. Therefore, an injection method was required which could inject the particles, whilst still maintaining a closed system. The injection system also needed to maintain an even distribution of particles to be injected, so that mean injection concentrations were predictable.

Particle injections were achieved using forced displacement. A regulated air supply was connected to a syringe pump (EFD Ultimus I); this unit could then control the air supply to provide air at a specific pressure for a specific length of time. This pressure could be controlled, and was varied to test the performance of the experimental rig.

The output from the syringe pump was then connected to a Drechsel bottle containing the mixture of water and hydrogel particles which was to be injected. When air was supplied to the Drechsel bottle by the syringe pump, it would force the solution through the outlet of the bottle head, which was connected to clear plastic tubing. This subsequently led to a Swagelok tubing section, which allowed this solution to pass through the tank above the measurement section, and down through the flow straightener in the measurement section. The swagelok pipe was adjusted during these investigations to have its outlet aligned with the flow straightener, at which points particles would freely enter the flow.

In order to provide an even suspension of particles, the Drechsel bottle was placed on a magnetic stirrer, and a magnetic bar was added to the solution. By stirring at 300-400 rpm, sufficient agitation occurred to give an even suspension, leading to consistent results from the injection events.

5.2.3. Particle filtration

Without a particle filter at the pump inlet or storage tank outlet, hydrogel particles would continue to circulate around the rig. This may either cause blockages in the pump and other instrumentation, or possibly cause extra particles to be observed in the experiments which were not the direct result of an injection event, but rather a recirculation.

Since this was not desirable, a filter was placed at the outlet of the storage tank. This is shown in the CAD model in Fig 5.5. The filter was designed to block any particles larger than 200µm, which comfortably blocks the hydrogel particles used with an average diameter of 500µm. The filter was also designed to prevent these hydrogel particle captures blocking the water flow, to avoid cavitation to the pump. However, the filter was able to allow passage of water and PIV tracer particles, therefore allowing ready circulation of the continuous phase.



Figure 5.5: CAD drawing of frame construction for tank outlet filter

Fig 5.6 shows a schematic of this filter, indicating the fluid flow through it. Fluid and particles are free to enter the filter at any point through the top; however, they are guided around the outside of the filter by a central insert. Once fluid and particles have reached the floor of the filter, they may then be dragged upwards, towards the tank outlet. This was extended upwards into the storage tank, and fitted with 200µm filter material, which formed a flat surface across the outlet.

In this way, water flowing out of the storage tank was not impeded. Particles dragged all the way through the filter were unlikely to reach the top of the extended outlet, and those which did were blocked by the filter material, and easily removed by the bulk flow. This filter also allowed for easy extraction of injected particles, which collected on the floor of the filter.



Figure 5.6: Schematic of fluid flow through the particle filter attached to the tank outlet. The direction of fluid flow is indicated using a cross-section of the designed filter. The black dashed lines indicate the attachment of the insert, which was permeable to water

5.3.Experimental operation and parameters

5.3.1. Calibration grid

A regular calibration grid was used to calibrate the scale factors required in these investigations. This grid is identical to the one seen in Fig. 4.28, which was used to map between different camera viewpoints. This grid was placed in a dummy test section filled with water and seeding particles, in order to account for distortions. The camera system was then translated to the appropriate position on the test section, so that the distance between the two was unaltered.



Figure 5.7: Example calibration image, captured from a single camera investigation. PIV tracer particles in the flow can be observed in front of the calibration grid.

In addition to providing a scale factor measurement for distance calibration, these grids also provide an array of translation points for image mapping in two camera investigations. These inputs are used, as described in the Data Analysis chapter, to calculate the transformations required to translate between images on different cameras. For these experiments the level of warping in the images is seen to be relatively low, due to the flat faces of the test section which the camera observes.

5.3.2. Data capture – LabVIEW

The capture of flow rate information was aided by using a LabVIEW code which relied upon an external trigger from the PIV system. This was taken from the first trigger pulse produced by the laser, which signals the first laser to fire. A buffer was pre-allocated, which was then filled with readings from the flow meter after the rising edge of the trigger had been received by the cDAQ. The buffer size and frequency were set to match the timescale of the laser measurements being made.

The flow meter output was recorded as 4-20mA. The calibration curve was given as:

$$Q = 1.14 \times I - 3.6$$

Equation 5.6: Calibration relationship between flow rate and output current

Where the flow rate Q is given in $m^{3}h^{-1}$ and the output current I is given in mA.

Prior to each experiment, output values from the flow meter were checked to determine that the flow rate in the channel was as required. These were then recorded with each experiment so that they could be verified using the LabVIEW data afterwards.

5.3.3. Pump stability

The pump was tested at a variety of inverter settings. The integrated inverter utilised a digital control, and therefore delivery could only be fine-tuned in these experiments by use of the bypass valve. The flow rate delivery was characterised for these settings, as shown in Fig 5.8.



Figure 5.8: Characterisation of flow rate delivery according to digital inverter settings. Data points represent flow rate at each digital inverter setting. Flow rates recorded are accurate to 1%, and therefore error bars are not shown.

This flow rate curve shows that the pump chosen for these experiments was more than capable of delivering the flow rates specified in Section 5.2.1, in the range of 3-7m³h⁻¹. However, there was also a requirement to use the bypass valve to achieve these flow rates, due to the fixed control settings of the motor inverter.

In order to assess the stability of the pump, data was recorded for single phase flows at fixed flow rates to determine the fluctuation in flow rate delivery. An example of this fluctuation is shown in Fig 5.9. Generally, the largest fluctuations for any time series were found to be 0.02 m³h⁻¹. Even at the lowest flow rates used, this corresponds to fluctuations

of less than 1% of the experimental value, and therefore the flow rates were determined to be reliable throughout the investigations.



Figure 5.9: Fluctuation of flow rate delivery to test section over time. Mean and standard deviations for this time series are overlaid as solid and dashed lines, respectively

5.3.4. Experimental operation

It was decided to record two types of data with respect to the injection of hydrogel particles into the channel flow. The first recording method was to space the PIV double frames widely, in order to give a reasonable separation between consecutive frames. This would therefore reduce the correlation between these frames, and lead to a greater degree of statistical independence. These experimental runs were then used to gather statistical data.

For a typical recorded image, the span in the vertical direction is 60mm. This is fixed by the need to capture the channel width in the horizontal direction, and the camera aspect ratio. For particles which are travelling on the order of 1ms⁻¹, it is therefore anticipated that they will pass through the recording region within a time of 60ms. The recording frequency for these 'statistical' experiments was usually 250Hz, which gives a time of 4ms between consecutive double-frames; as a result, it should be anticipated that each particle will only appear in 15 frames of data on average, and therefore these images would be unsuitable for building long hydrogel particle tracks. This is particularly true given that, between each of

these consecutive frames, a particle will travel around 50 pixels in distance, which would result in prohibitively slow tracking algorithms for particle numbers above 10³ per image.

For these experiments, typical sample sizes were around 2000 double-frame images. This resulted in recording times of 8s, which was taken from a sub-section of what would be a 20s hydrogel injection event.

The second recording method was to space the PIV double frames as closely as was allowed. This mitigates the problems associated with slow tracking algorithms, by allowing the linking section of the algorithm only a small search parameter. The recording frequency therefore varied with streamwise velocity, and was usually selected to just allow two camera captures to fit within the required time separation. For example, if a pulse separation, Δt , of 700µs is required due to the experimental conditions, then the highest recording frequency possible is 1428 Hz, without considering the non-zero pulse duration. In reality, a frequency of 1250Hz was used for experiments where $\Delta t = 700\mu s$, and frequencies ranged between 1250-2500Hz in these types of investigations.

It isn't possible to simply decrease Δt to very low values to allow for extremely high recording frequencies because there exists an optimal value of Δt for different experimental conditions. When considering the PIV analysis, it is desirable to have a frame separation which allows tracer particles to travel at least 8 pixels in distance, on average. If this displacement is too low, then significant peak locking can occur, and the accuracy of the measurement is reduced. In order to guarantee this, during capture of single phase information the value of Δt was varied to experimentally determine when this condition was met.

Values of pulse separation, Δt , can also be calculated if the mean streamwise velocity and scale factor, s.f., are known, by applying Eqn. 5.7. The scale factor is a measure of pixel distance to actual distance (i.e. the calibration measure from the grid measurement).

$$\tau = \frac{8 \times s. f.}{\bar{v}}$$

Equation 5.7: Derivation of time delay between PIV frames

The values found via this relationship were used as a starting guide before adjusting experimental values during data capture.

5.4.Experimental parameters

As explained previously, the Perspex channel allowed for measurements up to ninety channel heights downstream of the particle injection point. It was decided to perform measurements at three different locations along the channel. Firstly, measurements were made immediately after the injection point, in order to provide data about particle entry to the channel. In this region, we might hope to characterise the more chaotic particle behaviour as particles enter the bulk flow. Secondly, measurements were made at forty channel heights downstream. In this region, it would be expected that the flow would be partially developed, and the mean flow profile should be converging to its final value, whilst other higher order parameters may still not be settled to their final values. The final sets of measurements were made at a distance of eighty channel heights downstream of the injection point; in this region, flow properties should have converged to final values, and therefore these sets of data will represent fully-developed flow.

As already stated in this chapter, flow rates were varied to give Reynolds numbers within a desired range. Five Reynolds numbers were investigated, all of which should represent a fully turbulent flow regime. These are shown in Table 5.2.

Q, flow rate (m ³ h ⁻¹)	V, streamwise velocity (ms ⁻¹)	Reynolds number, Re	
3.26	0.91	16400	
4.14	1.15	20800	
5.43	1.51	27300	
6.28	1.74	31600	
7.02	1.95	35300	

Table 5.2: M	atrix of Reyn	olds numbers	s tested within	the vertical	channel	experiment
--------------	---------------	--------------	-----------------	--------------	---------	------------

Stokes numbers are defined for each of these flow rates. Stokes number changes due to the change in the fluid behaviour, which results in smaller fluid time scales at higher Reynolds numbers. These are shown in Table 5.3.

Reynolds number, Re	Stokes number, St
16400	1.82
20800	2.59
27300	3.89
31600	4.83
35300	5.72

Table 5.3: Stokes numbers for experiments conducted within the vertical channel experiment

Also under investigation was the injection rate of the particles, which could be varied by changing the pressure differential of the syringe pump. In Section 5.5, the effects of injecting single phase material (i.e. water and PIV tracer particles) at varying pressures are investigated.

Finally, concentrations were varied during experiments. Typically, it was possible to achieve hydrogel concentrations of up to 10% volume fraction in the plane of the laser light sheet. This is possible because there are very few particles in parallel planes through which light must travel, and can be achieved by pumping highly concentrated slurry through the entrance point of the channel. During these investigations, concentrations were varied between 1-10%, thus ensuring that all results are representative of four-way coupling flows, where particle-particle interactions are of importance to flow behaviour. However, with increased concentration, a much higher frequency of particle-particle collisions may be anticipated.

5.5.Injection rate experiments – single phase characterisation

PIV measurements were made using only single phase flows in order to characterise the performance and flow characteristics of the rig. In the first investigation, measurements were taken at a distance of 80 channel heights downstream of the injection point. These were made at constant flow rates of Re = 20800, whilst the injection system described in Section 5.2.2 introduced a mixture of water and PIV tracer particles to the flow.

Using these PIV measurements, the typical flow profile for these conditions was discovered. These are replicated in Fig. 5.10, with variation of the injecting pressure of the Drechsel bottle shown between different profiles. The asymmetry is a result of light scattering from

one side of the channel, due to the PIV laser configuration. As a result, it was not possible to obtain reliable results in this section of the channel.



Figure 5.10: Comparison of mean streamwise velocity profiles for differing injection parameters, under single flow conditions, at a distance of 80 channel heights downstream of the injection point and at a fixed Re = 20800

The measurements of singe phase profile shown are compiled from a set of 2000 doubleframe PIV images. If individual velocity measurements are assumed to have errors of up to 10%, the standard error of the mean velocity for the ensemble is therefore given as 0.22%. Additional structure observed on the plots shown is therefore outside of the margins of error. These fluctuations in streamwise velocity with crosswise position are indicative of vortex shedding, most probably from the particle injection point.

For the lowest strength of injection (0.2 bar injection pressure) the profile appears to show a wake at the centreline position. This is in comparison to the standard turbulent streamwise velocity profile, which is flat in shape across the majority of a channel flow. In this case, there is a velocity deficit immediately behind the injection point, in comparison to those velocities either side in the crosswise direction.

As the injection pressure is increased to 0.4 and 0.6 bar, this wake deficit becomes less evident, and the flow profile resembles a standard turbulent flow profile much more strongly. This presents a case that, for this Reynolds number, injection pressures in this range match the surrounding fluid in terms of the incoming velocity, thus 'filling in' the velocity discrepancy.

At the highest two injection pressures, 0.8 and 1.0 bar, the central region of the channel exhibits a much-enhanced velocity. For 1.0 bar injection pressure in particular, the enhancement of the peak velocity is 9% over the expected baseline velocity. These flow profiles are also greatly distorted from the anticipated form; despite the Reynolds number clearly being in the turbulent range, the profiles appear more akin to laminar flow profiles due to the added distortion. Therefore, these higher injection pressures contribute an added momentum along the channel centreline.



Figure 5.11: Comparison of mean streamwise velocity profiles at a distance of 40 channel heights downstream of the channel entrance, for injections at 0.4 bar and without any injected material. These are reproduced for a) Re = 20800 b) Re = 27300 c) Re = 31600

Further experiments were conducted at a distance of 40 channel heights downstream of the injection point. Fig. 5.11 shows comparisons between mean streamwise velocity profiles without material being injected into the channel and with injection of single phase material at 0.4 bar. The first case is taken at the identical Reynolds number to those experiments shown in Fig. 5.9, and therefore direct comparison can be made. In the less developed

region of 40 channel heights, the velocity profile hasn't completely flattened. Whilst it appears to be tending towards this shape, there is more complex behaviour in the channel centre, such as a number of local peaks. Additionally, the velocities in this region are around 0.04 ms⁻¹ higher than those in the more well-developed region of 80 channel heights downstream.

The velocity profile without any material injection shows the presence of a well-established wake region, superimposed onto a roughly turbulent flow profile. The velocity loss in this central region is 0.08 ms⁻¹, or 7% of the anticipated value, and is caused by the mechanism discussed previously. When material is injected in this instance, there actually appears to be an overcompensation, as opposed to the flattened profile observed in the well-developed region of 80 channel heights which indicates a matching of velocity.

The graphs from the two higher Reynolds numbers show similar behaviour, but with some differences. Firstly, the wake is still visible in the velocity profiles for single phase injection. This is particularly true for the higher Reynolds number of 31600. Therefore, the injection is less able to compensate for velocity loss in the central region. Instead, the injection of material appears to raise the overall velocity profile, across the whole width of the channel. This is particularly noticeable at the channel edges, where a 0.4 bar injection pressure at higher Reynolds numbers results in larger differences in velocity.

This indicates that, for higher Reynolds numbers, injection does not 'fill in' the wake profile, but instead contributes to fluid momentum across the full crosswise direction. This gives an indication of how the added fluid enters the channel when the flow conditions are faster and more turbulent. Greater mixing results in a spread of the additional downward momentum, resulting in a raising of the velocity values for the overall profile.

5.6.Conclusion

This chapter has presented the design decisions and creation of the vertical channel rig, along with implementation and initial testing. Dimensions of the test section were justified with respect to literature and CFD simulations being created as a comparison. Equipment was sourced in order to satisfy the experimental conditions required. Following on from this, details of PIV calibration, and the data capture from the flow meter were given. The test campaign for this experimental rig was then outlined, including the range of

concentrations, injection pressures, Reynolds numbers, particle Stokes numbers and the different distances downstream of the injection point at which measurements were made.

Finally, the performance of the rig was characterised by conducting single-phase only experiments, either by running the rig without material injection, or by injecting a mixture of water and PIV tracer particles to simulate full operation. In the well-developed region, injection pressures of 0.4 and 0.6 bar were found to reasonably compensate for the wake profile created by the injection system, thus creating mean streamwise velocity profiles as expected from turbulent flows. Particle injections were tested for a variety of injection pressures in the following chapter, but are usually injected at 0.4 bar unless otherwise stated within Chapter 7.

Further results were presented from the developing region of the channel flow. These showed more complicated forms of the mean streamwise velocity profile, and also indicated that the wake is more persistent for higher Reynolds numbers.

In the next chapter, the specific analysis routines required to characterise clustering of particles will be described, and verified using simulated particle distributions.

Chapter 6. Analysis techniques

In this chapter, the various techniques used to analyse clustering in particle distributions are outlined and validated using simulated particle distributions. This is tackled as a separate issue to provide some context of the algorithms used before showing the experimental results.

Initially, the simulated particle distributions used to validate these algorithms will be presented. Varying conditions were introduced to these distributions, which display important behaviour about some of the algorithms, and points to behaviour which might be expected experimentally.

Algorithms and validation are presented for three differing methods of quantifying preferential concentration in particles. Two of these were discussed in the literature review, and have been used by other researchers. These are known as the "box-counting method" and the "Voronoi cell analysis method". Following on from this, another method devised during this project, known as the "Particle pair potential" method, will also be presented.

Finally, the algorithm used to create the radial distribution function data shown in the results chapters is presented, along with some discussion of the Monte Carlo method which it relies upon.

6.1.Simulated particle distributions

6.1.1. Distribution descriptions

In order to validate the algorithms used to quantify particle distributions, test cases of simulated particles were developed. Some were developed to have well-known properties with respect to these algorithms, enabling the codes to be validated. Others were developed in order to test the behaviour of these codes with respect to simple changes in particle behaviour.

Some of the simulated particle distributions were created so that each particle was given an excluded radius. This then imposes a condition that no other particle can be placed within this distance of that particle. This corresponds to a real-life case, where particles cannot occupy the same volume.

For all of the simulated cases described, numbers will be quoted as the 'pixel values' or numerical co-ordinates, and units will be stated as arbitrary given that there is no physical significance to these particular results. These are denoted by A.U., short for Arbitrary Units.

The types of particle distributions considered are summarised in Table 6.1

Table 6.1: Description of the four types of simulated particle sets used to test the particle distribution codes

X-distribution	Y-distribution	Excluded radius (A.U.)
Random uniform	Random uniform	0
Random Gaussian	Random uniform	2.5
Random uniform	Random uniform	0
Random Gaussian	Random uniform	2.5

Each of these particle sets were simulated in a rectangular box with an area 4.8×10^5 A.U.². In addition, particle sets were varied in concentration for each of the types listed in the table, between 200-2200 particles per image. Images sets were created to give 200 distinct 'frames' of data to test on for each concentration.

6.1.2. Distribution properties

In all cases, a uniform distribution was chosen in the y-direction. This replicates a degree of axial uniformity which would likely be found in the experiments conducted within this thesis. For the uniform distributions, this was replicated in the x-direction. For the Gaussian distributions, particle positions were generated based on a Gaussian profile, chosen with mean value at the centre of the experimental volume and a sigma value which was half of the total width in the x-direction.

An example of these distributions is shown in Fig. 6.1, from a Gaussian profile with no excluded radius. As a result of sharp cut-offs due to the boundaries of the measurement area, the Gaussian profile is prematurely terminated. This means that the overall distribution is technically non-Gaussian, but serves as an approximation. The fitted curve gives parameter values which were given at the point of particle generation, and so the only difference is the abrupt cut-off.

As stated, some particle distributions were tested with an excluded radius. This means that, when creating the distribution, it should not be possible to place particles within a certain distance of each other, simulating the actual behaviour of particles with a finite volume. The

effect of this is to create a total excluded volume (or area from a two-dimensional point of view) which further particles cannot be added to within the measurement volume. As concentration increases, so will this excluded volume, leading to fewer degrees of freedom on particle placement. It is of interest to see how these excluded volume effects impact upon observed statistics for particle spatial distribution, since there are obvious parallels to experimental situations.



Figure 6.1: (Left) Plot of probability distribution in the X-position coordinate for the Gaussian profiles (Right) Plot of probability distribution in the Y-position for the corresponding particle distribution, with the mean probability superimposed (dashed line).

6.1.3. Particle concentration

For simulated particle cases where no excluded radius is used, it makes little sense to talk about concentration in the traditional sense. Since there is no volume or area attached to an infinitesimal point, a volume fraction cannot be calculated in these instances. However, for the excluded radius cases, it is possible to make such a calculation. For each instance, a particle radius of 2.5 A.U. has been assumed, and therefore particles may not appear within this distance of each other.

Equation 6.1 describe the volumes, or in the 2D case, area fraction of the particles.

$$\alpha = \frac{N \times (\pi \times r^2)}{V_T} \times 100$$

Equation 6.1: Area fraction calculation

In this case, r is the excluded radius, V_T is the measurement volume, and N is the number of particles. Table 6.2 shows the area fraction taken up by particles for each concentration of simulated particles:

Table 6.2: Area particle fractions associated with	different particle numbers per image frame in excluded
volume cases	

Particle number per frame	Area fraction (%)
200	0.82
400	1.64
600	2.45
800	3.27
1000	4.09
1200	4.91
1400	5.73
1600	6.55
1800	7.36
2200	8.18
2400	9.00

One factor which is affected by the addition of an excluded volume is the spread of the Gaussian distribution. As more particles are added in an excluded volume case, less space is available in which to place the particles. Therefore, the statistics are affected, as shown in Fig 6.2. When volume exclusion is considered, the spread of particle x-position increases as additional particles are added. This is a result of particles being added unnaturally to the image to avoid exclusions, forcing a greater positional spread as excluded volume becomes more influential.

109

Analysis techniques



Figure 6.2: Relationship between standard deviation of simulated particle position in x-direction, and number of particles, N, in each image, shown for both non volume excluded and volume excluded cases.

6.2.Box counting method

6.2.1. Overview of box counting method

One method of quantifying preferential particle concentration is to use a box counting method. This method relies upon overlaying a particle image with a regular grid. Particle occupation of each grid cell is counted, and the results across all cells in the grid can be expressed as a PDF for likelihood of particle occupation. The cell size may be changed arbitrarily in this analysis, and so multiple cell sizes can been investigated to determine effects of length scale.

For an entirely random process of particle placement with a uniform distribution, the statistics of particle occupation within a defined cell volume should obey a Poisson process. Therefore, the PDF in each case would take the form:

Analysis techniques

$$P(N) = \frac{\lambda^N e^{-\lambda}}{N!}$$

Equation 6.2: Equation of Poisson distribution

N represents the particle occupation of a cell, and λ represents the sample mean number of particles per cell, and would therefore be expressed as:

$$\lambda = \frac{N_T}{Cells_T}$$

Equation 6.3: Definition of Poisson mean for box counting method

In Equation 6.3, N_T is the total number of particles observed, and Cells_T is the total number of discrete volumes/areas used to compute the Poisson distribution statistics.

