Particle Image Velocimetry and its application in the study of gas velocities in the freeboard above bubbling fluidised beds

Steven J. L. Rix

Ph.D. thesis, October 1998

The University of Edinburgh

(Department of Chemical Engineering)



Abstract

Particle Image Velocimetry (PIV) has been used to study flow patterns in the **freeboard region** above bubbling **fluidised beds**. Most research was carried out using a 5 m-tall bed, with a square (0.3 m × 0.3 m) cross-section.

A simple, two-dimensional form of PIV was used, to illuminate thin vertical sections in the freeboard. Each photograph allowed instantaneous **gas velocity** measurements to be made across the whole of the field of view, at one moment in time. PIV thus allows transient flow structures to be seen (*e.g.* short-lived vortices and jets of rapid gas flow). Most previous research on flow in the freeboard has used single-point techniques such as Laser Doppler Anemometry, which yield a time series of velocity measurements at a given point.

In this work, PIV recordings have been processed using a digital analysis method. This technique is assessed in detail, using Monte Carlo simulations. An optimum image diameter has been obtained. Recommendations are also given for the minimum and maximum separations of consecutive particle images. Keane & Adrian (1990, 1991) gave recommendations for a different analysis system (using circular interrogation spots, rather than square-sided interrogation areas). The new recommendations are the more appropriate for the PIV method used here.

It was found that relatively large errors can arise in PIV measurements due to autocorrelation wrap-around effects. The analysis parameters and image separations which cause these errors have been investigated. Problems are most likely to occur if image-shifting is used and the digital interrogation areas are autocorrelated using unpadded FFTs. A model was developed, which describes these wrap-around errors very well.

The application of PIV has allowed velocities to be measured in the freeboard, at various heights above the bed surface. Time-average velocities have also been determined, by averaging each series of instantaneous vector maps. Both of these are of interest when carrying out **elutriation** calculations, for the carry-over of fine particles from fluidised beds. Velocities from the individual vector maps are relevant if particles are elutriated rapidly, whereas the time-average velocities might be more appropriate for particles with a longer freeboard residence time.

In the experiments, the greatest upward velocities were sometimes found near the centre of the freeboard, with lower velocities obtained near the walls. This was similar to the results of Horio *et al.* (1980). Dip profiles were also found, however, in which the greatest velocities occurred near the walls, as had previously been reported by Levy & Lockwood (1983). *Both* of these flow regimes could be found above the fluidised bed, according to the conditions of fluidisation. It is not clear what criteria determine the nature of the velocity profile.

Acknowledgements

In the beginning, this was a three-year fluidisation project which happened to use PIV. I would like to thank Denise McCluskey, Tom Bruce and Frank Morris for helping me to get started in the lab, over in the Department of Physics.

After a year or two, a large, unused room in the Sanderson Building was invaded and converted into a fluids lab. Tom Bruce, Bill Easson and (especially) Alastair '*Don't* use Blu-tack' Martin deserve thanks for getting this off the ground.

A larger-scale fluidised bed was designed and built in this lab. Matthew Rea was helpful in the design stage. From the Chemical Engineering Workshop, Kenny Fee, Bobby Hogg and Rab Kilgour have all provided invaluable technical support at various stages during this project (and also good-natured banter about the latest football news). Also from the workshop, Tommy Murray deserves special thanks for constructing the rig.

PIV photographs are of limited use if they cannot be analysed. I would like to thank Andy Hind for allowing the use of his excellent PIV software (both for analysis and afterwards, when examining the results under MATLAB).

Another program by Andy Hind allowed PIV images to be simulated under carefully-controlled conditions. The originally-intended simulation or two soon became a series of simulations, or two, or ten... This eventually accounted for over 100 pages of work in this thesis, none of which was part of the original brief. It is hoped that this work will be useful to other researchers, when setting parameters for PIV experiments. Andy deserves further thanks, both for encouraging this use of PIVSIM and for explaining how best to carry out the simulations.

Dave Laurenson's knowledge of signal processing was a great help, when trying to understand FFT autocorrelation and wrap-around phenomena. His services rendered as a LaTeX guru were also most appreciated. I would like to thank Dave for this, and also for being such an ideal flatmate for five years.

Finally, I would like to thank my supervisors, Don Glass and Clive Greated, for their advice, encouragement and tolerance throughout the rather extended duration of this project.

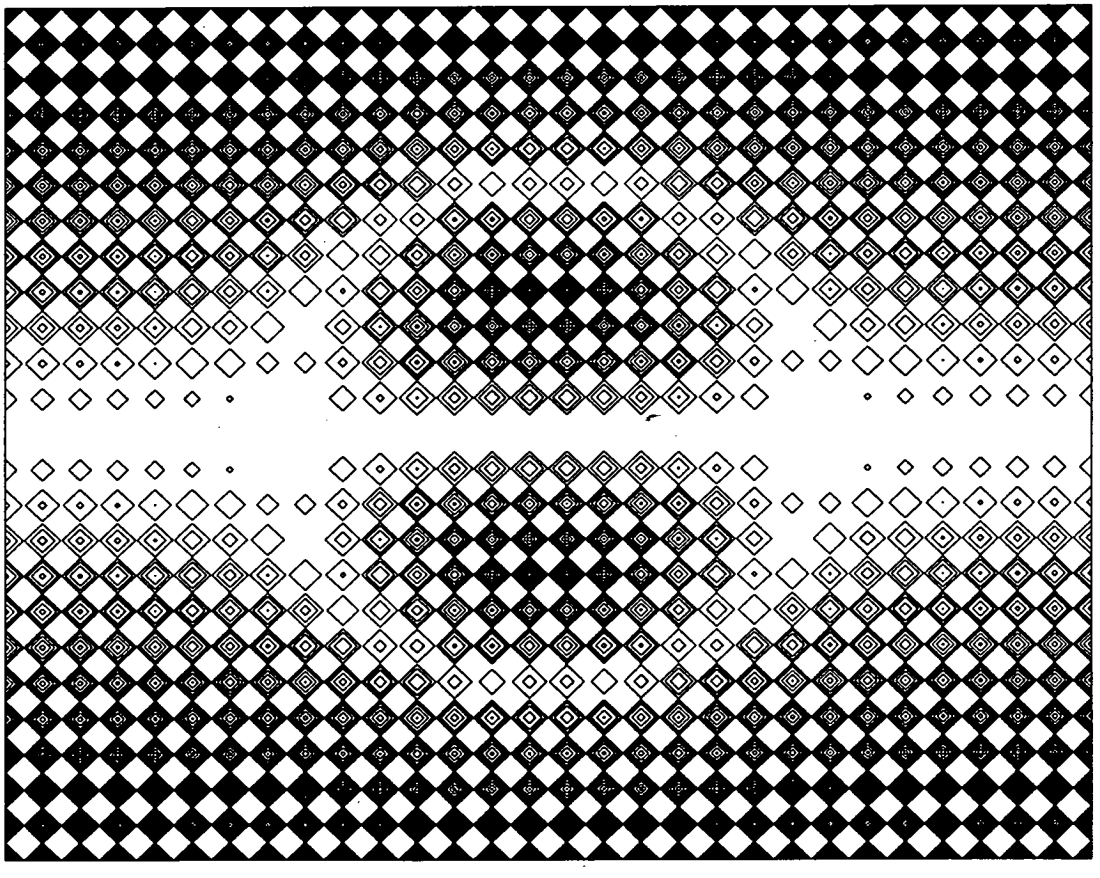
Declaration

This is thesis is based on my own research, unless otherwise attributed. Apart from my supervisors, who made many useful comments along the way, no other person was involved in writing up this research.

No work contained within this thesis has been submitted, in whole or in part, towards any degree of any description at this (or indeed any other) University.

Some of the work in this thesis has already appeared in print, in papers by McCluskey *et al.* (1994) and Rix *et al.* (1996). These papers are appended.

S. J. L. Rix



Contents

1 Introduction 1

2 Fluidisation 7

2.1 The basic principles of fluidisation 7

2.1.1 Apparatus 7

2.1.2 Low-velocity fluid flow 8

2.1.3 Mild fluidisation 10

2.1.4 Particulate fluidisation 11

2.1.5 Aggregative fluidisation and two-phase theory 12

2.1.6 Faster fluidisation 14

2.1.7 Some advantages and applications of fluidised beds 15

2.2 Previous research on fluidisation 16

2.2.1 Minimum fluidising velocity 16

2.2.2 Terminal velocity 18

2.2.3 Bubble diameter 19

2.2.4 Bubble rise velocity 21

2.2.5 Problems with the two-phase theory 24

2.3 Elutriation 25

2.3.1 Particle behaviour in a bubbling fluidised bed 25

2.3.2 The importance of the gas velocity in the freeboard 28

2.4 Selected previous research on elutriation 32

2.4.1 Bubble eruptions 32

2.4.2 Transport Disengaging height 34

2.4.3 Elutriation from tall beds, with fine particles 36

2.4.4 Elutriation from other beds 39

2.4.5 Velocity profiles in the freeboard region 43

2.4.6 Laminar flow in a square-sided conduit 54

3 Stream functions for bubbles in gas-fluidised beds **59**
3.1 Preamble 59
3.1.1 Frames of reference 59
3.1.2 Co-ordinate system 60
3.1.3 Stream functions and velocities 61
3.2 Stream functions for fluidisation bubbles 62
3.2.1 Two-dimensional bubbles, constant volume 63
3.2.2 Three-dimensional bubbles, no growth 70
3.2.3 Growing two-dimensional bubbles 74
3.2.4 Growing three-dimensional bubbles 84
3.3 Ghost bubble vector maps in the laboratory frame 88
3.4 Alternative stream functions 94
3.5 Experiments on fluidisation bubbles 97

4 Methods of flow measurement **99**
4.1 General comments 99
4.1.1 The aims 99
4.1.2 Point measurement techniques 100
4.1.3 Flow visualisation in liquids 101
4.1.4 Constraints on the choice of seeding 102
4.1.5 Particulate seeding in air 104
4.2 Light sheets 104
4.2.1 The expanded beam method 104
4.2.2 A scanning-beam light sheet 106
4.2.3 The geometry of a scanning beam system 107
4.2.4 The intensity profile in the x-direction 109
4.2.5 The intensity profile in the y-direction 110
4.2.6 The intensity profile in the z-direction 111
4.2.7 The number of rotating mirror facets 111
4.2.8 The brightness of scanning-beam and expanded light sheets 112
4.3 Streak photography 113
4.4 Particle Image Velocimetry 116
4.4.1 The advantages of a pulsing light sheet 116
4.4.2 PIV illumination methods 117
4.4.3 Mean particle displacements 119
4.4.4 Young's fringes techniques 120
4.4.5 Digital processing 122
4.4.6 Autocorrelation 124

4.4.7 Maxima of the autocorrelation function 128
4.4.8 Windowing 133
4.5 Setting PIV parameters 134
4.5.1 Laser power 134
4.5.2 The choice of film 135
4.5.3 Magnification 135
4.5.4 Image size 137
4.5.5 Illumination time interval 138
4.5.6 The exposure time 142
4.5.7 Seeding concentration 144
4.5.8 The dynamic range of a PIV system 145
4.6 Image-shifting 146
4.6.1 Problems of vector alignment and dynamic range 146
4.6.2 The use of image-shifting to enhance PIV 148
4.6.3 The pros and cons of image-shifting 150
4.6.4 Minimising errors 156
4.7 Other refinements to PIV 159

5 Assessing the PIV analysis software **161**
5.1 Background 161
5.2 Simulating PIV images 164
5.3 Analysing digital PIV images 165
5.4 Errors in digital PIV analysis when using unpadded FFTs 168
5.5 Results with padded FFTs 176
5.6 The probability of obtaining a valid vector 179
5.7 Other parameters 183
5.7.1 Small vertical displacements 183
5.7.2 Larger vertical displacements 186
5.7.3 Peak location method 190
5.7.4 Image diameter: no zero-padding 197
5.7.5 Image diameter: with full zero-padding 204
5.7.6 Magnification/resolution 205
5.7.7 Interrogation area size 211
5.7.8 Seeding particle concentration 213
5.8 Other simulations using $d_\tau = 2$ pixels 214
5.9 Conclusions 215

6 Loss of pairs and wrap-around errors **219**
6.1 Useful images from horizontal flow records 221
6.2 Second harmonic displacement peaks 230
6.3 The magnitude of errors due to wrap-around 233
6.3.1 The ratio of wrap-around peaks to correct peaks 234
6.3.2 Shapes of autocorrelation peaks 238
6.3.3 Conical peaks 245
6.3.4 Parabolic peaks 247
6.3.5 Gaussian peaks 253
6.4 Useful images from other flow records 259

7 Experimental apparatus **263**
7.1 Initial small-scale work 264
7.2 Later, larger-scale work 270

8 Experimental results with the smaller-scale apparatus **277**
8.1 Fluidised beds of AG grade ballotini 278
8.2 Fluidised beds of Grade 10 ballotini 285
8.2.1 Bed depths of 3 cm 285
8.2.2 Bed depths of 5 cm 298
8.2.3 Ensemble averages 305

9 Experimental results with the larger-scale apparatus **311**
9.1 31 cm bed depth, centre of freeboard 314
9.1.1 Viewing heights of 220 cm 314
9.1.2 Viewing heights of about 140 cm 333
9.1.3 Viewing heights of about 40 cm 344
9.2 31 cm bed depth, quarter of way in 347
9.2.1 Viewing heights of about 220 cm 347
9.2.2 Viewing heights of about 130 cm 352
9.2.3 Viewing heights of about 40 cm 357
9.3 60 cm bed, centre of freeboard 358
9.3.1 Viewing heights of about 190 cm 358
9.3.2 Viewing heights of about 130 cm 365
9.3.3 Viewing heights of about 40 cm 371
9.4 60 cm bed, quarter of way in 380
9.4.1 Viewing heights of about 190 cm 381
9.4.2 Viewing heights of about 130 cm 382
9.4.3 Viewing heights of about 40 cm 382
9.5 Ensemble averages 388

10 Conclusions **397**
10.1 PIV and image-shifting 397
10.2 Fluidisation results 401

11 Recommendations for future work **409**
11.1 PIV work 409
11.2 General fluidisation considerations 410
11.3 The natural convection pattern 411
11.4 Dip profiles and 'ghost bubbles' 411
11.5 Analysing existing results 412

12 Symbols **413**

13 References **419**

List of Figures

2.1 A simple fluidised bed can be formed if a fluid is passed upwards at a superficial velocity U through a bed of particles, of depth H. This flow produces a pressure drop, ΔP. 7

2.2 The pressure drop through a typical fluidised bed, plotted as a function of superficial gas velocity. The two dotted lines intersect at $(0.070, 7.80)$. 10

2.3 The particle concentration profiles in the freeboard of a fluidised bed, as measured by Lewis *et al.* (1962). In this example, freeboard A has a height $H_{f,A}$ which is greater than the TDH, whereas $H_{f,B} < z_{TDH}$. 41

2.4 The main features of the freeboard particle transport model, developed by Kunii & Levenspiel (1991). 42

2.5 The time-averaged vertical particle velocity profile across the freeboard above a cylindrical fluidised bed, measured by Horio *et al.* (1980). 44

2.6 The time-averaged vertical gas velocity profile across the freeboard above a rectangular fluidised bed, measured by Levy & Lockwood (1983). 46

2.7 Sketches of the flow sequence observed by Yórquez & Duursma (1999), when an injected bubble bursts through the surface of an incipiently-fluidised bed. 51

2.8 The fully-developed laminar velocity profile for flow in a square duct with a vertical axis, calculated using the result quoted by Knudsen & Katz (1958). Left: Contours of equal vertical velocity, for points in a 100×100 horizontal cross-section through the duct. Right: The velocity profiles along various lines in the left-hand figure, namely $z = \pm 38$ (dash-dot curve), $z = \pm 25$ (dashed curve), $z = \pm 12$ (dotted curve) and $z = 0$ (continuous curve). 56

3.1 The co-ordinate system around a bubble of radius a, showing the direction of the velocity components v_r and v_θ (left); three closely-spaced points A, B and C (right). 61

3.2 Velocity vectors and streamlines for a constant-volume 2D bubble 67

3.3 The vertical velocity along horizontal sections through a 2-D bubble. These sections run through the bubble centre ($x = 0$, solid curve), half a diameter above or below the bubble centre ($x = \pm 5$, dash-dot curve), just missing the top or the bottom of the bubble ($x = \pm 10$, dashed curve) or two bubble diameters above/below the origin ($x = \pm 20$, dotted curve). 68

3.4 Contours of stream function (dotted curves) around a two-dimensional bubble (inner continuous circle), plotted for various bubble rise velocities: from left to right and from top to bottom, $U_b/U_o = 0.25$, 0.5, 1.01, 1.5, 4 or 8 units s^{-1}, with $U_o = 1$ unit/s and $U_{mf} = 0.5$ units s^{-1} in all cases. 69

3.5 Velocity vectors and streamlines for a constant-volume 3D bubble 72

3.6 Contours of stream function (dotted curves) around a three-dimensional bubble (inner continuous circle), plotted for various bubble velocities. From left to right and from top to bottom, $U_b/U_o = 0.25$, 0.5, 1.01, 1.5, 4 or 8 units s^{-1}. $U_o = 1$ units s^{-1} and $U_{mf} = 0.5$ units s^{-1} in all cases. 73

3.7 Volume balance over an element AA' outside the bubble, which is bounded by the surface BB'. 75

3.8 Left: the various functions $f(r)$: a Bessel function ('+' signs, equation 3.35), a sine function (dotted line, equation 3.36) and a quadratic (continuous line, equation 3.37); Right: the ratio of the Bessel function to the sine function. 78

3.9 The modified sine function $y = f_1(r) = [\sin(\pi\, r/2\, a)]^n$, for various values of n. From top left to bottom right, $n = 1$ (continuous curve), 2 (dash-dot curve), 4 (dashed curve) or 8 (dotted curve). The dotted line plotted is $y = x$. 81

3.10 Streamlines for a growing 2D bubble, using a sine-squared function $f(r)$ inside the bubble and $\gamma = 0.1$. 82

3.11 Velocity vectors and streamlines for a growing 2D bubble, using a sine-squared function inside the bubble and $\gamma = 0.1$. 85

3.12 Velocity vectors and streamlines for a growing 2D bubble, using a sine-squared function inside the bubble and $\gamma = 0.2$. 86

3.13 Velocity vectors and streamlines for a growing 3D bubble ($\gamma = 0.1$) 89

3.14 Velocity vectors and streamlines for a growing 3D bubble ($\gamma = 0.2$) 90

3.15 Velocity vectors and streamlines for a growing 3D bubble ($\gamma = 0.4$) 91

3.16 Velocity vectors for a 3D fluidisation bubble rising without growth ($\gamma = 0$), after a vertical velocity U_b has been added to all vectors. 92

3.17 Velocity vectors for a 3D bubble rising without growth, with a vertical velocity of $0.75\, U_b$ (top) or $0.50\, U_b$ (bottom) added to all vectors. 93

4.1 A light sheet from a cylindrical lens. 105

4.2 A light sheet from a combination of mirrors. 106

4.3 Analysis of the scanning beam method. 107

4.4 The brightness of a light sheet produced by the scanning beam and expanded beam methods. 112

4.5 A comparison between a short streak and a PIV record. 116

4.6 A typical PIV record. 120

4.7 How the extent of particle image overlap changes as a function of increasing autocorrelation displacement. 127

4.8 The self-correlation of Figure 4.6. 128

4.9 Image overlaps at (a) the fundamental velocity peak and (b) the second harmonic peak. It can be seen that these produce eight image overlaps and five image overlaps, respectively. 130

4.10 The complete autocorrelation surface from the analysis of Figure 4.6 and a cross-section of this surface along $\Delta x = 0$ (from bottom left to top right). 132

4.11 The size distribution of displacement magnitudes (in pixels per illumination interval) from a manually-validated PIV record. The array of displacement vectors was 85 × 55, hence a total of 4675 vectors were generated by the analysis. 4521 valid vectors were found (97% of the total). 141

4.12 The size distribution of displacement magnitudes expected in a PIV measurement of laminar and turbulent flow. The mean flow displacement is 17.5 units in magnitude, in both cases. 142

4.13 The maximum and minimum displacements of a particle image: displacements which lie within the shaded region should yield a valid measurement. 145

4.14 The locus of image displacements from two hypothetical experiments. Circle A is well placed, so a good yield of accurate displacement vectors should be obtained. The position of Circle B suggests that there may be regions of low flow velocity, which might not yield satisfactory measurements. 146