However, if particles are preferentially accumulating, then the number of cells with high particle occupation will be higher than Poisson prediction would allow. Furthermore, voids would be left in other cells, which would result in a high N = 0 count. As a result, for strong preferential accumulation, we would expect to see a larger standard deviation in the cell occupation PDF than would be seen for a Poisson process.

The standard method of quantifying this deviation from the Poisson PDF is found in Wood et al (Wood, Hwang et al., 2005):

$$D = (\sigma - \sigma_P)/\lambda$$

Equation 6.4: Definition of Poisson deviation as defined for box counting method

In Equation 6.4, σ and σ_P represent the experimental and Poisson standard deviations respectively, and λ represents the sample mean cell occupation number. For experiments with no particle clustering, the value of D should remain at zero, because distribution of particles will essentially be random; as clustering levels increase, so will D.

In the next section, varying methods of implementing the box counting procedure will be discussed. These will be examined firstly for a uniform random distribution, and then for the other three simulated particle distributions to determine the effects of excluded volume and a Gaussian profile upon the behaviour for D.

Analysis techniques

6.2.2. Box counting methodology

For these validation tests, cell size was incremented in steps of five pixels, ranging from five to two hundred pixels. Since the total spatial range did not always divide by each cell size to make an integer number, the number of cells on the grid in each direction was rounded downwards so that the grid could fit entirely within the particle domain (a larger grid would have included vacant areas outside of the particle domain, which would have produced unbalanced statistics).

The next step to be determined was the placement of the grid on top of the particle domain. Given that the overall grid size is smaller than the particle domain, there is some freedom in this placement. The initial method tried was a central placement, whereby the centre of the grid was aligned with the centre of the particle domain. An example of this is shown in Fig. 6.3.



Figure 6.3: Centralised gridding method, shown for a randomly-generated particle set. In this example, the grid size is 55 pixels, with the generated grid being 10 x 14 in size. This method leaves gaps of 25 either side in the x-direction, and 15 either side in the y-direction

For a random set of particles, a full ensemble of two hundred images was analysed using the gridding method. The results are shown for a few values of N (particles per image) in Fig. 6.4. Values of D do not exceed a magnitude of 0.015, with values typically being negative. For higher concentrations of particles, the magnitude of D is lower, which indicates the convergence with higher statistical significance towards values of zero. Therefore, these distributions appear to conform to Poisson statistics within a percentage deviation of 1.5% in the most extreme case.

Analysis techniques



Figure 6.4: Comparison of D against cell size, for multiple values of N, particles per image frame. Shown for reference is the expected value of a Poisson process for this type of analysis (black dotted line)

Appendix B contains further information about other methods which may be used for a box counting method of quantifying clustering. The method used here was a 'centralised' grid, with a summation approach to averaging histogram statistics (statistics are combined across all available images). Appendix B investigates the possibilities of using a 'corner-anchored' grid and an alternative averaging method which quantifies clustering on a frame-by-frame basis.

It was determined that a centralised grid approach, using ensemble statistics, was preferable towards providing accurate results. The full breakdown of this decision is presented in Appendix B.

6.2.3. Further particle distributions

In this section, results for other types of simulated particle distributions are shown, using the chosen method.



Figure 6.5: Comparison of D against cell size, for three values of N, the number of particles per image frame, for a uniformly distributed set of particles with an excluded volume. Shown for reference is the zero line, the Poisson reference for a completely random set of particles

Fig. 6.5 shows the relationships of D against cell size when excluded particle volumes are included. As expected, this restriction lowers the standard deviation of the experimental distribution, which results in negative D values. The lowest D values occur for the highest concentrations, since these are the conditions under which the excluded volume has the most effect.

This deviation is particularly noticeable for the lower cell sizes. Observations on these length scales are comparable to the length scale of exclusion, and so capture this effect the most clearly. In comparison, those measurements made on length scales comparable to the full spatial domain do not notice these small scale effects, and so the deviation values become close to zero again.

Fig. 6.6 shows trends in D for various calculated particle area fractions, using the excluded radius of 2.5 A.U. The plot on the right hand side shows that, if the D values are normalised by the square root of the area fraction, α , there is a reasonable collapse of all the data.
Therefore, this scaling may be utilised when making comparisons to experimental data in which excluded particle volumes are important.



Figure 6.6: (Left) Plots of D against cell size for varying simulated particle area fractions, as calculated using the excluded area (Right) Normalisation of curves using the square root of volume fraction, α, showing collapse of data

Analysis techniques



Figure 6.7: Comparison of D against cell size, for three values of N, the number of particles per image frame, for a set of particles with Gaussian distribution. Shown for reference is the zero line, the Poisson reference for a completely random set of particles.

D values appear to follow an overall upward trend when considering a Gaussian distribution of particles. Therefore, as the cell size used increases, so does the level of preferential accumulation observed. This is consistent with expectations, since the deviation from random distribution is global in this distribution. Therefore, it would not be expected to pick up on these differences using a small box size, but it is understandable that it would be noticed by box sizes which are comparable to the size of the spatial domain.

A similar relationship is shown for the Gaussian distributions of particles with an excluded volume effect, as shown in Fig. 6.8. There is also a drop in D value with the inclusion of excluded volume, as seen with the uniformly distributed cases, when compared to Gaussian distributions with no excluded volume. Therefore, the trends start with negative values, before increasing at the same rate as those distributions with no volume exclusion. There is also some noticeable fluctuation of values within these graphs.

Analysis techniques



Figure 6.8: Comparison of D against cell size, for three values of N, the number of particles per image frame, for a set of particles with Gaussian distribution and an excluded volume. Shown for reference is the zero line, the Poisson reference for a completely random set of particles

6.2.4. Process of calculation

The implementation of the box counting method is shown in Fig 6.9. The code requires two

loops; one over the images in the data set and one over the cell sizes to be investigated.



Figure 6.9: Schematic for operation of the box counting routine.

6.3.Voronoi cell analysis method

6.3.1. Voronoi diagram formulation and statistics

For any set of points in a two-dimensional space, it is possible to construct a Voronoi diagram. These diagrams divide up the space into discrete cells, each of which contains only one of the points, and additionally all co-ordinates in the physical space which are closer to this point than any other. This diagram appears as a series of bisection lines between all the points in an image.

These diagrams can then be interpreted in terms of particle concentration. In an area of low concentration, neighbouring particles will be spaced at large distances, and henceforth the

local cells created in the Voronoi diagrams will be larger. Thus, the local cell area, A, can be considered the inverse of the local particle population density.



Figure 6.10: Example of a Voronoi diagram, constructed using a single frame of data from the uniform random distribution.

As with the box counting method for quantification of particle clustering, a comparison case is required in order to draw any meaningful comparisons from the experimental data. In this case again, there is a well-known comparison case. The most often cited comparisons in literature are those of a random Poisson process (RPP)(Monchaux, Bourgoin et al. 2010), which in our particle case studies refers to a random uniform distribution with no volume exclusion. The properties of the Voronoi diagrams for these distributions have been studied (Jarai-Szabo and Neda 2007). For a construction of the normalised Voronoi area PDF (V = A/<A>), one finds that the standard deviation of this distribution, σ^{RPP} , equals 0.53. There are no analytical solutions for this distribution available, although most are fitted using a Gamma distribution. Therefore, for the purposes of comparison with experimental results in these investigations, a set of Voronoi statistics was assembled for the simulated particle distributions illustrated earlier in this chapter, including the RPP case. Properties of the PDFs from these distributions will be discussed in this section, before using these as comparison cases for experimental results.

6.3.2. Standard deviations of cell-size distributions for simulated particles

As with the box counting method, the degree of particle clustering can be linked to the standard deviation of the experimentally acquired PDFs. If a large degree of clustering occurs, there will be regions of the physical domain with higher than average concentration, and regions which particles have vacated which will be very sparsely populated. In terms of the Voronoi analysis, this corresponds to an abundance of areas with very low and very high normalised Voronoi cell areas. Thus it would be expected that when clustering occurs, the experimental PDFs of normalised Voronoi cell area will exhibit a higher standard deviation than those generated by randomly simulated particle distributions.

In Fig. 6.11, it can be seen that for the non-volume excluded cases, distribution standard deviation has a weak dependency upon the particle concentration per image. Naturally for Gaussian distributed data, we would expect a higher σ value, due to large numbers of small Voronoi cell areas along the vertical centreline, and large numbers of larger Voronoi cell numbers towards the edges. The RPP data appears to confirm the assertion that $\sigma^{RPP} = 0.53$. The gradient of best fit for the Gaussian distributed data indicates a weak positive trend towards greater clustering at higher concentration. There appears to be some settling to a final value as N increases, so the upwards trend is likely due to settling towards statistical equilibrium.

For the volume excluded methods, negative gradients can be observed in the fits of standard deviation to N. As particle number, N, is increased, the space remaining to fill with further particles is reduced, thus providing less freedom. If we were to imagine an extreme example of this where all space is occupied by particles as fully as possible, we would observe a close-packed regular particle bed, for which the Voronoi PDF would clearly be single-valued. In this instance, the standard deviation would therefore be zero. This gives an

indication as to why the standard deviation drops with increases in particle concentration, and so this may be relevant when we are constructing comparison statistics.



Figure 6.11: : Standard deviations for normalised Voronoi cell PDFs created from four different simulated particle distributions, varying with simulated particle number N.

6.3.3. PDF shape

There is no analytical shape currently known for PDF of normalised Voronoi cell area in a random Poisson process. An attempt has been made to fit a shape using a Gamma function (Jarai-Szabo and Neda 2007), which is defined by Equation 6.5 for a two-dimensional problem.

$$f(y) = \frac{(\frac{7}{2})^{(\frac{7}{2})}}{\Gamma(7/2)} y^{5/2} \exp\left(-\frac{7}{2}y\right)$$

Equation 6.5: Analytical fit for PDF of normalised Voronoi cell area distribution for a random process

An attempt was made to fit this function to the data produced from the uniform distribution test case, which is shown in Fig 6.12.



Figure 6.12: Comparison of PDF of normalised Voronoi cell areas taken from uniform random particle distributions with the fitted Gamma equation.

From the regular scale plot, there does not appear to be a large difference between the predicted model and the data acquired. However, when observing on a log scale, it can be observed that the prediction actually is too low for the lowest normalised cell areas, with divergence taking place at around A/<A> = 0.1. This is also noted in the paper from which the model is derived, and so it should only be considered as a guide. Within the context of this investigation, it is preferable to use the raw data values taken from simulated particle distributions as a comparison.

Fig 6.13 shows the comparison of the Voronoi cell size PDFs for Uniform and Gaussian distributed points, at varying concentrations. The minimal differences in sigma values shown above are indeed reflected here, since most of the curves collapse on top of each other in these instances.



Figure 6.13: Comparison of Voronoi cell area PDFs for uniform and Gaussian distributions with no volume exclusion, over a range of concentration

Fig 6.14 shows the comparison between a uniformly distributed set of particles and those with a Gaussian distribution, for the same concentration of N = 2000 per image. The larger standard deviation of the Gaussian distribution is evident in the tails of the distribution where the probability values are larger, thus resulting in a flatter distribution.

For the Gaussian distribution, it can be seen that Equation 6.5 is an even less suitable fit, with significant divergence in both tails of the distribution.



Figure 6.14: Comparison between Voronoi cell area PDFs for uniformly and Gaussian distributed sets of particles at N = 2000. Also shown is the Gamma fit from the Equation 6.5

The PDFs were also calculated for the particle distributions which take volume exclusion effects into account. These are shown in Fig 6.15. As can be seen from the standard deviation values, these distributions have a lower spread, and hence the tails of these distributions typically sit below the Gamma-fit line. This is most noticeable in the low Voronoi cell area region, which would account for particle clustering in an experimental case. The effect is also exaggerated with increased concentration, since particle placement becomes more restrictive in these instances.



Figure 6.15: Comparison between Voronoi cell area PDFs for uniform and Gaussian particle distributions with volume exclusion at varying concentration. Also shown is the Gamma fit.

6.3.4. Voronoi code structure

Using these sample distributions, it is possible to make a comparison with experimental data. Fig 6.16 shows a flow of information within the Voronoi analysis routine, which will be explained in this section.



Figure 6.16: Flow of information in Voronoi analysis method

Once an ensemble PDF has been created for the Voronoi cell areas in an experimental run, it is compared to a sample PDF of the same particle concentration from a simulated particle distribution. It was felt that the most representative data set to use was the uniformly random, volume-excluded data, since this most closely represents the information we are able to extract from images. Given that this is the case, the instantaneous concentration is also taken into account for the experimental measurements, so that fluctuations in time are represented in the comparisons made. This is necessary due to the changing shape of PDFs with concentration in volume-excluded cases.

The comparison is made as shown in Fig 6.17. The experimental data is a sample, taken from the vertical channel results in Chapter 7, for illustrative purposes. The cross over points are calculated by using an initial estimate, followed by a linear regression. These cross over points, V_c and V_v , are then used as definitions of the upper limit for cluster cell size and lower limit of void cell size respectively. These are marked on Fig 6.17 by a solid and dashed line respectively. Once these have been found, the original Voronoi diagrams can be analysed again, with assignments made to each cell of cluster, void or neutral status.



Figure 6.17: Comparison of experimental (vertical channel, Chapter 7) and simulated data, with definitions of Vc and Vv determined by crossing over points for the two distributions. This splits the Voronoi cells found into three regions, depending upon size and relation to crossover points

Fig. 6.18 shows an assignment of cell status using these definitions. Once this has been achieved, local regions can be formed from adjacent, self-similar cells. This is achieved by looking for cells which share both common edges and the same designated status (e.g. void). These regions are then stored, with information such as vertex co-ordinates, area, perimeter and the number of particles contained within the region. This information is evaluated within the experimental results sections for this analysis technique for varying operational conditions.



Figure 6.18: Section of Voronoi diagram taken from sample experimental data, where cells have been assigned void (blue), neutral (white) or cluster (red)

6.4.Particle-pair potential method

This method uses only the nearest neighbour particles in order to calculate the levels of local clustering. The space around each particle is divided into angular regions. The closest particle within each region is identified, and the distances between each of these particle pairs is calculated. An example of this is shown in Fig. 6.19.



Figure 6.19: Identification of nearest neighbour particles (blue) around each central particle (red) is performed, for each angular region. Distances, r_i, are found between each of these pairs. The formulation thereby uses an assumption of shadowing, so that particles further away in the same angular region (black) do not contribute to the calculation

The level of clustering, Σ , is thereby quantified by using a potential formulation, using Equation 6.6, where n is the number of neighbouring particles used.

$$\Sigma = \frac{1}{n} \sum_{i=1}^{n} \frac{1}{r_i}$$

Equation 6.6: Calculation of local particle potential

This means that each of the individual potentials contributes to an average. The angular form of these local clusters can also be determined, if so required. Some examples of this angular distribution are analysed in Chapters 7 and 9.

The statistics of the potentials calculated for each particle within the experimental flow can be compiled, creating a PDF which may then be analysed. Initially, the potentials are normalised, in order to adjust for concentration effects. This is achieved by using the potential calculated from a lattice formation using an identical concentration of particles. In this model, the potential of each particle would be, using 12 nearest neighbours:

$$D = \sqrt{\frac{N}{A}} \left(\frac{3 + 2\sqrt{3}}{9}\right)$$

Equation 6.7: Calculation of normalisation distance

For three-dimensions, this value should be adjusted accordingly to reflect the extra degree of positional freedom. In this equation, N represents the number of particles imaged and A

the imaging area.



Figure 6.20: Comparison of PDFs of particle potential for an experimental distribution and a random simulated particle distribution. In this case, a 12-particle nearest neighbour formulation is used, thus giving the normalisation in Equation 6.7.

The PDF of normalised particle potential is shown in Fig. 6.20, for experimental data from the vertical channel experiment, in contrast to a simulated set of random particles. It is clear that clustering exists in comparison to the random placement, since the experimental distribution is wider. There are therefore more instances of low particle potential (i.e. few close local neighbours) which represents void particles (particle occupying sparse regions), and also more instances of high particle potential (i.e. many close local neighbours) which represents cluster particles (particles occupying highly populated regions).

As with the Voronoi formulation, the crossovers of these PDFs can be used to define relevant critical potentials regarding differing behaviour within the flow. Therefore, those particles with potentials below the first cut-off are denoted as being void particles. Those particles with potentials above the second cut-off are denoted as clustering particles i.e. those drawn into cluster regions preferentially. Examples of these are shown in Fig. 6.21.



Figure 6.21: Example of labelled particles, where red particles represent those considered to be clustered. Particles cluster along strands in this image, which bears similarities to particles clustering along lines of strain and vacating vortices, as described in Chapter 2

6.5.Radial distribution functions

Radial Distribution Functions can be used in a particle distribution to quantify how particle density varies as a function of distance from a chosen reference particle. This is done by analysing a series of annular elements around the reference particle, and calculating the number of secondary particles which fall within this area. This particle density is then compared to the total density of particle pairs within the whole measurement volume; any value of this ratio which exceeds 1 indicates clustering of particles.

The layout of the RDF code is shown in Fig. 6.22. The first step is to reduce the quantity of data to restrict computational time by pre-selecting camera frames to be analysed, as discussed in Appendix C. A range of radial elements are set up, and the corresponding full annular element areas are calculated.

Following on from this, a loop is established to consider different primary particles. In a secondary loop, displacements are calculated to all secondary particles in the frame data. Naturally, to avoid double counting, a primary particle which has been considered already will be omitted from any further secondary particle lists. Also calculated for each primary particle are the partial annular elements, calculated using the given boundary conditions and the Monte Carlo method.

Following collection of data, the annular element areas and overall measurement volume areas are normalised by the total pair density within the measurement volume, N/V. The RDF is then calculated according to Equation 6.8.

$$g(r_i) = \frac{\frac{N_i}{\Delta V_i}}{\frac{N}{V_V}}$$

Equation 6.8: Definition of Radial Distribution Function

In this equation, the terms ΔV_i and V represent the averaged annular-volumes and measurement volume described above, and N is calculated from the summation of all possible particle pairs in the data set analysed.



Figure 6.22: Schematic of the RDF code used, starting with the raw positional data

In the following section, the performance of the Monte Carlo codes required in each secondary loop are examined. Due to their frequent usage, and potentially high computational cost, an assessment is made to determine the best trade-off in terms of accuracy and computational time.

6.5.1. Monte Carlo methods

In order to accurately assess the Radial Distribution Functions (RDFs) of particle positions, it was necessary to consider edge effects, which would occur when the annular elements

surrounding the central particle cross the boundaries of the particle-containing area. This is particularly likely if we are considering the RDF at larger displacements, and therefore if this is not factored into the analysis, measurements will only be possible for the smallest displacements.

These edge effects can be accounted for by calculating the percentage of the annulus which is included in the particle containing region. In other words, the contribution to average area from each annular element is weighted by the percentage of which is contained in the measurement area. For example, in the case illustrated in Fig. 6.23, the weighting should be 0.25. This weighting must be applied to the area used in the overall calculation, since the remaining 75% is not a representative sample of the measurement volume.



Figure 6.23: A particle located at (0,0), surrounded by an annular element with Rmax = 2 and Rmin = 1. The measurement region is specified as being bounded by x = 0 and y = 0, which means the particle is located at the corner of the measurement region. In this case, the annular element contained within the measurement region is only 25% of the total area, and so this weighting is applied to the RDF calculation.

In order to calculate the areas included and therefore the weighting required, the Monte Carlo method is used as a means of area estimation. Using the inputs of particle positions and the annular radii being considered, a random set of particles may be generated on this annular element, as shown in Fig. 6.24.



Figure 6.24: Monte Carlo scheme applied to the situation illustrated in fig. 1. The annular element is populated randomly, and then the selection criteria for the measurement area is applied to leave only the points shown in red.

Using the definition of the measurement area, points lying outside of this region can be excluded. The area of the annular element contained within the measurement area is then calculated by:

$$Fractional area = \frac{Points remaining after selection}{Points initially used for seeding}$$

Equation 6.9: Calculation of fractional area within the RDF calculation

In the case of Fig. 6.24, we would therefore expect this number to be 0.25. Naturally, this heavily depends upon the level of seeding we apply, since random noise can produce some anomalous results in this analysis. Therefore, we would ideally like to sample with a very high level of seeding in order to mitigate for these noise effects. Typically, the data being analysed contains around 500 particles per average image, with experimental runs extending to 2000 individual image pairs. This results in a total of 10⁶ sets of calculations

Analysis techniques

being required to compute a complete Radial Distribution Function for each experimental run. It is imperative that the run time for the analysis is not overly long given the number of calculations required. Thus, a payoff must be sought between accuracy and computational expense for this analysis technique.

Appendix C analyses in detail the performance of Monte Carlo algorithms to estimate fractional areas, weighing up the additional computational expense of adding further samples versus the accuracy gained. For these investigations, sample numbers of 10⁴ per angular area were used, which typically gives errors of 1% per data point. A strategy was identified to use a subset of the available data with temporal spacing, thus reducing systematic error and giving samples across a wide time-frame. This approach also gave workable data sets, which could be analysed with sufficient accuracy within the available time.

6.5.2. Radial Distribution Functions of simulated particle distributions

The Radial Distribution Function effectively quantifies the ratio between local pair separation density, and pair separation density in the entire measurement volume. It would therefore be expected that, for all radial separations, in a randomly distributed set of particles this ratio should collapse to one, since there will be no preferable radial separation.

This was tested using the completely random, non-volume-excluded, sets of generated particles. Fig. 6.25 shows a comparison for image sets of 200 and 2200 particles per image generated. It can be seen that the fit for the larger population of particles is flatter, with a gradient which is two orders of magnitude smaller. Additionally, the line of best fit in this instance lies much closer to the $g(r_i) = 1$ line which we would expect. It should also be noted that, for the very smallest radii tested, discrepancies are more noticeable due to fewer data points lying within this range.



Figure 6.25: Comparison of Radial Distribution Function, g(r_i), at varying radial distances, R_i, for generated sets of random particle images.

These linear fits were then computed for all eleven particle per image values used, in order to assess the behaviour of the gradients and intercepts as N increases. This behaviour is shown in Fig. 6.26.



Figure 6.26: (top) Change in linear fitted gradient with particles per simulated image. (bottom) Change in linear fitted intercept with particles per simulated image

The equations used to fit the data shown in Fig. 6.26 are:

$$Gradient(N) = 4.2 \times 10^{-5} (e^{-2 \times 10^{-3}N}) + 1.3 \times 10^{-7} (e^{1 \times 10^{-3}N})$$

Equation 6.10: Equation for gradient of best fit line with particle variation

$$Intercept(N) = 9.96 \times 10^{-1} (e^{1.2 \times 10^{-6}N}) - 2.5 \times 10^{-2} (e^{4.5 \times 10^{-3}N})$$

Equation 6.11: Equation for y-intercept with particle variation

As particles per image, N, is increased, the gradients tends towards zero, and the intercept also tends towards 1. Therefore, as the number of data points is being increased, behaviour of the RDF code is tending to produce uniform RDFs of value 1. Therefore, the code can be considered validated by these test studies. Other types of particle distributions can also be studied using this code, to verify behavioural types. The same analysis was performed, incorporating a volume of exclusion in the particle behaviour. We can see that, for most separation radii, $g(r_i)$ collapses to 1. Indeed, the spread of the data remains consistent with that of the previous case. However, for the lowest radial values, the $g(r_i)$ collapses quickly to zero, since particles are unable to occupy the same space. This is a type of behaviour we should expect to see within the experiments conducted.



Figure 6.27: Comparison for three different simulated particle concentrations of Radial Distribution Function with radial separation, for an excluded volume set of simulated particles

This code can be further applied to the Gaussian profiled set of quasi random particles, with both non-excluded and excluded volumes. As can be seen from Fig. 6.28, neither of these particle distributions obey the standard $g(r_i) = 1$ relationship. For the non-volume excluded case, as N is increased, there is a convergence of g(0) to 1.2-1.25. These profiles monotonically drop off, until at distances of greater than 300 A.U., the value of $g(r_i)$ is actually lower than 1. The final value reached for values of 400 A.U. is within the range 0.92-0.94. For the volume excluded case, the same behaviour occurs as for truly randomly distributed particles. For the lowest radial distances, $g(r_i)$ drops to zero, since particles are unable to occupy the same space. Following on from this, the profile rises, albeit with a reduced peak value compared to the non-volume excluded case. However, the differences between these two distributions are minimal.



Figure 6.28: (top) Radial distribution functions for three particle concentrations with Gaussian x-dimension profiles. (bottom) corresponding data for inclusion of an excluded volume

This type of behaviour should be expected. By enforcing a Gaussian profile in one direction, we move more particles into the central region, at the expense of the outer region. Therefore, there will be a larger number of particles within small radial separations than expected. Correspondingly, in the x-direction, there will be a reduction in particles with larger radial separations. The points where the RDFs usually cross the $g(r_i) = 1$ line is around 300 A.U., which corresponds to two sigma values of the Gaussian distribution.

Analysis techniques

6.6.Conclusion

This chapter has presented in detail some of the analysis methods used in order to present data describing the distribution of particles in the experiments conducted. Initially, randomised sets of virtual particles were created in order to test known properties of these methods. Four distinct types of particle distribution were created, with variation between uniform and Gaussian distributions, and also inclusion of an excluded volume. Excluded volume cases enforced separation between particles, simulating experimental behaviour of particles with real-volume. For each type of distribution, the number density of particles is also altered to give a wide range of data for comparison with experiments.

Firstly, the box counting method for quantifying preferential accumulation is tested. For a random distribution of particles, the procedure should return values of σ - σ_P/λ which are close to zero. This was found at all number densities and for all box sizes used, with some deviation due to lack of statistical convergence.

Other particle distributions were tested using the box counting method. For uniform distributions with an excluded volume, values of D are downshifted in comparison to a non-volume excluded case. All calculated values are found to be negative, with an increase in magnitude as the volume fraction of virtual particles increases. The presence of an excluded volume impacts the order of particle placement, acting to create a distribution which is more ordered than random Poisson placement. The curves of D against box size are found to roughly collapse when dividing by the square root of the particle volume fraction. Gaussian distributions were also tested. There is some observable fluctuation as box size is changed, and values no longer hold close to zero due to the large-scale structure of the particle distribution.

Secondly, the verification statistics of the Voronoi analysis method were calculated using the various simulated particle distributions. PDFs of normalised cell size area were created for each particle distribution. Standard deviations for these PDFs were found to hold roughly constant for non-volume excluded distributions. For distributions with excluded volumes there is a linear decrease in standard deviation with particle volume fraction. Again, this indicates the imposition of a more ordered structure with particle addition under these conditions.

PDFs calculated from random distributions were found to be in close match to Gamma function relationships posited in literature. Those with excluded volumes typically are narrower, which is mostly accounted for by a lowering of probability in the smallest cell areas, which correspond to cluster cell areas. Since these most closely represent the PDFs found experimentally, these volume excluded PDFs are used, with appropriate matching of volume fraction. Following on from this, there is an illustration of how cells which have been labelled as clusters or voids may be combined to create larger-scale structures.

Another method of quantifying local clustering was formulated, which relies upon calculating local particle potential based on the nearest neighbours to the central particle. This method allows for probing of angular variations in mean separation distance, and acts only at the local level to determine behaviour.

Finally, the mechanisms for computing radial distribution functions of particle positions are considered. This algorithm has a high dependency upon the Monte Carlo routine, which has a high computational expense when used many times. Radial distribution functions were calculated for simulated particle distributions. For random particle distributions, g(r_i) is close to 1 at all radial separations r_i. Deviations in this trend are tracked with changes in particle density, showing the convergence with increase in data points used. RDFs are also calculated for other particle distributions. For volume-excluded uniform distributions, there is a decrease to 0 at the lowest radial separations. For Gaussian distributions particle sets, there is a preference to find particles at closer distances, which monotonically decreases to the largest radial separations measured.

The results outlined in this chapter are of great consequence for the experimental results shown from here on. In some cases, such as the Voronoi method, these statistics are necessary in order to carry out the method. In some cases, the validation provided by using known statistics provides confidence that results obtained are true results, and not an artefact of the algorithms used. Additionally, some of the results displayed using a controlled test in these instances may shed some light on the experimental results discussed later.

Chapter 7. Vertical channel results

This chapter will present experimental results taken from the vertical channel experiment, described in Chapter 5. This experiment investigated the injection of 500 µm hydrogel particles into a downward moving bulk flow of water. These injections were captured using a PIV experimental set-up; the dispersed phase particles were analysed using PTV algorithms described in Chapter 4. The principle focus is upon the concentration of particles being injected, although some results shown will focus upon the distance travelled from the particle injection point.

The majority of results within this chapter aim to describe the particle distribution within these types of flows. This is done by a variety of methods, including the Box counting method, Voronoi analysis, and Particle Potential method. Additionally, Particle Pair statistics such as Radial Distribution Functions and collision probabilities will be quantified, which can provide useful inputs towards CFD models where four-way coupling of particles/droplets is of importance.

7.1.Dispersed phase distribution

7.1.1. Box counting method

Figure 7.1 shows results of the deviation factor, $(\sigma - \sigma_P)/\lambda$, calculated for varying particulate phase fractions. These results are a representation of behaviour 40 channel heights downstream of the injection point, and were calculated for 50 different box sizes, varying up to 37 particle diameters (18.5mm) in size. For the very smallest box sizes, the value of D is lower than zero. This is a sign of a distribution with more imposed order than a random Poisson distribution, and is caused by the excluded volume of particles, as described in Chapter 6.

Subsequently, as box size increases, there is a marked increase in D. Therefore, for additional increases in box size, preferential accumulation is easier to observe. Typically, the increase of D is larger for a higher concentration of hydrogels; in other words, clustering is more exaggerated when the particle density is higher. This indicates that the presence of additional particles accelerates the coupling mechanisms which encourage clustering.



Figure 7.1: Deviations from Poisson histogram statistics of particle positions, as computed via the box counting method, at a distance of 40H from the injection point for a variety of particulate volume fractions

There is additional structure on these plots, most evidently for the higher particulate phase fractions, where fluctuations are typically on the order of D = 0.05-0.07. Examples of this were seen in Chapter 6, where a Gaussian distribution of particles was used to validate the box counting algorithm. Therefore, this may be explained by the Gaussian distribution which develops in the crosswise direction within the channel, as the particles begin the spread outwards from the injection point.

In addition to being able to quantify the levels of clustering for each experimental condition using the peak value of D, it is also possible to quantify the length scale of maximal clustering. These were found by calculating the peak co-ordinates for the plots shown in Fig. 7.1. This peak co-ordinate analysis was performed for varying particulate volume fraction at all three distances from the injection point which were recorded.

The plot on the left hand side of Fig 7.2 shows this peak clustering length scale, attained by extracting the x-coordinate of the peaks in Fig 7.1. For distances of OH and 40H, there is a strong linear relationship, indicating a decrease in characteristic clustering length scale when particle volume fraction increases. The same overall trend holds true for 80H, but with greater deviations from a linear trend. As expected, the peak clustering length scales are

larger with increasing distance from the injection point, as a consequence of the particles spreading out in the crosswise direction as they progress down the vertical channel.



Figure 7.2: (left) Relationship between length scale of peak clustering, taken from the x co-ordinate of the data in Figure 7.1, and particle volume fraction, for varying distance from the injection point (right) Relationship between peak clustering value observed, obtained from the y co-ordinate of Figure 7.1, and particulate volume fraction, for varying distance from the injection point

This behaviour therefore suggests that, as more particles are added to a flow under otherwise identical conditions, the structure of the clustering changes. With more particles in the flow, the most common size of the clusters observed is smaller. This behaviour is particularly strong at the injection point in this case; from the left hand plot, it can be observed that the gradient of change is much larger than for more developed flows.

The addition of particles may act to exclude other particles in the flow from cluster structures. This will result in a distribution of cluster sizes whereby there are a greater proportion of smaller scale clusters.

The right hand plot of Fig 7.2 shows the peak clustering values which can be expected for these flow conditions. The clustering levels for the well-developed flow regions are very similar, with peak values between 0.25-0.6. As shown in Fig. 7.1, clustering levels tend to

increase with addition of extra particles; therefore, it seems reasonable to conclude that addition of more particles encourages the mechanisms which cause clustering to occur.

In the case of particles which have just been injected, the behaviour is very much different, with a sharp decrease in levels of clustering. This behaviour is not strictly linear; as can be seen in Fig. 7.2, there is significant deviation from a linear trend. For low particulate concentrations (2-4%), deviations from random Poisson statistics are very high, indicating a great deal of disorder. Following on from this, there is a sharp decrease, until at higher particulate concentrations (8-9%), values of D are more comparable with those in well-developed flow regions.

Given that the time the particles have been in the flow at this point is negligible, it is unlikely that this behaviour would be caused by sampling of continuous phase motions. It is more likely that this behaviour is determined by the injection conditions. Particles within the injection system are confined to a tight space, making collisions with both each other and the wall more likely. As a result, the velocity distribution of these particles has a large width upon entry. This further enhances the probability of inter-particle interactions, which may explain the large degree of disorder seen. However, for higher concentrations, particles are more likely to be tightly packed into the injection system. This may act to regulate the velocity distributions, and introduce more order to the particle distribution as they enter the channel. These results may provide a useful input for CFD models, in which the entry conditions of particles into a flow must be specified.

Some example images used to quantify particle dispersions are shown in Fig. 7.3. These were taken at OH, just behind the injection point, and at 80H, in the most well-developed region of the flow that was measured.







Figure 7.3: (Left) Still image of particle dispersion at OH, where H is channel height (right) Still image of particle dispersion at 80H

7.1.2. Voronoi analysis method

Initially, the Voronoi analysis method was used simply to quantify the degree of clustering within the particle distributions observed. This was achieved by calculating the standard deviation of the experimental cell size distribution, σ_{EXP} and subtracting the standard deviation of the simulated particle distribution, σ_{REF} , the derivation of which is explained in Chapter 6. Typically, an excluded volume reference case of particles provided a better fit to experimental results. These were therefore used, with statistics calculated for different mean concentrations to account for changes due to excluded volume effects.

Results are shown in Fig. 7.4, for varying phase fraction at different distances downstream from the injection point. These results are therefore the analogue of the right hand plot of Fig. 7.2 for this method, and there are similarities and differences to observe.

Firstly it should be noted that, for every experimental condition, there is a positive value of σ_{EXP} - σ_{REF} . This indicates that there is clustering of particles under all conditions, since experimental cell size distributions are broader than simulated particle counterpart distributions and therefore exhibit greater probabilities of cluster and void particles.

Furthermore, the trend in the well-developed flow regions of 40 and 80 channel heights is an upwards linear one, as with the box counting method. This corroborates the idea that, under these conditions, additional particles do lead to higher levels of preferential accumulation. Finally, results seem to hold true which indicate a stronger preference for clustering nearer to the injection point.

There are some differences between the two figures. Firstly, from the Voronoi analysis method, it can be suggested that levels of clustering are actually higher for D = 40H than for D = 80H. This is not the same for the box counting method, where values of the deviation factor suggest roughly comparable clustering levels at these two positions. Additionally, the gradients differ markedly in the Voronoi method. The Voronoi method suggests that the level of clustering is more heavily influenced by addition of particles as the point of observation moves towards the injection point.





Finally, the trend of clustering behaviour at the injection point is the most starkly contrasted difference between the two analysis methods. For the Voronoi method, there is a general upwards trend. This is complicated by two local turning points in the line of best fit (a third-

order polynomial provided the best fit). However, there is some evidence of the first of these two-turning points in the box counting method. Depending upon interpretation, these methods could be seen to give a similar trend up until hydrogel concentrations of 6%, at which point the Voronoi results have a second turning point, and box counting results continue to show a downwards trend.

The data at the injection point (OH) provides information about the entrance conditions of particles into the vertical channel. This profile shape may therefore be able to identify a particular set of injection conditions, and the clustering that results just after injection, versus another set of entrance conditions which induce a different effect. Furthermore, information about injection conditions is useful for validation of CFD codes, where boundary conditions are often of utmost importance.

These two methods therefore produce some differing results. The box counting method relies upon assignment of particles to a physical space, whereas the Voronoi analysis method relies upon assignment of physical space to each particle. Therefore, comparisons between the two methods must be made with care.

Nevertheless, in this instance, some conclusions can be drawn from both sets of data. Preferential accumulation is inherent to this system, even when observing in a fullydeveloped flow region. In general, the levels of clustering observed increase nearer to the injection point, and also increase with addition of particles when considering the welldeveloped flow region.

The following section will consider an extension of the Voronoi technique, which allows assignment of cluster and void regions to the physical space under consideration. This is the main advantage over the box counting method, which must really be considered a "wholefield" technique.

7.1.3. Further Voronoi analysis

Once an experimental Voronoi cell size distribution has been ascertained, individual particles can be assigned as voids, neutral or cluster regions, as described in Chapter 6. From this point onwards, it is possible to join self-similar regions which share edges in order to determine larger scale structure of clusters. The geometrical properties of these cluster

structures was analysed by using the data at 80H from the injection point, and is shown in Fig. 7.5.



Figure 7.5: (Left) Plot of structure perimeter, P_c, against the square root of the structure area, A_c^{1/2}, for an ensemble of experimental results. (Right) From this, two differing gradients can be extracted from two regions of the plot, which are plotted against particle concentration. Also shown are examples of structures from four plot regions

The relationship between cluster structure perimeter and cluster structure area was determined, in order to find information about the geometrical properties of these regions. Plots of these are shown on the left of Fig. 7.5, where a variety of concentrations were used.

For simple shapes, there would appear to be a gradient of 1 between P_c and $A_c^{1/2}$ (for example, when considering the scaling of a circle or square). Therefore, any gradient over this value is indicative of more complicated cluster shapes, typically analogous to fractal structures. The relationship between perimeter and area will scale as follows:

$$P_C = \left(A_C^{1/2}\right)^m + c$$

Equation 7.1: Relationship between cluster structure perimeter and area

Therefore, in the graph on the left of Fig. 7.5, m would be the calculated gradient of the data points, since a logarithmic scale has been chosen to represent the large range of scales present.

As can be seen from the plot of P_c against $A_c^{1/2}$, there are two distinct regions, between which the plot gradient changes. These regions are delineated by the black dashed line. This change in gradient indicates that there is a difference in scaling properties between smaller cluster structures on the left hand side, and larger ones on the right hand side.

Each of these two regions may also be sub-divided into two, depending upon whether the data point is found at the top or bottom of the spread. In the left hand side of the plot, cluster structures found on the upper side of the data points, denoted as A, are typically representative of small isolated structures, containing two or three cluster cells. Those structures found on the lower side of the data points, denoted as C, are characteristic of isolated single cluster cells. The differing behaviour of these types of cluster structures can be seen from how the data points extrude from the bottom of the plot in Fig. 7.5.

It should be noted that, in general, it is not correct to call the particle at the centre of an isolated cluster cell a cluster structure. Instead, the surrounding particles should be considered as contributing to the structure, since these are the particles which enforce the small size of the central cluster cell. These are therefore small isolated clusters, which do not extend beyond the nearest neighbours of the central particle. However, this method does identify the centre of the clusters and the approximate size; if the neighbouring cells were included, this could include large areas not associated with the cluster structure. Thus, it is better to slightly underestimate, than to vastly overestimate the cluster structure size.

Upon inspection, there appears to be a good collapse of data on the left hand side of the plot with differences in particle concentration. However, there is a small linear increase in the gradient m_1 as particulate concentration is increased. This is mostly attributable to the cluster structures in section C, and how they change with concentration.

For the lowest particle concentrations, the band of data points in 'Region C' is wide, covering a large range of structure areas. As particle concentration is increased, there is a suppression of the range of areas over which these types of cluster structures occur. For the highest particle concentrations, only the smallest structure areas are found as a 'Region C' cluster structure.
These changes therefore account for the small change in gradient m_1 of the left hand plot with changes in concentration. As concentration increases and 'Region C' is suppressed, the weighting of the data curve moves towards the lower left, thus altering the gradient calculated for these data points.

This is indicative of greater connectivity occurring in these experiments as concentration increases, since there are much fewer isolated cluster structures. Therefore, these types of structures become more like those found in region A as more particles are added.

The right hand side of the in Fig. 7.5 can again be subdivided into two different types of cluster structure, depending upon the ratio between perimeter and area. Those structures in Region B, along the top of the data curve, have a 'fractal' nature. These are typically complicated structures with many inlets and thin connections between similar regions. In contrast, those structures in Region D are generally rounded in shape, as seen in the relevant example in Fig. 7.5. In these instances, cluster particles do not have many non cluster particle neighbours; only at the edges of the structure is there any differing behaviour.

In this region, there is a larger change in gradient with respect to change in concentration than is seen in the left hand plot region. The values placed upon this gradient can be thought of as a measure of complexity; as the value of m increases, we observe clustering features that have a more fractal structure. An example of this can be seen in Fig. 7.5, showing the various strands that form together to create a larger cluster structure. This is also indicative of greater mixing, since there is a greater surface area in common between the differing void and clustering areas of the flow field.

7.1.4. Local cluster method

In this section, particle distributions from the vertical channel rig are described using the local cluster method, as described in Section 6.4. This method uses the distances to local particle neighbours in different angular sections, and compares to a mean particle potential, calculated assuming an even distribution of particles with the same concentration.

Fig. 7.6 shows the particle potentials for sets of particles at varying concentrations, with the distribution normalised by the mean particle potential. The formula for calculating this

factor is described in Section 6.4. Naturally, as particle concentration increases, the average separation decreases, resulting in larger potential values on average, so this normalisation is performed to compare the distributions at different concentrations.



Figure 7.6: PDFs of individual normalised particle potential, using a 12-neighbour formulation, for varying particulate concentrations, using experimental data from the vertical channel rig, at Re = 20800

Typically it appears that the lower tail of the distributions, corresponding to non-cluster particles, is similar for all particulate concentrations. This is probably most reflective of similar behaviour at the edges of the particle jet in all cases, where these conditions are most likely to be met.

On the higher side of these distributions, there is relatively little distinction between particle concentrations of around 1-3%; these PDF values mostly peak at a value of 1.2-1.4, and have a similar width. The most distinct shapes occur for higher concentration distributions, with noticeably narrower shapes and higher peaks. This is indicative of some ordering of local cluster behaviour; therefore, it is more likely to see replicated cluster structures throughout the particle jet in these situations.

Using this local cluster formulation with 12 nearest neighbours, it is possible to analyse the typical shape that these clusters exhibit. For example, within each angular section, it is possible to characterise the mean and standard deviations of the particle-particle separation, under varying conditions.

The angular separation, θ , is defined as the angle formed from the positive x-axis, as experiments were recorded. In the vertical channel rig, this results in the angle being defined as shown in Fig 7.7.



Figure 7.7: Illustration of angular separation in the vertical channel experiment. Angles within A are 0° to 90°, within B are 90° to 180°, within C are -90° to 0° and within D are -180° to 90°



Figure 7.8: Comparison of variation in nearest neighbour at different angular separations, at 9% particulate concentration. These plots are shown for the full distribution, saturated particles (those with neighbours in each angular region, and those which are classified as clusters

Fig. 7.8 shows the variation in local particle separation over the full set of particles, looking at each angular section individually. Initially, this calculation was performed for a full set of particles, including those which did not have a full complement of neighbours. Therefore, in these cases the sample size varied with the angle being considered due to 'voids' in the data.

For calculations of this type, the standard deviations computed are relatively large. They are also highly variable depending upon the angle of separation being considered. For instance, at values of around 0°, the variation is very high. This is likely due to particles on the edge of the injected mass, which have no immediate neighbour in this direction, giving uneven results. Additionally, unpredictable behaviour at the edges of the jet may result in more randomised particle distributions.

Further sets of particles considered were those in which all angular sectors contained a nearest neighbour of some sort, called saturated particles, and those which were

considered cluster particles using the particle potential method. In this case, a simple threshold was assigned, relating to the mean expectation potential; those particles with a larger local potential were denoted as belonging to clusters.

Both of these more selective distributions resulted in flatter profiles of variability. This is not surprising, since behaviour considered only within the bulk will be much more homogeneous. Additionally, the separation variation was lower, due to the lack of anomalous behaviour from edge particles.

There is however, still some variation with respect to angular displacement. Typically, as seen with the first profile, there is a rise in variability to 0° , with lower variability around $\pm 90^{\circ}$. Since there are no edge effects to consider anymore within these distributions, there must be other effects which are influencing the variability at different angles, such as wake dynamics which draw some particles closer and push others away.

As a result of the high variability in the full particle set, the following results will consider only those particles in the bulk (i.e. saturated condition) and those which are denoted as being cluster particles, in order to look at differences between them.



Figure 7.9: Mean nearest neighbour particle separation as a function of angular separation, for varying particulate concentration, for cluster particles only



Figure 7.10: Mean nearest neighbour particle separation as a function of angular separation, for varying particulate concentration, for saturated particles only

Figures 7.9 and 7.10 show the mean nearest neighbour particle separation changes with angle, at various bulk concentrations for both the cluster and saturated conditions. Mean separations are typically lower for the cluster particle cases at all concentrations, due to the filter being imposed on the data - if all particles chosen are designated as belonging to clusters, then data will typically show very closely packed particles, whereas saturated particles require less stringent selection conditions which could result in more sparsely packed regions of the flow being represented.