4.15 The locus of image displacements obtained from a hypothetical flow system, photographed directly with an illumination time interval of 10 ms. 148

4.16 The locus of particle image displacements obtained from a hypothetical flow system: (a) after the superposition of an image-shifting velocity such that the added displacement $|\boldsymbol{d}_s| = 24$ pixels to the left on the original basis ($\Delta t = 10$ ms); (b) $|\boldsymbol{d}_s| = 24$ pixels to the left on the original basis, but with the illumination time interval now reduced to $\Delta t = 5.56$ ms. 149

4.17 A common arrangement for horizontal image-shifting using a rotating mirror, viewed from above. A camera of focal length F views the light sheet (which is vertical and located a total distance B away) via a mirror placed at 45° to its optical axis. Rotating this mirror causes a stationary seeding particle to move to a different position in the image plane (*i.e.* on the photographic film). 152

4.18 A plot of the theoretical image-shifting displacement on the film, using the work of Raffel & Kompenhans (1995). Six vertical marker vectors have been added to the top left-hand corner of this vector map, for reference. Their magnitudes are 50, 100, 150, 175, 200 and 225 μm. The second-largest vector is similar in magnitude to the mean horizontal displacement. 155

4.19 The deviation of the theoretical image-shifting displacement on the film, as calculated by Raffel & Kompenhans (1995), from the central value. 156

5.1 The average error in the PIVANAL displacement, calculated without padding and plotted as a function of the simulated displacement, for the case of two images of each particle. 168

5.2 An example of the relative intensities of autocorrelation values (shown as degrees of blackness) around a displacement peak at the point (10, 10), which would satisfy the validity criteria imposed by PIVANAL. 169

5.3 The average error in the PIVANAL displacement, calculated without padding for the case of three images of each particle. . . . 171

5.4 Tiling of the data contained in each interrogation area occurs if the data is autocorrelated using Fourier transforms. Copies of the interrogation area in question (numbered 5 and including pixels 1–64 in each direction) are placed around the interrogation area in a 3×3 array. 172

5.5 A sample interrogation area, where particle images are shown as black circles, being autocorrelated using a 64×64 pixel unpadded FFT. Points near the right-hand edge of the interrogation area must be correlated with points from a 'tiled' copy (lightly-shaded images). For this displacement, the particle images A3, B3 and D2 are correlated with the regions shown as unfilled circles. 173

5.6 The error in the displacement calculated without padding, using four or five images of each particle, plotted against the simulated displacement. The dotted vertical lines mark simulated displacements of 64/5, 64/4 and 64/3 pixels. 175

5.7 A plot of the most common output displacement calculated by PIVANAL against simulated displacement. The continuous lines are plots of ordinate = abscissa and ordinate = 64 − abscissa, shown for comparison. 176

5.8 The average error when using fully-padded FFTs as a function of the simulated horizontal displacement, for the cases of two, three, four and five particle images. 177

5.9 The difference in the average horizontal measurements with and without padding for the five-image case, plotted against the simulated horizontal displacement. Data obtained using padded FFTs was subtracted from data obtained without padding. Dotted lines mark values of 64/3, 64/4 and 64/5 pixels. 178

5.10 The fraction of valid vectors determined by the analysis (without padding) of simulated PIV images, plotted against the simulated displacement. Plots are given for simulations using 2–5 images of each seeding particle. 180

5.11 The fraction of valid vectors determined by the analysis (with full zero-padding) of simulated PIV images, plotted against the simulated displacement. Plots are given for simulations using 2–5 images of each seeding particle. 181

5.12 The average error in the measured horizontal displacement, plotted against the simulated horizontal displacement. In these simulations, a constant downward displacement of 0.4, 0.7, 1.0, 1.4, 1.7 or 2.0 pixels was added to a range of horizontal displacements. Three images of each seeding particle were simulated and the autocorrelation analysis took place without padding the FFTs. The vertical dotted lines mark a simulated horizontal displacement of 64/3 pixels. 184

5.13 The average vertical displacement calculated by PIVANAL, plotted against the simulated horizontal displacement, when a downward displacement of 0.0, 0.4, 0.7, 1.0, 1.4 or 1.7 pixels was added to a range of horizontal displacements. Three images of each seeding particle were created and the autocorrelation analysis used unpadded FFTs. The vertical dotted lines mark a simulated horizontal displacement of 64/3 pixels. 185

5.14 The yield of valid vectors obtained by PIVANAL, plotted against the simulated horizontal displacement, when various downward displacements were added to a range of horizontal displacements. Three images of each seeding particle were simulated and the autocorrelation analysis took place with fully-padded (left) or unpadded (right) FFTs. From top to bottom, the curves shown in each graph are for simulated downward displacements of 0 pixels (uppermost solid line), 10 pixels (dotted line) or 20 pixels (dashed line); the lowermost solid curve is for diagonal displacements. 187

5.15 For one quadrant of the displacement plane, these graphs show plots of points at which the yield of valid vectors was 80% (circles) or 90% ($\times$ signs), for the cases of analysis with (left) or without (right) FFT padding. 189

5.16 Analyses using Whittaker's reconstruction (left) or volume moments (right) to determine the location of the peak value of the autocorrelation function. These analyses were obtained using unpadded FFTs (top row) or fully-padded FFTs (central row). The yields obtained when padding was used are shown in the bottom row. 191

5.17 Thirty evenly-spaced simulated horizontal displacements in the range 14.6–17.5 pixels are shown in the middle row, using circles to plot these points on a number line. The displacements calculated by PIVANAL from these simulated displacements are also shown, for analysis with unpadded FFTs using Gaussian interpolation (× characters) or Whittaker's reconstruction ('+' signs). 193

5.18 Plots of the raw data produced by PIVANAL using fully-padded FFTs, for simulated horizontal displacements of 24.5 pixels (left) or 24.6 pixels (right). Either Gaussian interpolation (top) or Whittaker's reconstruction (bottom) was used, in order to determine the location of peaks to sub-pixel accuracy. 194

5.19 The average error in the displacements measured by PIVANAL using unpadded FFTs, plotted against the simulated displacement, for a range of entirely-horizontal displacements. From left to right and top to bottom, the e^{-2} image diameter of each particle image was 1.0, 1.1, 1.2, 1.5, 2.0 or 3.0 pixels. The minimum image separation for accurate results is indicated by an asterisk. 199

5.20 The magnitude of the variable error in the displacement measured by PIVANAL (using unpadded FFTs and 63 × 63 interrogation areas), plotted as a function of the simulated e^{-2} image diameter. A full-scale plot is shown on the left, with a close-up of the lower part of the curve shown on the right. 200

5.21 For measurements calculated by PIVANAL using unpadded FFTs to analyse simulations in which three images of each particle were simulated: Left: the average systematic error in the PIV measurement, plotted as a function of the simulated e^{-2} image diameter; Right: The minimum resolvable displacement. 201

5.22 The possible appearance of a digital PIV recording, when images are 2.0 pixels or 4.0 pixels in diameter and one image diameter apart. 203

5.23 The 'corrected' average error in the measured displacement (calculated with fully-padded FFTs for the case of five images of each 4.0 pixel-diameter particle), plotted against the simulated displacement. The measurements were corrected either by multiplying the average measured displacement by $(1 + d_\tau^2/d_I^2)$ (left) or by adding 0.07 pixels to each average measured displacement (right). 203

5.24 For measurements with either no zero-padding ('+' characters; continuous curve) or full zero-padding ('×' characters; dotted curve), using three images of each particle: Left: the average offset error in the PIV measurement, plotted against the simulated e^{-2} image diameter; Right: The variable error. 204

5.25 Plots of half the average error against half of the simulated horizontal displacement, obtained by simulating PIV images at double the magnification and analysing using 127×127 pixel interrogation areas (without padding the FFTs). 207

5.26 For measurements using unpadded FFTs to analyse simulations in which three images of each particle were created at two different magnifications: Left: the variable error in the average measurement, plotted against the simulated diameter at the lower magnification; Right: the average systematic offset error in the PIV measurement, plotted as a function of the simulated e^{-2} image diameter, again at the lower magnification. The points were obtained using 63×63 pixel interrogation areas ('+' characters; solid lines) or double the recording resolution then 127×127 pixel interrogation areas ('o' characters; dotted lines). The latter results have been halved, due to their doubled magnification. 208

5.27 Left: the systematic offset error in each measurement, plotted against the actual simulated image diameter; Right: The minimum resolvable image displacement, plotted against the actual simulated image diameter. The analyses used unpadded FFTs to analyse simulations featuring three images of each particle. Either 63×63 pixel interrogation areas ('+' characters if zero-padding was not used or '$\times$' characters if it was) or 127×127 pixel interrogation areas without zero-padding ('o' characters) were used. 209

5.28 The behaviour of the analysis program PIVANAL when 128×128 pixel interrogation areas are used to analyse three-image particle tracks, without zero-padding. The simulated image diameter was 4 pixels and the particle concentration was lower than in previous runs. Left: the plot of measurement error against image displacement; Right: the yield of valid vectors. 212

5.29 The error in the horizontal measurement of displacement, calculated by PIVANAL using unpadded FFTs and $d_\tau = 2$ pixels, plotted against the simulated horizontal displacement, for simulations using three (left) or five (right) images of each particle. . . . 214

5.30 Contours of the yield of valid vectors, as measured by PIVANAL using padded (left) or unpadded (right) FFTs and $d_\tau = 2$ pixels, plotted on a graph of the simulated image displacement. The contours are for yields of 90% (solid lines, $\times$ signs) and 80% (dashed lines, 'o' symbols). 215

5.31 If a PIV recording captures three images of each particle, each image being 4 pixels in diameter (top), analysis with fully-padded FFTs (left) gives yields of 90% (solid lines) if the displacement $\boldsymbol{d} = (d_x, d_y)$ is greater than 6 pixels in magnitude and satisfies $d_x + d_y < 36$ pixels or thereabouts. On the other hand, analysis with unpadded FFTs (right) introduces wrap-around errors at many points; these 'no-go' areas are coloured black in these diagrams. The region of safest operation is $6 < |\boldsymbol{d}| < 19$ pixels (dotted circle), which gives yields of over 90% and no wrap-around errors, at the expense of a limited dynamic range. The peak detection algorithm required that $-29.5 < d_x, d_y < 29.5$ pixels (dotted square) for unpadded analyses. Yields improve if smaller seeding particles are simulated (bottom). In both cases, padding the FFTs increases the yields and avoids wrap-around errors, significantly increasing the dynamic range at the expense of requiring more analysis time. 218

6.1 For the case of three images of each particle separated by a horizontal displacement 10 pixels to the right, there are a number of ways for one or more images from a particle track to lie within a randomly-placed interrogation area. 220

6.2 For three images of each particle, displacements to the right less than 32 pixels in magnitude and 64×64 pixel interrogation areas, consider the position of the first image in a random particle track, relative to the position of an interrogation area containing at least one image of that particle. There are five possibilities to consider, represented by the five lightly-shaded rectangles. 222

6.3 For the case of three images of each particle, separated by displacements to the right more than 32 pixels in magnitude, with 64×64 pixel interrogation areas, consider the location of the first image of a particle track (lightly-shaded rectangles), relative to the position of an interrogation area containing at least one image of that particle. 224

6.4 Plots of the fraction of useful images within an average interrogation area, as a function of the image separation, for the cases of two ($*$), three ($+$), four ($\circ$) and five ($\times$) images of each particle. . 226

6.5 For the case of three images of each particle and 24 particle images per 64×64 interrogation area, the number of image pairs (solid line, '+' symbols, images a distance d apart) and the number of harmonic pairs (dash-dot line, '×' symbols, images $2d$ apart) are plotted against the image separation. The dotted line represents correlations between different particle tracks. 233

6.6 For the case of five images of each particle, separated by displacements to the right which are $d_I/3$ in magnitude, with interrogation areas measuring $d_I \times d_I$, consider the location of the first image of a random particle track, relative to the position of an interrogation area containing at least one image of that particle. The various starting locations are represented by lightly-shaded rectangles. . . 234

6.7 Square images of uniform intensity: profile in the image field, contours in the image field, profile of the autocorrelation plane, contours in the autocorrelation plane, cross-section through the autocorrelation plane along $\Delta y = 0$ pixels and a Gaussian fit to discrete data points 1 pixel apart. 240

6.8 Circular images of uniform intensity: profile in the image field, contours in the image field, profile of the autocorrelation plane, contours in the autocorrelation plane, cross-section through the autocorrelation plane along $\Delta y = 0$ pixels and a Gaussian fit to data points 1 pixel apart. 241

6.9 Images with a conical intensity profile: profile of the image intensity, contours of the image intensity, profile of the autocorrelation plane, contours in the autocorrelation plane, cross-section through the autocorrelation plane along $\Delta y = 0$ pixels and a Gaussian fit to discrete data points 1 pixel apart. 242

6.10 Images with a Gaussian intensity profile: profile in the image field, contours in the image field, profile of the autocorrelation plane, contours in the autocorrelation plane, cross-section through the autocorrelation plane along $\Delta y = 0$ pixels and a Gaussian fit to discrete data points 1 pixel apart. 243

6.11 Images with a Gaussian intensity, but saturated at a value of 62.5% of the possible maximum: profile of the image intensity, contours of the image intensity, profile of the autocorrelation plane, contours in the autocorrelation plane, cross-section through the autocorrelation plane along $\Delta y = 0$ pixels and a Gaussian fit to discrete data points 1 pixel apart. 244

6.12 For the case of three images of each particle, separated by entirely horizontal displacements, analysed using interrogation areas measuring 64×64 pixels, consider the effect of combining a conical fundamental correlation peak of magnitude 4 units (solid lines) with a conical wrap-around peak of magnitude 1 unit (dash-dot lines). A summation function can be obtained (dotted lines). The vertical lines indicate the location of the maximum of each function. . . . 246

6.13 As in the previous figure, except this time a parabolic fundamental correlation peak (solid lines) at $d_1 = 20.961$ pixels is added to a parabolic wrap-around peak (dash-dot lines) to obtain a summation peak (dotted lines). The magnitude of the summation peak has been increased by 0.05 pixels for clarity. 247

6.14 As in the previous figure, except the simulated displacement has been increased to $d_1 = 21.25$ pixels (left) or 21.706 pixels (right). . 249

6.15 A comparison of simulated PIV results against the predictions which result from assuming parabolic autocorrelation peaks. Simulations were carried out using three images of each particle and analysis both with and without FFT padding, for particles with an e^{-2} image diameter of 4 pixels. The padded results have been subtracted from the unpadded results, to produce a plot of the measured error against the simulated displacement (dash-dot lines, $\times$ symbols). The predicted wrap-around error for parabolic peaks is plotted using a continuous line. 252

6.16 A section of the autocorrelation plane along $\Delta y = 0$, plotted in order to assess the predictions which result from assuming that autocorrelation peaks are Gaussian. Three images of each particle were simulated, 21.0 pixels apart. The fundamental displacement peak (solid curve) has four times the amplitude of the wrap-around peak (dash-dot curve). The summation peak (dotted curve) has a maximum at 21.10 pixels. 254

6.17 A comparison of simulated PIV results against the predicted behaviour of Gaussian autocorrelation peaks. Simulations were analysed both with and without FFT padding. The padded results plus 0.01 pixels were subtracted from the corresponding unpadded results, then the wrap-around error was plotted against the simulated displacement ($\times$ symbols). The numerically-predicted wrap-around error for Gaussian peaks is plotted using a continuous line, with the result for parabolic autocorrelation peaks also given for comparison (dash-dot line). Top: three images of each particle; Bottom: five images of each particle. 255

6.18 Gaussian fits to the predicted autocorrelation peak data from Figure 6.16, where the true image displacement was 21.0 pixels (dotted curves, $\times$ symbols). Left: the predicted autocorrelation plane values are fitted by a Gaussian function with a maximum value at the global maximum point (21.1 pixels) predicted by the Gaussian peak model. Right: the predicted data are fitted by a broader Gaussian function, whose maximum value occurs at 21.15 pixels. 258

6.19 For the measurement of a diagonal displacement (16, 16) using 64×64 pixel interrogation areas, five-images tracks can cause between 1 and 4 particle images to be found in the interrogation area (thick square). Left: sample places in which the first image of the track (darker circles) can lie, in order that at least one image from the track lies within the interrogation area; Right: the number in each square indicates the number of images which will be found within the given interrogation area, if the first image of the five-image track lies in that square. 260

6.20 For the measurement of a displacement (12, 8) using 64 × 64 pixel interrogation areas, tracks which are five images long can cause between 1 and 5 particle images to lie in the interrogation area. Left: some sample places in which the first image of a particle track (darker circles) can be situated, in order that at least one image from the track lies within the interrogation area (dark square); Right: the number within each square indicates the number of images from a track of five images which will lie within the given interrogation area, if the first image of the track is found in the given square. 261

7.1 A schematic diagram of the apparatus used in the initial work. . . 263

7.2 Pressure drop characteristics of Grade 10 ballotini. The two dashed lines intersect at a value $U_{mf} = 0.0715\,\mathrm{m\,s^{-1}}$ (dotted line). 266

7.3 Schematic diagram of the apparatus used in the larger-scale work. 271

7.4 Normal and log-normal fits to size distribution data for 180–300 μm impact beads. The measurements were made using four separate runs on a Malvern particle size analyser. The median size is indicated using a dashed line. 273

8.1 PIV photographs for the flow about 10 cm above a bed of AG grade glass particles (film 3, frames 4 (top) and 5 (bottom). $U/U_{mf} \approx 30$. 279

8.2 PIV photographs for the flow about 10 cm above a bed of AG grade glass particles (f3f6 and f3f7). $U/U_{mf} \approx 30$. 280

8.3 PIV photographs for the flow about 10 cm above a bed of AG grade glass particles. Top: f3f8, for which $U/U_{mf} \approx 30$; bottom: f3f11, where U/U_{mf} was about 19. 281

8.4 PIV vector maps for the flow about 18 cm above a 3 cm bed of particles. $U \approx 1.0\,U_{mf}$ 287

8.5 PIV photograph of the flow about 18 cm above a 3 cm bed of ballotini (film 39, frame 23). 288

8.6 PIV vector maps for the flow about 18 cm above a 3 cm bed of particles. $U \approx 1.0\,U_{mf}$ 289

8.7 Film 39 frame 23, analysed using 64 × 64 interrogation areas. . . . 293

8.8 PIV vector maps for the flow about 33 cm above a 3 cm bed of particles. $U \approx 1.0\,U_{mf}$ 294

8.9 PIV vector map (top) and photograph (bottom) for the flow about 33 cm above a 3 cm bed of particles. $U \approx 1.0\,U_{mf}$. 295

8.10 PIV vector maps for the flow about 17 cm above a 5 cm bed of particles. $U \approx 1.0\,U_{mf}$ 300

8.11 PIV vector maps for the flow about 17 cm above a 5 cm bed of particles. $U \approx 1.0\,U_{mf}$ 301

8.12 PIV vector maps for the flow about 33 cm above a 5 cm bed of particles. $U \approx 1.0\,U_{mf}$ 302

8.13 PIV vector maps for the flow about 33 cm above a 5 cm bed of particles. $U \approx 1.0\,U_{mf}$ 303

8.14 Ensemble-averaged PIV vector maps for the flow about 33 cm (top) and about 18 cm (bottom) above a 3 cm bed of particles. $U \approx 1.0\,U_{mf}$; Re = 350. 307

8.15 Ensemble-averaged PIV vector maps for the flow about 33 cm (top) and about 17 cm (bottom) above a 5 cm bed of particles. $U \approx 1.0\,U_{mf}$; Re = 350. 308

8.16 Vertical velocity profiles across the freeboard of the smaller fluidised bed, at viewing heights of 33 cm (top) and about 17 cm (bottom) above beds of particles which were 3 cm (left) or 5 cm (right) deep. The curves represent the laminar velocity profile based on the given flowrate (continuous curves) and 25% more than the given flowrate (dotted curves). $U \approx 1.0\,U_{mf}$; Re = 350. 309

9.1 PIV vector maps for the flow about 220 cm above a 31 cm bed of particles. $U \approx 1.7\,U_{mf}$. The marker vectors in the top left corner are 0.1–0.5 m s^{-1} in magnitude. 316

9.2 Top: PIV photograph from frame 11 of film 88; Bottom: velocity values for this film, using $U = 0.118\,\mathrm{m\,s^{-1}} = 1.69\,U_{mf}$. The bed depth was 31 cm and the mean viewing height was about 220 cm above the unfluidised bed surface. 317