It can be seen that, for both sets of particles, the separation profile flattens as a result of increasing concentration. This means that local cluster structure is more isotropic under highly concentrated conditions. For lower concentration, it is typical to observe lower mean separations at angles of ±90°, which suggests that particles are more likely to be dragged in the streamline of the particle into closer proximity i.e. drafting. As further packing of the channel occurs, this effect is less noticeable due to particles being forced to occupy closer spaces, in particular perpendicular to the streamline being followed.



Figure 7.11: Mean separation variability for varying angular separation and particulate concentration, for cluster particles



Figure 7.12: Nearest neighbour separation variability for varying angular separation and particulate concentration, for saturated particles

Variability of nearest neighbour separation is shown in Figures 7.11 and 7.12 for both cluster and saturated particles. Variability becomes more consistent with angular separation as

concentration increases, due to the more isotropic nature of the clusters being formed. There is lower variability at low concentrations for angles of $\pm 90^{\circ}$, which is along the particle streamline. As noted earlier, along this direction the mean separation is typically lower, due to drafting of particles. As concentration increases, the variability and mean parameters for nearest neighbour separation even out with angle.

For higher concentrations, there is symmetry in the variability plots. This shows similar behaviour upstream and downstream of a particle, and also between each crosswise direction of the particle. This symmetry is somewhat broken for lower concentrations. This inconsistent variability with respect to direction is most apparent in the crosswise direction, which suggests that small numbers of edge particles are having an influence over the statistics.

For saturated particle sets, there is no clear reduction in variability until higher particle concentration (3-5%), and variability is much stronger with angular separation. This is due to the effects of particles in less sparsely populated regions, with exhibit different behaviours which affect the overall statistics.

7.1.4.1. Non-dimensionalising

As shown in Figures 7.9 and 7.10, mean separation between local neighbours reduces with increasing particle concentration, as would be anticipated. However, it was of interest to determine whether these changes fell in line with expectation. Since the majority of particles lie along the PIV measurement plane in these experiments, an appropriate scaling to non-dimensionalise was $\alpha^{1/2}$, where α represents bulk particle concentration.

The results of applying this scaling are shown in Figs. 7.13 and 7.14 for both saturated particles and cluster particles. The variability of the values plotted compared to the means are 7.5% and 13.3% respectively. These results show that the attempts to non-dimensionalise do not entirely collapse the curves of mean particle separation; there are therefore other mechanisms which determine this value than just the bulk concentration.



Figure 7.13: Mean separation profiles for saturated particles, <r>, multiplied by the square root of the bulk particle concentration, $\alpha^{0.5}$



Figure 7.14: Mean separation profiles for cluster particles, <r>, multiplied by the square root of the bulk particle concentration, $\alpha^{0.5}$

It is worth considering the errors which might be present in these plots to determine how good the collapse of data is. The random error in assessing the distance to neighbouring particles was typically lower than 1/8D_P, and so even for the highest particle concentrations would typically be 10%. Combined with the number of samples for each angular region taken to produce these samples, the error in the mean would therefore be 0.1%.

The total error is therefore likely dominated by the error in measuring particle concentration. This is due to large fluctuations, typically on the order of 10%. This would result in a total error which is also 10%. As can be seen from Figs 7.13 and Figs 7.14, this means that the data for saturated particles collapses within the margins of error, whereas the data for cluster particles does not.

7.1.4.2. Connectivity

Using full sets of cluster data, it is possible to assess self-similarity of neighbouring particles within the flow being observed. For example, all central cluster particles can be assessed to determine how many cluster particles also comprise the local cluster of the nearest 12 neighbours. The probability of N neighbours is shown in the Fig 7.15, with changes to particulate concentration.



Figure 7.15: Plots of probability for a central cluster particle to have N local cluster neighbours, for varying particulate concentration. Fits achieved using a smoothing spline to indicate general trend.

Each of these plots shows a similar shape, with the exception of the data for 3.3%. In each of the other cases, it is particularly rare for a cluster particle to neighbour with precisely zero cluster particles. There is then typically a rise to an intermediate number of particles, the peak of which varies according to concentration. There is then typically a return to zero probability for 12 cluster neighbours i.e. a central particle completely surrounded by other cluster particles.

Such a probability function should be evident given how dense particles behave in turbulent flows. Typically, particles in these types of flows cluster along high strain regions, which will result in 'strings' of particles forming along these narrow regions. Therefore, it is most likely that particles will experience local crowding in two specific spatial directions, and sparsely populated regions in the other directions. This would therefore not lead to 12 local cluster neighbours, but more likely an intermediate number dependent upon the precise geometry of the strain region.

The anomalous result in this case is for particle concentrations of 3.3%, where there is a steady rise in probability to a full complement of 12 local cluster neighbours. This would suggest cluster behaviour focused around compact circles of cluster particles, whereby the majority of cluster particles are buried within the local structure. It seems more likely that this is an anomalous result, possibly caused by an unusual particle distribution (or perhaps some temporary clumping of particles), since understanding of typical particle behaviour would determine this to be unlikely.

For those particle concentrations which display the typical profile, there is a shift in the peak neighbour number. This indicates that, for higher particulate concentrations, there is a higher probability of more cluster neighbours. In this case, this means that the cluster structures, and therefore the strain regions, are disrupted by addition of further particles, resulting in a broader spread of the structure.



Figure 7.16: Percentage of cluster particles with 12 local neighbours with changes to bulk particulate concentration

Fig. 7.16 shows the probability of central cluster particles having 12 nearest neighbour particles of any type. This can therefore be interpreted as the probability of a cluster particle being contained within the bulk flow, and experiencing a particle interaction in any angular direction. Clearly, the probabilities of this are usually quite high, with values of >88% being observed at all concentrations, since clusters are most likely to form with the centre of the particle jet. There is a small drop in probability initially, followed by a rise with concentration. This indicates that, for the highest concentrations achieved, cluster particles are increasingly likely to be found in the bulk flow, as opposed to being on the edges of the particle jet.

163



Figure 7.17: Percentage of cluster particles with N local neighbours, for varying bulk particle concentration and values of N.

These probabilities can also be calculated for lower numbers of local neighbours, representing cases where cluster particles are found at the outer extremities of the particle plume. As expected, probabilities drop with decreasing values of N, reflecting the decreasing likelihood of finding a cluster particle near to the edges of a plume. Typically, for values of N which are 7 or lower, the probabilities calculated are vanishingly small, and therefore not considered here. This is unsurprising since, even for a particle which is on the very edge of the jet, it should be possible to find at least 6 neighbours, since a full 180° angle will still encapsulate those particles in the bulk.

There is a small probability rise to intermediate concentrations, followed by a drop. This is observed uniformly for all values of N. Therefore, the probability loss shown for full saturated cluster particles is split fairly evenly amongst other cases of N local neighbours.

These particles are most likely clustering in small groups at the edges of the plume. Therefore, whilst there are vacancies in particular angular regions, some particle neighbours are in such close proximity that the overall particle potential is still greater than that of an ordered structure. These statistics may therefore give some insight into clustering

probabilities towards the edges of the jet, and the likelihood of cluster formation given separation for the jet flow.

7.2.Particle pair statistics

Radial Distribution Functions were calculated for various flow conditions. Generally, the concentration and particle Reynolds numbers were found to have little effect upon the peak value at short radial separations. The main difference, as expected for a spreading particle jet, is found for changes in distances downstream of the injection point. These are shown in Fig. 7.18; the peaks for well-developed regions only correspond to about double the probability of random placement. For particles leaving the injection point, this probability is more than tripled.



Figure 7.18: RDFs for a range of separation distances, for three distances downstream of the injection point. Fits achieved using smoothing spline fit, show a peak at low radial separation, with collapse to zero for infinitesimal radial separation

Using the relative velocity and position vectors between particle pairs, it is possible to define a closing time, in which particles would meet if continuing along the same trajectory. A PDF of this quantity, t*, is shown in Fig. 7.19. t* is defined as the distance between the

two particles in a pair, divided by the relative velocity which is superimposed onto this vector.

All negative values of t* indicate particle pairs which have a 'closing' relative velocity, and therefore some chance of colliding. These are shown in the left hand side of Fig. 7.19. Those with positive values of t* indicate 'opening' relative velocities i.e. those particle pairs which are separating. These are shown in the right hand side of Fig. 7.19.

Due to the chaotic nature of these flows, only those within the shortest timescales could realistically be considered candidates of collisions. To say that all particle pairs represented as 'closing pairs' will collide would be incorrect; if this collision time t* is large, then other events will likely cause one or both of the particles to deviate from their original course.



Figure 7.19: Sample PDF of separation timescales, t*, with closing and opening particle pairs labelled. Closing particles are defined as those which have the potential to collide, whereas opening particles are those which are considering to be moving away from each other.

By integrating under the curve in Fig. 7.19 for different values of t*, it is possible to determine the collision probabilities for the full particle set. These results are shown in Fig.

7.20 for varying particle concentrations, for only the most relevant timescales. In the timescale of 0.01s, probabilities of collisions range between 0.5-3% in this concentration range. There is also a good collapse of data when normalising by concentration, indicating that collision probability can be well-predicted by this factor.





7.3.Conclusions

Data from the vertical channel experiments has been analysed, with a particular focus on particulate concentration effects. This included analysing whole-field clustering and cluster structure, using a variety of different approaches.

Application of the box counting method showed that clustering of particles occurred for all experimental conditions. This was observed in the standard relationship between deviation from Poisson statistics and box size, in which a peak value suggests a preferential clustering size and strength.

Data from different distances from the injection point was characterised. At all distances from the injection point, there exists a negative linear correlation between the

characteristic size of peak clustering, and the particulate concentration. As particle concentration increases, clusters form into tighter groups. The strength of this peak clustering was characterised as growing with particulate concentration in most welldeveloped conditions, but decreasing with particulate concentration just downstream from the injection point. This difference in behaviour is likely caused by the randomisation of particle velocity vectors caused by the injection mechanism. As particles enter the free flow, they have just been released from a series of physical interactions with other particles and the wall of the injection pipe. Therefore, the cause of this correlation is likely to rely more upon the injection mechanism than the underlying fluid mechanics.

The Voronoi analysis technique was initially applied as a competing method for assessing the global clustering parameters. There were some common features between both sets of results; in both cases, the levels of clustering increasing with particle concentration in the more well-developed regions. However, the key difference was seen in results from just downstream of the injection point, where clustering observes a generally upwards trend with added particle concentration. This effect was not linear as with the other distances downstream, and it is therefore still distinct, but there is no agreement between the two methods used on this set of results.

The Voronoi technique was further used to analyse the structure of clusters in a more local manner. Using the thresholds defined by intersecting probability density functions, it was possible to join similar and neighbouring regions in the flow. These were characterised by using their geometrical properties, which resulted in the identification of four different types of cluster shape. It was then shown that, whilst the relationship remains roughly similar across particle concentrations, there are differences in local cluster structure, which can be broken down into regions of large and small clusters.

With regards to particle clustering, the final technique employed was a nearest-neighbour based formulation which assigned potentials to each particle based on separation distance. The probability distribution functions of potential suggest that behaviour at the jet edges is very similar regardless of concentration. Behaviour differs most at high local particle potentials, which will mostly occur in the bulk of the flow. For higher concentrations, it was

noticeable that local particle potential appeared more uniform; therefore, the immediate cluster structures may be very similar across the flow domain in these instances.

Profiles of mean separation with angle were also calculated. These indicate a preference towards angles which suggest drafting of particles within wakes. There was also a reasonable collapse of data with bulk particle concentration, more so for the general bulk population than for just cluster particles. Therefore, in the bulk the typical length separation determines most of these factors; in the centre of particle clusters, more complicated length scales such as those dictated by the underlying fluid may come into play.

This formulation could also be used to determine the connectivity of cluster particles with other particle types. Generally, it is found that there is a peak at the intermediate number of cluster neighbours; this is indicative of our typical understanding of heavier-than-fluid clustering, whereby particles form 'strands' in the high strain regions. As expected, the majority of cluster particles were found in the bulk with a full complement of nearest neighbours; with increasing concentration, this becomes increasingly likely, so most clusters are drawn towards the bulk.

Radial distribution functions were found to depend mainly upon distance from the injection point, and no other factors. Typical enhanced probabilities ranged between 1.5-3 times that of a random placement. Closing times were found between particle pairs, which indicated immediate probabilities of collisions to be 0.5-3%, a value which was well-predicted by overall concentration.

The following chapter will describe the design, building and initial testing of a subsequent experimental rig, known as the Circulating Dispersions Rig. This rig was designed to test similar phenomena to those found in this chapter, and so the results will be described in Chapter 9.

However, there were some key differences to this rig. Firstly, a fully three-dimensional flow was enabled. Secondly, the Reynolds numbers were chosen to be higher, with streamwise velocities up to three times larger than those obtained in the vertical channel rig. As a consequence of the increased optical path through mixed states, there was also a need to employ refractive index matching and consider lower concentrations. Chapter 8 will begin

by describing attempts to refractive index match the hydrogels used in these experiments, and a judgement of their suitability for these types of experiments, before going on the address other key design criteria for the Circulating Dispersions Rig.

Chapter 8. Circulating dispersions rig - design

As discussed in Chapter 1, collaborative partners within the TMF consortium planned to test new codes using a simple cylindrical pipe geometry. In order to provide experimental data which can be used to validate these simulations, an experimental facility was designed from scratch, using the various criteria required to match up flow regimes.

The CFD codes being developed are designed to deal with "liquid-liquid" dispersions. Since these generally occur when the carrier phase superficial velocity is large, the experimental rig needed a capability of high Reynolds number in order to make comparisons meaningful. Additionally, the dispersed phase needed to occupy a length scale range which would allow it to interact with the turbulent eddies in the continuous phase. This generally limited allowed particle diameters to 100-500 μ m.

In this chapter, the hydrogel particles used as the dispersed phase are reviewed to determine their suitability for an experiment requiring a high optical visibility. Measurements were made, as part of an associated Masters' thesis, to determine the level of scattering and opacity one could expect from an even dispersion of these particles. Following on from this is a review of literature of other dispersed phase systems, concluding with the selection made for use in these experiments.

The second part of this chapter will discuss the design and building of the "Circulating Dispersions Rig", with consideration of the brief outlined above. This will include an overview of the equipment used, the operation of the facility, and the collection of data, including any calibration procedures which were required. Finally, characterisation of the rig using an aqueous single phase will be presented.

8.1.Optical transparency of hydrogel-water mixtures

To provide supporting data for this project, a Masters' project was devised to characterise the optical transparency of mixtures which could potentially be used in the "Circulating Dispersions Rig". Therefore, all work contained within Section 8.1 is credited to Ramsey Dao as the principal investigator (Dao 2015), whilst the author provided a supervisory role in the experimental work. The two principal systems investigated in this work were water and undyed hydrogel particles, and sodium iodide solution with PMMA particles. The latter system was chosen as a comparison due to its frequent use in such applications (Wiederseiner, Andreini et al. 2010). For both of these systems, suspensions of varying concentrations were agitated inside a cylindrical test chamber using a magnetic stirrer. In addition, this stirrer also doubled up as a hotplate, allowing for temperature control of the contents. This temperature was measured using a T-type thermocouple, chosen due to its suitability for low-temperature applications.

Transparency of each solution was measured by using a "laser transmittance" method. The laser used in these investigations has a very similar wavelength (532nm) to those of the PIV laser used in the main investigation. Therefore, the refractive index behaviour is expected to be very similar. The original beam was passed through a beam-splitter, with one arm of the split beam sent through the air to a photodiode. This signal measured the baseline intensity of the laser, known as I₀. The second arm of the split beam was directed through the suspension, and then on to a second photodiode. The attenuation of the light signal through the suspension, I, could then be measured relative to the first signal. The transmittance is then given by:

$$T_{\lambda} = \frac{I}{I_0}$$

Equation 8.1: Definition of light transmission

In these experiments, PMMA matching with NaI solution represents the comparison case, since these particles have been refractive index matched much more extensively in literature, and optical behaviour of these suspensions is therefore well-characterised. Fig. 8.1 represents a scan of the agitated mixture over a temperature range of 16-40 degrees Celsius. The results show that there exists an optimal matching temperature of 21.5 degrees Celsius. At this temperature, for suspensions of 5% particle concentration, there was a transmission of 0.4. This is in contrast to a transmission of 0.625 with a 0% particle concentration (i.e. pure NaI solution), and represents at 36% loss in laser intensity.



Figure 8.1: The temperature and transmittance over time of a 5% PMMA and Nal suspension at the optical path length of 49mm

Having found the optimal matching temperature of this suspension, it was then possible to investigate the effects of optical path length. Fig. 8.2 shows the peak transmission values (i.e. taken at 21.5 degrees Celsius) with changes in optical path length, L, and for differing suspension concentrations. L is defined as the total path length light must traverse to pass through the suspension. As expected, for changes in optical path length, the transmission roughly obeys the Beer-Lambert law:

$$T_{\lambda} = e^{-\alpha L}$$

Equation 8.2: Beer-Lambert law

Where α is the attenuation coefficient and L is the optical path length through the solution.

As concentration increases, the number of interfaces the laser beam must cross increases, thus reducing the transmission of the suspension. At the lowest optical path length of 30mm, transmission is just above 50%. For 15% concentration, the transmission drops to around 30%; for 25% concentrations, there is a further drop to around 15% transmission.



Figure 8.2: Plots of peak transmittance values for NaI/PMMA suspensions, with variations in optical path length and the particle concentration

Using these results as a baseline, comparisons could then be made to measurements using water and hydrogel suspensions. Fig. 8.3 shows the corresponding scan with increasing temperature for a 5% concentration of suspension at a path length of 49 mm. In contrast to the previous test, temperature was only ramped up to 32.5 degrees Celsius, to avoid the hydrogel protein structure denaturing.



Figure 8.3: The temperature and transmittance over time of a 5% Hydrogel and water suspension at the optical path length of 49mm

There is no observable preference for temperature in transmission. This suggests that the dominant factor for opacity in these suspensions is not refractive index mismatching, but rather some other mechanism. Furthermore, the transmission is very low in comparison to the PMMA/NaI suspension tests, averaging only 9%. This is in contrast to the baseline transmission, conducted using only water, in which transmission was measured at 82%. Therefore, the reduction is much larger compared to the control test when adding particles.

Further experiments with higher concentrations found that the transmission was not measureable above the noise; therefore, suspensions of hydrogels and water higher than 5% must be considered opaque.

Optical path length was varied for 5% suspensions, with transmission values dropping from 15% to 8% with an increase in path length from 30mm to 55mm.

It should be noted that there was evidence that scattering was particularly large in some instances. For example, some photographs show that light was scattering around the perimeter of the photodiode. It might have been possible to refocus this light using another lens; however, the deviation of the light is still of importance when making PIV measurements. Thus, any deviation was taken into account.

A further point of note from these results was the importance of refractive index matching the experiment to the surroundings. The baseline results of "zero concentration suspensions" showed very different transmission percentages, dropping from 82% to 62.5% from water to Nal. This illustrates that the importance of additionally refractive index matching the materials containing the experiment, in order to reduce signal losses at the experimental boundaries.

In order to provide an explanation for the differences in behaviour between these two suspensions, microscopic images were taken of samples of both types of particles. Fig. 8.4 shows a sample of PMMA particles, with a magnification of 100. The image shows that the particles have a very regular structure, albeit with a few surface impurities. This explains why, despite the refractive index match between the two phases, there is a drop in transmission, since scattering will occur at these impurities. These scattering events will become compounded as the number of interfaces to cross increases, thereby reducing the laser intensity at the receiving photodiode.



Figure 8.4: 100x magnification image of PMMA particles used in the refractive index matching experiments Fig. 8.5 shows a microscope image of the hydrogel particles acquired from GeniaLab, using a magnification of 20. The image includes the solution used to keep the hydrogels hydrated, and therefore an aqueous film is visible throughout the image. It can be noted that, due to the flexible shapes of the particles, their shape is less regular than those of the PMMA particles. This may be an influencing factor for the lower laser transmission, since refraction events at irregular boundaries may be more influential for slight refractive index mismatches. Additionally, there is a clear secondary structure to these particles, in the form of an outer shell surrounding the core of each particle. Due to the proprietary nature of the manufacture of these particles, it is not clear why this additional structure is present, although it may be due to the option to include a fluorescent dye into only the central part of the particle, as was done in the vertical channel experiments.

The more complex structure of these hydrogel particles may explain the lower transmission of water-hydrogel suspensions. If there is a refractive index mismatch between the inner and outer regions of these particles, for instance due to differences in composition, then traversal of light through these particles will require crossing of four boundaries instead of two. Combined with the less regular shape of these particles, this may then result in many deviations of the beam due to refraction events. As a result of this, the beam transmission as seen by the receiving photodiode will be significantly reduced.



Figure 8.5: 20x magnification of agarose hydrogel particles used in the refractive index matching experiments

As a result of these experiments, other suppliers and manufacturing methods for hydrogel particles were considered. In general, most existing hydrogel production uses large-scale moulds to create structures. An application for this is the creation of flow boundaries to

investigate complex water flows using non-intrusive optical techniques. (Weitzman, Samuel et al. 2014). However, there are few suppliers of small scale hydrogel particles. This is most probably a result of the long production time for each particle, which puts production timescales to around 2-3 months for any reasonable quantity i.e. larger than 1kg. The same problems were found when attempting to source similar particles from academic institutions with the relevant facilities.

Within the existing literature, there are not many examples of hydrogel particles being used in multiphase flow experiments. In their paper, Byron and Variano produced two types of hydrogel particle, using both agarose and polyacrylamide (Byron and Variano 2012). These particles were typically 20mm in diameter, and were doped with seeding particle in the manufacturing process. These were then suspended in homogeneous and isotropic turbulence, with the macroparticle tethered to keep it in the image area.

Using illumination of a 532nm laser and commercial PIV software, these particles and the surrounding fluid were tracked, with velocity maps produced and particle rotation quantified. Whilst this shows potential for the use of hydrogels in analysis of two-phase flows, and the injection moulding recipes are included in the paper, a number of limitations are noted.

Firstly it should be noted that agarose particles, whilst safer and cheaper to mass produce, exhibit a poor performance due to increased scattering of light. The hydrogels procured for experiments in this thesis were produced from agarose. This may explain the poor performance of these mass-produced particles.

Furthermore, Byron and Variano note that their manufacturing process becomes challenging when attempting to mould features of around 1mm in size. As seen in the macroscopic images above, the manufacture of particles of this size may result in less regularity in the particle population, which will result in lower optical transmission through a suspension.

Finally, the paper notes that the manufacturing process used would be inappropriate for large scale production of particles. Presently, compromises on material quality must therefore be made to produce large numbers of particles of this type. In addition, the

experiments presented only required imaging of one particle, whereas the attempts made in this thesis required many thousands of particles being measured simultaneously.

Given the current state of manufacturing process for hydrogel particles, it was therefore decided that they would be inappropriate for use in the circulating dispersions rig. Whilst they worked adequately for experiments requiring lower depth of field, such as in the vertical channel experiment, the refractive index matching of a homogeneous suspension was felt to be unrealistic at the concentrations required.

In the next section, an overview of the common refractive index matched pairs is given, along with the choice that was made for use in the circulating dispersions rig.

8.2.Optically matched liquid-solid pairs

Many of the alternative refractive index matched pairs are listed in a review of the field by Wiederseiner et al. (Wiederseiner, Andreini et al. 2010) This paper provided the experimental method used to characterise the optical transmittance of the water-hydrogel mixtures, as described above. This paper lists many examples of particles, in particular plastics, which can be used to match the refractive indices of particular fluids, as well as reproducing many examples which have been used in the literature.

Within the candidate materials to match low RI fluids, there are some which can be used with water. Examples of these include Fluorinated ethylene-propylene (FEP) and Tetrafluoroethylene (TFE), both of which have RI values of 1.34. However, researchers who have attempted to use FEP found the optical transparency to be poor, making it ill-suited for these measurements (Hassan and Dominguez-Ontiveros 2008).

Higher RI plastic particles can be used with the correct fluid. For instance, PMMA is commonly used (as above for the comparison case) (Yuki, Okumura et al. 2008). This has a stated RI of 1.49. Additionally, fused quartz represents an intermediate possibility, with an RI of 1.458. 4.7mm PMMA beads were used to refractive index match p-cymene (Hassan and Dominguez-Ontiveros 2008). This system was then utilised to find the fluid velocity field in between pebbles.

PTV visualisation has been achieved using 8mm PMMA spheres and Sodium Iodide (NaI) solution (Satake, Aoyagi et al. 2014). This is the same system which was used in the

counterfactual experiments for the hydrogel-water testing. In these experiments, threedimensional velocity fields were reconstructed around the particle positions.

Another possibility is to use glass particles (Mehta 2007). However, the RI values of these types of particles are broad-ranging, usually between 1.47 and 1.65. Therefore, the behaviour can be unpredictable, which does not make these an ideal candidate.

Use of pyrex appeared to be particularly promising. Pyrex is a brand of borosilicate glass, which has a well-defined refractive index of 1.473. This is in contrast to other types of glass, which tend to exhibit a larger range of refractive index values. There are numerous examples of matching fluids for these types of particles, including but not limited to:

- Potassium thiocyanate (KSCN) solution (Gijsen 1996)
- Aqueous zinc iodide solution
- Kerosine and oil mixtures
- Various Dow Corning fluid mixtures
- Demineralised water and zinc chloride mixtures
- Ethyl and benzyl alcohols
- Sodium iodide solution and glycerin

Broadly speaking, the matching fluids could be placed into two categories; organic molecule based solutions, such as mixes of alcohols and oils, or aqueous salt solutions, such as the sodium iodide solution used in the matching experiment above. It was felt that, for these experiments, it would be much easier to operate with aqueous salt solutions in the experimental rig.

Numerous relationships can be found in literature which attempt to quantify variations of refractive index with respect to temperature, concentration and wavelength for these solutions. Papers by Patil and Liburdy (Patil and Liburdy 2010, Patil and Liburdy 2012) investigated the refractive index matching of ammonium thiocyanate and pyrex beads of 6mm diameter. They discern that the beads act as converging or diverging lenses within the solution, dependent upon the direction of any refractive index mismatch. Therefore, focal distances shift occurring to bead size, number of phase boundaries along a light path, and the total path length. Decrease in transmittance was found to be symmetrical around the peak value, and totalled one-half as the refractive index mismatch reached 0.0035.

Comparisons have also been made between sodium iodide solutions, as used in Section 8.1, and potassium thiocyanate. Potassium thiocyanate is a more attractive material to use experimentally; unlike sodium iodide, it does not discolour when it comes into contact with air. It is also regarded as being lower in toxicity, which is a benefit when considering the loading and unloading of an experimental system.

The refractive index of sodium iodide has been characterised in literature by a number of sources. Bai and Katz (Bai and Katz 2014) measured refractive index at 23°C and 589.3nm, and developed the following relationship, where c is the concentration of NaI salt by mass in the solution:

 $n_{\text{NaI}} = 0.2425c^2 + 0.09511c + 1.335$

Equation 8.3: Refractive index variation of Nal, reported by Bai and Katz

Narrow et al (Narrow 2000) used a smaller concentration range, but also investigated the effects of temperature and wavelength on the optical behaviour of the solutions. They reported the following relationship:

 $n_{Nal} = 1.252 - (2.91 \times 10^{-4} c^{-1})T + (0.365)c + 5542 \times \lambda^{-2}$

Equation 8.4: Refractive index variation of Nal, reported by Narrow et al.

Agrawal et al (Agrawal, Shokri et al. 2015) confirmed these relationships, in addition to making comparisons using potassium thiocyanate solutions. At a wavelength of 589.3nm, they reported the relationship below, with c as the mass percentage of KSCN in the solution:

$$n_{KSCN} = 0.0774c^2 - 0.00015T(c+1) + 0.1745c + 1.338$$

Equation 8.5: Refractive index variation of KSCN, reported by Agrawal et al.

They determined from these equations that the temperature sensitivity of both solutions is low, and roughly of the same value, leaving concentration fluctuations as the most dominant mechanism affecting refractive index. A separate investigation of refractive index around that of borosilicate glass returned linear relationships of the form, for a temperature of 25°C:

$$n_{NaI} = 0.3982c + 1.2427$$

Equation 8.6: Refractive index variation of NaI, accounting only for concentration

$n_{KSCN} = 0.2605c + 1.3082$

Equation 8.7: Refractive index variation of KSCN, accounting only for concentration

These relationships demonstrate the lower sensitivity of potassium thiocyanate on concentration, since the gradient of change is 1.53 times greater for sodium iodide solutions. As a result, it will be easier to maintain refractive index matching for pyrex and KSCN mixtures.

As a result of these advantages, potassium thiocyanate was chosen as the matching fluid for the pyrex particles in these experiments. In order to match with the refractive index of pyrex, according to the relationship above, it would be desirable to use KSCN with a concentration of 63.3%. In separate experiments, Agrawal et al (Agrawal, Shokri et al. 2015) demonstrated transmittance through a packed bed of 3mm pyrex beads and KSCN solution at differing concentrations. The best match was shown to be at c = 62.4%, close to the theoretical value. However, since batches of pyrex may vary in refractive index, it may be that the concentration to use varies depending upon the particles being matched to.

Finally, the temperature variations at the optimal concentration were assessed. It was determined that the refractive index obeys the following relationship:

 $n_{KSCN} = -0.0003T + 1.478$

Equation 8.8: Temperature dependency, T, of refractive index of KSCN at optimum conditions

This relationship indicates that, for 5°C changes, the refractive index of the liquid is affected by 0.0015. Whilst this isn't a large change compared to the effects of concentration changes, it does indicate the importance of temperature control in maintaining optical visibility.

8.3.Rig design considerations

To provide experimental results which served as a good comparison to CFD investigations within TMF, various parameters were pre-determined in order to inform the experimental design. Experiments were required to fall within the highly turbulent regime, and therefore a Reynolds number range of 25000-100000 was specified. To investigate this range, a method of flow rate control was required in the rig design.

Furthermore, to provide relevant results on phase interactions, it was necessary for the particles to be of a similar lengthscale to the underlying turbulent eddies. This therefore placed a constraint on particle diameters of less than 1mm, with a preference of particles which measured under 0.5mm. These particles were to be pumped with the liquid in a continuous flow, rather than introduced separately as with the vertical channel experiments. Therefore, the pumping mechanism chosen was to be suitable for processing of fine particulates.

As discussed in Chapter 1, results of interest are those at high dispersed phase concentrations, so that four-way coupling effects are very prominent. In most available literature, particle phase fractions of 1% are considered the limit of what can be achieved, due to issues with optical visibility in more dense suspensions. This is considered comfortably within the range of the four-way coupled regime. Experiments within this rig were designed in order to maximise dispersed phase fraction up until the limits of optical visibility; given previous attempts, this was anticipated to be of the order of 1%.

As discussed previously, the contents of the rig were to be refractive index matched. In order to facilitate this, the temperature of the contents was to be controlled by using a temperature measurement and heat exchanger feedback loop.

Finally, there was a requirement to make measurements in a well-developed region of the flow. These length scales were discussed in the context of the vertical channel rig, and for the purposes of this experiment a distance of 70-80 pipe diameters was considered sufficient to satisfy this constraint.

Bearing these factors in mind, a basic design for the building of this rig was conceived, as shown Figure 8.6.



Figure 8.6: Schematic of design for the circulating dispersions rig

The pipe to be used was 50mm internal diameter. This allowed the range of Reynolds numbers specified to be achieved, whilst reducing the optical path length required. The PIV test section was located 3880mm upstream of the pump outlet, corresponding to a distance of 76 pipe diameters, enabling measurements of well-developed flow. Furthermore, the pipeline leading up to the measurement section was entirely straight, without any bends; this was to eliminate swirl in the flow, to encourage convergence on well-developed flow. Finally, a distance of 1000mm to the top bend allows measurements to be independent of effects from liquid traversing this corner.

Valves V1 and V2 were located inside the laboratory, allowing for control and manipulation of flow rates through the main and bypass lines. Valves V3 and V4 were only used for shut down and start up procedures, and therefore access was not necessary during experimental work. Valve V5 represents control of cooling water, dictated by pump flow rate and LabVIEW control. Valves V1 and V2 were chosen to reduce damage to particles, which could be caused by tight constrictions. Therefore, butterfly valves were chosen to minimise these effects. In order to make measurements at a distance of 76 pipe diameters from the pump, an existing platform was used, with the new experimental apparatus built around this existing structure. The first stage in construction was the build of a laser safety cabin. This is a box which surrounds the experiment and is light impenetrable, thereby confining high intensity laser light to a closed area, which may then be considered in isolation for health and safety purposes. A light above the entrance to the cabin is linked to the laser, warning personnel outside that the laser is in operation. Additionally, switches were attached to the doors of the cabin, which were linked to a laser interlock box. This provided an additional measure of safety against anyone wishing to enter the cabin who had not observed the warning light. If the doors are opened in such an eventuality, the circuit is broken, and the laser light source is terminated.





Following the installation of this cabin, the experimental apparatus was then constructed. In the following subsections, the requirements and design of the principal components will be discussed. Additionally, the instrumentation required, along with its calibration and operation, will be covered.

8.3.1. Equipment

The first consideration when finding equipment for this experimental rig was to find a suitable pumping mechanism for particles in a slurry. It is undesirable for any pumping mechanism to impart large amounts of shear stress to particles within the pumped solution, since this may damage the particles significantly. The clear choice was therefore a progressive cavity pump. These are also known as positive displacement pumps. They act via

the mechanism of a turning part within the barrel of the pump, much like a screw. This part contains several small, discrete cavities, and as the rotor mechanism turns, the fluid in these cavities progresses along the length of the barrel, until it is expelled. These are often design with overlaps between neighbouring cavities so that the arrival of a cavity at the outlet doesn't cause flow pulsing. These pumps therefore deliver very reliable and steady flow rates, which is another consideration for their choice in this application.

Further considerations were: the choice of media of an aqueous-based solution, guidelines for variation in the flow rate delivery in order to provide Reynolds numbers between 25000-100000, and a pumping height of 5.28m, and the need for a pumping mechanism which would be largely unaffected by reasonable particulate volume. The pump chosen is shown in Fig. 8.8. It was selected with an integrated motor inverter, which allowed for variation of the output flow rates between 5-25 m³/hr. In terms of this experimental rig, this corresponds to a Reynolds number range of 22000-109000. This mostly covers the desired range, with the option to achieve the lower ranges by utilising the bypass valve shown in the design schematic in Fig. 8.6.



Figure 8.8: Progressive cavity pump used in the circulating dispersions rig. This was fitted with the additional safety mechanism provided, which was an overpressure switch

The storage tank was sourced with a few criteria in mind:

- With at least 40 litres of pipeline volume required for this experimental rig, a tank volume of 200+ litres was required to provide sufficient reserve.
- The tank used would require an impeller mechanism. This would act to prevent uneven dispersed phase fractions in start-up of the experimental apparatus, which could be a safety concern if the pump were overloading with a large amount of particulates.
- Fittings for heat exchanger input and output, either by a heat jacket or an internal exchanger, were also considered desirable.
- Sufficient inlet and outlets for operation, including ease-of-loading both the liquid and solid phases throughout the experimental procedure.

In order to optimise "off-bottom" clearance, the shape of the vessel used is of importance. Ideally, a liquid height of Z/T - 1.15, where T is the tank diameter, is considered most optimal (Bao, Lu et al. 2018). Additionally, a dished bottom vessel can aid the mixing process. These vessels require less power than flat-bottomed tanks, where dead zones can develop in the corners resulting in trapping of particles.
Another consideration is the geometry and size of the impeller used. Different impellers produce differing flow patterns; in the limits, these can be described as axial and radial flow patterns (Albright 2009), as shown in Fig. 8.9. Axial flow patterns are considered desirable for suspension of particles, and so design choices of the impeller should reflect this where possible. The size of the impeller is usually chosen between 0.2-0.6T, depending upon the specific application, with T/3 most often being chosen for agitation investigations and applications (Jafari, Tanguy et al. 2012). Usually, it is also desirable to operate the impeller at low rotational speeds (200-600 rpm) to avoid significant damage occurring to the entrained particles.



Figure 8.9: Illustration of typical radial and axial flow patterns in an agitated vessel. For suspension of solids, creating an axial flow pattern is preferable.

The tank purchased for this experiment is shown in Fig. 8.10, and it satisfies many of the criteria outlined above. The tank has a dish-bottomed surface, which aids particle suspension. Two sets of fittings were provided to a heat jacket, consisting of a half-coil, which could then be used in conjunction with a heat exchanger. Furthermore, the impeller is a marine-type propeller, which is known to cause axial flow patterns which encourage particle suspension (Kumaresan and Joshi 2006).

Further information about the storage tank is presented in Table 8.1.

Parameter	Value
Z, Tank filling height (m)	0.95
T, Tank interior diameter (m)	0.72
D, Impeller diameter (m)	0.30
Ω , Rotational motor speed (rpm)	1420

Table 8.1: Physical parameters of storage tank used in circulating dispersions rig

These parameters give ratios of D/T = 0.42 and Z/T = 1.32, which are in good agreement with the desirable values outlined above. However, the motor supplied with the tank has a rotational rate much higher than required. As a result, the motor was fitted with an external inverter, which then allowed a slowing of the rotational speed, and finer control of the particle agitation. It was found that a rotational speed of 400-500 rpm was sufficient to achieve off-bottom clearance of the particles used, and therefore this was used throughout the experiments.



Figure 8.10: (Left) Photograph of storage tank used in circulating dispersions rig; A is the stirrer motor, and B is the heat jacket for use with the heat exchanger (Right) Photograph of the interior of the tank; C is the impeller used to agitate the particulate-liquid mixture

A transparent measurement section was built to provide optical access for PIV measurements. Ideally, it would be best to match the refractive index exactly with the choice of material for this section. However, facilities only exist to manufacture such equipment from certain material. A compromise was made, using Perspex (PMMA) for the pipe material. This material has a refractive index of 1.49, which represents a slight mismatch with the pipe contents. However, this shouldn't have a large effect, since there are only two boundaries to traverse.



Figure 8.11: (A) Central measurement section (B) measurement section including rectangular box to reduce distortions at boundaries

An integral part of such measurement sections is a surrounding box containing the same fluid as the pipe. Without this inclusion, large distortions occur at the curved boundary walls, making the PIV signals captured by the camera scattered and therefore useless. By creating a straight boundary edge, this can be avoided. This is shown in the Fig. 8.12, taken by half filling the surrounding box with fluid. The filling line is marked in red on this image. Below the filling line, the section of the measurement section entirely filled with liquid presents an entirely flat surface, allowing optical visualisation and clear images of tracer particles. In contrast, above the filling line the camera has to image from a curved surface, and cannot focus at one distance. This creates a distorted image which is dominated by noise, as shown in the upper half of Fig. 8.12.



Figure 8.12: Demonstration of functionality of the surrounding box in the measurement section, achieved by filling the box to half way.

8.3.2. Instrumentation and rig control

A flowmeter was chosen for this rig to minimise disturbance to the flow on the mainline. The best choice given this requirement is for an electromagnetic flowmeter. These devices operate by applying a magnetic field to the pipeline under scrutiny, and then measuring the potential difference generated by moving charges within the fluid. This potential difference is proportional to the velocity of the flow perpendicular to the magnetic field direction, and therefore allows for calculation of the flow rate, if the pipe diameter is taken into account.

The flowmeter is shown in Fig. 8.13. This was a suitable choice for this application, since it is designed to operate on both water flows and in chemical industry applications, measuring acidic and alkaline solutions. It was possible to calibrate this flowmeter based on the potassium thiocyanate solution used specifically for this experiment, in advance of the main experimental run. The device was capable of measuring in a flow range of 0.03-12 ms⁻¹ to an accuracy of 0.25%, comfortably covering the required range. Results were obtained using the 4-20mA output, whilst the on-screen display was used to perform checks on performance throughout the experimental run.



Figure 8.13: Photograph of the electromagnetic flowmeter used on the circulating dispersions rig

To monitor the temperature of the experimental rig, T-type thermocouples were used. Since the expected range of temperatures was only between 15-30°C in the most extreme of circumstances, the range provided by a K-type thermocouple was unnecessary. The cold junction compensation for these thermocouples was provided by the cDAQ used to make the measurements. The coupling used to hold these thermocouples within the flow is shown in Fig. 8.14.



Figure 8.14: T-type thermocouple, held in coupling to maintain position within the fluid flow

The differential pressure across the measurement section was measured using an Omni PD-23 pressure transducer. This unit was capable of measuring in the range of 0-500 mBar, which comfortably covered the expected range of differential pressures across the 1m section of pipe under measurement (typically, the average would be expected to vary between 150-200 mBar, dependent upon the theoretical assumptions made). Measurements were made at an accuracy of ±2% of the experimental value., and with a response time of 1 kHz, which is comparable to the frequency at which laser measurements are obtained. As with the magnetic flowmeter, measurements were acquired using a 4-20mA output.

Connections were made using flexible PVC tube, which connect the G 1/4"female ports on the transducer to identical fittings on the pipe, which were placed at a distance of 0.3m from the edges of the measurement section. As per the manufacturer's instructions, these were filled with water during operation.



Figure 8.15: Photograph of differential pressure cell used across the measurement section of the circulating dispersions rig

Calibration of the flow meter was performed using varying particulate concentrations and flow rates, corresponding to those used in the experimental work. Particles were not found to have a large effect upon the flow rate readings produced, both in overall value and stability. A collection procedure was used to record the throughput of liquid phase and the corresponding current output. These values could then be programmed into the flow meter, to give both digital and display readouts which matched.

The typical heat input due to pumping action was characterised by measuring over a standard running time of around two hours, without any temperature control, using water as the liquid phase to utilise its well-known properties. The results of this experiment are shown in Fig. 8.16. It should be noted that the results from the other thermocouple were very similar; typically these values were around 0.1°C larger, possibly due to the slightly warmer temperature inside the cabin. Nevertheless, the results followed a very similar trend, and therefore are not replicated here.



Figure 8.16: Temperature measured at pump outlet over standard running time of the experimental rig without any temperature control of the pipe contents, using water as the liquid phase for a test case.

The data shows a variable rate of temperature increase across the time measurements were made. This might be partly due to temperature increases taking place in the laboratory due to the time of day. Nevertheless, these environmental factors must also be considered when attempting to compensate for temperature increases within the rig. The irregular increase in temperature may also be caused by factors such as heating or air conditioning turning on in the laboratory and the changing temperature gradient affecting the rate of cooling.

A linear fit was made using the data, to extract a value for the average heating of the rig contents within a period of two hours. The gradient found was 0.0268 ± 0.0018 °C/min. When this level of heating is considered over the period of one hour, this equates to an increase of 1.6°C. With a filling level of 300 litres, this corresponds to a heat input of 0.56 kW. This is likely a combination of the pumping mechanism, friction losses due to movement through the pipelines, the agitation from the tank impeller, and the environmental factors such as heat exchange with the environment.

The temperature was controlled by using the half-coil heat jacket surrounding the storage tank. This heat jacket consists of 20.1m of 2" half pipe, corresponding to an instantaneous coolant volume of 19.8 litres. The coolant used was domestic cold water, typically at a

temperature of 15°C. An actuator was fitted to the input valve of the heat jacket, allowing for control of the coolant flow rate using feedback from the fitted thermocouples.

As shown in the Fig. 8.16, the typical heating rate of the contents without any temperature control was 0.56kW. The simplest equation to describe heat transfer is given below:

$$\frac{\Delta Q}{\Delta t} = UA\Delta T$$

Equation 8.9: Equation for heating of a liquid using a coolant

where Q is the heat transferred, t is the time elapsed, A is the area available for flow of heat, ΔT is the different in temperature between the processing fluid and coolant, and U is the heat transfer coefficient. Using this relationship, we can determine that Q must be equal to 0.56kW in order to extract the heat added during operation.

For this tank, A can be calculated as 1.26 m³. Typically, to maintain the temperature of the rig contents at 20°C, at equilibrium $\Delta T = 5$ °C. This means that the heat transfer coefficient must be larger than 88.9 Wm⁻²K. This is lower than the typical values cited for most heat exchangers using jacketed vessels, and therefore was within the capability of this set of equipment.

8.3.3. Parameter space and data acquisition

Acquisition of flowmeter, thermocouple and differential pressure cell signals were achieved simultaneously using a National Instruments cDAQ-9178 containing a 4-channel NI-9219 Analog Input Module. The capture of these signals was triggered in each experiment by the first rising edge of the laser trigger signal, thus ensuring that the data captured was synchronised to the images taken. These measurements were taken at 1 kHz and for a duration of 10 seconds, in order to capture data during and for a large period after the images were captured.

The dispersion to be considered, as stated in Section 8.2, was a liquid phase of Potassium Thiocyanate with a solid phase of Pyrex particles. These particles were provided by a bulk supplier, and were given with a size specification of 250 μ m, with a standard deviation of 20 μ m.

Experiments were conducted to investigate variations in liquid flow rate and particle concentration. In total, seven flow rates and five particle concentrations were chosen, with measurements taken for each combination. The flow rates are as shown in Table 8.2.

 Table 8.2: Liquid flow rate characteristics relevant to the experiments conducted within the circulating dispersions rig

Q, flow rate $(m^{3}h^{-1})$	V, streamwise velocity (ms ⁻¹)	Re, Reynolds number
8.7	1.23	3.86×10^4
11.0	1.56	4.89×10^4
13.4	1.90	5.96×10^4
15.2	2.15	6.74 x 10 ⁴
17.2	2.43	7.62 x 10 ⁴
19.4	2.75	8.63 x 10 ⁴
21.7	3.07	9.63 x 10 ⁴

For each of these liquid Reynolds numbers, the corresponding particle Stokes number is

detailed in Table 8.3.