9.3 PIV vector maps for the flow about 220 cm above a 31 cm bed of particles. $U \approx 2.25\,U_{mf}$. 318

9.4 Histograms of the distribution of vertical velocities in two sample photographs from film 88 ($U \approx 1.69\,U_{mf}$, left) and film 86 ($U \approx 2.25\,U_{mf}$, right). 320

9.5 Histograms of the velocity distributions obtained from frame 10 of film 90 ($U \approx 1.16\,U_{mf}$). Left: a histogram of the vertical velocity, in m s^{-1}; Right: a histogram of the magnitude of each displacement vector, in pixels. 321

9.6 PIV vector maps for the flow about 220 cm above a 31 cm bed of particles. $U \approx 1.2\,U_{mf}$. 322

9.7 PIV vector maps for the flow about 220 cm above a 31 cm bed of particles, with image-shifting used. $U \approx 1.27\,U_{mf}$. 323

9.8 PIV vector maps for the flow about 220 cm above a 31 cm bed of particles. Image-shifting was used. $U \approx 1.27\,U_{mf}$. 324

9.9 PIV vector maps for the flow about 220 cm above a 31 cm bed of particles, using image-shifting. $U \approx 1.27\,U_{mf}$. Top: the full width of the test section; Bottom: a close-up view near the centre of the full vector map. 325

9.10 The scattering of the light sheet as it enters or leaves the test section. Images can be viewed either directly (continuous-line rays) or after reflection in the back wall of the test section (dashed lines). . 326

9.11 PIV photograph of the flow recorded in film 105, frame 31. 327

9.12 Histograms of the distribution of the vertical (left) and horizontal (right) velocity information obtained from frame 33 of film 105 ($U \approx 1.27\,U_{mf}$). 328

9.13 A PIV vector map for the flow about 220 cm above a 31 cm bed of particles. Image-shifting was used, then the smallest vectors were deleted, to illustrate the gains which can be made by using image-shifting. $U \approx 1.27\,U_{mf}$. 329

9.14 Histograms of the vertical (left) and horizontal (right) velocity distributions from film 104, frame 14. 330

9.15 PIV vector map for the flow about 220 cm above a 31 cm bed of particles, with image-shifting used. 331

9.16 PIV vector maps for the flow about 150 cm above a 31 cm bed of particles. $U \approx 2.67\,U_{mf}$. 334

9.17 Top: Histograms of the distribution of the vertical (left) and horizontal (right) velocity information obtained from frame 10 of film 77. Bottom: Velocity statistics for films 77 and 78. In both runs, $U = 0.187\,\mathrm{m\,s^{-1}} \approx 2.67\,U_{mf}$, the bed depth was 31 cm and the mean viewing height was 150 cm above the bed surface. 335

9.18 PIV vector maps for the flow about 150 cm above a 31 cm bed of particles. $U \approx 2.67\,U_{mf}$. 337

9.19 PIV vector maps for the flow about 150 cm above a 31 cm bed of particles. $U \approx 1.69\,U_{mf}$. 338

9.20 Direct PIV vector maps for the flow about 140 cm above a 31 cm bed of particles. $U \approx 1.27\,U_{mf}$. 340

9.21 PIV vector maps for the flow about 130 cm above a 31 cm bed of particles, after subtracting the image-shifting velocity. $U \approx 1.27\,U_{mf}$. 342

9.22 A PIV vector map for the flow about 130 cm above a 31 cm bed: film 133 frame 24, but with the smallest vectors deleted. $U \approx 1.27\,U_{mf}$. 344

9.23 PIV vector maps for the flow about 40 cm above a 31 cm bed of particles. $U \approx 1.55\,U_{mf}$. 345

9.24 Results for film 102, where the superficial velocity was $0.0889\,\mathrm{m\,s^{-1}}$, or about $1.27\,U_{mf}$. The bed depth was 31 cm and the mean viewing height was about 220 cm above the unfluidised bed surface, one quarter of the way into the test section. Image-shifting was used. Top: the vector map from frame 18 of this film; Bottom: the velocity statistics. 349

9.25 Results from film 93, where $U = 0.108\,\mathrm{m\,s^{-1}} \approx 1.55\,U_{mf}$. The bed depth was 31 cm and the mean viewing height was about 220 cm above the unfluidised bed surface, one quarter of the way into the freeboard. Top: Vector map of the velocity information obtained from frame 15; Bottom: velocity statistics. 350

9.26 Results from film 91, where the superficial velocity was $0.136\,\mathrm{m\,s^{-1}}$ ($1.94\,U_{mf}$). The bed depth was 31 cm. The mean viewing height was about 220 cm above the bed surface, one quarter of the way into the freeboard. Image-shifting was not used. Top: vector map for frame 13; Bottom: some velocity statistics. 351

9.27 Results for film 132, where the superficial velocity was $0.0787\,\mathrm{m\,s^{-1}}$, or about $1.12\,U_{mf}$. The bed depth was 31 cm and the mean viewing height was about 130 cm above the unfluidised bed surface (one quarter of the way into the freeboard). Image-shifting was used. Top: vector map of the velocity information obtained from frame 22; Bottom: some velocity statistics. 353

9.28 Results obtained from film 84, with $U = 0.125\,\mathrm{m\,s^{-1}}$, or about $1.79\,U_{mf}$. The bed depth was 31 cm and the mean viewing height was about 140 cm above the unfluidised bed surface (one quarter of the way into the freeboard). Image-shifting was not used here. Top: vector map of the velocity information obtained from frame 31; Bottom: the velocity statistics for this run. 354

9.29 Summary of results for film 85. Conditions were the same as in film 84, but with an increased superficial velocity $U = 0.157\,\mathrm{m\,s^{-1}}$ ($2.25\,U_{mf}$). Top: the vector map from frame 24; Bottom: some velocity statistics for the run. 356

9.30 Vector map of the velocity information obtained from frame 10 of film 108, where $U = 0.0852\,\mathrm{m\,s^{-1}} \approx 1.22\,U_{mf}$. Image-shifting was used. 359

9.31 The vector map obtained from frame 15 of film 107, with $U = 0.122\,\mathrm{m\,s^{-1}}$ (about $1.75\,U_{mf}$). The bed depth was 60 cm. The mean viewing height was about 190 cm above the bed surface and through the centre of the freeboard. Image-shifting was not used. 361

9.32 Vector map of the velocity information obtained from frame 11 of film 118, where $U = 0.165\,\mathrm{m\,s^{-1}} \approx 2.35\,U_{mf}$. Image-shifting was not used. 363

9.33 Vector map of the velocity information obtained from frame 12 of film 121, where $U = 0.0926\,\mathrm{m\,s^{-1}} \approx 1.32\,U_{mf}$. Image-shifting was used. 366

9.34 Distributions of the vertical (left) and horizontal (right) velocities in frame 17 of film 124. $U = 0.1296\,\mathrm{m\,s^{-1}} \approx 1.85\,U_{mf}$. Image-shifting was used. 367

9.35 Top: Vector map of the velocities for frame 17 of film 124, where $U = 0.1296\,\mathrm{m\,s^{-1}} \approx 1.85\,U_{mf}$; Bottom: Summary of velocity values for film 124. The mean viewing height was about 130 cm above the unfluidised bed surface and the light sheet was directed at the centre of the bed. Image-shifting was used. 368

9.36 Vector map of the velocities measured in frame 10 of film 119, where $U = 0.172\,\mathrm{m\,s^{-1}}$ ($2.46\,U_{mf}$). Image-shifting was not used. 370

9.37 PIV vector maps for the flow about 40 cm above a 60 cm-deep bed of particles, after subtracting the image-shifting velocity. $U \approx 1.31\,U_{mf}$. 372

9.38 PIV vector maps for the flow about 40 cm above a 60 cm-deep bed of particles, after subtracting the image-shifting velocity. $U \approx 1.31\,U_{mf}$. 373

9.39 Histograms of the vertical (left) and horizontal (right) distributions of velocity in frame 12 of film 140, where $U = 0.0917\,\mathrm{m\,s^{-1}} \approx 1.31\,U_{mf}$. 374

9.40 PIV vector maps for the flow about 40 cm above a 60 cm bed of particles, after subtracting the image-shifting velocity. $U \approx 1.75\,U_{mf}$. 377

9.41 PIV vector maps for the flow about 40 cm above a 60 cm bed of particles, after subtracting the image-shifting velocity. $U \approx 1.75\,U_{mf}$. 378

9.42 Histograms of the distribution of the vertical (left) and horizontal (right) velocities from film 140 frame 21, where $U = 0.122\,\mathrm{m\,s^{-1}} \approx 1.75\,U_{mf}$. 379

9.43 PIV vector maps for the flow about 40 cm above a 60 cm bed of particles, after subtracting the image-shifting velocity. $U \approx 1.31\,U_{mf}$. 384

9.44 PIV vector maps for the flow about 40 cm above a 60 cm bed of particles, after subtracting the image-shifting velocity. $U \approx 1.31\,U_{mf}$. 385

9.45 PIV vector maps about 40 cm above a 60 cm bed of particles using image-shifting, one quarter of the way into the freeboard. $U \approx 1.75\,U_{mf}$. 387

9.46 Ensemble-averaged vector maps for the flow 1.3 m above a 31 cm fluidised bed, looking through the centre of the freeboard. $U = 0.0889\,\mathrm{m\,s^{-1}} = 1.27\,U_{mf}$ in the upper vector map and $0.118\,\mathrm{m\,s^{-1}}$ ($1.69\,U_{mf}$) in the lower one. 390

9.47 Profiles of vertical velocities along horizontal sections through the centre of the freeboard above a 31 cm deep fluidised bed, at heights of 2.2 m (top) and 1.3 m (bottom) above the bed surface. The superficial gas velocity was $0.0889\,\mathrm{m\,s^{-1}}$ in the left-hand figures and $0.118\,\mathrm{m\,s^{-1}}$ on the right. 392

9.48 Profiles of vertical velocities along horizontal lines through the centre of the freeboard above a 60 cm deep fluidised bed, at heights of 1.9 m (top), 1.3 m (middle) and 0.4 m (bottom) above the bed surface. The superficial velocity was roughly $0.09\,\mathrm{m\,s^{-1}}$ (left) or $0.125\,\mathrm{m\,s^{-1}}$ (right). The continuous curves are profiles of the velocity under fully-developed turbulent flow, based on a one-seventh power law for cylindrical pipes; hopefully these will be reasonable estimates of turbulent flow in a square duct. 393

9.49 Vertical velocity profiles along horizontal lines one quarter of the way into the freeboard space above a 31 cm deep fluidised bed, at heights of 2.2 m (top) and about 1.35 m (bottom) above the bed surface. The superficial velocity was about $0.083\,\mathrm{m\,s^{-1}}$ (left) or about $0.130\,\mathrm{m\,s^{-1}}$ (right). The expected profile for turbulent flow is plotted using a continuous curve. 394

9.50 Further vertical velocity profiles along horizontal lines through the freeboard. Left: one quarter of the way into the section and 2.2 m above a 31 cm bed; Right: in the centre of the freeboard, 1.3 m above a 60 cm bed. 395

List of Tables

4.1 The properties of various scanning beam boxes, each providing a 1 m wide light sheet 112

5.1 The seeding concentrations in each simulation. 165
5.2 The simulated horizontal displacement at which the average yield of valid vectors dropped to 90% and then 80%, when using fully-padded FFTs. 182
5.3 Horizontal and vertical errors, at various vertical displacements. . 186
5.4 The horizontal displacements which produce given fractional yields of valid vectors, when analysing a range of two-dimensional displacements. 188
5.5 For PIV analyses in which the centre of the autocorrelation peak is calculated to sub-pixel accuracy using Whittaker's reconstruction, this table summarises the regions in which the measured vectors might lie. 196

6.1 Calculations of the fraction of paired images. 225
6.2 The average ratios of paired images to unpaired images. 228
6.3 Calculated ratios of wrap-around correlations to valid correlations between paired images. 235

8.1 Summary of films taken using PIV to record the flow through the centre of the freeboard space above the smaller fluidised bed. . . . 277
8.2 Summary of velocity values for film 39, where the bed depth was 3 cm and the mean viewing height was 18 cm above the unfluidised bed surface. 291
8.3 Summary of velocity values for film 37, where the bed depth was 3 cm and the mean viewing height was 33 cm above the unfluidised bed surface. 296
8.4 Summary of velocity values for film 41, where the bed depth was 5 cm and the mean viewing height was 17 cm above the unfluidised bed surface. 298
8.5 Summary of velocity values for film 33, where the bed depth was 5 cm and the mean viewing height was 33 cm above the unfluidised bed surface. 304

8.6 Summary of average velocities for the films in this chapter. 'Mean v_y' is the mean of the various section-average vertical velocities from the analysed frames of each film, whilst 'Typical max v_y' refers to the average of each frame's value of greatest upward vertical velocity, *etc.* 304

9.1 Bubble diameters (Rowe, 1976) and rise velocities (Rowe & Partridge, 1965) in the larger fluidised bed. 312

9.2 Transport Disengaging Heights predicted for the larger fluidised bed by Horio *et al.* (1980), Baron *et al.* (1988) and Fournol *et al.* (1973). 313

9.3 Summary of films taken using PIV to record the flow through the centre of the freeboard space above a 31 cm fluidised bed. 314

9.4 Summary of velocity values for film 86, using $U = 0.157\,\mathrm{m\,s^{-1}} = 2.25\,U_{mf}$. The bed depth was 31 cm and the mean viewing height was about 220 cm above the position of the bed surface (measured with the air off). 319

9.5 Summary of velocity values for film 105, where the superficial velocity U was $0.0889\,\mathrm{m\,s^{-1}} = 1.27\,U_{mf}$ and image-shifting was used. The bed depth was 31 cm and the mean viewing height was about 220 cm above the unfluidised bed surface. Yields were measured before interpolation. 330

9.6 Summary of velocity values for film 80, where the superficial velocity was $0.119\,\mathrm{m\,s^{-1}}$, or about $1.69\,U_{mf}$. The bed depth was 31 cm and the mean viewing height was about 150 cm above the unfluidised bed surface. 339

9.7 Summary of velocities for film 133, where the superficial velocity U was $0.0889\,\mathrm{m\,s^{-1}}$ (about $1.27\,U_{mf}$). The bed depth was 31 cm and the mean viewing height was about 130 cm above the bed surface. Image-shifting was used. 343

9.8 Summary of velocity values for film 71, where the superficial velocity was $0.108\,\mathrm{m\,s^{-1}}$, or about $1.55\,U_{mf}$. The bed depth was 31 cm and the mean viewing height was about 40 cm above the unfluidised bed surface. 346

9.9 A list of films taken using PIV to record the flow in a section 7.5 cm into the freeboard space above a 31 cm fluidised bed. 347

9.10 Details of films taken using PIV to record the flow through the centre of the freeboard space, when the bed depth was 60 cm. . . 358

9.11 Summary of velocity values for film 108, where the superficial velocity U was $0.0851\,\mathrm{m\,s^{-1}}$, or about $1.22\,U_{mf}$. The bed depth was 60 cm. The mean viewing height was about 190 cm above the unfluidised bed surface and through the centre of the bed. Image-shifting was used. 360

9.12 Results for film 107, with $U = 0.122\,\mathrm{m\,s^{-1}}$ (about $1.75\,U_{mf}$). The bed depth was 60 cm. The mean viewing height was about 190 cm above the bed surface and through the centre of the bed. Image-shifting was not used. Top: the vector map obtained from frame 15; Bottom: some velocity statistics. 362

9.13 Summary of velocities for film 118: $U = 0.165\,\mathrm{m\,s^{-1}} = 2.35\,U_{mf}$. The bed depth was 60 cm. The mean viewing height was about 190 cm above the unfluidised bed surface and through the centre of the bed. 364

9.14 Summary of velocity values for film 121. The superficial velocity was $0.0926\,\mathrm{m\,s^{-1}}$ ($1.32\,U_{mf}$). The bed depth was 60 cm. The test section was 130 cm above the bed surface, and through the bed's centre. Image-shifting was used. 365

9.15 Velocity statistics for film 119, where $U = 0.172\,\mathrm{m\,s^{-1}}$ ($2.46\,U_{mf}$). The bed depth was 60 cm. The mean viewing height was 130 cm above the bed surface and through the centre of the bed. Image-shifting was not used. 369

9.16 Summary of velocity values for the first part of film 140, where the superficial velocity U was $0.0917\,\mathrm{m\,s^{-1}}$ ($1.31\,U_{mf}$). The bed depth was 60 cm and the mean viewing height was 40 cm above the bed. Image-shifting was used. 375

9.17 Summary of velocities for the second part of film 140, where the superficial velocity $U = 0.122\,\mathrm{m\,s^{-1}} = 1.75\,U_{mf}$. The bed depth was 60 cm and the mean viewing height was 40 cm above the bed surface. Image-shifting was used. 379

9.18 Details of films taken using PIV to record the flow one quarter of the way into the freeboard space, when the bed depth was 60 cm. 381

9.19 Summary of velocities for the first part of film 136, where the superficial velocity $U = 0.0917\,\mathrm{m\,s^{-1}}$ ($1.31\,U_{mf}$). The bed depth was 60 cm. The mean viewing height was 40 cm above the bed surface, one quarter of the way into the freeboard. Image-shifting was used. 383

9.20 Summary of velocities for the second part of film 136, where the superficial velocity $U = 0.122\,\mathrm{m\,s^{-1}} = 1.75\,U_{mf}$. The bed depth was 60 cm and the mean viewing height was 40 cm above the bed surface, one quarter of the way into the freeboard. Image-shifting was used. 386

9.21 Summary of velocity statistics for the main runs on the larger fluidised bed. The letter 'C' after the bed depth indicates a photograph of the centre of the test section, whilst 'Q' denotes a run one quarter of the way into the test section. 389

Chapter 1

Introduction

When a fluid is pumped upwards through a bed of solid particles, the pressure drop across the bed helps to balance the 'buoyant weight' of the particles (*i.e.* the net weight of the particles, after subtracting the buoyancy force due to their immersion in a fluid). At a certain velocity known as the 'minimum fluidising velocity', this pressure drop is large enough for the particles to become suspended in the fluid stream, without any need to rest on top of one another. When this happens, the bed of particles behaves like a fluid and is said to form a 'fluidised bed'. Fluidised beds have many industrial applications, *e.g.* when a solid catalyst encourages gaseous hydrocarbons to pyrolyse in a 'Fluid Catalytic Cracking' unit.

If the fluidising material is a gas, it can flow upwards through the bed in one of two ways. The first mechanism is percolation between the particles in the bed; this is always possible. A second mechanism comes into play only if the fluidising gas is supplied at a larger flowrate than that required to achieve incipient fluidisation: gas can flow up through the bed as a particle-free bubble phase, in much the same way that bubbles of air rise through a liquid. The bubbling action means that some gas bypasses the denser particle/gas phase and helps to avoid creating a pressure drop across the bed in excess of the buoyant weight of the particles.

The bubbling action in this type of fluidisation has a number of consequences.

One of these is that the bubbles stir up the particles, helping to achieve good mixing. This enables a constant temperature to be maintained within the bed and allows good heat transfer to take place. Another result is that particles are thrown up above the bed when bubbles burst through its upper surface, into the space above the bed surface but below the gas exit pipe (the 'freeboard region').

Above the bed, the particles experience an upward drag force, due to the upward flow of gas. If the particles are small, they can be carried out of the system by this gas flow and onwards into downstream processing equipment. One industrial application of fluidised bed technology is for coal combustion; pulverised coal is mixed into a bed of sand, at a mean bed temperature of about 800 °C. Coal particles burn in the fluidising gas (air) and therefore become smaller the longer they remain in the bed. If ejected into the freeboard region, these particles could well be carried downstream by the hot gases and might cause problems (*e.g.* a turbine might be damaged).

The rate of particle entrainment from such a system is a useful thing to know, as it will enable an engineer to add a suitable cyclone to the system (so as to remove small particles from the exit gases). Alternatively, an engineer might wish to assess the benefits of building a taller freeboard region, so as to give the particles a longer residence time in the freeboard (hopefully allowing complete combustion to be achieved before the particles can leave the system).

The interaction between spherical particles and gas in the freeboard is complex. A theoretical calculation of the progress of an ejected particle is possible, if the following properties are known:

- the mass and diameter of the ejected particle and/or its terminal velocity
- the initial velocity of the particle
- the region of space in which the ejection occurs
- the rate of change of the particle diameter, as a result of combustion

- the mean gas velocities in the freeboard at the time of the ejection, as a function of position

- the degree of turbulence in the mean gas velocity at each point in the freeboard, and

- the height of the freeboard region, in case the maximum height of rise brings the particle close to the exit pipe (whereupon end effects might arise).