 Table 8.3: Particle Stokes number for each of the experiments conducted within the circulating dispersions

 rig

Reynolds number, Re	Stokes number, St
38600	0.34
48900	0.48
59600	0.65
67400	0.78
76200	0.94
86300	1.13
96300	1.33

Five particulate concentrations were studied. Initial tests were conducted which determined an upper limit of around 3% before images begin to deteriorate in quality, making analysis difficult. Using this upper limit, five particle concentrations, equally spaced between 0.56% and 2.76% were chosen to span the range of phase fractions.

All PIV measurements during this set of experiments were taken at a frequency of 3135 kHz. For each experimental condition, 4000 sets of doubles images were taken, giving a total record time of 1.28s per data set. The recording frequency chosen was sufficiently high to ensure accurate particle tracking, by linking frame 2 of t_n to frame 1 of $t_{(n+1)}$. As is standard for PIV measurements, the time between frames was chosen to optimise the performance of the autocorrelation algorithm, and was thus varied in accordance with the average streamwise velocity of the particles.

8.4.Single phase characterisation

Single phase operation of the experimental rig was recorded at all Reynolds numbers used by means of Particle Image Velocimetry (PIV). The ensembles of measurements created were 2000 images in size; errors on mean velocities are therefore assessed as being 0.045%. This characterisation is described within the section that follows.

In order to maximise optical visibility, measurements were taken in a plane at a distance of 17mm from the pipe centre. This is represented in Fig. 8.17.





As a result, the radial distances measured in these experiments are between 17-25.5mm, which corresponds to 0.66-1R.

The mean axial velocity profiles for all Reynolds numbers used are shown in Fig. 8.18 for this range of radii. Naturally, the velocity profiles appear to be flattening around 0.66R, as expected for a turbulent flow profile. Measurements are not precise enough at the wall to measure the no-slip conditions. This is due to measurements taking place in a finite volume which encapsulates some bulk flow, and also some light scatter from the pipe wall affecting the ability to resolve in this region. The average velocity measurement at the 'pipe wall' is around 0.7 times the value of the central velocity.

It should be noted that these profiles do not collapse, despite being divided by the anticipated average fluid velocity. This could be a result of flow fluctuations, which result in a marginally incorrect normalisation.



Figure 8.18: Axial velocity profiles for varying Reynolds numbers in single phase flow, normalised using average axial velocity calculated from measured flow rate

Velocities perpendicular to the axial velocity in the plane of measurements were also calculated. The mean profiles of these velocities are shown in the Fig. 8.19. These velocities indicate the tendency for fluid elements to move perpendicular to the overall fluid velocity i.e. across the pipe. As anticipated, the crosswise velocities are low compared with the streamwise components, typically around 3-4% of the streamwise velocity values. As can be observed in Fig. 8.19, for all Reynolds numbers there is a lower crosswise movement towards the circumference of the pipe. Therefore, these profiles show how much larger the crosswise redistribution is at the pipe wall. All of these profiles appear to be flattening towards the centre of the pipe.



Figure 8.19: Crosswise velocity profiles for varying Reynolds numbers in single phase flow, with varying normalised Radial Distance

Turbulence intensity in the streamwise direction was calculated from this ensemble average, by dividing the root mean square of the streamwise velocity component by the mean streamwise velocity. At the circumference of the pipe, the turbulence intensity varies between 11-12%. Profiles at all Reynolds numbers show a decrease in turbulent intensity towards the pipe centre, with values of 5-6% at 0.66R. Typically, the profiles flatten towards the centre of the pipe. However, in the case of the highest and lowest Reynolds numbers, the profiles are raised between 0.75-0.9R, so this behaviour is not evident.



Figure 8.20: Turbulent intensity profiles, u_{rms}/<u>, in streamwise direction

Fig 8.21 shows the calculated turbulent velocity fluctuations, v', for the crosswise component of the two-dimensional velocity. There is typically a peak in this quantity between 0.9-1.0R, with a decrease towards the centre of the pipe, where turbulence levels are lower. The fluctuations in this direction are comparable to the velocity fluctuations in the streamwise direction, but generally lower.



Figure 8.21: Turbulent velocity fluctuations in the crosswise direction, v', for varying normalised Radial Distance, at different Reynolds numbers

The development of the flow was assessed by comparing streamwise velocity profiles at consecutive positions along the pipe, using the PIV velocity maps obtained. This is shown in the left hand side of Fig. 8.22 in the form of a raster scan. Raster scans are produced by reading values along one row, and then again along subsequent rows. The repeating form of the velocity profile is shown for two Reynolds numbers. From these, it appears that the lower Reynolds number case exhibits a lower drift in profile. This is shown in the plots of fractional deviation, whereby deviation is shown from the 'reference' profile for the subsequent velocity profiles, as shown in the right hand side of Fig. 8.22. For the lower Reynolds number case, the deviations are typically lower than 0.5%, with a maximum of 1%. Additionally, these values are most likely to occur at the wall, where higher turbulence intensities have an impact. For the higher Reynolds number, variations are typically twice as large, and appear to be more systematic.



Figure 8.22: (left) Raster scan presentation of mean streamwise velocities, for Re = 3.86×10^4 (top) and Re = 9.63×10^4 (bottom) (right) Deviation in streamwise velocity profiles at subsequent distances along the pipe, for Re = 3.86×10^4 (top) and Re = 9.63×10^4 (bottom)

8.5.Conclusion

This chapter began by investigating the suitability of hydrogel particles for use in the circulating dispersions rig. Experimental results were obtained from an associated Master's project, supervised during the PhD by the author. Whilst these particles are theoretically ideal for refractive index matching with water, the construction methods of these particles appeared to prohibit their use in a dense suspension, since the test case of PMMA in sodium iodide produced more transparent suspensions.

Following on from this, there is a small review of the types of particle-liquid pairs which have commonly been used in this field of research. A mixture of pyrex and potassium thiocyanate was chosen to move forward onto the circulating dispersions rig experiments. Pyrex particles are considered to have good optical quality, and can be purchased in large quantities, making them a suitable dispersed phase candidate. Potassium thiocyanate was considered superior to sodium iodide; its refractive index has a lower concentration dependence, it discolours less easily in oxygen, and it carries a lower health and safety risk. Following on from this, the design of the circulating dispersions rig is discussed. Given the initial design specifications and the available lab area, a schematic was created in which the measurement area was enclosed by a laser cabin on top of a raised platform. A pump was chosen to provide reliable and safe delivery of liquid mixed with particles, and the tank was designed in order to provide adequate suspension of particles. Other factors, such as instrumentation, the measurement section, and the heating of the experimental rig contents were discussed. An experimental parameter matrix was then determined, consisting of varying Reynolds numbers and particle volume fractions.

Finally, a single phase characterisation of the experimental facility is presented, using data sets of 2000 images. In order to maximise optical visibility, the laser light sheet was placed off-centre of the pipe, so that the distance for light to travel was reduced. Therefore, results are only presented for 0.66-1.0R.

Axial velocity profiles do not show a decent collapse when normalising by the average fluid velocity. This may be a consequence of unaccounted for fluctuations in the flow. The turbulence intensity in the axial direction is calculated, and appears to be converging towards 5% in the bulk of the flow; at the wall, values reach as high as 11%. Finally, there is a comparison between flow profiles for consecutive measurements along the pipe. At the lowest Reynolds number, the deviations are typically around 0.5%, and are randomly fluctuating either side of the initial distribution. For the higher Reynolds number, these fluctuations are twice as large, and typically seem to be more systematic. These single phase results can be used to make comparisons with corresponding CFD simulations which may be performed.

Chapter 9. Circulating dispersions rig - results

This chapter presents the multiphase flow results of the circulating dispersions rig, as described in the previous chapter. Whilst there are not many direct overlaps with the results of the vertical channel rig, it is possible to draw comparisons between the two data sets in specific areas, and speculate as to what may cause any discrepancies. As with the vertical channel, the results will focus on the distribution statistics of the dispersed phase and the relative velocities of particle pairs in motion.

The result shown in this chapter focus upon the effect of concentration changes, due to this being the operating area of research. Further results were also recorded during the project to account for changes in Reynolds number. The concentrations used ranged from 0.75-3.0% by volume; at higher concentrations, results could no longer be achieved due to visibility issues. A continuous phase of KSCN and a dispersed phase of 250µm Pyrex particles were used throughout these experiments.

9.1.Box counting method

Fig. 9.1 shows the results of applying the box counting method to the observed particle positions in these experiments. For each concentration, there is an obvious peak in clusters, which typically occurs at a cell size, as normalised by the particle diameter, of around 2-4. This is much lower than for the vertical channel, for which the peak value usually occurred in the range of 12-20 particle diameters, in well-developed flow conditions.



Figure 9.1: Deviations, D, from Poisson histogram statistics of particle position distributions, as computed via a Box Counting Method for a variety of particulate volume fractions, at varying length scales of observation

These typical clustering length scales are more similar to those of the injection point in the vertical channel. Nevertheless, we should consider that, for these experiments, we were typically observing at much higher liquid Reynolds numbers. As a result, the underlying eddies which cause clustering may have a lower length scale, which could be reflected here.

Another point of note is that, for the most part, the peak clustering values are very similar to the well-developed flows in the vertical channel, falling between a range of 0.25-0.5. This is actually a reasonable standard value to compute from calculations of these types. This proves that, in both experiments, the clustering observed is likely a real phenomenon, the scale of which remains reasonably constant across these experiments.

However, another key deviation from the previous experiments is the correlation between peak clustering values and particulate concentration, which no longer seems to exist here. Considering the lack of pattern in these results, it is highly likely that concentration is not a dominating factor in clustering strength for these types of flows.

9.2.Particle-particle closing times

Fig. 9.2 shows data concerning the closing times of particle pairs, calculated by considering the magnitude of radial separation and the resolved relative velocity in the direction of separation. Taken together, these components may give the time of closing/collision between the two flow elements. Integrating over time, this may then be used to predict the probability of particle-particle collisions within a given time frame.





For these plots, there is evidently a lower probability of particle-particle collisions at 0.75%, the lowest concentration. For all other concentrations, these probability values at a given timeframe are somewhat similar, although perhaps surprisingly probabilities for 2.25% concentration are lower than would be anticipated.

In the interests of providing context, it should also be noted that particles were viewed in these experiments for a travelling length of 0.061m. At typical travelling velocities of 2.36ms⁻¹, there is therefore only a time on screen for these particles of 0.025s. In this range,

probabilities of collisions typically range between 0.15-0.2, which is suggestive of a large amount of particle-particle interactions, compared with those observed in the vertical channel experiments.

9.3.Particle-particle Radial Distribution Functions

The radial distributions functions shown in Fig. 9.3 represent the relative probability of two particles being within a specified distance, relative to the entirely random case. All curves take the shape that would be anticipated; however, image resolution issues mean that the curves do not collapse to a value of zero for a particle diameter equal to 1.



Figure 9.3: Radial Distribution Functions of Particle-Particle separation, shown for varying particulate concentrations.

The peaks are illustrative of a peak in radial separation; for these experiments, this typically occurs at a very low value of around 2 particle diameters, which is symptomatic of particle-particle pairs forming within the flow. Generally, any increase in probability over a random distribution is curtailed for a separation of more than six particle diameters. This indicates that particles from very tight-knit clusters in these types of flows; this is also represented by

the low characteristic length scale observed in the box counting measurements, and the relatively high collision probabilities seen in the relative velocity data.

However, as with the box counting data, the concentration seems to have little correlation with the strength of the clustering which is observed. For all conditions, this peak is reasonably high, taking values between 7-17. The only common characteristic with the box counting data is that the particulate concentration experiments of 2.25% exhibit the lowest levels of clustering by the metric used.

9.4. Local cluster method

This section describes the application of the local cluster method to particle distributions observed within the circulating dispersions rig. This analysis applies potentials to particles, based on the distances of the nearest neighbours in each case. These are normalised by the mean particle potential, a distance based upon the bulk particle concentration.



Figure 9.4: Plots of individual normalised particle potential, using a 12-neighbour formulation, for varying particulate concentrations.

Fig. 9.4 shows the local particle potentials for sets of particles at varying concentrations, whereby the distribution has been normalised by a threshold potential related to an even

spacing with a three-dimensional scaling. Whilst all of these potential curves observe a similar shape, there are a few subtle differences. What is noticeable is that the peak value is higher for lower concentrations, and vice versa, indicating that for high concentrations, less particles have a threshold which corresponds to an evenly spaced distribution.

For the highest particulate concentration, it can easily be seen that particles are more likely to be found at higher particle potentials i.e. in clustering conditions. Whilst this is a small difference in probability, it is nevertheless present and systematic for these higher concentrations. There are also some differences in the tails of the distributions, which also indicate that particles are less likely to be found in void conditions for such high concentrations. Generally, it appears that there is a small but perceptible shift to the right hand side as concentration increases, indicating a transfer of particles from void regions to cluster regions. The largest transfer however occurs at the highest concentration, whereby many particles are found in cluster region and are reduced in the central values of particle potential.

Similar analysis of local cluster shape with angular separation is performed, as seen in Chapter 7. However, it is worth noting that the angles of separation correspond to different situations, due to the orientation of recording equipment that was used on the Circulating Dispersions Rig. The angle of separation is therefore defined as shown in Fig. 9.5.



Figure 9.5: Definition of angular separation for particles in the circulating dispersion rig. Angles within A correspond to 0° to 90°, within B are 90° to 180°, within C are 0° to -90°, and within D are -180° to -90°



Figure 9.6: Comparison of variation in nearest neighbour at different angular separations, at 3% particulate concentration. These plots are shown for the full distribution, saturated particles (those with neighbours in each angular region), and those which are classified as clusters

Fig. 9.6 shows the variation calculated from samples of particles obtained using different conditions. Unlike the vertical channel, there is actually very little difference in the statistics recorded. This is likely due to a more even distribution of particles. In the vertical channel, particles could easily be delineated into bulk and edge particles making up a jet, whereas in this scenario the only particles which experience different conditions are those at the pipe wall. This area is typically very difficult to measure, and so particles like this do not skew the statistics as much.



Figure 9.7: Mean nearest neighbour particle separation as a function of angular separation, for varying particulate concentration, for cluster particles only

Fig 9.7 shows the profile of mean particle separation as a function of angular separation for the experiments conducted. The most notably different profile is for 0.75%, which as expected features some very large separations, particularly at 0°, which is downstream of the particle under consideration. This profile also features a very prominent minima at around -90° to -100°, which is roughly equivalent to a particle tracking alongside. In fact, for this angular displacement, the mean separation is almost identical to those of other concentrations.

For other particulate concentrations, there is much less distinction between the mean separation profiles. Somewhat surprisingly, the separations for the highest concentration are marginally higher than for the intermediate concentrations. There also doesn't appear to be an obvious trend with regards to angle, which suggests that local cluster patterns do not persist with the change in concentration. This could be a result of very small clusters being dominant, resulting in a lack of obvious cluster shape.

The same trends can broadly be observed in Fig. 9.8 showing mean separation profiles of saturated particles (those particles with a full complement of neighbours). Whilst the profiles can certainly not be described as isotropic, there doesn't seem to be any persistent pattern across the whole concentration range. Therefore, the mechanisms by which particles locally cluster may depend quite heavily upon the flow conditions, particularly when three-dimensional swirling and higher Reynolds numbers become a factor.



Figure 9.8: Mean nearest neighbour particle separation as a function of angular separation, for varying particulate concentration, for saturated particles only.

9.4.1. Local cluster method discrepancies

Given the plots for both saturated and cluster conditioned particles, it is not reasonable to perform an attempt at non-dimensionalising, as was done with the results from the vertical channel. It also clearly seems counter-intuitive that mean separations should not correlate readily with particulate concentration.

There are a couple of reasons which could account for this lack of correlation, which also give some insight into the difficulty of performing this type of experiment. Firstly it should be noted that any attempt to analyse a two-dimensional projection of a three-dimensional data set, as is the case with planar PIV, will create difficulties. It is entirely possible that there are many extra particles in parallel planes to the light sheet used in the experiment, which could reduce the measured mean separation values of the highest concentration experiment. Therefore in order to truly characterise such a local cluster, it would be preferable to obtain either a volumetric-PIV, or multi-plane PIV set of results, which are capable of determining these out-of-plane interactions.

Another issue which may have affected these results is the scatter of light intensity upon further addition of particles. Even with information loss due to the three-dimensional nature of the experiment, it would still be expected to observe a higher concentration in the plane of viewing as overall particle concentration was increased. Therefore, although the results would not be fully representative of the cluster shapes forming, there should still be a decrease in mean separation.

However, it is noticeable that, for higher particulate concentrations, the number of particles returned by analysis did not increase as expected. This is not due to trapping of particles in the rig, since this was never observed throughout the experiments. Instead, it is likely due to the scattering of light intensity by an ever-increasing number of particles. This scattering results in a loss of image focus, making the automatic detection of particle outlines very difficult, because such detection relies heavily upon contrast and image gradients. Therefore, it is the author's opinion that, to achieve measurements at higher concentrations, it is necessary to reduce this scatter in a selection of particles. This topic will be discussed in the further work section of this thesis.

9.4.2. Local cluster connectivity

As with the vertical channel particle data, the self-similarity of central cluster particles and their nearest neighbours has been characterised. The probabilities of each having N local cluster neighbours is shown in Fig 9.9.



Figure 9.9: Plots of probability for a central cluster particle to have N local cluster neighbours, for varying particulate concentration

Firstly, these plots all exhibit the same shape, with a peak to an intermediate number of local cluster neighbours, followed by a rapid drop of probability when considering higher numbers. This gives further credence to the idea that the result obtained for 3.3% particle concentration in the vertical channel rig was anomalous, and these results correspond well with the anticipated formation of particle clusters for heavier-than-fluid particles.

In general, the most common number of local cluster neighbours to have varies between seven and eight. This lies in the middle of the range established by results from the vertical channel, but with much less variation. There is no strong correlation or change which is observable with changes in concentration. At 3.0%, the probability of higher numbers of local cluster neighbours is greater, but the second highest probabilities are for 0.75%. Therefore, there doesn't appear to be any causality between concentration and the probability of cluster particles grouping together.



Figure 9.10: Percentage of cluster particles with N local neighbours, for varying bulk particle concentration and values of N.

Finally, the number of local neighbours is considered for all central cluster particles. There is a large deviation from the vertical channel results when considering the probability of a full complement of particles. In those experiments, it was uncommon to observe lower than 90% of cluster particles which fulfilled this requirement; in these experiments, values typically range between 20-50%, depending upon the concentration in question. This illustrates the situations which often arise in a truly three-dimensional flow, which mean that whilst a particle may form part of a cluster, there are many more angular displacements to consider, making the possibility of a void section much higher.

The remaining values of N local neighbours from N = 11 to N = 8 are split reasonably evenly across all concentrations, typically ranging from 10-20%. Only for N = 7 is an obvious drop in

probability seen. Therefore, it is roughly equally likely that there are 1-4 voids within this formulation, which again shows the nature of considering two-dimensional projections of three-dimensional cluster structures.

9.5. Conclusions

This chapter has presented the results of data analysed from the circulating dispersions rig. As with Chapter 7, the focus has been on particle concentration effects, since this the issue of most interest.

Many of the results reported here differ to Chapter 7, and effects of concentration are generally not the same. This is not to say that concentration does not affect the results at all, but rather that there do not appear to be any obvious correlations, as were found in the channel rig flows.

The box counting method shows peak clustering strengths which are comparable with other experiments. However, the peak clustering lengthscale is much smaller than in the vertical channel experiment, typically falling between 2-4 particle diameters, which is a factor of five smaller than the previous experiments showed. This is most likely due to the higher Reynolds numbers in these experiments, which promote fluid length scales of different magnitudes, which then create the clustering structures. There is no correlation of peak clustering with concentration, although there are changes in values reported between experiments.

The closing times between particle pairs were reported in terms of probabilities of collision. For the time on screen of 0.025s, the probability of collisions 0.15-0.2. These are reasonably high values, indicating that multiple collisions would be observed in these experiments. Practically speaking though, it may be more prudent to consider collision probabilities for lower times than this, since the particle path will quickly diverge from its anticipated trajectory. These results also found that, for concentration of 0.75%, collision probabilities were low; for all other concentrations, probabilities remain roughly constant, and appear to have reached a plateau.

Radial distribution functions appear to confirm results from the box counting method. There are high values for g(r) at the lowest values of r, which drop of rapidly to indicate no

preference in radial separation for values of r > 6D. These results indicate a low characteristic clustering length scale, which is actually reasonably strong given the number of particles which can be involved. As with the box counting data, there is little obvious correlation between the strength of clustering in particle pairs and the bulk concentration.

Local particle potential results indicate that there is less variation in mean separation profiles depending upon the particle concentration. For concentration of 0.75%, mean separations are larger, but there a fewer deviations for higher concentrations. Again, this is reminiscent of the collision probability results. There was therefore little point in attempting a normalisation with characteristic separation, since length scale in this case is most likely determined by the fluid motions.

Connectivity analysis showed that, as anticipated for particle clusters forming along strain regions, there is a preference towards an intermediate number of cluster neighbours. There is also a higher probability of cluster particles still having void regions; this is likely to the additional spread of particles due to the three-dimensional nature of this flow.

These experiments were reliant upon two-dimensional projections of a three-dimensional flow, and therefore loss of information was always going to be a possibility. For example, given the equipment provided, it would not be possible to construct a three-dimensional cluster shape, and even if it were the number of angular elements required would be far higher. Additionally, the depth of flow made experiments much more challenging, due to a larger optical path length through unmatched solids and liquid. The scattering generated along this path length contributed towards a decrease in contrast of image intensities, a factor which may have resulted in incomplete capture of particles by automated algorithms. Therefore, some loss of particle information may explain a lack of clear correlation observed in, for example, the separation length scales observed.

The final chapter of this thesis will present an overview of the entire body of work. Following on from this, there will be a discussion of the limitations of these experiments, in particular the circulating dispersions rig, as alluded to in this chapter. Finally, there will be an analysis about what extensions to the experimental techniques could be applied to these experimental rigs, what extra data could be extracted, and what the overall contribution has been from this work in terms of deliverables.

Chapter 10. Conclusions and future work

10.1. Thesis conclusions

The work presented in this thesis forms part of a larger body of work within the Transient Multiphase Flows (TMF), which is aimed at solving the problem of poor dispersed phase predictions within existing CFD solvers. Using current models, flow behaviour simulations are unable to provide an accurate representation of instrumentation such as flow separators; often, flow conditions which should give rise to dispersed or bubbly flows produce only stratified profiles. The aim of this work was to provide experimental data for validation of a new CFD code, designed to specifically cope with these issues. If successful, a solver such as this could have great economic and environmental benefits, by improving the efficiency of equipment design in large, deep-sea oil fields.

The objectives were therefore outlined as follows:

- There was a requirement to investigate flows in a four-way coupling regime. This is largely understood to be the given by the condition that the dispersed phase fraction is greater than 10⁻³. The aim in these experiments was to achieve a phase fraction of at least 10⁻², and higher if permissible.
- Experiments were to be made using solid-liquid dispersions, with the aim of providing comparison points for CFD solvers aimed at liquid-liquid dispersions.
- To investigate highly turbulent dispersed phase flows in experimental rigs. One such rig, a pipeline flow containing freely-dispersed particles, was pre-specified, whilst the vertical channel was agreed upon during the project timeline.

10.1.1. Literature survey

The literature surveyed identified some of the main topics of interest in dispersed turbulent flows. Firstly, the clustering of inertial particles is considered of great importance to describing fluid mechanics in these types of flows. Secondly, the settling velocity of such particles (and also the effective drag coefficients) can be very difficult to predict, being either enhanced or suppressed depending upon the dominant mechanism. Finally, the effect of particles upon turbulent kinetic energy is of great interest, since it defines the phase boundary conditions which must be imposed to perform accurate CFD modelling.

10.1.2. Experimental methodology

An experimental methodology of high speed imaging was chosen for these investigations. In order to gain insight into whole-field behaviour of dispersed phase flows, with spatially resolved information as required by the objectives of the project, this was the only suitable technique. Following on from a discussion of the key aspects of PIV and the processing of its results, the anticipated theoretical error for each PIV calculated vector was discussed. Typically, these errors are found to be 1-2% F.S., or 0.1 pixel in terms of the cross-correlation displacement. Investigations have determined that standard commercial PIV algorithms achieve an error which is an order of magnitude smaller using synthetic images, and the most likely reason for the additional error is the uneven profile of the laser light sheet, coupled with perpendicular movement of tracer particles. Due to additional error on an experimental PIV measurement may be as high as 10%.

An extension of the technique, using two cameras with different optical filters and one flow constituent tagged with a fluorescent dye, was discussed with regards providing some means of phase discrimination. This technique was used for some of the experiments in the vertical channel, because hydrogel particles are constructed to allow for easy encapsulation of a fluorescent dye which is suitable for this purpose.

10.1.3. Particle tracking methods

Particle identification and tracking methods were investigated subsequently. This data was vital for creating particle dispersion and collision data. Initial information can be discovered by performing a watershed transformation on a dispersed phase image. This algorithm attempts to find "catchment basins" in the pixel intensity, thus giving an initial approximation of particle positions. Following on from this, particle centroids can be found using this initial information; firstly by attempting to find local maxima, and then by using a local mask to weight an initial estimate followed by determination of the particle centre with sub-pixel level accuracy. The application of the watershed algorithm prior to this process was found to improve instances of over and under prediction of particles.

Following on from this, two different candidates for tracking algorithms were examined in order to quantify their performance. The Hungarian performance time scales as N³ and the

nearest-neighbour algorithm scales as N². Experimental PDFs of particle velocity were created by both algorithms. These showed that the nearest-neighbour algorithm typically over-predicted extreme velocities due to poor matching. However, this can also be rectified by limiting the area of search, which can be achieved using a priori knowledge of the particle images. This approach can be applied to both algorithms, thus reducing the number of partners to be searched over (and therefore time of computation) and resulting in less poor matches.

Both algorithms were also tested using virtual particle images, in which particle velocity PDFs were given the form of Gaussian distributions. The Hungarian algorithm performed very well at replicating the input distribution in all situations. The nearest-neighbour algorithm struggled unless suitable constraints of searching distance were assigned initially.

As a result of these investigations, the Hungarian algorithm is recommended for more accurate extraction of particle velocities. However, in instances where time of computation is prohibitive, the nearest neighbour algorithm can be used to give reasonable answers, if reasonable restrictions are applied.

10.1.4. Masking and geometric transforms

Discrimination in mixed-phase images was performed using a mask based on particle size, after the image had been thresholded. Structure elements were used to perform an erosion of this mask, followed by a dilation to recover the elements lost from the dispersed phase mask. The effects of different structure elements were discussed with respect to different types of particle image. Typically, a structure element of radius 2 was used to perform the erosion, whilst a structure element of radius 3 was used as the basis for the dilation.

The procedure for matching images taken from two cameras simultaneously was performed by determining geometric transforms between the image planes. This combines with the method of particle discrimination which uses optical filters attached to the camera. Geometric transformations were performed in MatLab using control points, determined using the calibration grid. Projective transformations were found to provide lower errors in matching particle centroids, due to their inclusion of tilt transformations.

10.1.5. Vertical channel rig

The initial experimental apparatus outlined by this thesis is known as the vertical channel rig. This was designed as a simple experiment, which could be easily compared with CFD simulations. In this experiment, the continuous phase runs in a loop, with dispersed phase introduced to the flow at the entrance to the channel. The effects of this were then captured by synchronising the injection event with PIV measurements.

The dispersed phase system under investigation for these experiments was water and hydrogels tagged with a fluorescent dye. These two phases of water, given their properties, provide a very closely-matched system in terms of density, and presented an opportunity to test hydrogel particles for use in further experiments.

The experimental facility was designed to allow optical access at any point along the channel, allowing for variation of the measurement position. Further variations could then be made to the particulate concentration, the Reynolds number of the channel, and injection velocity of the particles.

The performance of the rig was characterised by conducting single-phase only experiments, either by running the rig without material injection, or by injecting a mixture of water and PIV tracer particles to simulate full operation. In the well-developed region, injection pressures of 0.4 and 0.6 bar were found to reasonably compensate for the wake profile created by the injection system, thus creating streamwise velocity profiles as anticipated from turbulent flows. Further results were presented from the developing region of the channel flow. These showed more complicated forms of the streamwise velocity profile, and also indicated that the wake is more persistent for higher Reynolds numbers.

10.1.6. Investigation of analysis techniques

Simulated particle distributions were created in order to validate and provide comparison data for the box counting method, Voronoi analysis method, and Radial Distribution Functions. These were created using uniform and Gaussian distributions, and using non volume excluded and volume excluded formulations.

The box counting method for quantifying preferential accumulation should return values of σ - σ_P/λ which are close to zero for a random set of particles. This was found at all number

densities and for all box sizes used, with some deviation due to lack of statistical convergence. The grid placement and averaging method were tested to determine the most accurate method, with a centralised grid and combined statistical PDF approach found to be the most appropriate.

Other particle distributions were tested using the box counting method. For uniform distributions with an excluded volume, values of the statistical deviation, D, are downshifted in comparison to a non volume excluded case. All calculated values are found to be negative, with an increase in magnitude as the volume fraction of virtual particles increases. Clearly, the presence of an excluded volume impacts the order of particle placement, acting to create a distribution which is more ordered than random Poisson placement. The curves of D against box size are found to roughly collapse when dividing by the square root of the particle volume fraction. Gaussian distributions were also tested. There is some observable fluctuation as box size is changed, and values no longer hold close to zero due to the large-scale structure of the particle distribution.

Verification statistics for the Voronoi analysis method were calculated using the various simulated particle distributions. PDFs of normalised cell size area were created for each particle distributions. Standard deviations for these PDFs were found to hold roughly constant for non volume excluded distributions. For distributions with excluded volumes there is a linear decrease in standard deviation with particle volume fraction. Again, this indicates the imposition of a more ordered structure with particle addition under these conditions.

PDFs calculated from random distributions were found to be in close match to Gamma function relationships posited in literature. Those with excluded volumes typically are narrower, which is mostly accounted for by a lowering of probability in the smallest cell areas, which correspond to cluster cell areas. Since these most closely represent the PDFs found experimentally, these volume excluded PDFs are used, with appropriate matching of volume fraction.

The mechanisms for computing radial distribution functions of particle positions are considered. Radial distribution functions were calculated for simulated particle distributions. For random particle distributions, g(r_i) is close to 1 at all radial separations r_i.
RDFs are also calculated for other particle distributions. For volume-excluded uniform distributions, there is a decrease to 0 at the lowest radial separations. For Gaussian distributions particle sets, there is a preference to find particles at closer distances, which monotonically decreases to the largest radial separations measured.

10.1.7. Key experimental results - two dimensional channel rig

Data from the vertical channel experiments was analysed, with a focus on particulate concentration effects. This included analysing whole-field clustering and cluster structure, using a variety of different approaches.

The box counting method showed that clustering of particles occurred for all experimental conditions. Data from different distances from the injection point was characterised. For all distances from the injection point, there exists a negative linear correlation between the characteristic size of peak clustering, and the particulate concentration. Therefore, it appears that, as particles are added to the flow, clusters form into tighter groups. The strength of this peak clustering was characterised as growing with particulate concentration just downstream from the injection point. This difference in behaviour is likely caused by the randomisation of particle velocity vectors caused by the injection mechanism.

The Voronoi analysis technique was initially applied as a competing method for assessing the global clustering parameters. The key difference to the box counting method was seen in results from just downstream of the injection point, where clustering observes a generally upwards trend with increasing particle concentration. The Voronoi technique was able to characterise the structure of clusters by using their geometrical properties, which resulted in the identification of four different types of cluster shape, which were identified based on the relationship between perimeter and area. It was then shown that there are differences in local cluster structure, which can be broken down into regions of large and small clusters.

With regards to particle clustering, the final technique employed was a nearest-neighbour based formulation which assigned potentials to each particle based on separation distance. Profiles of mean separation with angle were also calculated. These indicated a preference towards angles which suggest drafting of particles within wakes. In the particle bulk the typical length separation determines mean separation most accurately; in the centre of

particle clusters, more complicated length scales such as those dictated by the underlying fluid may come into play.

It was found that there is a peak at the intermediate number of cluster neighbours for a central cluster particle; this is indicative of our typical understanding of heavier-than-fluid clustering, whereby particles form 'strands' in the high strain regions. As expected, the majority of cluster particles were found in the bulk with a full complement of nearest neighbours; with increasing concentration, this becomes increasingly likely, so most clusters are drawn towards the bulk.

Times of collision were found to heavily depend on concentration of particles, with realistic imminent values ranging between 0.5-3%. RDF peak values varied between 1.5-3, depending mostly upon the distance from the injection point.

10.1.8. Hydrogels and refractive index matching

The suitability of hydrogel particles for use in the circulating dispersions rig was investigated by using an optical transmission experiment. These particles are theoretically ideally refractive index matched with water, and should therefore provide an ideal candidate dispersed phase system. However, the optical transmission of these suspensions was found to be poor, especially in comparison to the PMMA in sodium iodide used as a comparison. Microscopic images of the hydrogel particles used indicate that there was additional structure to the interior of these particles, which may have formed some part in increasing the level of scatter.

Particle-liquid pairs which have commonly been used in this field of research were reviewed. A mixture of pyrex and potassium thiocyanate was chosen to as the best candidate for the circulating dispersions rig. Pyrex particles are considered to have good optical quality, and can be purchased in large quantities, making them a suitable dispersed phase candidate. Potassium thiocyanate was considered superior to sodium iodide; its refractive index has a lower concentration dependence, it discolours less easily in oxygen, and it carries a lower health and safety risk. Empirical relationships were outlined for the refractive index of KSCN, which enabled concentration and temperature to be adjusted to maximise refractive index matching.

10.1.9. Circulating dispersions rig

For the circulating dispersions rig, a pump was chosen to provide reliable and safe delivery of liquid mixed with particles, and the tank was designed in order to provide adequate suspension of particles. Experimentally, investigations were made over a range of Reynolds numbers and particle volume fractions.

A single phase characterisation of the experimental facility was performed. In order to maximise optical visibility, the laser light sheet was placed off-centre of the pipe, so that the distance for light to travel was reduced. Therefore, results are only presented for 0.66-1.0R.

Axial velocity profiles do not show a decent collapse when normalising by the average fluid velocity. This may be a consequence of unaccounted for fluctuations in the flow. The turbulence intensity in the axial direction is calculated, and appears to be converging towards 5% in the bulk of the flow; at the wall, values reach as high as 11%. Finally, there is a comparison between flow profiles for consecutive measurements along the pipe. At the lowest Reynolds number, the deviations are typically around 0.5%, and are randomly fluctuating either side of the initial distribution. For the higher Reynolds number, these fluctuations are twice as large, and typically seem to be more systematic.

10.1.10. Key experimental results - circulating dispersions rig

Many of the results from this experiment differ to those reported in Chapter 7. In particular, the effects of concentration are generally not commutable.

The box counting method showed peak clustering strengths which are comparable with other experiments. Peak clustering lengthscale is much smaller than in the vertical channel experiments, typically falling between 2-4 particle diameters. This is most likely due to the higher Reynolds numbers in these experiments, which promote fluid length scales of different magnitudes, which act to create small and focused clusters. There is no correlation of peak clustering with concentration, although there are changes in values reported between experiments.

The closing times between particle pairs were reported in terms of probabilities of collision. For the time on screen of 0.025s, the probability of collisions 15-20%. These are reasonably high values, indicating that multiple collisions would be observed in these experiments.

These results also found that, for particle concentration of 0.75%, collision probabilities were low; for all other concentrations, probabilities appear to have reached a plateau.

Radial distribution functions appear to confirm results from the box counting method. There are high values for g(r) at the lowest values of radial separation, r, which drop off rapidly to indicate no preference in radial separation for values of r > 6D, where D is the particle diameter. These results indicate a low characteristic clustering length scale, which is actually reasonably strong given the number of particles which can be involved. As with the box counting data, there is little obvious correlation between the strength of clustering in particle pairs and the bulk concentration.

Local particle potential results indicate that there is less variation in mean separation profile depending upon the particle concentration. For concentrations of 0.75%, mean separations are larger, but there a fewer deviations for higher concentrations. Again, this is reminiscent of the collision probability results. Normalisation with characteristic separation was not attempted, since length scale in this case is most likely determined by the fluid motions.

Connectivity analysis showed that, as anticipated for clusters forming along strain regions, there is a preference towards an intermediate number of cluster neighbours. There is also a higher probability of cluster particles still having void regions, due to the spread of particles into parallel planes.

The three-dimensional nature of these experiments may have contributed to some surprising results; loss of particles to other planes of the flow, and the poor optical quality of the mixture, were all factors which may have resulted in incomplete collection of information.

10.2. Future work

10.2.1. Achieving higher particle phase fractions

Whilst any information on four-way coupled dispersed phase flows is valuable, further interest in this field will lie in pushing the boundaries of what can be achieved in terms of dispersed phase fraction. Up to date, most research has only been able to achieve dispersed phase fractions on the order of 1%; in this thesis, the highest measured suspension was 3%.

It would be desirable to increase this value, in order to observe suspensions behaviour as particle-particle interactions become increasingly common and influential.

In particular, this will be useful for making comparisons of very bubbly flows, where dispersed phase fraction can reach up to 30% in industrial components. Whilst it is no more difficult to model these situations in CFD simulations, except perhaps for time of computation, experimentally this is going to be difficult to achieve.

One possibility lies in creating experiments which use liquid-liquid or gas-liquid dispersed phase systems to create this validation data. Certainly the match between mechanics would be more exact; for example, polydisperse 'particulate' size and break-up/combination of 'particles' will be included in the experimental data.

An additional experimental advantage is that it will not be difficult to control the optical quality of each phase. Since most components in these systems can be made to suit in the laboratory where the tests are conducted, one can ensure there are few impurities, and that the refractive index is tuned exactly as required. Additionally, there is a reasonable guarantee that one fluid droplet with have the same properties as another.

However, there are disadvantages to these types of experiments. Due to the fluctuating boundaries inherent in such flows, light scattering can be very large, depending upon the conditions for total internal reflection. Therefore, images taken may contain areas of unexpectedly high intensity, which can cause problems when attempting to automate testing.

A further practical issue is caused by the break-up/coalescence of particles. The tracking codes in this thesis do not require any assumption of such events, because they are impossible. However, if this was to be considered, each image frame would have to be analysed to determine sizes of each 'particle', to determine there is consistency between frames. This may add significant computational time to an already lengthy process.

The other route to be considered is continuing with solid-liquid dispersions. As can be seen, there is some promise in taking this route, but a few changes may need to be made in order to further increase optical visibility.

One perceived problem with this experiment was that, due to continual pumping and agitation, the beads used were being damaged. Such damage causes dents and scratches on the surface, lower the optical visibility. In order to mitigate for this, an experiment with a gentler pumping mechanism would have be devised. Alternatively, the possibility of regular replacement of the dispersed phase would have to be considered.

Another issue is the optical quality of the batch being sourced for these experiments. Generally, to attain a supply of many kilograms of particles, it is necessary that there are some differences in properties across the particle in the batch. Included in this may be the refractive index of the particles, which are usually only quoted to the second decimal place by any manufacturer. If there is non-zero width in the distribution of refractive indices, this means that it is impossible to refractive index match every particle to the underlying fluid, regardless of temperature control. Those particles which are not matched will scatter light significantly even with RI differences of around 0.002, and at a critical quantity this scatter will reduce the contrast of the captured images, making particle detection too difficult to achieve.

However, it may be possible to develop a manufacturing technique for these types of particles which narrows down the range of refractive indices. If this can be achieved, a source of particles such as this could be used much more successfully in a refractive index matched campaign such as this one.

A final consideration is the size of particle used. To achieve a fixed volume fraction of particles at a smaller particle diameter requires a greater number of particles. Therefore, making measurements of finer particles in suspensions carries with it the necessary requirement that transmitted light must traverse more boundaries. However, as discussed particles cannot simply be made very large, because they may not interact or respond to the fluid motions. Any future experiments should take into account this trade-off before considering the parameters of the particle size distribution to be used.

10.2.2. Further data analysis

One area of interest would be to use dispersed phase data, such as Voronoi cell assignments, taken with highly resolved temporal resolution, to track cluster structures. Combined with the simultaneously captured continuous phase velocity maps, correlations

between clusters and the continuous phase motions could be analysed. Furthermore, it could be seen how long clusters persist for within the fluid motion. This work was not attempted due to the intensive computational effort which would be required in not only identifying flow structures, but also tracking them across multiple frames. Nevertheless, it could easily be within the scope of a dedicated project.

It would also be of interest to conduct some of these experiments using a high resolution lens. Whilst this would not capture information from the whole of the fluid flow, a high temporal resolution could be used to observe particles as they pass the point of measurement. With such images, it would be possible to see detailed motions of fluids around the phase boundaries, and with some manual identification it would also be possible to present detailed breakdowns of particle-particle collisions. This would need to be undertaken as a separate campaign, because the analysis and acquisition of data is much different in this scenario to those presented in this thesis.

10.3. Overall conclusions

This body of work provides an insight into some of the clustering phenomena which govern particles in turbulent dispersed flows, particularly at a high dispersed phase fraction at which particle-particle collisions become more commonplace (four-way coupling). Experimental results of this type, and particularly at concentrations of greater than 1%, are very hard to come by, but are of utmost importance for validating models in which dispersed phase elements are likely to collide and coalesce.

The work has produced a library of analysis tools and routines which will enable future researchers to continue working in this field. In particular, tracking algorithms and phase discrimination techniques have been characterised for their performance, and developed to allow for use with other sets of results.

Additionally, analysis techniques specific to this field have either been investigated and developed, or entirely devised from the ground up. With this library of tools, it should therefore be possible to perform further analysis of particle behaviour.

An output of the work is a large library of particle and fluid velocity data, which can be used in a multitude of different ways. The results presented in this thesis only scratch the surface

of what might be possible with these results, and these can therefore be provided in tabulated or raw image form to any other researchers wishing to extract further data which suits their own needs.

In particular, it will be possible to apply this data to relevant CFD models, as per the project aims. Clustering data will provide useful validation of CFD simulations, and particle-pair data will provide inputs towards collision modelling. This will enable their performance to improve by providing validation inputs, and therefore create more powerful tools for understanding a variety of environmental and industrial fluid flows.

References

Agrawal, Y. K., et al. (2015). "Turbulence statistics in solid-liquid flow using a refractive index matching and particle image velocimetry technique." <u>10th Pacific Symposium on Flow Visualization and Image Processing</u>.

Albright, L. F. (2009). "Albright's Chemical Engineering Handbook."

Aliseda, A., et al. (2002). "Effect of preferential concentration on the settling velocity of heavy particles in homogeneous isotropic turbulence." <u>Journal of Fluid Mechanics</u> 468.

Anselmet, F., et al. (2009). "Axial development of the mean flow in the entrance region of turbulent pipe and duct flows." <u>Comptes Rendus Mécanique</u> 337(8): 573-584.

Bagchi, P. and S. Balachandar (2003). "Effect of turbulence on the drag and lift of a particle." <u>Physics of Fluids</u> 15(11): 3496-3513.

Bagchi, P. and S. Balachandar (2004). "Response of the wake of an isolated particle to an isotropic turbulent flow." <u>Journal of Fluid Mechanics</u> 518: 95-123.

Bai, K. and J. Katz (2014). "On the refractive index of sodium iodide solutions for index matching in PIV." <u>Experiments in Fluids</u> 55(4).

Balachandar, S. and J. K. Eaton (2010). "Turbulent Dispersed Multiphase Flow." <u>Annual</u> <u>Review of Fluid Mechanics</u> 42(1): 111-133.

Banitalebi, B., Amiri, H. (2008). "An Improved Nearest Neighbor Data Association Method for Underwater Multi-Target Tracking." <u>IEEE Workshop</u>.

Bao, Y., et al. (2018). "Effects of rotational speed and fill level on particle mixing in a stirred tank with different impellers." <u>Chinese Journal of Chemical Engineering</u> 26(6): 1383-1391.

Birzer, C. H., et al. (2012). "The influences of particle mass loading on mean and instantaneous particle distributions in precessing jet flows." <u>International Journal of Multiphase Flow</u> 41: 13-22.

Bosse, T., et al. (2006). "Small particles in homogeneous turbulence: Settling velocity enhancement by two-way coupling." <u>Physics of Fluids</u> 18(2): 027102.

Brucato, A., et al. (1998). "Particle drag coefficients in turbulent fluids." <u>Chemical</u> <u>Engineering Science</u> 53(18): 3295-3314.

Burton, T. M. and J. K. Eaton (2005). "Fully resolved simulations of particle-turbulence interaction." <u>Journal of Fluid Mechanics</u> 545(-1): 67.

Byron, M. L. and E. A. Variano (2012). "Refractive-index-matched hydrogel materials for measuring flow-structure interactions."

Chung, J. N. and T. R. Troutt (1988). "Simultation of particle dispersion in an axisymmetric jet." Journal of Fluid Mechanics 186: 199-222.

Crowe, C. T. e. a. (2006). "Multiphase Flow Handbook."