The size of the ejected particles will depend on the size distribution function for the pulverised coal feed and the mixing properties within the fluidised bed. For example, particles which have just been fed into the bed might circulate within the bed for a certain time, before reaching the surface and being ejected into the freeboard. Particle terminal velocities can be calculated, from the size and mass of the ejected particles.

The ejection velocity is likely to depend on the bubbling properties of the bed, *e.g.* the size, rise velocity and frequency of bubbles reaching the surface. When each bubble bursts, a number of particles are ejected into the freeboard, at a range of different speeds and in a range of directions. If necessary, a worst-case scenario could be considered in order to determine the greatest possible height reached in the freeboard, but a whole range of other possibilities will need to be considered if the total mass flowrate out of the system is to be calculated.

Some other properties of interest concern the fluid dynamics of the gas in the freeboard region. The nature of the gas velocity profile is open to question. Some researchers (Horio *et al.*, 1980) have found that the gas flows upwards mainly in a central core region, with velocities somewhat greater than the 'superficial' mean velocity (based on the volumetric flowrate and the total pipe area); outside this core, gas and particles return to the bed in a downward-flowing annular region. Other researchers (*e.g.* Levy & Lockwood, 1983) have found that the greatest velocities are actually found close to the wall, with rather low velocities observed close to the axis of the freeboard. In both cases, it is useful to know the maximum

vertical gas velocity in the freeboard region, *e.g.* in order to seek ways of achieving a flatter profile.

Chapters 2 and 3 of this thesis outline the theory behind fluidisation and the 'elutriation' of fine particles in the gas flow above a bubbling fluidised bed.

Most of the previous research on gas flow in the freeboard has used 'point measurement' techniques. For example, a hot wire anemometer measures velocity indirectly, by measuring the rate of cooling of a hot wire placed into the flow. These techniques can be used to determine a time-history of velocity at a given point, but record no information about the flow elsewhere at these times. After recording a time-history at one point, the probe can then be moved to another point to record another time-history. The technique works best if the flow is steady, but the absence of information about the structure of the flow at any time can be a problem in some flows (*e.g.* if unsteady vortices are present). It might be possible to use several probes in a number of different places so as to record simultaneous velocity information, but the more probes are used, the more invasive the experimental technique becomes.

Particle Image Velocimetry (PIV) is a non-invasive experimental technique, which allows velocities to be determined simultaneously at thousands of points in space, at one particular moment in time. PIV is, therefore, a 'whole field' velocity measuring technique. A vortex can clearly be seen using PIV, because simultaneous velocities will be measured at many points in and around the vortex. It is also possible to capture a series of PIV recordings at short time intervals, in order to obtain information about the development of the flow with time.

PIV involves illuminating a thin section of a flow with a stroboscopic light source. Small 'seeding particles' are introduced into the flow, chosen to be small enough to follow even minor fluctuations in the local gas velocity. The seeding particles scatter light each time they are illuminated, allowing their position to be recorded on a few occasions using photography, then the photographic negatives are developed. It is possible to examine small regions of these negatives, in order

to determine the average particle displacement between successive illumination pulses (and hence the local flow velocity). The basic principles behind PIV, with or without a refinement known as 'image-shifting', are set out in Chapter 4.

As explained above, PIV usually involves introducing a number of small particles into the flow. The technique is, therefore, highly suitable for measuring flows which already contain particles, such as the flows in the freeboard region above a bubbling fluidised bed. Various models have been proposed for flows above a bubbling fluidised bed, following research using point measurement techniques. Whole field velocity data would be useful, in order to confirm or to question the existence of vortex structures in the freeboard region.

The PIV technique itself was also tested. The analysis software used here was developed by Hind (1997), as part of his research work in this department. Hind also wrote a PIV simulation program, which allows a PIV flow record to be simulated under a number of variable conditions. For example, by keeping PIV parameters such as the image concentration, image size, image intensity, number of images of each particle, analysis parameters, *etc.*, constant over a number of simulations whilst varying only the separation of successive images, it was possible to assess how successful the PIV system was at measuring a range of image displacements. This success was assessed in terms of accuracy (*i.e.* how close the average measured displacement was to the corresponding simulated displacement) and yield (*i.e.* the probability that the measurement at one point in the flow will lie within a given tolerance of the mean measured displacement). The original aim was simply to determine the least and greatest image displacements which gave satisfactory measurements, but a number of other factors were also considered. These simulations are described in Chapter 5.

Chapter 5 also considered the danger of obtaining spurious results, if the fastest analysis technique is used to measure certain image displacements. These errors are due to autocorrelation wrap-around effects when using unpadded Fast Fourier Transforms (FFTs). The causes of these errors are outlined in Chapter 5, along with some schemes for their removal.

An attempt is made in Chapter 6 to explain the general reduction in the yield of valid vectors measured by a typical PIV system, as the simulated image displacement increased. A model was also developed, in an attempt to calculate a theoretical value for the errors due to wrap-around correlations. This model, based on the superposition of Gaussian autocorrelation peaks of prescribed magnitudes, gave a good fit to the data from the simulations presented in Chapter 5 and was therefore deemed successful.

PIV experiments have been undertaken, using two different square-sided fluidised beds. The apparatus used in this work is described in Chapter 7. Experimental results using the smaller-scale apparatus are presented in Chapter 8, then results from the larger scale apparatus are summarised in Chapter 9. The larger fluidised bed measured $0.30\,\mathrm{m} \times 0.30\,\mathrm{m}$ in horizontal cross-section and had several metres of freeboard space, in which the gas flow could be studied.

It was found that the velocity profiles above the smaller fluidised bed were 'domed', with the greatest vertical velocities found near the bed axis (as expected). In contrast, the largest vertical velocities in the larger rig were usually found near the walls when the flowrate was low, with the velocities near the centre of the section often very small or even negative. As the flowrate through the larger bed was increased, the velocity profiles levelled out progressively before beginning to show signs of bulging up into a dome.

Chapter 2

Fluidisation

2.1 The basic principles of fluidisation

2.1.1 Apparatus

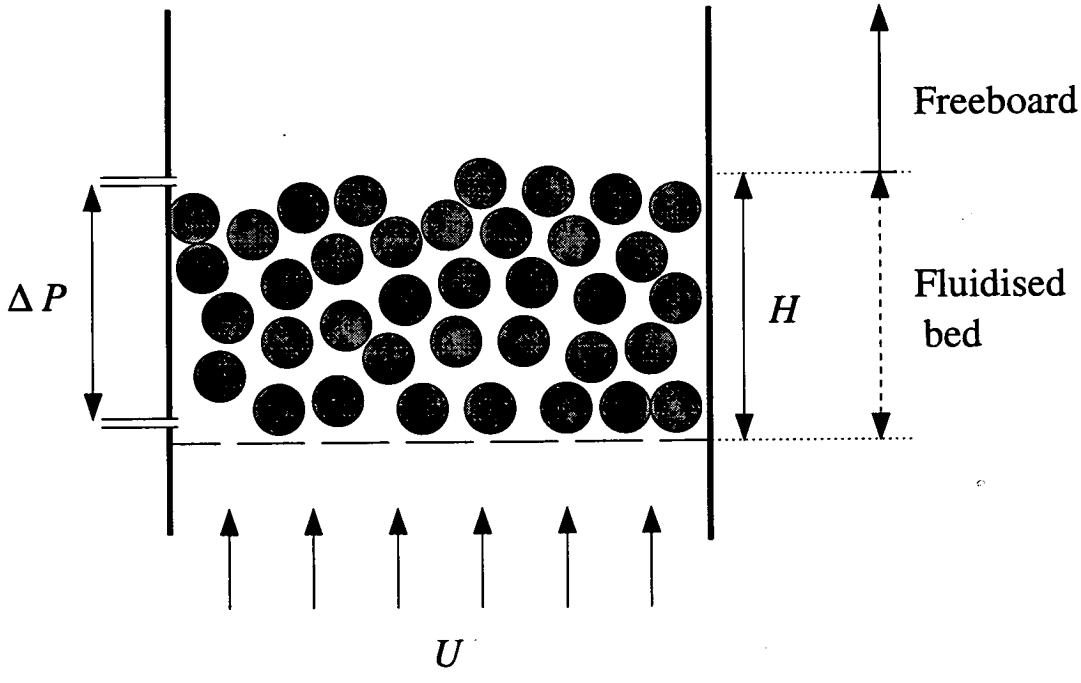


Figure 2.1: A simple fluidised bed can be formed if a fluid is passed upwards at a superficial velocity U through a bed of particles, of depth H. This flow produces a pressure drop, ΔP.

The basic details of fluidisation are shown in Figure 2.1, which shows a vertical section through a simple fluidised bed. Particles, of diameter d_p and solid density ρ_s, are held in place by a porous support (the 'distributor') and are filled to a depth H. When poured into such a system, particles do not occupy all of the bed volume. The fraction of free space within the bed, ε, is usually about 0.4 to 0.6; correspondingly, the solids fraction in the bed is $(1-\varepsilon)$. A fluid is passed upwards through the bed at a superficial velocity U (based on the volumetric flowrate and the pipe cross-sectional area, A, *i.e.* neglecting the particles). The fluid has a density ρ_f and a viscosity μ.

2.1.2 Low-velocity fluid flow

At low fluid flowrates, the fluid experiences a small downward drag force as it flows through the network of narrow gaps between particles; this means that a pressure drop (ΔP) is required to maintain a steady flow. Correspondingly, the particles experience an equal and opposite upward force. The other forces on the particles are their downward weight, an upward buoyancy force (which is equal to the weight of fluid displaced by the particles) and an upward reaction force from the distributor. If the drag force from the fluid is small, the particles will rest on top of one another and will not be able to move.

If the fluid velocity increases, the pressure drop also increases. For a flow through a fixed packed bed of spherical particles at low Reynolds number, the semi-empirical Carman-Kozeny equation gives:

$$\Delta P = \frac{C\,\mu\,H\,(1-\varepsilon)^2}{d_p^2\,\varepsilon^3} U, \tag{2.1}$$

where C is a dimensionless constant, found to be equal to about 180. At higher Reynolds number, momentum effects come into play. The Ergun equation for fluid flow through packed columns can then be used (equation 2.4).

Thus the upward force on the particle bed, due to viscous fluid flow at low Reynolds number, is essentially proportional to the upward fluid velocity. As the fluid velocity increases, there comes a point where the upward force due to the fluid flow exactly balances the downward weight of the particles (less the upward buoyancy force). At this point, the particles are supported by the fluid and no longer need to rest on top of each other; they are also free to re-arrange their configuration (*e.g.* by rotation) or to move about within the bed.

A bed in this state is said to be incipiently fluidised. The fluid velocity, bed height, void fraction, *etc.*, are given the subscript 'mf' to denote the particular values of these parameters found at this point of minimum fluidisation.

A fluidised bed behaves more like a liquid than a bed of solids, *e.g.* the bed is easy to stir, its upper surface becomes horizontal and immersed objects either float or sink, according to their density. Also, the levels in connected fluidised beds equalise out and, if a small hole is made in the side of a bed, a jet of solids will gush out (making fluidised beds relatively easy to empty).

A force balance on the bed material, at incipient fluidisation, gives:

$$A\Delta P_{mf} = (1 - \varepsilon_{mf})\, A\, H\, (\rho_s - \rho_f)\, g, \tag{2.2}$$

where g is the acceleration due to gravity. Substituting for ΔP using equation 2.1 and rearranging:

$$U_{mf} = \frac{(\rho_s - \rho_f)\, g\, d_p^2\, \varepsilon_{mf}^3}{180\, \mu\, (1 - \varepsilon_{mf})}. \tag{2.3}$$

This simple force-balance theory allows U_{mf} to be estimated. The equation is found to be valid for small, spherical particles. The result suggests that the value of U_{mf} is quite sensitive to changes in particle diameter and extremely sensitive to changes in void fraction.

2.1.3 Mild fluidisation

If the fluid velocity is increased beyond U_{mf}, then the Carman-Kozeny equation predicts that ΔP will also increase. The bed weight cannot increase, however, and that is the only downward term in the force balance. The buoyancy term cannot be reduced, as that depends only on the total particle volume (which is constant). Thus, if the previous equations continue to apply, it appears that a net upward force on the bed will arise if $U > U_{mf}$. If the bed was a rigid porous object instead of a loose pile of particles, it would be lifted up off the distributor as a plug and carried downstream by the fluid.

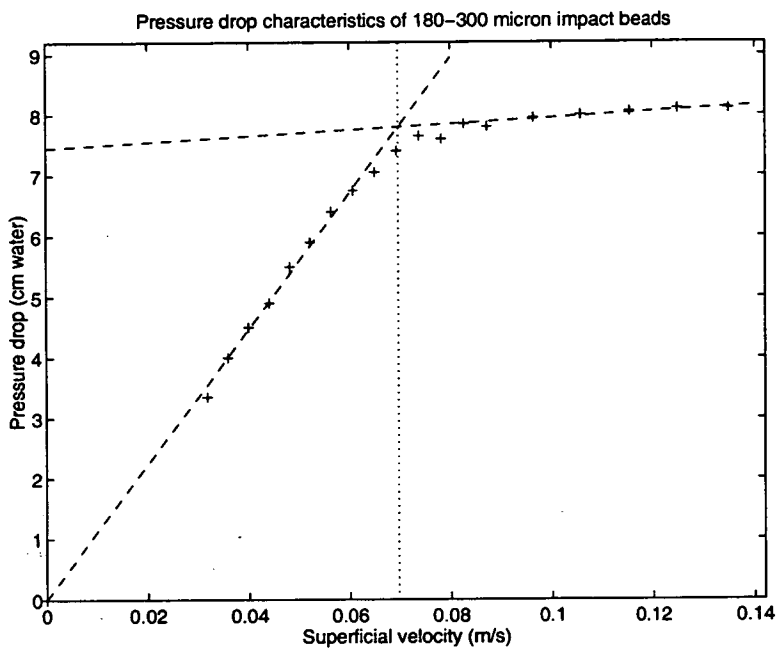


Figure 2.2: The pressure drop through a typical fluidised bed, plotted as a function of superficial gas velocity. The two dotted lines intersect at $(0.070, 7.80)$.

This does not usually occur in fluidisation, because interparticle forces are not usually large. The observed behaviour of ΔP as a function of U is shown in Figure 2.2, for the case of particles 180–300 μm in diameter, fluidised by air. At low gas velocities, ΔP is proportional to U, as predicted by the Carman-Kozeny equation. For $U > 0.08\,\mathrm{m\,s^{-1}}$, however, only a minimal increase in ΔP is found.

The value of U_{mf} is taken to be the value of U at the intersection of the two lines, which is $0.070\,\mathrm{m\,s^{-1}}$ in this case. The slight decrease in the measured ΔP

when $U \approx U_{mf}$ is a result of small regions within the bed being fluidised, allowing some particles to re-arrange themselves into a configuration which offers reduced resistance to fluid flow. This behaviour is often observed when the particles do not have exactly the same diameter.

The behaviour illustrated in Figure 2.2 is that of an 'unlocked' bed; the points plotted were obtained in a random order. If instead the bed had been compacted before use (*e.g.* by gently tapping its walls) and readings had been taken in order of increasing gas velocity, a steeper line through the origin would have been obtained for $U < U_{mf}$, owing to the reduced voidage (equation 2.1). If the bed particles were 'locked' in close contact with the wall due to this compaction, frictional forces at the walls could come into play, allowing pressure drops slightly greater than ΔP_{mf} to be obtained. When the gas velocity is increased so that $U \approx U_{mf}$, a locked bed expands to a voidage of ε_{mf}. In this looser state, no frictional forces are transmitted from the walls and so the pressure drop decreases to ΔP_{mf}. When the fluid flowrate to an already-fluidised bed is decreased, the pressure drops are usually similar to those shown in Figure 2.2 because the bed will not get 'locked' without outside intervention.

When the superficial velocity U exceeds U_{mf}, the way in which a fluidised bed avoids creating a pressure drop greater than ΔP_{mf} (as shown in Figure 2.2) depends on the nature of the fluid-particle system. One key factor is the fluid, which could be a liquid or a gas.

2.1.4 Particulate fluidisation

If the fluid and the particles have a relatively small density difference, *e.g.* when a bed of sand is fluidised by water, then 'particulate' fluidisation is usually obtained. For these systems, the bed expands if $U > U_{mf}$, acquiring a uniform, larger void fraction. This increase in void fraction balances the increase in U, so that the value of ΔP predicted by the Carman-Kozeny equation remains close to that at incipient fluidisation (assuming that the equation remains valid). The relationship

between U and ε was correlated by Richardson & Zaki (1954):

$$\frac{U}{u_t} = \varepsilon^n,$$

where u_t is the terminal velocity of an individual particle and n is a correlated function of a Reynolds number ($2.39 \leq n \leq 4.65$).

For example, a hypothetical incipiently-fluidised bed might have $\varepsilon_{mf} = 0.4$ at $U = U_{mf}$. If the liquid velocity is increased to $U = 2.76\,U_{mf}$, the effect of this in the Carman-Kozeny equation can be cancelled out if the void fraction increases to 0.5 (all other things being equal). This would in fact be the expected voidage according to the Richardson-Zaki equation, if the exponent $n = 4.55$ (this is a reasonable value for small particles, for which the Carman-Kozeny equation is most likely to be valid).

At mild fluidisation conditions, therefore, the Richardson-Zaki correlation appears to be consistent with the bed expansion calculated by setting $\Delta P = \Delta P_{mf}$ (*i.e.* a constant) in the Carman-Kozeny equation and then observing the behaviour of U as a function of ε. At higher fluid velocities, the bed simply expands further (although not necessarily with the same exponent n, as this will decrease as the flow becomes increasingly turbulent).

A review of particulate fluidisation has been carried out by Di Felice (1995).

2.1.5 Aggregative fluidisation and two-phase theory

If the density difference between the fluid and the particles is relatively large, *e.g.* when the fluid is a gas or the particles are unusually dense, 'aggregative' fluidisation occurs. In more detail, Jackson (1963a) showed that systems with a Froude number given by $U_{mf}^2/g\,d_p > 1$ allow small local disturbances to an initial state of uniform fluidisation to grow rapidly within the bed, leading to the formation of particle-free regions. The actual behaviour of such systems, however, depends on the nature of the particles.

According to the simple two-phase theory of aggregative fluidisation, a typical gas-fluidised bed has a 'dense' phase and a 'lean' phase. In the dense phase, fluid flows between the particles in the bed with upward velocities close to those found at incipient fluidisation, with the voidage in the dense phase remaining close to ε_{mf}. Beyond a certain point, any gas flow in excess of a minimum value (at which the superficial gas velocity is U_{mb}) bypasses the dense phase and flows through the bed as a stream of particle-free gas bubbles ($\varepsilon = 1$ in this lean phase). Bubbles in aggregative fluidisation form at (or close to) the distributor. An analogy is to the sparging of gas into a liquid.

When the feed to a bubbling fluidised bed is increased so that $U > U_{mb}$, the bubbling rate increases. If the dense phase remains at a voidage close to ε_{mb}, its total volume depends on the volume of particles present in the bed, which does not change. Thus the volume of dense phase material is constant. The increased rate of bubbling creates more and larger bubbles, increasing the volume of the lean phase and thus the total volume occupied by the fluidised bed (which includes the dense phase and the lean phase). The position of the upper bed surface varies as a function of time, due to short-term variations in the hold-up of lean phase gas in the bed. The bed surface approaches maxima just before large bubbles reach the bed surface and visits minima just after such bubbles burst.

If a bubbling fluidised bed is tall enough, bubbles can coalesce to form larger bubbles. In the limit, the bubbles eventually grow large enough to become influenced by the geometry of the fluidised bed, forming slugs. Slug-like characteristics develop when the 'bubble' diameter reaches one third of the bed diameter (D_t). Slugs can grow until they occupy almost the entire width of the bed.

Geldart (1973) classified aggregative fluidisation systems into four main types, according to the physical properties of the particles:

- Class A particles are usually relatively light ($\rho_s < 1400\,\mathrm{kg\,m^{-3}}$) and relatively small ($30\,\mu\mathrm{m} < d_p < 150\,\mu\mathrm{m}$). For these powders, the key property is that $U_{mb} > U_{mf}$, *i.e.* bubbling only begins when U is somewhat larger

than U_{mf}. In the intermediate region where $U_{mf} < U < U_{mb}$, this type of fluidised bed expands uniformly, as found in particulate fluidisation.

- Class B particles ($40\,\mu\text{m} < d_p < 500\,\mu\text{m}$) can be larger than those in Class A and are also somewhat denser ($1500\,\text{kg}\,\text{m}^{-3} < \rho_s < 4000\,\text{kg}\,\text{m}^{-3}$). Their feature is that $U_{mb} \approx U_{mf}$, *i.e.* a bed of Class B particles begins to bubble at (or only very slightly above) the minimum fluidising velocity.

- Most Class C particles have a mean size of less than about $30\,\mu\text{m}$ and tend to be cohesive. Interparticle forces are significant, making uniform fluidisation difficult to achieve. When $U > U_{mf}$, clumps of particles within the bed can be moved aside as a plug, leaving behind channels through which gas can flow with little resistance.

- Class D particles are larger and/or more dense than Class B particles. Bubbles in these fluidised beds rise more slowly than the interstitial gas and act mainly as void regions, through which gas can easily flow.

Plots of $(\rho_s - \rho_f)$ against d_p can be used to create regime maps, indicating the regions in which each Geldart particle class can be obtained. As is often the case with regime maps, however, the transitions are not always clearly defined.

2.1.6 Faster fluidisation

Kunii & Levenspiel (1991, ch. 3) summarised the regimes found in aggregative fluidisation at high gas flowrates.

Starting from $U = U_{mf}$ or U_{mb}, the rate of bubbling increases if the gas velocity is increased. Each bubble throws a number of particles up into the 'freeboard' region, which is the region of the system above the bed surface but below the gas exit pipe. As $U \to u_t$, isolated particles in the freeboard tend to be carried out of the system, so a recovery and recycle system is necessary if the solids inventory is to remain constant. The rate of particle carryover is limited by the bubbling

properties of the bed: if most of the fluidisation gas bypasses the dense phase by means of bubbles, particles in the dense phase can stay there for some time. As a result, fluidised beds can be operated well beyond $U = u_t$.

The magnitude of fluctuations in the pressure drop across the bed increases with the gas velocity, until a peak is reached. At higher gas velocities, the pressure fluctuations decrease (either gradually or suddenly) until they reach a reasonably constant value. This transition, marking the onset of 'turbulent fluidisation', was found by Yerushalmi & Cankurt (1979) to take place at gas velocities between $8u_t$ and $13u_t$. In turbulent fluidisation, the void fraction of the bed increases to about 0.8–0.9. The bed surface becomes hard to discern, with unstable clusters or strands of particles moving around rapidly both in the bed and in the freeboard. The 'terminal velocity' of these clusters can be similar to the superficial gas velocity, because a cluster of particles is larger and heavier than a single particle.

At even higher gas velocities, *e.g.* $U > 20\,u_t$, 'fast fluidisation' takes over. The behaviour of the bed is influenced by the solids feed rate: at higher feed rates, a zone of relatively high solids concentration is found close to the solids inlet, as the incoming solids are accelerated upwards by the gas feed. The solids fraction decreases with height above the distributor. The type of gas exit has an effect: a smoothly-tapered exit allows the solids fraction to decrease right up to the exit, whereas a more abrupt exit (for example, a capped top with a horizontal outlet at the side) means that solids accumulate at the top of the freeboard (Grace, 1990).

If the mass flowrate ratio of solids to gas reaches about 1:20, pneumatic conveying is found. The void fraction in this regime can exceed 0.98.

2.1.7 Some advantages and applications of fluidised beds

Compared to fixed beds, fluidised beds are very well mixed. If an exothermic chemical reaction is being carried out between a gas and a solid catalyst, this natural mixing helps to control chemical changes by avoiding the occurrence of runaway reactions in 'hot spots'.

In addition, the liquid-like nature of a fluidised bed allows the bed particles to be manipulated in various ways. These may allow a fluidised bed process to be operated continuously, in situations where a fixed bed (*e.g.* with downward fluid flow) could only be operated batchwise.

For example, in a fluidised bed catalytic cracker, catalyst particles can become coated by carbon deposits and thus lose effectiveness, if they have been in the bed for some time. It is possible to remove particles from the bed continuously, *e.g.* via an overflow pipe near the bed surface, from a cyclone on the gas exit stream if the particles are elutriable or through a deliberately-defluidised region at the bottom of the bed. These particles can be pneumatically conveyed by air into a regenerator unit, where any undesirable surface coatings are burnt off, before the revitalised catalyst particles are carried back into the reactor by the hydrocarbon feed stream. Schematic diagrams of such processes are given in textbooks by Coulson & Richardson (1978) and Kunii & Levenspiel (1991, ch. 2).

2.2 Previous research on fluidisation

2.2.1 Minimum fluidising velocity

If the more general Ergun equation is used to estimate ΔP, rather than the Carman-Kozeny equation, the expression for the pressure drop across the fluidised bed no longer assumes that viscous flow takes place. The results, therefore, need not be restricted to small particle diameters. Ergun's equation is:

$$\Delta P = 150 \frac{\mu H (1-\varepsilon)^2}{d_p^2 \varepsilon^3} U + 1.75 \frac{\rho_f H (1-\varepsilon)}{d_p \varepsilon^3} U^2. \tag{2.4}$$

After substituting this equation into the force balance equation, a quadratic equation for U_{mf} is obtained. This can be expressed most compactly in terms of a particle Reynolds number based on the particle size and fluid properties at minimum fluidisation ($\mathrm{Re}_{p,mf} = \rho_f U_{mf} d_p / \mu$) and an Archimedes number ($\mathrm{Ar} = \rho_f (\rho_s - \rho_f) g d_p^3 / \mu^2$):

$$\frac{1.75}{\varepsilon_{mf}^2\,\phi_p}\mathrm{Re}_{p,mf}^2 + \frac{150\,(1-\varepsilon_{mf})}{\varepsilon_{mf}^3\,\phi_p^2}\mathrm{Re}_{p,mf} = \mathrm{Ar},$$

where ϕ_p is the sphericity of the particles (allowing a wider variety of particles to be described). The Archimedes number described above is proportional to the product of a drag coefficient and the square of a Reynolds number (with all groups based on particles fluidised at $U = U_{mf}$).

For 'small' particles which are not necessarily spherical, this simplifies to:

$$U_{mf} = \frac{(\rho_s - \rho_f)\; g\; \phi_p^2\; d_p^2\; \varepsilon_{mf}^3}{150\;\mu\;(1-\varepsilon_{mf})}. \tag{2.5}$$

This equation is valid if $\mathrm{Re}_{p,mf}$ is less than 20, although it might perhaps be better to use a numerical constant of 180 instead of 150 if the flow is assumed to be laminar after all (*cf* the Carman-Kozeny equation). For glass spheres 200 μm in diameter which have $U_{mf} = 0.07\,\mathrm{m\,s^{-1}}$ in air, $\mathrm{Re}_{p,mf} \approx 1$ so it appears that the particles do not need to be particularly small for the equation to be valid. Glass spheres with a diameter of about 500 μm would be close to the limit of validity, when fluidised by air.

Wen & Yu (1966) correlated 284 published U_{mf} values and found:

$$\mathrm{Re}_{p,mf} = (33.7^2 + 0.0408\,\mathrm{Ar})^{0.5} - 33.7.$$

Rowe & Nienow (1975) investigated the effect of mixing two powders, each of which has a known minimum fluidising velocity. Let these velocities be $U_{mf,A}$ and $U_{mf,B}$ (with the former being smaller than the latter) and let the weight fraction of the larger particles be x_B. Provided that the larger set of particles is no more than three times bigger than the smaller, an empirical correlation was given for the minimum fluidising velocity ($U_{mf,C}$) of such a mixture:

$$U_{mf,C} = U_{mf,A}\left(\frac{U_{mf,B}}{U_{mf,A}}\right)^{x_B^2}.$$

Geldart & Abrahamsen (1980) gave U_{mf} correlations for Class A materials. They also investigated the minimum bubbling velocity for these powders, which was strongly dependent on the mass fraction of fine particles which were less than 45 μm in diameter (F_{45}). Their rather inelegant correlation for U_{mb} was:

$$\frac{U_{mb}}{U_{mf}} = \frac{2300\ \rho_f^{0.126}\ \mu^{0.523}\ e^{0.716\,F_{45}}}{(\rho_s - \rho_f)^{0.934}\ g^{0.934}\ d_p^{0.8}}.$$

2.2.2 Terminal velocity

If the particles are spherical and sufficiently small, their terminal velocity can be found using Stokes' Law:

$$u_t = \frac{(\rho_s - \rho_f)\ g\ d_p^2}{18\,\mu}. \tag{2.6}$$

The limit to this equation is $\mathrm{Re} = \rho_f\ u_t\ d_p/\mu < 2$. In air, glass particles ($\rho_s = 2500\,\mathrm{kg\,m^{-3}}$) which are 73 μm in diameter have $u_t = 0.40\,\mathrm{m\,s^{-1}}$ and $\mathrm{Re} = 1.96$. These particles are therefore on the limit of validity of the above expression; the drag force for larger particles is somewhat greater than that predicted by Stokes' Law, which will therefore overestimate u_t.

Another correlation for u_t, which is valid for $2 < \mathrm{Re} < 200$, was quoted by Leva (1959):

$$u_t = \frac{0.152\ g^{0.714}\ (\rho_s - \rho_f)^{0.714}\ d_p^{1.14}}{\rho_f^{0.285}\ \mu^{0.428}}.$$

Other, more general correlations are available, covering a wider range of particle sizes and sphericities. These can be used directly or via charts. The most convenient correlations are those in which only one dimensionless group is a function of the terminal velocity, so an iterative solution is not required.

2.2.3 Bubble diameter

Various researchers have made observations about the bubble diameter in a bubbling fluidised bed, D_b. An early review, including a number of instructive photographs, was given by Rowe (1971).

Bubbles form at the distributor (*e.g.* a porous plate or a grid of orifices) and rise up towards the bed surface, coalescing with other bubbles along the way. Photographs of two bubbles coalescing in a liquid have been published by Manga & Stone (1993).

Even when bubbles are not limited in size by the dimensions of the fluidised bed (*e.g.* the bubbles are in a tall, wide bed), an apparent maximum bubble size can be observed in tall beds of Geldart A materials. An equilibrium situation can be reached, in which separate bubbles coalesce and large bubbles 'knife' themselves in two at identical rates; these phenomena were observed using X-rays by Rowe & Partridge (1965). As a result, bubbles are usually smaller in beds of fine powders.

Various correlations have been developed for bubble diameter as a function of height above the distributor (z). Bubbles are considered to form with a diameter D_{b0} a distance z_0 above the distributor. Other bubble sizes of note are D_{be}, the equilibrium bubble size when bubble splitting and coalescence both occur, and D_{bm}, the maximum obtainable bubble diameter if only coalescence occurs. A number of other parameters are relevant, *e.g.* the bed diameter (D_t) and the bubble rise velocity (U_b).

By considering only bubble coalescence, Mori & Wen (1975) correlated the bubble diameter using:

$$\frac{D_{bm} - D_b}{D_{bm} - D_{b0}} = e^{-0.3(z-z_0)/D_t}. \tag{2.7}$$

For fluidised beds with a porous plate distributor:

$$
\begin{aligned}
D_{b0} &= 3.77\,(U - U_{mf})^2/g \\
D_{bm} &= 2.59\,[(U - U_{mf})\,A]^{0.4}\,g^{-0.2} \\
\text{and } z_0 &= 0.
\end{aligned}
$$

For Geldart B and D materials, this result has been validated to an accuracy of $\pm 50\%$ over a wide range of operating conditions. Mori & Wen (1975) recommended that this correlation should be used if $0.005\,\mathrm{m\,s^{-1}} < U_{mf} < 0.2\,\mathrm{m\,s^{-1}}$, $D_t < 1.3\,\mathrm{m}$, $60\,\mu\mathrm{m} < d_p < 450\,\mu\mathrm{m}$ and $(U - U_{mf}) < 0.48\,\mathrm{m\,s^{-1}}$.

Rowe (1976) presented a simpler equation for Geldart B materials, forming fluidisation bubbles of up to 0.7 m in diameter:

$$
D_b = (U - U_{mf})^{0.5}\,(z + z_0)^{0.75}\,g^{-0.25}, \tag{2.8}
$$

assuming that no limitations on bubble size are imposed by the finite dimensions of the apparatus. Here, the constant z_0 is a characteristic of the distributor which is expected to equal zero for porous plate distributors.

Darton *et al.* (1977) gave a similar expression:

$$
D_e = 0.54(U - U_{mf})^{0.4}\left(z + 4\sqrt{A_0}\right)^{0.8} g^{-0.2},
$$

where D_e is an equivalent bubble diameter and A_0 is a catchment area for the bubble stream the distributor plate ('the area of plate per orifice').

Horio & Nonaka (1987) developed a more general model than the above, by taking both bubble coalescence and bubble splitting into account. When bubble splitting is not important, *e.g.* for Class B and D materials, their model reduces to that of Mori & Wen (1975) and gives similar results to those predicted by Rowe (1976) and Darton *et al.* (1977). When Class A materials are fluidised, the more general result of Horio & Nonaka (1987) is better: the other correlations significantly overestimate D_b because they fail to take bubble-splitting mechanisms

into account. As this work concerns only Geldart B materials, however, the simpler results will be preferred to the more general correlation of Horio & Nonaka.

2.2.4 Bubble rise velocity

There are many similarities between bubble properties in a fluidised bed and those in a liquid. Bubble rise velocities in fluidised beds are usually independent of the size and density of the fluidised particles, which suggests that the dense phase in a bubbling fluidised bed behaves like a low-viscosity liquid with zero surface tension (Davidson *et al.*, 1959).

For a single bubble injected into a fluidised bed held in an infinitely-wide vessel, Rowe & Partridge (1965) suggested that the rise velocity was given by:

$$U_{b,\infty} = K\sqrt{g\,D_b/2}, \tag{2.9}$$

for some constant in the range $0.8 < K < 1.2$, although the experimental data were inconclusive. Rowe & Partridge found that values of K were larger for fluidised beds containing small or aspherical particles. For particles satisfying $60\,\mu\text{m} < d_p < 550\,\mu\text{m}$, they correlated K to $\pm 10\%$ by the equation:

$$K = 1.34\,\varepsilon_{mf} + 0.31.$$

Rowe & Yacono (1976) preferred to correlate K against d_p. These values of K can be compared to results for the rise velocity of bubbles in a liquid, which are given by the formula:

$$U_{b,\infty} = 0.71\sqrt{g\,D_e}$$

according to Davidson & Harrison (1963), taking D_e to be a volume-based equivalent bubble diameter. This book drew on results published by Davidson *et al.* (1959) and Harrison & Leung (1962).

Substituting a typical void fraction $\varepsilon_{mf} = 0.52$ into Rowe & Partridge's results for fluidised beds gives $K \approx 1$ and an overall numerical constant of $1/\sqrt{2} \approx 0.71$, when their equation is phrased in the same way as that of Davidson & Harrison. Similar results are found when a particle diameter of about $200\,\mu m$ is used in the correlation of Rowe & Yacono (1976).

Rowe & Matsuno (1971) carried out experiments on lead glass ballotini using two different techniques. In all experiments, the ballotini were fluidised vigorously for a few minutes before the flowrate was reduced and the first reading taken. With 'settled' beds, the flowrate was then reduced progressively between subsequent readings, but for 'pre-agitated' beds the flowrate was increased between readings, so that the particles were fluidised vigorously immediately before every reading. Rowe & Matsuno found that the pressure drop was greater in the pre-agitated bed when $U < U_{mf}$, suggesting that its void fraction was lower. Correlations revealed that the bubble rise in a pre-agitated bed ($K = 0.86$) was about 10% greater than that in a settled bed ($K = 0.79$), although there was a considerable amount of scatter in the data. Fluidised beds must be extremely complicated devices if this difference in the experimental technique has such an effect.

The above equations for isolated bubbles are only valid if $D_b \ll D_t$. When bubbles are large enough to have an equivalent diameter close to that of their container, an equation for slug flow must be used. Dumitrescu (in 1943) and Davies & Taylor (1950) found that:

$$U_{slug} = 0.35\sqrt{g\,D_t}. \tag{2.10}$$

Davidson *et al.* (1977) reviewed the liquid-like behaviour of fluidised beds. They presented a plot of $U_b/\sqrt{g\,D_t}$ against $\sqrt{D_e/D_t}$, from which it appears that the slug flow equation should be used when $D_e/D_t > 0.3$. The scatter of the experimental data suggests that $U_b/\sqrt{g\,D_e}$ is between 0.52 and 0.73. The line $U_b/\sqrt{g\,D_e} = 0.62$ fits most of the data to an accuracy of about $\pm 17\%$. A compromise line of this slope is plotted by Davidson *et al.* and attributed to Rowe & Matsuno (1971).

For a series of bubbles rising in a fluidised bed, the absolute rise velocity of each bubble is given, if simple two-phase theory applies, by:

$$U_b = U_{b,\infty} + (U - U_{mf}), \tag{2.11}$$

as stated by Davidson & Harrison (1963). In the case of a Geldart Class A material, U_{mf} should be replaced by U_{mb}. An explanation of this is that the dense phase around a rising bubble in the upper part of the fluidised bed is pushed up at a velocity of $(U - U_{mf})$ due to the continuous feed of lean-phase gas into the bed via the distributor; the bubble rises relative to the rising dense phase.

The model used by Davidson & Harrison (1963) assumed that the flow of gas through isolated bubbles was dominated by viscous forces. According to Glicksman *et al.* (1987), this means that the particles in a bed fluidised by air at ambient conditions should really be less than 400 μm in diameter. However, their experiments showed that Davidson & Harrison's formula for the average bubble rise velocity still gave good results with particles as large as 1 mm.

Grace & Clift (1974), however, reported that 'there is ample evidence ... that the original two-phase theory seriously overestimates the visible bubble flowrate in most cases'. Various models have been devised to correct this, *e.g.* Davidson and Harrison (1966) proposed that:

$$\frac{Q_b}{A} = U - k\,U_{mf},$$

where Q_b is the total volumetric flowrate of bubbles in the bed and k is a constant, possibly depending on the mean bubble-phase volume fraction (Lockett *et al.*, 1967). This constant could reflect an increase in gas flow through the dense phase when $U > U_{mf}$ (Pyle & Harrison, 1967) or the flow of gas through rising bubbles, which present a low-resistance path up the bed (Grace & Harrison, 1969).

Grace & Clift (1974) reviewed the literature and found a wide range of values for k, which varied between 0.7 and 18; a value of $k = 1.5$ was common.

2.2.5 Problems with the two-phase theory

Kunii & Levenspiel (1991, ch. 6) report that the assumptions made in the simple two-phase model of fluidisation are not always satisfied. In particular, the void fraction in the denser emulsion phase rises above ε_{mf} if $U > U_{mf}$ and the superficial gas velocity through the emulsion phase, U_o, can exceed U_{mf}. Werther (1983) investigated Geldart B and D particles in three-dimensional beds and found that:

$$\frac{U_o - U_{mf}}{U - U_{mf}} = \frac{1}{3}.$$

Considering the gas flow in excess of that required to achieve the state of minimum fluidisation, *i.e.* an excess air flowrate of $A(U - U_{mf})$: simple two-phase theory states that all of this gas should pass through the bed as bubbles, but Werther's result suggests that only two-thirds of the excess gas actually does so in practice. Hilligardt & Werther (1986) found that this fraction increased towards 1 if the ratio of fluidised bed depth to diameter exceeded two. In all cases, the remaining excess gas must pass through the dense phase, which presumably expands to cope with the extra flow (perhaps following the Richardson-Zaki correlation).

Further, the dense emulsion phase of a fluidised bed is not necessarily stagnant. Various types of bulk circulation (or 'gulf streaming') have been observed, depending on the particular conditions (Lin *et al.*, 1985). For example, in a 'deep' bed with depth/diameter > 1 and $U > U_{mf}$, there will be a tendency for large bubbles to form by coalescence. These bubbles tend to rise close to the bed axis and the rapidly-rising wake behind each bubble drags some of the particulate phase up with it. In this case, the dense phase in the upper part of the bed will circulate in a toroidal pattern, rising up near the bed axis and descending near the walls. A toroidal circulation in the opposite direction can also be obtained, if the gas flowrate is closer to that at incipient fluidisation and if the height/diameter ratio is about 0.95. In deeper beds, these circulation patterns can exist simultaneously, with the former 'up in the centre, down at the wall' circulation sitting on top of the 'down in the centre, up at the wall' pattern (which remains close to the distributor). This suggests that most bubbling occurs near the walls at shallow

depths, before bubble coalescence moves the bubbling activity towards the centre of the bed at greater heights above the distributor.