Dao, R. D., Yates, M., Hann, D.B. (2015). "Refractive Index Matching of Varying Concentrations and Temperatures of Liquid-Particles Suspensions and a Review of Their Safety and Cost."

de Jong, J., et al. (2010). "Measurement of inertial particle clustering and relative velocity statistics in isotropic turbulence using holographic imaging." <u>International Journal of</u> <u>Multiphase Flow</u> 36(4): 324-332.

Dean, R. B. (1978). "Reynolds number dependence of skin friction and other bulk flow variables in two-dimensional rectangular duct flow." J. Fluids Engineering 100(215-233).

Doherty, J., et al. (2007). "The development of turbulent pipe flow." <u>16th Australasian</u> <u>Fluid Mechanics Conference</u>.

Elghobashi, S. (1994). "On predicting particle-laden turbulent flows." <u>Applied Scientific</u> <u>Research</u> 52(4): 309-329.

Fessler, J. R., et al. (1994). "Preferential concentration of heavy particles in a turbulent channel flow." <u>Physics of Fluids</u> 6(11): 3742-3749.

García-Villalba, M., et al. (2012). "DNS of vertical plane channel flow with finite-size particles: Voronoi analysis, acceleration statistics and particle-conditioned averaging." <u>International Journal of Multiphase Flow</u> 46: 54-74.

Geiss, S., et al. (2004). "Investigation of turbulence modification in a non-reactive twophase flow." <u>Experiments in Fluids</u> 36(2): 344-354.

Gijsen, F. J. H. e. a. (1996). "Analysis of the Axial Flow Field in Stenosed Cartoid Artery Bifurcation Models - LDA Experiments." <u>Journal of Biomechanics</u> 29(11): 1483-1489.

Good, G. H., et al. (2014). "Settling regimes of inertial particles in isotropic turbulence." Journal of Fluid Mechanics 759.

Gualtieri, P., et al. (2012). "Statistics of particle pair relative velocity in the homogeneous shear flow." <u>Physica D: Nonlinear Phenomena</u> 241(3): 245-250.

Hassan, Y. A. and E. E. Dominguez-Ontiveros (2008). "Flow visualization in a pebble bed reactor experiment using PIV and refractive index matching techniques." <u>Nuclear</u> <u>Engineering and Design</u> 238(11): 3080-3085. Holtzer, G. L. and L. R. Collins (2002). "Relationship between the intrinsic radial distribution function for an isotropic field of particles and lower-dimensional measurements." Journal of Fluid Mechanics 459.

Jafari, R., et al. (2012). "Characterization of Minimum Impeller Speed for Suspension of Solids in Liquid at High Solid Concentration, Using Gamma-Ray Densitometry." <u>International Journal of Chemical Engineering</u> 2012: 1-15.

Jarai-Szabo, F. and Z. Neda (2007). "On the size-distributon of Poisson Voronoi cells." <u>Physics A</u> 385: 518-526.

Kajishima, T., Takiguchi, S. (2002). "Interaction between particle clusters and particleinduced turbulence." <u>International Journal of Heat and Fluid Flow</u> 23: 639-646.

Kiger, K. T. and C. Pan (1999). "Two-Phase PIV for Particle-Laden Flows." <u>Turbulence on</u> <u>Shear Flow-1</u>: 843-848.

Kuhn, H. W. (2009). "The Hungarian Method for the Assignment Problem." <u>50 Years of</u> <u>Integer Programming 1958-2008</u>.

Kulick, J. D., et al. (1994). "Particle response and turbulence modification in fully developed channel flow." <u>Journal of Fluid Mechanics</u> 277: 109-134.

Kumaresan, T. and J. B. Joshi (2006). "Effect of impeller design on the flow pattern and mixing in stirred tanks." <u>Chemical Engineering Journal</u> 115(3): 173-193.

Kussin, J. and M. Sommerfeld (2002). "Experimental studies on particle behaviour and turbulence modification in horizontal channel flow with different wall roughness." <u>Experiments in Fluids</u> 33(1): 143-159.

Lau, T. C. W. and G. J. Nathan (2017). "A method for identifying and characterising particle clusters in a two-phase turbulent jet." <u>International Journal of Multiphase Flow</u> 88: 191-204.

Laufer, J. (1953). "The structure of turbulence in fully developed pipe flow." <u>NACA</u>.

Lázaro, B. J. and J. C. Lasheras (1989). "Particle dispersion in a turbulent, plane, free shear layer." <u>Physics of Fluids A: Fluid Dynamics</u> 1(6): 1035-1044.

Lecordier, B. and M. Trinite (2003). "Advanced PIV algorithms with image distortion -Validation and comparison from synthetic images of turbulent flow." <u>5th International</u> <u>Symposium on Particle Image Velocimetry</u>.

Lien, K., et al. (2004). "The Entrace Length for Fully Developed Turbulent Channel Flow." <u>15th Australian Fluid Mechanics Conference</u>.

Liu, X., et al. (2017). "Distribution characteristics of inertial sediment particles in the turbulent boundary layer of an open channel flow determined using Voronoï analysis." <u>International Journal of Sediment Research</u> 32(3): 401-409.

Longmire, E. K. and J. K. Eaton (1992). "Structure of a particle-laden round jet." <u>Journal of Fluid Mechanics</u> 236: 217-257.

Lutteke, F., Zhang, X., Franke, J. (2012). "Implementation of the Hungarian Method for object tracking on a camera monitored transportation system." <u>Robotik 2012</u>.

Mazzaferri, J., et al. (2015). "Adaptive settings for the nearest-neighbor particle tracking algorithm." <u>Bioinformatics</u> 31(8): 1279-1285.

Mehta, M. e. a. (2007). "Particle velocities in the rotating impeller of a slurry pump." <u>5th</u> <u>Joint ASME/JSME conference proceedings</u> 1: 369-378.

Melling, A. (1997). "Tracer particles and seeding for particle image velocimetry." <u>Meas.</u> <u>Sci. Technol.</u> 8: 1406-1416.

Merle, A., et al. (2005). "Forces on a high-Reynolds-number spherical bubble in a turbulent flow." Journal of Fluid Mechanics 532: 53-62.

Meyer, F. (1993). "Topographic distance and watershed lines." <u>Signal Processing</u> 38: 113-125.

Miessner, U., et al. (2006). "Experimental investigation of an oil droplet colliding with an oil-water interface." <u>13th Int. Symp. on Applications of Laser Techniques to Fluid</u> <u>Mechanics</u>.

Monchaux, R., et al. (2010). "Preferential concentration of heavy particles: A Voronoï analysis." <u>Physics of Fluids</u> 22(10): 103304.

Montante, G., et al. (2012). "Particle settling velocity in slurry stirred vessels." <u>16th Int.</u> <u>Symp on Applicatons of Laser Techniques to Fluid Mechanics</u>.

Narrow, T. L. Y., M.; Abdel-Khalik, S.I. (2000). "A simple model for the refractive index of sodium iodide aqueous solutions." <u>Experiments in Fluids</u> 28: 282-283.

Nobach, H. and E. Bodenschatz (2009). "Limitations of accuracy in PIV due to individual variations of particle image intensities." <u>Experiments in Fluids</u> 47(1): 27-38.

Nobach, H. and M. Honkanen (2005). "Two-dimensional Gaussian regression for sub-pixel displacement estimation in particle image velocimetry or particle position estimation in particle tracking velocimetry." <u>Experiments in Fluids</u> 38(4): 511-515.

Obligado, M., et al. (2015). "Experimental detection of superclusters of water droplets in homogeneous isotropic turbulence." <u>EPL (Europhysics Letters)</u> 112(5): 54004.

Obligado, M., et al. (2014). "Preferential concentration of heavy particles in turbulence." Journal of Turbulence 15(5): 293-310.

Patil, V. A. and J. A. Liburdy (2010). "Refractive index matching with distortion measurements in a bed of irregularly packed spheres." <u>Proceedings of the AMSE 2010</u> Joint US-European Fluids Engineering Summer Meeting.

Patil, V. A. and J. A. Liburdy (2012). "Optical measurement uncertainties due to refractive index mismatch for flow in porous media." <u>Experiments in Fluids</u> 53(5): 1453-1468.

Perry, A. E. and C. J. Abell (1978). "Scaling laws for pipe-flow turbulence." <u>Journal of Fluid</u> <u>Mechanics</u> 67: 257-271.

Poelma, C., et al. (2007). "Particle–fluid interactions in grid-generated turbulence." Journal of Fluid Mechanics 589.

Rabencov, B. and R. van Hout (2015). "Voronoi analysis of beads suspended in a turbulent square channel flow." <u>International Journal of Multiphase Flow</u> 68: 10-13.

Rogers, C. B. and J. K. Eaton (1991). "The effect of small particles on fluid turbulence in a flat-plate, turbulent boundary layer in air." <u>Physics of Fluids A</u> 3.

Rudoff, R. C. and W. D. Bachalo (1988). "Measurement of droplet drag coefficients in polydispersed turbulent flow field." <u>Aerosp. Sci. Meet.</u> 29th.

Sahbani, B., Adiprawita, W. (2016). "Kalman Filter and Iterative-Hungarian Algorithm Implementation for Low Complexity Point Tracking as Part of FastMultiple Object Tracking System." <u>IEEE 6th International Conference on System Engineering and Technology</u>.

Satake, S., et al. (2014). "Three-dimensional flow measurement of a sphere-packed pipe by a digital hologram and refractive index-matching method." <u>Fusion Engineering and</u> <u>Design</u> 89(7-8): 1064-1067.

Sawford, B. (2001). "Turbulent relative dispersion." <u>Annual Review of Fluid Mechanics</u> 33: 289-317.

Scott, S. J., et al. (2009). "On the quantification of preferential accumulation." International Journal of Heat and Fluid Flow 30(4): 789-795.

Shindler, L., et al. (2010). "Spatial-temporal improvements of a two-frame particle-tracking algorithm." <u>Measurement Science and Technology</u> 21(11): 115401.

Shokri, R., et al. (2017). "Investigation of particle-laden turbulent pipe flow at high-Reynolds-number using particle image/tracking velocimetry (PIV/PTV)." <u>International</u> <u>Journal of Multiphase Flow</u> 89: 136-149. Squires, K. D. and J. K. Eaton (1990). "Particle response and turbulence modification in isotropic turbulence." <u>Physics of Fluids A: Fluid Dynamics</u> 2(7): 1191-1203.

Squires, K. D. and J. K. Eaton (1991). "Preferential concentration of particles by turbulence." <u>Physics of Fluids A: Fluid Dynamics</u> 3(5): 1169-1178.

Sumbekova, S., et al. (2017). "Preferential concentration of inertial sub-Kolmogorov particles: The roles of mass loading of particles, Stokes numbers, and Reynolds numbers." <u>Physical Review Fluids</u> 2(2).

Sundaram, S. and L. R. Collins (1997). "Collision statistics in an isotropic particle-laden turbulent suspension. Part 1. Direct numerical simulations." <u>Journal of Fluid Mechanics</u> 335: 75-109.

Suzuki, Y., et al. (2000). "Simultaneous measurement of fluid and dispersed phases in a particle-laden turbulent channel flow with the aid of 3-D PTV." <u>Experiments in Fluids</u>: 185-193.

Tagawa, Y., et al. (2012). "Three-dimensional Lagrangian Voronoi analysis for clustering of particles and bubbles in turbulence." <u>Journal of Fluid Mechanics</u> 693: 201-215.

Tanaka, T. and J. K. Eaton (2010). "Sub-Kolmogorov resolution partical image velocimetry measurements of particle-laden forced turbulence." Journal of Fluid Mechanics 643: 177.

Traczyk, M. (2016). "Numerical computations of liquid-liquid dispersions at high phase fractions - PhD Thesis."

Vinuesa, R., et al. (2014). "New insight into flow development and two dimensionality of turbulent channel flows." <u>Experiments in Fluids</u> 55(6).

Wang, L. P. and M. R. Maxey (1993). "Settling velocity and concentration distribution of heavy particles in homogeneous isotropic turbulence." <u>Journal of Fluid Mechanics</u> 256: 27-68.

Weitzman, J. S., et al. (2014). "On the use of refractive-index-matched hydrogel for fluid velocity measurement within and around geometrically complex solid obstructions." <u>Experiments in Fluids</u> 55(12).

Westerweel, J., et al. (2013). "Particle Image Velocimetry for Complex and Turbulent Flows." <u>Annual Review of Fluid Mechanics</u> 45(1): 409-436.

Westerweel, J. and F. Scarano (2005). "Universal outlier detection for PIV data." <u>Experiments in Fluids</u> 39(6): 1096-1100.

White, F. M. (2003). "Fluid Mechanics 5th Edition."

Wicker, R. B. and J. K. Eaton (2001). "Structure of a swirling, recirculating coaxial free jet and its effect on particle motion." <u>International Journal of Multiphase Flow</u> 27: 949-970.

Wiederseiner, S., et al. (2010). "Refractive-index and density matching in concentrated particle suspensions: a review." <u>Experiments in Fluids</u> 50(5): 1183-1206.

Wood, A. M., et al. (2005). "Preferential concentration of particles in homogeneous and isotropic turbulence." <u>International Journal of Multiphase Flow</u> 31(10-11): 1220-1230.

Yang, T. S. and S. S. Shy (2003). "The settling velocity of heavy particles in an aqueous near-isotropic turbulence." <u>Physics of Fluids</u> 15(4): 868-880.

Yuki, K., et al. (2008). "Flow Visualization and Heat Transfer Characteristics for Sphere-Packed Pipes." Journal of Thermophysics and Heat Transfer 22(4): 632-648.

Zanoun, E. S., et al. (2009). "A Study on Flow Transition and Development in Circular and Rectangular Ducts." Journal of Fluids Engineering 131(6): 061204.

Zeng, L., et al. (2007). "Interactions of a stationary finite-sized particle with wall turbulence." <u>Journal of Fluid Mechanics</u> 594.

Zeng, L., et al. (2010). "Wake response of a stationary finite-sized particle in a turbulent channel flow." <u>International Journal of Multiphase Flow</u> 36(5): 406-422.

Appendix A

Conference Presentations

- 68th Annual Meeting of the APS Division of Fluid Dynamics, 2015
- 9th International Conference on Multiphase Flow, 2016

Conference paper

 Investigation of particle cluster in two-dimensional particulate channel flows. Yates,
M., Hann, D.B., Hewakandamby, B., Azzopardi, B., Rhyne, L. 9th World Conference on Experimental Heat Transfer, Fluid Mechanics and Thermodynamics.

Appendix B. Box counting method

An alternative approach to centralised grid placement is to utilise a grid which uses the same x and y start points for each calculation. This will be referred to as a corner-anchored gridding method. An example of this is shown in Fig B.1.



Figure B.1: Corner-anchored gridding method, shown for a randomly-generated particle set. In this example, the grid size is 55 pixels, with the generated grid being 10 x 14 in size. This method leaves gaps of 50 on one side in the x-direction, and 30 on one side in the y-direction

An ensemble of two hundred images was used to gather the statistics for this method of gridding. These are shown for a few values of particle number per image frame, N, in Fig. B.2. These results are similar to those using the centralised method of grid placement, so this aspect does not affect the statistics, provided that behaviour is identical across the region being analysed.



Figure B.2: Comparison of deviation factor, D against cell size, for multiple values of N, particles per image frame. Shown for reference is the expected value of a Poisson process for this type of analysis (black dotted line)



Figure B.3: Comparison between two gridding methods using identical data, for N = 2000 particles per image frame

There are two available options for performing the calculation to obtain D. The optioned used in this thesis is to sum all the quantities across all of the frames to be analysed. This means that particles accounted for, grids used, and histogram data are all collected together from an ensemble of particle images. These may then be used to calculate the Poisson predicted properties of the whole sample, and the variance of the experimental histogram. At this point, one calculation is performed in order to find D.

The second method is to perform an individual calculation of deviation factor, D, for each frame of data. Therefore, for each set of particles, a histogram is calculated, with the predictive Poisson values and experimental standard deviation calculated individually for that frame. Therefore, a value of D can be calculated for each frame. These are then used together at the end to find the mean value.

The difference between these two methods was investigated, for both of the gridding methods described above. Firstly, some comparisons for the centralised grid method are presented. For N = 200, there is a small difference between the values of D obtained by these methods, with method 1, called the summation method, producing slightly larger values. However, once concentration has been increased to N = 2000, the two methods appear to converge.



Figure B.4: Comparison of averaging methods to obtain D, for two different values of N, using a centralised gridding method.

The same analysis was performed for the corner-anchored gridding. In both instances, the same trends are followed as the cell size increases. However, there is a widening gap between the two values of D as cell size increases, and the summation method is generally closer to zero out of the two approaches.



Figure B.5: Comparison of averaging methods to obtain D, for two different values of N, using a corneranchored gridding method

Given the information found so far, it was decided that a summation method should be used, since these values were always closer to zero. With very little difference between the grid placements, it was felt that the best results for experimental distributions would be found by centrally placing the grid, as this would capture the main behaviour of the flow and avoid spurious edge effects.

Appendix C. Monte Carlo performance

Monte Carlo predictions

The quarter circle described in Section 6.5.1 was used to test both accuracy and speed of Monte Carlo computations. The advantage of this situation is that it has a well-known value to which the computations can be compared.

An initial test was performed with the inner radius fixed at 1 A.U., whilst the outer radius was varied. The seeding density was also varied by changing the number of points used in each computation. For each condition, 100 simulations were performed in order to capture means and also standard deviations associated with Monte Carlo computations.

The change in outer radius size was found to have no effect upon the predicted mean and standard deviation values of the calculation, and therefore was not an important consideration for the remainder of the investigation.

Fig. C.1 shows the convergence of the predicted fractional area as the number of points used in the Monte Carlo method is increased. Typically, we see that the predicted value settles to the actual value in the region of 10^3 - 10^4 points used. This behaviour replicates for all of the outer radii used. There is evidently some oscillatory behaviour within this plot; this may have been damped out by using a larger sample size than one hundred annular elements.



Figure C.1: Progression of predicted fractional area for increases in the number of simulated points, N, used per annular element, for R_{max} = 2

Fig. C.2 shows a similar progression, but for a percentage deviation, defined by:

$$\delta = \frac{100 \times (Predicted - 0.25)}{0.25}$$

Equation C.1: Percentage deviation for RDF test case

Fig. C.2 shows behaviour for the two most extreme outer radii values used. As stated, there is repeatability between the two radii, which is evident across the range that was used. An equation for each case can be defined by taking logarithms such that:

$$\log(\langle \delta \rangle) = m \times \log(N) + c$$

Equation C.2: Linear line fit of percentage deviations

Between the two cases, we can see minimal differences between the calculated gradients, and thus the predicted deviation can be ascertained using the relationships derived. However, the oscillatory behaviour, as a result of randomly simulating particles, means that there is some spread of data, and so this should only be used as a guide.



Figure C.2: Variation of δ , the percentage difference between actual and predicted values of included area, with the number of points N used per annular element. The equation for linear line of best fit in each case is displayed on the graphs, with comparison between $R_{max} = 2$ and $R_{max} = 10$.

In all the cases above, a sample of one hundred Monte Carlo simulations have been used to create an ensemble. This will increase the accuracy, particularly for the low density seeding cases, in comparison to using just one estimation. However, in all likelihood, only one pass would be desirable when analysing experimental data, and so the sigma values of these data sets must also be considered. Of interest is the standard deviation with respect to simulated

points, N. If it is possible to guarantee a satisfactory spread of predictive values for a desired value of N, results can then be stated with reasonable confidence intervals.



Figure C.3: Standard deviations from Monte Carlo simulations using varying number of points, N, per annular element, for R_{max} = 2

For the lowest number of simulated points per annular regions, N = 10, the spread of data σ is around 50% of the actual fractional value, 0.25. The trend can be fitted with a logarithmic relationship. In general, for values of N = 10^3 - 10^4 , errors vary between 1-5%, and this would be considered a lower bound for seeding numbers to provide a reliable estimation of area.

Computational expense of Monte Carlo routines

The relationship between number of simulated points and time of computation is shown in Fig. C.4, for the minimum outer radius used.



Figure C.4: Plot of average time per Monte Carlo computation for variations in number of simulated points per annular element, for R_{min} = 1 A.U. and R_{max} = 2 A.U.

Up until 200 samples per annular element, there is a slow increase in time per computation. As sample number is increased from this, a sharper rise in time of computation can be seen. Linear relationships were fitted to fit the following equation in each case, and these are shown separately in the two different regions of Fig. C.4

$$\log(< T >) = m \times \log(N) + c$$

Equation C.3: Linear line fit to assess computational expense

In Eqn C.3, T represents time of computation, and N represents the number of data points used in each Monte Carlo process. The values for gradient, m, and intercept, c, are shown on the graph for each of the two different regions. For the higher point density region, the gradient is thirteen times larger, which represents a significantly larger growth in computational expense.

Implications with respect to data

To put these findings to use, they must be considered in the context of the experimental results obtained. In an average experimental run, as stated above, in the order of 10^6 particles may be captured. Needless to say, not all instances will require Monte Carlo calculations to be performed; in particular, at low radii of separation, only those particles near to the edge will need to be subjected to this analysis. Nevertheless, to obtain a significantly detailed RDF which captures the effects of radius in sufficient resolution, in the order of 50 annular elements must be applied to each particle position. This means that the total number of calculations required per experimental run could reach up to 5 x 10^7 , if we consider a total number of particles of 10^6 , as indicated previously.

As already stated, not all annular elements will require this level of analysis. Nevertheless, we can obtain an upper estimate for the computational time of the RDF from an experimental run via the equation below:

$T = m \times t \times f \times r$

Equation C.4: Estimation of computation time for complete data set

In Eqn C.4, T represents the total time of computation, m is an average number of particles per image, t is the time for an individual Monte Carlo computation, f is the number of images comprising a single experimental run, and r is the number of annular elements required to make up the total radial distribution function.

The highest number of simulated points per annular element trialled was 5×10^5 , which has a corresponding time of individual computation of 0.053s. If an entire experimental run was analysed using this level of fidelity, the total computational time would run up to 736 hours, which is clearly prohibitively expensive.

By only using 10000 simulated points per annular element (50 times fewer), only 1.8% of the time given above would be required for a full analysis. However, this is still 14 hours per experimental run. Clearly this is still too long for analysis, but further reduction of points per annular element could result in inaccurate measurements.

The other consideration, which may reduce the need for larger computations, is the large amount of data points created by this analysis. For an individual frame containing m

particles, up to ½*m(m-1) particle pairs can be obtained. Given the estimations above of m = 500, this could result in an average of 2500 particles per data bin being collected per frame, when using 50 different radial separations. Therefore, a more sensible collection strategy might be to collect data from a random selection of frames throughout the experimental regime. This would reduce the time of computation, whilst still producing sufficient quantities of data to create statistical convergence towards a defined radial distribution function. By selecting frames at random throughout the duration, any bias introduced by fluctuating flow conditions present at a particular moment in time would then be eliminated.