Circulation patterns in the dense phase are likely to influence, and be influenced by, the bubbling pattern. As explained above, bubbles rising along the axis of a fluidised bed can induce a circulation pattern in a previously-stagnant dense phase; similarly, density differences within the dense phase could give rise to a circulation pattern, which then alters the distribution of bubble-phase material.

2.3 Elutriation

2.3.1 Particle behaviour in a bubbling fluidised bed

Particles in a fluidised bed are usually of different sizes, unless the samples have been very carefully processed before use (*e.g.* by sieving). An industrial fluidised bed can be a rather dynamic system, in which the size distribution of particles in the bed could well change as a function of feed rates. The size of each individual particle could well change as a function of time. For example, repeated collisions between particles can cause fragmentation or agglomeration, whilst chemical reactions in a fluidised bed combustion unit will consume burning coal particles, causing each particle to shrink. For these and other reasons, the size distribution of particles in a real fluidised bed need not necessarily bear much resemblance to the size distributions of the feeds.

The behaviour of a fluidised bed can be complex if the particles have a broad range of sizes. The behaviour of such beds is dominated by the larger particles. If $U < U_{mf}$, the bed is not fluidised and the finer particles remain trapped within the bed, even if U exceeds their terminal velocity.

When fluidisation occurs, at $U > U_{mf}$, the larger particles become mobile. Rising bubbles stir up the dense-phase material and release the finer particles,

which can now escape from the fluidised bed and be carried away with the gas flow. This entrainment of fine particles is known as 'elutriation'.

Elutriation can cause a variety of problems with downstream process equipment. When coal particles are burnt, elutriation means that particles will not necessarily spend enough time in the fluidised bed to achieve complete combustion. Combustion will also take place in the freeboard region, before the gases leave the system. In order to avoid any possibly-undesirable carry-over of fine coal particles into downstream equipment, such as a turbine, the residence time of such particles in the freeboard needs to be long enough to ensure that their combustion is complete. If not, equipment will need to be installed in order to remove the particles (*e.g.* in a cyclone) for disposal or recycle. Knowledge of the elutriation rate is useful, because this allows the freeboard region or the cyclone to be sized.

Examining the freeboard region immediately above a gently-bubbling fluidised bed, the average concentration of particles varies as a function of particle diameter and height above the bed surface. Particles of all sizes are ejected upwards into the lowest part of the freeboard region (the 'splash zone') whenever a fluidisation bubble bursts at the surface. The larger particles have $u_t > U$, so they tend to fall back into the bed after rising only a short distance. Smaller bed particles tend to rise higher in the freeboard before falling back down, as their weight may only be slightly larger than the upward drag force they experience. The finest particles in the bed might be small enough for their terminal velocity to be less than the fluid velocity, in which case the upward flow of fluidisation gas will carry the particles out of the system, once they enter the freeboard.

The elutriation of fine particles from a gas-fluidised bed was shown by Leva (1951) to follow an approximately first-order concentration decay, when monitored as a function of time. This is consistent with a model of a well-mixed fluidised bed, where gas flow carries fines out of the bed at a rate proportional to the fines concentration in the bed. Osberg & Charlesworth (1951) found that this first-order decay was only true initially: as elutriation progressed, the rate of

elutriation fell below the first-order value. This might have been due to a change in the mean size or concentration of the fines still in the top part of the bed, or due to the existence of an equilibrium fines concentration.

A distance known as the 'transport disengaging height' (TDH) can be defined, which measures the height above the surface of a very tall fluidised bed, beyond which no change is found in the concentration of the particles present in the flow. All particles which remain associated with the fluidised bed, *i.e.* which are not elutriated, disengage from the gas flow before reaching this height and fall back down towards the bed. The only particles found above the TDH are those which are small enough to be elutriated.

Wen & Hashinger (1960) quote results which illustrate the effect of a TDH. Data of cumulative mass of elutriated fines versus time were collected, using a fluidisation system with varying heights of freeboard region (12, 24 or 36 inches). Considering the results for the intermediate freeboard height, the cumulative mass collected at any time was about half that collected when the shorter freeboard region was used, but about the same as that collected using the taller freeboard. This suggests that the TDH for the system was between 12 and 24 inches.

The TDH might well determine the size of freeboard region. There is little point building more freeboard space than is necessary, as a greater capital cost will be incurred without reducing elutriation. In Wen & Hashinger's example, 24 inches would be a cost-effective freeboard height if particle carry-over is not desirable and no other factors come into consideration. On the other hand, a fluidised bed with a taller freeboard might fit in better with other process equipment. For example, a more-elevated gas exit pipe might allow the exit gases to pass through a cyclone with a longer dip-leg for feeding the recovered fines back into the bed; problems will be encountered if this dip-leg is too short.

2.3.2 The importance of the gas velocity in the freeboard

Elutriation is likely to be affected by the bubbling properties of the fluidised bed. Rapidly-rising bubbles burst at the surface of the bed, ejecting particles into the freeboard and also creating a jet of rapid gas upflow. Thus, particles are ejected into the freeboard region with an initial upward velocity, and they are ejected into a region where the upward drag force from the gas is probably greater than average. Close to the surface of the bed, this encourages newly-ejected particles to continue to rise. Other nearby particles, outside regions of rapid gas upflow, fall back down due to gravity and rejoin the bed.

A long way above the bed, one would expect the gas flow to resemble normal pipe flow, if the solids loading is low. For example, if the fluid Reynolds number is greater than 5000, the flow in the pipe will be turbulent. If the fluidised bed is cylindrical, the velocity profile will be relatively flat and the greatest velocity will be about 1.2 times greater than the mean (or superficial) velocity. On the other hand, the greatest velocity will be twice the superficial velocity if laminar flow conditions prevail instead. Particles present in the flow are most likely to be elutriated if they lie close to the axis of symmetry of the fluidised bed, where the gas velocity (and hence the upward drag force on each particle) is greatest.

If the freeboard region is tall enough for a known velocity profile to be reached, if the particle concentration in the freeboard is dilute and if random fluctuations in the local gas velocity due to turbulence are short-lived and can be neglected, the condition for elutriation is that:

$$U_{max} \geq u_t,$$

where U_{max} is the maximum upward velocity found in a flow of mean upward velocity U and u_t is the terminal velocity of the largest particles which can be elutriated. The former is a function of the fluid Reynolds number in the freeboard, whilst the latter is a function of various physical properties of the particles and the fluid (*e.g.* the size and shape of the particles).

In the splash zone immediately above the surface of a bubbling fluidised bed, the kinematics of particle motion are likely to be dominated by bubbling properties and the gas flow is likely to be extremely turbulent. Further above the bed, the gas velocity profile should become less turbulent and elutriation calculations can be carried out. The exact nature of the velocity profile is of great importance, as this reveals the location and magnitude of the greatest upward gas velocity. Bed particles in this region are the most likely to be elutriated.

In some parts of the freeboard, however, the average vertical velocity profile might not be useful when attempting to calculate elutriation rates, if the average velocities pale into insignificance when compared to the velocities in regions affected by short-lived high-velocity 'jets'. It would be expected that such regions would be most likely to be found close to the bed surface, but some penetration into the freeboard space should also occur.

Particle circulation patterns in the freeboard region were measured by Horio *et al.* (1980). They found that the time-averaged particle velocity profile was upward and relatively uniform across various horizontal cross-sections through a cylindrical freeboard region, except in a thin annulus near the wall (where downward particle velocities were found). The magnitudes of the particle velocities in the bulk of the flow were usually larger than the superficial velocity. This suggested that a downward gas circulation was induced near the walls, due to the downward particle motion. Further, the presence of particles in the freeboard appears to result in a flatter gas velocity profile (as inferred from the particle velocity profile) than would be expected for the given Reynolds number, which suggests that the presence of particles in the freeboard region might have been affecting the gas velocity profile. Horio *et al.* proposed that particles are transported radially outwards from the centre of the freeboard to the wall and only begin to fall back down once they have reached the outer descending zone.

An interesting circulation pattern for gas in the freeboard was proposed by Pemberton & Davidson (1984). When bubbles rise in the bed relative to the dense-phase material, particle movement past the bubble surface creates a toroidal gas

flow pattern within each bubble. Considered in the bubble reference frame, the direction of this circulation is upwards at the bubble centre and downwards at the bubble surface. When large bubbles burst at the bed surface, this vortex might persist into the freeboard region. Using a water wave tank, Reynolds (1876) had observed that 'vortex rings' such as this could travel distances of several metres, slowing down along the way due to the entrainment of slow-moving fluid with which momentum has to be shared. In the context of a fluidised bed, Pemberton & Davidson (1984) called these vortex rings 'ghost bubbles'.

In the freeboard, 'ghost bubbles' would be expected to grow to the width of the freeboard, rising relatively slowly. Downward gas velocities would be expected near the walls, with strongly-upward velocities along the bed axis. In this model, a substantial part of the horizontal cross-sectional area in the freeboard would be occupied by recirculating pockets of gas and would not be available for bulk gas flow. If only a quarter of the bed cross-section is available for flow, a volume balance on the fluidising gas suggests that its average velocity, in the limited region in which net gas upflow can occur, must be four times greater than the superficial velocity. According to the ghost bubble model, the local gas velocity at some points in the freeboard could well be much larger than the superficial velocity. As the ghost bubbles themselves are thought to rise relatively slowly, this restriction on the area available for flow should last for some time. This means that the elutriated gas from a tall fluidised bed might contain somewhat larger particles than would be predicted by a model based on the greatest gas velocities found in orthodox turbulent ($U_{max} \approx 1.22\,U$) or laminar ($U_{max} \approx 2\,U$) pipe flow systems.

Vortex rings have been studied in liquids by Maxworthy (1972, 1974) and Didden (1979). Photographs of vortex rings slipping through each other have been published by Yamada & Matsui (1978). More recently, velocity measurements using PIV have been made, by Willert & Gharib (1991) and by Sridhar & Katz (1995). A review of vortex rings has been carried out by Shariff & Leonard (1992).

As ghost bubbles are a spatial feature of the flow, they should be easy to spot using PIV, in which thin slices of the flow are illuminated with a pulsing light sheet

and then photographed. Each PIV photograph is an instantaneous snapshot of the flow, allowing velocities to be determined at many points in space, at one moment in time. Whole-field velocity measurement techniques (such as PIV) make it easier to observe unsteady structures in the flow than is the case with point measurement techniques (such as hot wire anemometry).

Horio & Kuroki (1994) reported work using Fluid Catalytic Cracking particles with a mean diameter of 61 μm, which were fluidised in air. They used a *circulating* fluidised bed, so high flowrates were used (in the range $0.75\,u_t < U < 3\,u_t$). The FCC particles were in Class B according to Geldart's classification, although at the small and light end of the range (close to Class A). In their experiments, a light sheet from an expanded laser beam illuminated the flow, which was recorded on video. Clusters of particles were seen in the freeboard, along with particle-free gas pockets. The observed motion of particle clusters around rising gas pockets suggested that the pockets had vortices in both vertical and horizontal planes.

Horio & Kuroki therefore stated that ghost bubbles 'were clearly visualized'. This conclusion seems dubious, because bubbles do not exist in circulating fluidised beds. It is possible that circulation patterns were indeed induced in gas pockets, as they rose up through the freeboard past slow-moving particle clusters; these vortex rings do not originate from bubbles in the fluidised bed, however, as predicted by Pemberton & Davidson (1984). It is unlikely that these vortex rings will be found above bubbling fluidised beds with $U = 1\text{–}5\,U_{mf}$ or so, because particle clusters do not form in the freeboard under these conditions.

Potential flow analysis has been carried out on bubbles rising in gas-fluidised beds, starting with the work presented by Davidson in an early symposium (Various, 1960) and then continued by Pyle & Rose (1965) and Collins (1982, 1983). These analyses consider the flow around and inside an apparently-stationary fluidisation bubble, with or without bubble growth. As a result, the velocity profile of a newly-formed ghost bubble can be estimated and its likely appearance in the laboratory frame (after a change of reference frame) can be noted. This work will be presented in the following chapter.

2.4 Selected previous research on elutriation

2.4.1 Bubble eruptions

Kunii & Levenspiel (1991, ch. 7) state that solids are ejected into the freeboard when fluidisation gas pockets reach the surface, by various mechanisms. When the bed is slugging, the slug pressure is greater than the pressure in the freeboard, so the solids just above a bursting slug are exploded into the freeboard. In bubbling beds, the solids ahead of each bubble run down the side of the bubble and material from the bubble wake is ejected into the freeboard. Occasionally two bubbles coalesce just as the higher bubble breaks through the bed surface; when this happens, the ejection of material into the freeboard is particularly energetic.

In all cases, a representative sample of bed material is ejected into the freeboard region (*i.e.* not just the finer particles). Particle transport within the fluidised bed is needed if elutriation is to continue at a steady rate, in order to ensure that the fines concentration near the bed surface is not depleted (Leva & Wen, 1971). This could happen either by means of good mixing in the fluidised bed or by a process of stratification, in which the larger bed particles gradually work their way down to the bottom of the fluidised bed whilst the finer particles migrate towards the surface. Some evidence of stratification along these lines was tabulated by Osberg & Charlesworth (1951), who fluidised an initially-uniform bed for ten minutes. They found that the concentration of fine particles was greatest at the bed surface and slightly lower than average near the distributor.

On the other hand, Rowe *et al.* (1965) found that rising bubbles in fluidised beds draw up a wake of dense-phase material, which is continually exchanged with other solids in the bed. One of their experiments involved setting up a fluidised bed, with a thin lead-glass layer sandwiched between the distributor and a soda-glass bed fill. X-rays through the apparatus after the injection of a single bubble into an incipiently-fluidised bed showed that a spout of the radiographically-denser lead-glass particles was drawn up behind the bubble. A subsequent bed dissection revealed that many lead-glass particles were left behind in the bed as a 'finger',

after leaving the bubble wake at some stage during the bubble's ascent. This finger connected the remains of the initial lead-glass layer near the distributor to a newly-formed deposit on the bed surface. Thus, the particulate matter from a bubble/dense-phase 'cloud' in a fluidised bed appears to undergo continuous mass transfer with the other particles in the dense phase, during the rise through the bed. Particulate material in the wake region below each cloud appears to do likewise. Thus, particles ejected into the freeboard are likely to be representative of the bed as a whole, *e.g.* some fine particles should be ejected even if these happen to be most concentrated near the distributor.

Pemberton & Davidson (1986) considered two theoretical models for particle ejection into the freeboard. In the first, a dome forms at the bed surface as a (single) bubble approaches. The dome then thins and becomes hemispherical with a radius $D_e/2$, before bursting when its thickness falls to d_p. When this happens, the monolayer of particles in the dome is ejected into the freeboard. The mass of these particles could be determined. The two-phase theory allowed the bubble eruption frequency to be found. The mass flux at the foot of the freeboard was hence estimated to be $3\,\rho_s\,(U - U_{mf})\,(1 - \varepsilon_{mf})\,d_p/D_e$.

This model for ejection from the roof of a bubble gave fair agreement with elutriation mass flux data from the literature. Variations of up to one order of magnitude were found, when all previous results were considered. The expected error was a factor of about 4 due to uncertainties in the published data, *e.g.* in ε_{mf} and in the actual bubble-phase flowrate (0.5–$1.0\,(U - U_{mf})\,A$). Agreement was best if the particles were Class B (Geldart, 1973) and when $U/U_{mf} < 10$–15.

In Pemberton & Davidson's alternative model, bubbles were so close together that eruption was usually accompanied by coalescence. They assumed that 40% of the particles in the bubble wake were ejected into the freeboard and that the bubble wake occupied 25% of the bubble volume). The predicted particle mass flux of $0.1\,\rho_s\,(U - U_{mf})\,(1 - \varepsilon_{mf})$ gave reasonable results for Class A materials, and for Class B materials with $U/U_{mf} > 10$–15. In other cases, the mass flux could be over-estimated by as many as three orders of magnitude.

It can be seen that the ratio of the theoretical mass flux from wake ejection to that from roof ejection is equal to $D_e/(30\,d_p)$. The equivalent bubble diameter D_e increases as $(U - U_{mf})$ increases (equation 2.8). This ratio will therefore be large at high flowrates and/or for small particles (*e.g.* Class A materials). Under these conditions, the bubble wake model predicts the greater mass flux. In their experiments, Pemberton & Davidson (1986) measured mass fluxes between their two theoretical limits, so they concluded that a change in the particle ejection mechanism takes place as U increases.

2.4.2 Transport Disengaging height

Zenz & Weil (1958) correlated the TDH against fluidised bed diameter, for a range of gas velocities. Given a superficial velocity, it is therefore possible to estimate the TDH from their chart.

Fournol *et al.* (1973) correlated the elutriation rate (and the mean diameter of the elutriated particles) against an inverse Froude number, $g\,z/U^2$, where z represents the height above the bed surface. They found that little change was found in either elutriation property when this number exceeded 1000, for a range of gas velocities. Kunii & Levenspiel (1991) expected that this type of correlation would hold for other systems, although the exact value of the numerical constant might need to be determined for each system. If so, the method would allow the TDH (z_{TDH}) to be measured experimentally at one reference velocity and then estimated at any other plausible gas velocity.

However, both of the above results were based on experiments with fine Fluid Catalytic Cracking catalyst particles. The predictions for the TDH should be used only with caution when the bed material is not Class A under Geldart's classification or when the superficial velocity U is less than the terminal velocity of most of the particles (Kunii & Levenspiel, 1991).

Horio *et al.* (1980) estimated the TDH from the decay of RMS turbulence fluctuations in the freeboard above beds of some reasonably large particles (mean particle diameters of 160 μm to 1.35 mm were used). Their prediction was that the TDH would equal $14\sqrt{D_b/g}$, where D_b was the diameter of bubbles erupting at the bed surface. The '14' in this expression is a simplification of terms involving a velocity parameter, hence its units are $\mathrm{m\,s^{-1}}$.

Baron *et al.* (1988) considered the behaviour of particle clusters, after ejection into the freeboard from a bubbling fluidised bed. These clusters are larger than individual particles and can reach greater heights in the freeboard, because drag forces have less effect on larger objects. In the absence of drag, the greatest height reached by an object ejected upwards at a velocity U_{ej} is simply $U_{ej}^2/2\,g$, so the problem becomes one of determining U_{ej} for clusters ejected from a fluidised bed. Research by previous workers suggested that $U_{ej}/U_b = 2$ was reasonable, for eruptions involving coalescing bubbles. The TDH was therefore estimated to be $(2\,U_b)^2/2\,g$. It can be noted that this type of expression can be re-arranged into a form similar to that proposed by Fournol *et al.* (1973) but based on the bubble rise velocity instead of U, *i.e.* $g\,z_{TDH}/U_b^2 = 2$.

Comparisons were made between pre-existing experimental data and the predictions made by this formula; a reasonable fit was found. Some other researchers' correlations were also tested, *e.g.* the result of Horio *et al.* (1980), which appeared to overestimate the TDH in most cases. The correlation proposed by Fournol *et al.* (1973) gave very large errors, suggesting that the numerical value of the inverse Froude number at the TDH might vary significantly from system to system or that U_b should have been used instead of U in the inverse Froude number.

Baron *et al.* (1988) also correlated the TDH above slugging fluidised beds, where the ejection processes are completely different. Their best-fit correlation appears to be $z_{TDH} = (7.5\,U_b)^2/2\,g$. The authors also noted that it would be prudent to add 10% to their best-fit correlations for the TDH when designing the freeboard above a fluidised bed, so as to allow some margin for error.

2.4.3 Elutriation from tall beds, with fine particles

Zenz & Weil (1958) developed a theoretical expression for the greatest height reached by particles ejected into the freeboard, starting at the bed surface and working upwards. Particles were assumed to be ejected upwards at a velocity equal to the bubble rise velocity, into a freeboard region which had a constant gas velocity (U). The freeboard region had a total height H_f, which was larger than the TDH. The bed particles were divided up into a number of small particle size fractions. It was assumed that the drag coefficient for particles in each fraction followed a typical form for laminar flow ($C_D = k/\mathrm{Re}$, for a Reynolds number based on the velocity difference between the particles). For each size fraction, equations of motion could be integrated in order to calculate the final 'balancing velocity' of the particles in that fraction (upward if elutriable, downward if retained) and the maximum height reached. From these calculations, it was possible to determine the relative entrainment at various heights above the bed surface. The results from the method as described above were poor, with very short TDH values predicted.

Zenz & Weil (1958) then proceeded to estimate elutriation by an empirical method, starting at the top of a large hypothetical freeboard region and working down. It was assumed that the bubbling action in a fluidised bed causes the formation of intermittent high-velocity jets at the bed surface, so that the splash zone near the bed surface has a highly irregular velocity profile and a large solids concentration. Higher up in the freeboard, the velocity profile becomes more uniform and the solids concentration decreases, eventually reaching a limiting value. It was therefore assumed that an elutriating fluidised bed acts as a saturation feed device for a fine-particle transport system. Pneumatic-conveying research had previously suggested that each solid-gas system has a saturation particle-carrying capacity, which could be determined using correlated data. As a result, the solids carrying capacity at the TDH could be determined.

Agreement between these empirical and theoretical methods of Zenz & Weil (1958) was improved if, when moving up the freeboard, the gas velocity was allowed to equal the bubble velocity at the bed surface, before gradually decreasing

as a function of freeboard height (reaching the superficial velocity at the TDH). The combination of the two methods gave predictions which were in reasonable agreement with experimental data from a two-dimensional test rig, in which fine particles (50% smaller than 60 μm) were fluidised vigorously. It was found that the TDH was a function of the superficial gas velocity and the vessel diameter, but not the actual particle diameter or the bubbling properties.[1]

Zenz & Othmer (1960) stated that the TDH correlations given by Zenz & Weil (1958) are reasonable approximations if the mean particle diameter is between 40 μm and 200 μm and if the superficial gas velocity satisfies $0.15\,\mathrm{m\,s^{-1}} < U < 0.9\,\mathrm{m\,s^{-1}}$. Under these conditions, however, most of the bed particles are elutriable: Gugnoni & Zenz (1980) report that 90% of the cracking catalyst particles tested by Zenz & Weil (1958) were elutriable at the lowest superficial velocity tested, which seems reasonable given that Stokes' Law predicts that $u_t = 0.28\,\mathrm{m\,s^{-1}}$ for glass particles 60 μm in diameter in air[2]. The work of Zenz & Weil, therefore, is only relevant to Geldart Class A materials in fluidised beds taller than the TDH.

Wen & Hashinger (1960) obtained a correlation for the first-order elutriation rate constant, along the lines of work by Leva (1951). Their data included particles satisfying $40\,\mu\mathrm{m} < d_p < 150\,\mu\mathrm{m}$ and $1300\,\mathrm{kg\,m^{-3}} < \rho_s < 5000\,\mathrm{kg\,m^{-3}}$, gases with densities in the range $0.16\,\mathrm{kg\,m^{-3}} < \rho_f < 1.2\,\mathrm{kg\,m^{-3}}$, superficial velocities in the range $0.22\,\mathrm{m\,s^{-1}} < U < 1.3\,\mathrm{m\,s^{-1}}$ (converting from the original imperial units) and fines concentrations of up to 25% by mass. The equation used to describe the rate of mass loss was:

$$-\frac{1}{A}\frac{dw}{dt} = K\frac{w}{W},$$

where w is the mass of fines in the bed, W is the total mass of solid particles in the bed and K is an elutriation rate constant. When the elutriation mass loss is expressed in this form, K should be independent of the bed height or the total

[1]This was surprising. It was thought that the effect of frequent slowly-rising bubbles must somehow be similar to that of less frequent but more rapidly-rising bubbles.

[2]FCC particles may be less dense than glass and hence might have a lower terminal velocity.

particle mass. Leva & Wen (1971) revised Wen & Hashinger's correlation for K (in $\mathrm{kg\,m^{-2}\,s^{-1}}$) to:

$$\frac{K}{\rho_f(U-u_t)} = 1.7\times10^{-5}\left(\frac{(U-u_t)^2}{g\,d_p}\right)^{0.5}\left(\frac{\rho_f\,u_t\,d_p}{\mu}\right)^{0.725}\left(\frac{\rho_s-\rho_f}{\rho_f}\right)^{1.15}\left(\frac{U-u_t}{u_t}\right)^{0.1}$$

In Wen & Hashinger's original work, the numerical constant was actually 1.52×10^{-5} and the final term (involving only U and u_t) was neglected. If the fines concentration exceeds 25% by mass, a correction factor can be used. This correlation is only valid for beds with a freeboard height greater than the TDH.

This rate-constant method could be used after dividing the particles in the bed into suitable size fractions. For each particle diameter, u_t could be be calculated and used to estimate K for the particles in that size fraction. The value of K so determined gives the elutriation flux which would be measured if all the particles in the bed were the same size as the particles in the given size fraction. The total mass flux G_s of elutriated solids from a mixed-size bed is given by:

$$G_s = \sum_i \frac{w_i\,K_i}{W},$$

where w_i is the mass of particles in the i^{th} size fraction and K_i is the corresponding elutriation rate constant.

Kunii & Levenspiel (1991, ch. 7) summarised some of the correlations for K in the published literature. Beds of various shapes (circular or rectangular) and sizes (0.076 m to 0.92 m) had been used by previous workers, along with a wide range of bed-particle and fine-particle sizes. They then tabulated the results from each set of correlations in a 'worked example', for particles in the range 30 μm to 100 μm elutriated from a mixed-size fluidised bed by air at $U = 1.0\,\mathrm{m\,s^{-1}}$. A high degree of variation (*e.g.* about ±100%) was found in the predicted values of K.

2.4.4 Elutriation from other beds

Gugnoni & Zenz (1980) found that data obtained from a Geldart Class B powder were not at all consistent with the results from previous workers (including Zenz & Weil (1958) and Wen & Hashinger (1960)). These differences were attributed to the fact that only about 20% by weight of Gugnoni & Zenz's somewhat larger powder could be elutriated in the runs at their lowest flowrates, whereas the vast majority of particles were elutriable in the experiments carried out by the other authors. Gugnoni & Zenz noted Matsen's observations that some of the trends in the literature defied logic. For example, the correlation by Wen & Hashinger (1960) for the elutriation constant suggests that particles with a high solid density and a large diameter have a large K value and therefore appear to be easier to elutriate; however, increasing both of these particle properties will also increase u_t, which will offset some (or possibly all) of the increase.

By making an analogy between elutriation from a fluidised bed and the release of vapour from a boiling liquid, Gugnoni & Zenz (1980) proposed that better results could be obtained by correlating a dimensionless elutriation rate (mass elutriation rate/$(A\ U\ \rho_f)$) against the dimensionless excess velocity $(U - U_{mb})/\sqrt{g\ d_p}$. They assumed that the total fines elutriation rate is given by a correlation in the above form and that the *relative* elutriation rate of particles in a small size fraction (in the elutriable range) is given by the same functional relationship of dimensionless groups but with U_{mb} replaced by u_t, the terminal velocity for that size fraction. This allowed Gugnoni & Zenz to explain their elutriation data.

Two different correlations of the total elutriation rate were found, for elutriation from a highly-entrainable Class A powder at a high superficial velocity (Zenz & Weil, 1958) and for the elutriation of just the finest components from a mostly-unelutriable Class B material (Gugnoni & Zenz, 1980). Care should be taken to ensure that the correct curve is used.

Lewis *et al.* (1962) carried out elutriation experiments under a number of conditions, including varying freeboard heights. Beds taller than the TDH were

operated in a condition termed 'complete reflux', as all of the unelutriable solids had had the opportunity to disengage from the gas flow before reaching the top of the freeboard. Lewis *et al.* found that the particle bulk density in the freeboard ($\overline{\rho}$) decayed exponentially as a function of height in the freeboard above the bed surface, before reaching a constant value at the TDH:

$$\overline{\rho_R}(z) = \overline{\rho_{R,0}}\, e^{-a\,z},$$

where $\overline{\rho_R}(z)$ is the mean particle bulk density at a height z under conditions of complete reflux, $\overline{\rho_{R,0}}$ is the particle bulk density in the splash zone just above the bed surface when complete reflux occurs, and a is a constant. This expression can be used when $z < z_{TDH} < H_f$.

When runs were repeated using freeboard sections which were shorter than the TDH ($H_f < z_{TDH}$), lower particle bulk densities were found at all positions in the freeboard than had been found at the same height above the bed surface in the condition of complete reflux. It was noted that the solids density just above the bed surface was reduced, due to the fact that some of the particles which cannot escape when $H_f > z_{TDH}$ manage to do so when the freeboard is shorter, so the flux of solids returning to the bed is reduced. The solids density at a randomly-selected height (*e.g.* $z = z_{TDH}/2$) was also lower than when complete reflux applied. In detail, it was found that:

$$\overline{\rho_R}(z) - \overline{\rho}(z) = \text{constant},$$

i.e. the exponential decay curve for solids bulk density against height in the freeboard was translated a constant distance in the direction of decreasing $\overline{\rho}$ but truncated before $\overline{\rho}$ had a chance to decrease as far as $\overline{\rho_{R,0}}\, e^{-a.z_{TDH}}$, the particle bulk density at the TDH. The solids density in the exit stream was increased, however, due to the lowering of the exit system. As the entrainment rate is proportional to $\overline{\rho}$, it was found that entrainment was reduced exponentially as the freeboard height in the system increased towards z_{TDH}.

A sketch of the overall particle concentration profile found by Lewis *et al.* (1962) is given in Figure 2.3.

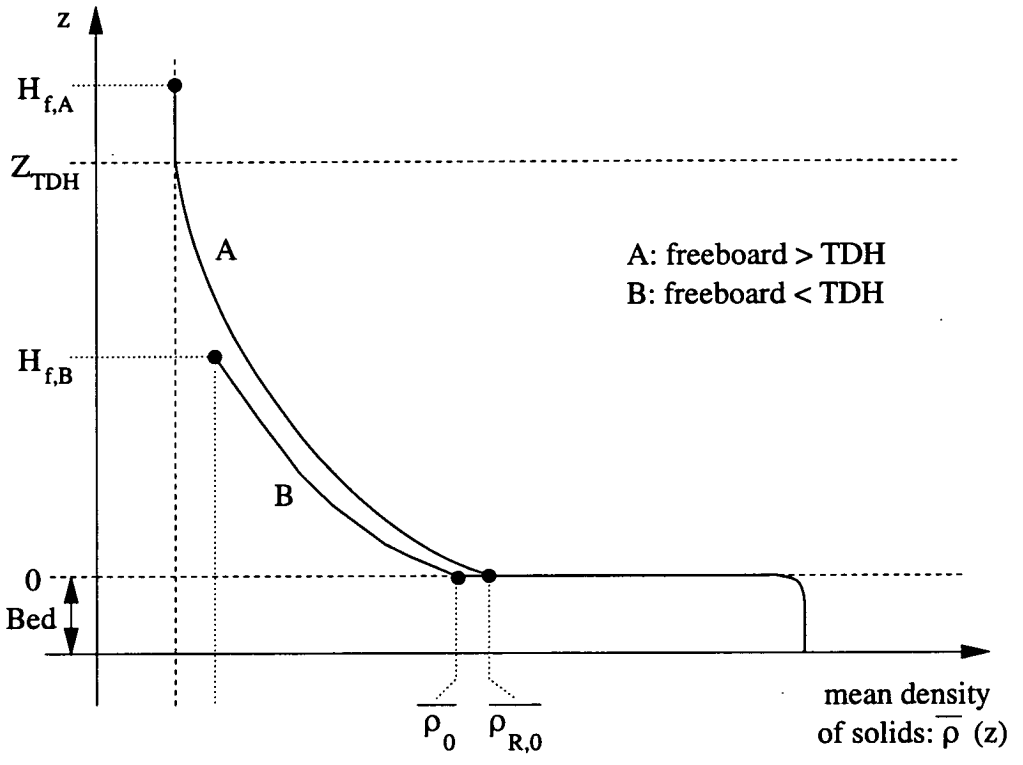


Figure 2.3: The particle concentration profiles in the freeboard of a fluidised bed, as measured by Lewis *et al.* (1962). In this example, freeboard A has a height $H_{f,A}$ which is greater than the TDH, whereas $H_{f,B} < z_{TDH}$.

In 1969, Kunii & Levenspiel developed a general model for entrainment from fluidised beds. The model was summarised in their later book (Kunii & Levenspiel, 1991). The basic details are that solids can be in one of three different phases in the freeboard, numbered as follows:

1. particles carried up by the gas, in a dispersed suspension;
2. particles rising in the freeboard as an agglomerate, after being ejected from the fluidised bed by the bubbling action; and
3. particles in a descending agglomerate cluster or wall layer.

Bulk densities, mass fluxes and velocities were introduced for each phase, *e.g.* ρ_1 was the bulk density of dispersed particles rising in suspension, at a velocity U_1

which remained constant throughout the freeboard. A mass balance revealed that the sum of the three particle mass fluxes must equal the total upward mass flux, which was also constant throughout the freeboard. It was assumed that particles could only enter the freeboard from the fluidised bed via the agglomerates of phase 2 (*i.e.* $\rho_1 = 0$ at $z = 0$) and that all particles reaching the top of the bed were elutriated (*i.e.* $\rho_3 = 0$ at $z = H_f$, as it was assumed that no particles which reach that height via phases 1 and 2 were able to return towards the bed via phase 3). Some features of this model are shown in Figure 2.4.

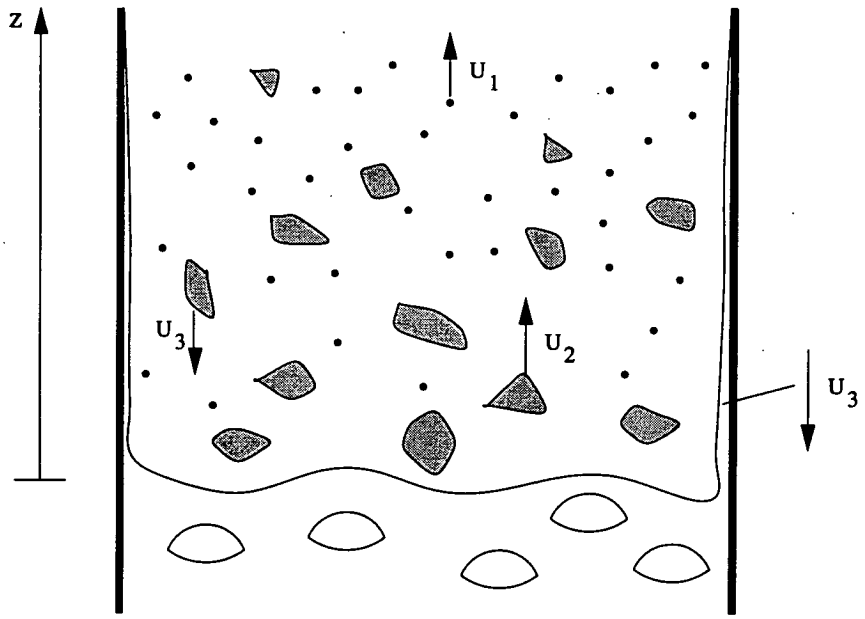


Figure 2.4: The main features of the freeboard particle transport model, developed by Kunii & Levenspiel (1991).

Some mass transfer between the phases was modelled, using differential equations. The transfer of mass into phase 1 at any height in the freeboard was proportional to the total agglomerate concentration at that height, *i.e.* the break-up of both of the agglomerate phases boosted phase 1 by an amount proportional to $K_1(\rho_2+\rho_3)$, for some constant transfer coefficient K_1. Transfer from the rising agglomerate phase to the falling agglomerate phase was also modelled, but no mass transfer out of the suspension phase (1) or from the falling agglomerate phase (3) to the rising agglomerate phase (2) was considered. The rising agglomerate density ρ_2 was greatest at the bed surface but was reduced at greater heights in the freeboard (due to irreversible mass transfer to phases 1 and 3).

On integration, expressions for all three bulk densities could be found, as functions of the transfer coefficients, the phase velocities, the total freeboard height H_f and the height variable z. It was found that ρ_2 and ρ_3 were greatest at the bed surface and tended towards zero at the top of tall fluidised beds ($H_f \to \infty$), whilst ρ_1 was zero at $z = 0$ and approached a limiting value at large freeboard heights. The predictions for the total mass flux and the total solids density were in agreement with the general form of the experimental results published by Lewis *et al.* (1962). Comparisons with other work were also made, *e.g.* a correlation was given for the freeboard decay constant, a.

Kunii & Levenspiel (1995) later extended their model to take into account exit effects. For example, the top of the freeboard might be blanked off using a flange, with air leaving the system through a horizontal exit pipe a short distance below the flange. In this case of an abrupt exit, not all particles in the freeboard will be able to change direction and leave the system on their first attempt. Some particles will rise further (possibly as far as the blanking flange) before changing direction and falling back down towards the bed. The revised model allowed mass transfer to take place from both of the rising phases (1 and 2) to the falling phase (3), but only at the very top of the column. As a result of this change, the particle concentration in the new model was allowed to increase as the top of the freeboard was approached, rather than having to decay smoothly as before. However, some brackets were omitted in key equations. The results agreed with experimental observations on systems with various exit arrangements (Grace, 1990).

2.4.5 Velocity profiles in the freeboard region

Horio *et al.* (1980) used an optical probe to measure particle velocities in the freeboard above an elutriating fluidised bed. They found that the time-average particle velocity profile was fairly flat across their 0.24 m-diameter freeboard, except for a region near the walls in which descending particle velocities were found. The velocity profile they measured is sketched in Figure 2.5. The thickness of the descending region was about 5 to 25 mm. At a given height above the bed

surface, this descending region became thicker if the gas flowrate was increased; at a given superficial velocity, the region became thinner as the height increased. Some downward particle velocities were observed in the centre of the freeboard, in other runs at lower flowrates.

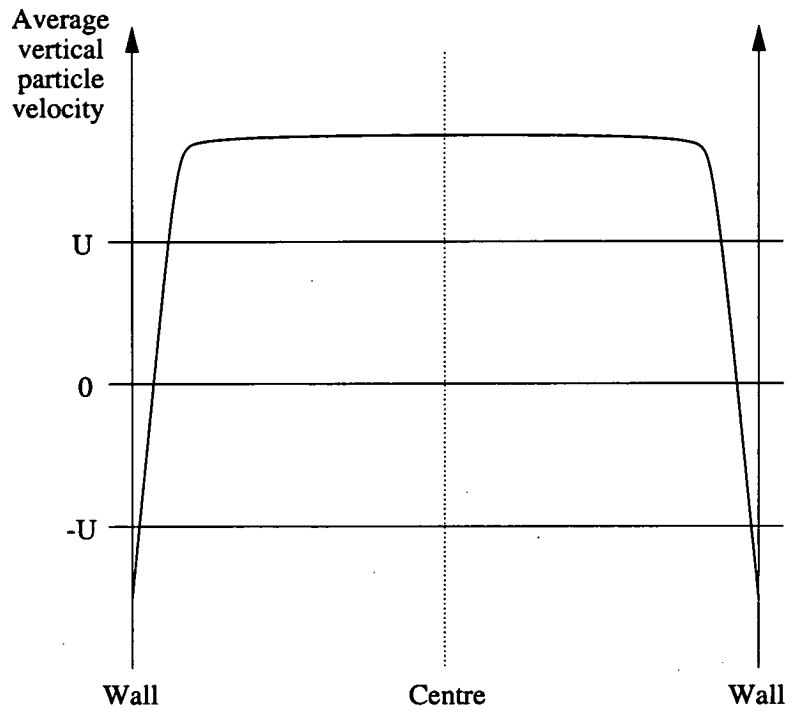


Figure 2.5: The time-averaged vertical particle velocity profile across the freeboard above a cylindrical fluidised bed, measured by Horio *et al.* (1980).

These findings are consistent with the model of Kunii & Levenspiel (1991). At relatively high superficial velocities, the bubbling action will be more vigorous, throwing more particles into the freeboard, more frequently. Some of these particles will make their way to the wall, where they will fall back down towards the bed surface. It seems entirely reasonable that the particle fluxes in the freeboard will increase if the flowrate is increased. Further, there will be no returning particles if $z > z_{TDH}$, so it would be expected that the layer of descending particles will be non-existent above the TDH. Lower down, the thickness of this layer should increase as the bed surface is approached: almost all particles which return to the bed will have to fall back down at the walls through a horizontal level 0.001 m above the bed surface, but only some of the particles will reach a height of 0.1 m above the bed surface and fewer still will make it higher than $z = 1$ m. The more particles are present, the thicker the descending layer will be.

The actual time-average *particle* velocities found by Horio *et al.* (1980) were always somewhat larger than the superficial *gas* velocity, except in the descending region. They reasoned that the descending region probably contained downward-flowing particles *and* downward-flowing gas. In one run using $U = 0.3\,\mathrm{m\,s^{-1}}$, the vertical particle velocities at a number of freeboard heights were uniformly about $0.5\,\mathrm{m\,s^{-1}}$ in an inner cylindrical 'core' region which was about 0.22 m in diameter, with typical downward velocities of about $0.4\,\mathrm{m\,s^{-1}}$ found in the outer annular region (which was 0.01 m thick). A volume balance based on these dimensions and *particle* velocities gives a result which is very close to the superficial *gas* velocity. It would therefore appear that particle and gas velocities were very similar in this run, even though the terminal velocity of the fine particles was by no means negligible ($u_t = 0.24\,\mathrm{m\,s^{-1}} = 0.8\,U$). Although the thickness of the descending region might appear small, a 0.01 m thick annulus on the inside of a 0.24 m diameter cylinder actually accounts for 16% of the total freeboard area (so this 'thin' region can actually make quite a difference).

Rather different velocity profiles were found by Levy & Lockwood (1983). Their experiments were carried out in a rectangular bed measuring 0.3 m × 0.6 m, in which sand particles either 0.4 mm or 1.0 mm in diameter were fluidised (along with some finer particles). Superficial velocities of about $2\,U_{mf}$ were typical. Using Laser Doppler Anemometry to measure time-averaged mean vertical gas velocities at various points in the freeboard, Levy & Lockwood were surprised to find that the velocities were greatest near the walls (Figure 2.6) even though their beds were reasonably deep. These results were obtained in runs made with both sets of particles; U and the viewing height z were also varied. Looking at the longer cross-section through the centre of their bed (spanning the 0.6 m distance between the two narrower walls), they found that the mean vertical velocity was usually slightly less than U in the central region less than 0.1 m away from the bed axis, gradually increased when moving further away from the bed axis (reaching a peak value of about $1.5\,U$ about 0.26 m away from the bed axis, *i.e.* 0.04 m away from either wall) and then decreased rapidly as the wall was approached. Thus, gas velocities somewhat larger than U were found, but not in the expected location.

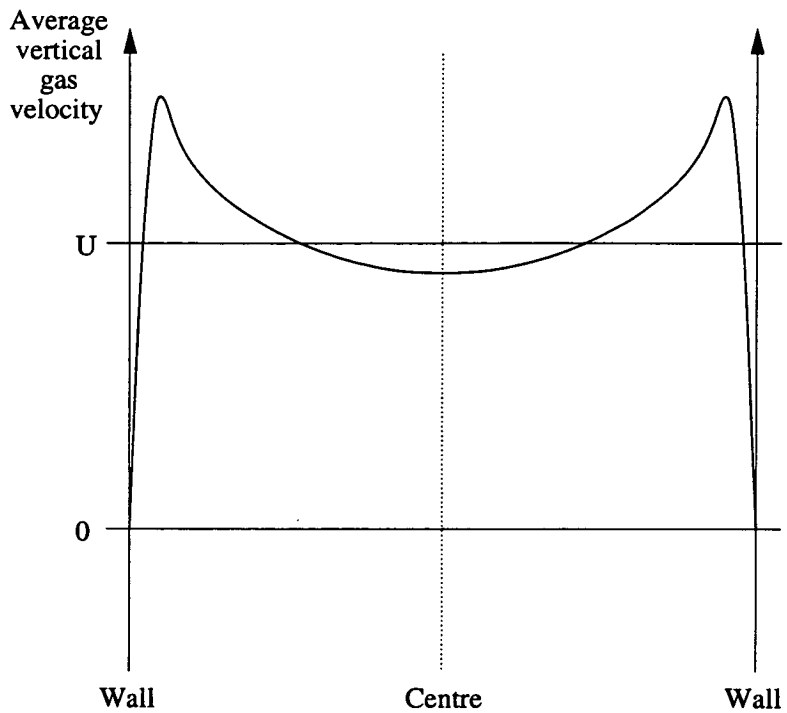


Figure 2.6: The time-averaged vertical gas velocity profile across the freeboard above a rectangular fluidised bed, measured by Levy & Lockwood (1983).

Levy & Lockwood (1983) also carried out bubble injection experiments. Single bubbles were injected into the centre of a fluidised bed, which was operated at $U = U_{mf} = 0.135\,\mathrm{m\,s^{-1}}$. Gas velocities were recorded, as a function of time, at various points in the centre of the freeboard (*e.g.* on the axis of the freeboard, at a height 0.060 m above the bed surface). They found that:

- Rapid upward gas and particle velocities were obtained at first (reaching over $1\,\mathrm{m\,s^{-1}}$).

- However, the velocities of both phases soon reversed their direction. Peak downward velocities were about $0.6\,\mathrm{m\,s^{-1}}$ for the gas and about $1.2\,\mathrm{m\,s^{-1}}$ for the particles.

- The series of particle velocity measurements stopped when all of the sand particles had returned to the bed surface, but measurements of gas velocity continued; this steadily became more upward, until a value close to the superficial velocity was reached.

- At greater freeboard heights (*e.g.* $z > 0.21\,\mathrm{m}$), the gas velocity remained close to zero after the bubble injection, suggesting that air upflow was taking place away from the bed axis.

Levy & Lockwood's explanation was that the bubble causes rapid upflow for only a short time after it bursts (about 0.2 s). After that, the particles fall back down towards the bed under gravity, dragging gas along too. This causes a vortex ring to be formed just above the bed surface, rotating downwards along the axis of the bed and upwards near the walls (*i.e.* in the opposite direction to the 'ghost bubbles' envisaged by Pemberton & Davidson (1984)). When single bubbles were injected, this 'down-in-centre, up-at-wall' vortex ring would be carried away by the fluidising gas. Smoke injection experiments were also carried out, as a method of visualising the flow; these confirmed the existence of the vortex ring.

Moving from these single-bubble injection experiments to their freely-bubbling bed, Levy & Lockwood (1983) suggested that bubbles rising in the centre of a fluidised bed might cause stop-start gas motion in the lowest part of the freeboard, setting up a down-at-centre, up-at-wall vortex ring. This would create strong upward gas flow near the walls, explaining their velocity profiles. Obviously these gas profiles were very different to the experimental results of Horio *et al.* (1980).

Lang (1986) investigated the flow in the freeboard region above rectangular fluidised beds, also using LDA. In his smaller apparatus ($0.075\,\mathrm{m} \times 0.075\,\mathrm{m}$), the vertical velocity profiles were usually rather flat. Glass ballotini particles (440–530 μm in diameter; $U_{mf} = 0.21\,\mathrm{m\,s^{-1}}$) were fluidised in air at $U = 0.3\,\mathrm{m\,s^{-1}}$ or $0.4\,\mathrm{m\,s^{-1}}$. The average vertical air velocity observed was usually about 20% greater than U, suggesting that the flow near the walls was only minimally upward or perhaps even downward (*cf* the results of Horio *et al.* (1980), as sketched earlier in Figure 2.5). The LDA system suffered from reflections at its perspex walls, which made it impossible to obtain velocity measurements any closer than 0.01 m away from the walls; as a result, gas downflow near the walls was suspected but could not be measured experimentally.

Lang's larger fluidised bed was 0.49 m × 0.28 m in cross-section. Round sand was used as the bed material; this had a broad size distribution function, with 50% by weight being less than about 120 μm. The minimum fluidising velocity was about 0.010 m s^{-1}. Air superficial velocities of 0.03 or 0.04 m s^{-1} were used. Viewing 1.2 m above a 0.60 m deep bed, Lang found that his mean gas velocities followed the 'dip' profiles reported by Levy & Lockwood (1983), with low vertical velocities found near the centre of the freeboard and large velocities observed near the walls (along the lines of Figure 2.6). Lang's measurements were taken by tracking along both of the lines which passed through the bed axis and ran parallel to the walls; the dip profile was observed in both cases.

Comparing the results of Levy & Lockwood (1983) to those of Lang (1986):

- Each mean air velocity when $U = 0.03$ m s^{-1} in Lang's work was less than the corresponding mean velocity at the same point when the higher air flowrate was used; Levy & Lockwood observed this too.
- In both sets of work, the greatest mean vertical velocity measured in any run was usually about $1.6\,U$.
- Lang's dip profile was rather more pronounced than Levy & Lockwood's. For example, Lang found some upward mean velocities of only $0.5\,U$ (or even lower), whereas Levy & Lockwood's lowest velocities were about $0.85\,U$. Further, Lang found that the mean vertical velocity was less than U at points in the centremost 70% of his cross-section, whereas Levy & Lockwood found this only for the centremost 30% of theirs.

The above results from Lang (1986) were taken using a 0.60 m deep fluidised bed, which was double the bed depth used by Levy & Lockwood (1983). Lang repeated his experiments using a lower bed depth of 0.30 m and halved his viewing height to $z = 0.60$ m (to maintain the same ratio of viewing height to bed depth). He did not, however, find a 'dip' profile under these conditions. At the same superficial air velocities as before, the measured vertical air profiles were relatively

flat (especially the run at $U = 0.04\,\mathrm{m\,s^{-1}}$), with the measured velocities oscillating about the superficial velocity (not some higher value).

Comparing the turbulence levels at each bed depth and superficial velocity, Lang noticed that more turbulence was found 0.60 m above the shallower bed when $U = 0.03\,\mathrm{m\,s^{-1}}$ than at the same height above the deeper bed, using the same velocity. This was the opposite of his observations in his smaller rig, where deeper beds produced more turbulence because they produced larger bubbles.

Using $U = 0.04\,\mathrm{m\,s^{-1}}$ instead, however, similar turbulence levels were found at both bed depths. For the shallower 0.30 m bed, the turbulence levels when $U = 0.03\,\mathrm{m\,s^{-1}}$ and $U = 0.04\,\mathrm{m\,s^{-1}}$ were not significantly different, but turbulence was greatly reduced when the superficial velocity was cut back to $0.02\,\mathrm{m\,s^{-1}}$.

Lang also found interesting results when he compared his time-averaged vertical gas velocity profiles at various heights above the bed surface. With a bed depth of 0.60 m and $U = 0.04\,\mathrm{m\,s^{-1}}$, all of the velocity profiles were of the dip pattern described above, for viewing heights of 1.2 m, 1.5 m and 1.8 m. However, much variation was found when runs with a bed depth of 0.30 m and a superficial velocity $U = 0.04\,\mathrm{m\,s^{-1}}$ were compared. No dip profile was found at viewing heights below 1 m, (*e.g.* as described above for the case $z = 0.60$ m). The profile was indeterminate when $z = 1.2$ m, but clear dip profiles were seen when the viewing height was increased to 1.5 m or 1.8 m. It would therefore appear that the dip profile found by Levy & Lockwood (1983) needs a certain amount of freeboard height in which to develop. At $z = 1.8$ m, the peak vertical gas velocity was about 2–3 U (near either wall), but the observed velocities were only 0.8–1.5 U at a viewing height of 0.9 m. The unexpected conclusion is that a 2 m-tall freeboard might produce larger vertical gas velocities (and therefore *more* entrainment, or entrainment of larger particles) than a 1 m freeboard.

Hamdullahpur & MacKay (1986) also used LDA to measure gas velocities above a bubbling fluidised bed. Their bed was rectangular, with a cross-section measuring 0.32 m × 0.18 m. Bed depths of 5 cm and 12 cm were used. The bed particles

were 300 μm sand, therefore a Geldart Class B material. The minimum fluidising velocity was not quoted, but was presumably about $0.1\,\mathrm{m\,s^{-1}}$.

Tracking across their apparatus at low freeboard heights, Hamdullahpur & MacKay found reasonably-constant time-average gas velocities near the bed axis, which then fell away linearly as the walls were approached. At intermediate heights, the width of the central plateau increased and dips became apparent; the average velocity had to decrease more sharply near the wall. At freeboard heights of about 30 cm when using $U = 0.20\,\mathrm{m\,s^{-1}}$ (or about 65 cm when $U = 0.45\,\mathrm{m\,s^{-1}}$), the velocity profiles were very flat (at a value close to U). The turbulence intensity levels were much reduced at these greater heights.

Lang (1986) found that, under certain conditions, his dip profiles became more pronounced at greater freeboard heights. In contrast, Hamdullahpur & MacKay (1986) found that their velocity profiles became much flatter at greater freeboard heights. Hamdullahpur & MacKay had about 4 m of freeboard space above the distributor, but only investigated the lowest 0.8 m of it.

Bed depths should also be considered. Levy & Lockwood (1983) used bed depths (*e.g.* 0.3 m) which were comparable to one of the side lengths of their rectangular sections (0.3 m × 0.6 m), as did Lang (1986). With these deep beds, it would be expected that large bubbles will form by coalescence and then rise near the bed axis (Lin *et al.*, 1985). In Levy & Lockwood's mechanism for the creation of a down-at-centre vortex ring for gas in the freeboard, these large bubbles eject a large mass of bed particles into the centre of the freeboard whenever they erupt. When the bed particles fall back down, a vortex ring is formed. Hamdullahpur & MacKay (1986) used a bed depth of only 5 cm in most of their published results, so their bed might well have produced many small bubbles throughout the cross-section, rather than fewer but larger bubbles near the centre. Levy & Lockwood's vortex ring might not have been able to develop.

Yórquez & Duursma (1999) are currently doing PIV research in the Department of Chemical Engineering at this university, investigating the flow in the freeboard

region above a 0.25 m-diameter fluidised bed. Yórquez has carried out similar single-bubble experiments to those described by Levy & Lockwood (1983), using a light sheet from a laser to illuminate a thin section of the freeboard region above a fluidised bed. Her bed material is sand, with some small hollow glass spheres added to 'seed' the flow. Large bed depths have been used, *e.g.* 0.5–1.0 m. The flow recorded with this PIV system has been difficult to analyse using simple PIV methods, *i.e.* photographic film to record instantaneous snapshots of the flow and digital analysis techniques to determine velocities at various points within the photographed regions. Nevertheless, the light sheet technique allows the flow to be visualised and recorded on video, as the particles illuminated by the laser are very much brighter than the other particles in the freeboard.

When single bubbles were injected into her fluidised bed, Yórquez's video shows that seeding particles in the fluidised bed do not penetrate very far into the centre of the freeboard after each bubble bursts. A reasonable time was allowed between bubbles, so most of the particles ejected into the freeboard by one bubble were removed before the next one arrived. If the models of jets into the freeboard (Zenz & Weil, 1958) or ghost bubbles (Pemberton & Davidson, 1984) applied, it would be expected that newly-ejected seeding particles would rise up the freeboard at great speed. This would have been easy to observe, because the centre of the freeboard was illuminated by the light sheet and was initially relatively free of particles. Instead, the mechanism of Levy & Lockwood (1983) seems more appropriate.

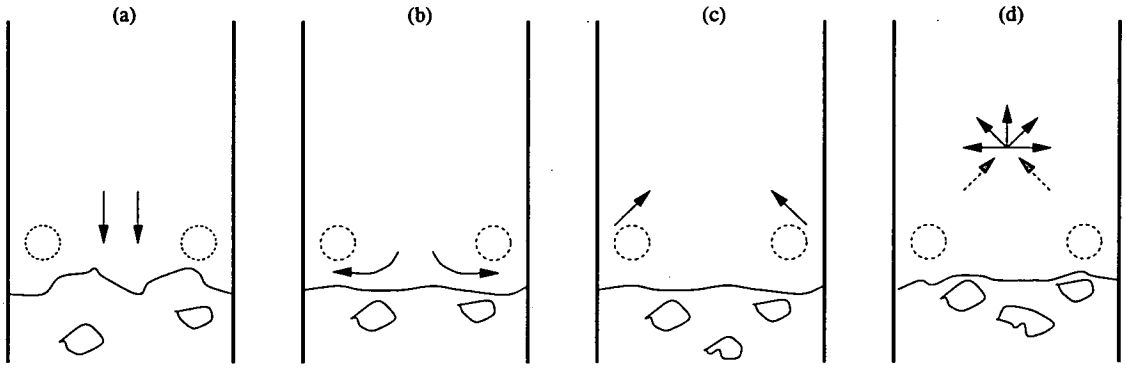


Figure 2.7: Sketches of the flow sequence observed by Yórquez & Duursma (1999), when an injected bubble bursts through the surface of an incipiently-fluidised bed.

EX BIBL. UNIV. EDINBURGH

The series of events observed on video, for a thin slice through the freeboard centre during a sequence of 'individual' bubbles, is sketched in Figure 2.7.

1. When a bubble was injected into the fluidised bed, close to the distributor, the bed surface (and objects in the freeboard) rose a constant distance, due to the extra gas within the bed. A short time then passed during which nothing much happened in the freeboard, whilst the bubble worked its way up through the fluidised bed.

2. As the bubble approached the surface, a bulge was formed. The bulge had a peak height of about 10 cm above the bed surface and a width of about 15–20 cm. The bubble itself would probably be smaller than these dimensions.

3. As each bulge reached its peak height, the sand particles from the bed had largely drained away from the top surface of the bubble or had begun to fall back. Seeding particles appeared from within the bubble, moving upwards.

4. The bubble 'burst' and the sand particles fell back down. Sand particles did not appear to be ejected any significant distance into the freeboard. Gas and seeding particles were sucked back down, from all points in the freeboard less than about 10 cm from the bubble centre (Figure 2.7, (a)).

5. The newly-ejected seeding particles bounced off the bed surface (close to its centre) and then moved radially outwards along the bed surface, at great speed. (Figure 2.7, (b)). This may or may not be related to the prior movement of bed particles in the initial bulge; some of these particles moved radially outwards when draining from the bulge, before it burst.

6. These seeding particle 'jets' hit the wall and bounced back towards the centre of the bed. They also started to move upwards (Figure 2.7, (c)). At any rate, 'jets' were seen using the two-dimensional light sheet, but in reality these will have been part of a continuous three-dimensional stream of material, hitting the flat bed surface by moving downwards after each bubble collapsed and then spreading out radially, moving approximately parallel to

the bed surface until it hit the wall. If a vortex ring structure exists just above the bed surface, then this material moves around it.

7. Some seeding particles from this bubble returned to the axis of the freeboard, after hitting the wall. Some of these particles did so after originally moving out of the plane of the light sheet, hitting an unilluminated part of the wall before bouncing back into the illuminated region.

8. The 'jets' bounced back off the walls and collided with 'each other' near the centre of the freeboard (Figure 2.7, (d)). A lot of turbulence was produced. Further 'jets' were created, some of which bounced back off unilluminated parts of the wall before returning to the centre of the freeboard (at heights of about 0–20 cm above the bed surface).

9. Some recirculating regions could be observed near the walls, after some of the bubbles. These vortices gradually rose up the freeboard, apparently clinging to the walls.

The general seeding particle movement was down in the centre of the freeboard and up at the walls. Asymmetries were present, however, *e.g.* the recirculating regions near a wall were not always paired with an equivalent vortex on the opposite site of the freeboard (thus, a full toroidal vortex was not always formed) and some bubbles gave rise to no vortices at all. Possibly because the bubbles rose slightly away from the bed axis, the 'jets' of seeding particles bouncing off the walls moved back towards the centre by differing amounts and collided with 'each other' in different places. Also, the freeboard region immediately above the bed surface was not necessarily uniform before each bubble burst, due to some turbulence (and seeding particles from previous bubbles) still being present. These extra seeding particles made it possible to observe the flow in more of the freeboard, but the turbulence meant that the initial condition of the freeboard was not reproducible.

Similar video work by Yórquez & Duursma (1999) has allowed the flow above a freely-bubbling fluidised bed to be videoed at $U = 2\,U_{mf}$. The continuous bubbling action meant that paired recirculation regions were found at the foot of

the freeboard almost all of the time, due to the jets of gas which were sucked down by descending bed particles after each bubble burst. The centres of these regions were located about 8 cm above the bed surface and about 5 cm in from the wall. The flow was reasonably well described by a model of a stationary down-in-centre, up-at-wall vortex, but the flow in the freeboard was also somewhat chaotic and stop-start (*e.g.* the flow in the centre of the freeboard took one step upwards then two or three steps downwards).

Overall, the model of Levy & Lockwood (1983) appears to have merit, although it must be recognised that it is a simplification. For example, whilst the vortex ring structure they describe undoubtedly exists, it seems to direct newly-ejected gas towards a point just above the centre of the bed surface (*i.e.* *not* directly upwards, along the wall). Results presented by Lang (1986) suggested that Levy & Lockwood's velocity profile only develops somewhat higher above the bed surface, so a vortex ring near the bed surface might not be the main cause of the dip velocity profiles observed by Levy & Lockwood.

2.4.6 Laminar flow in a square-sided conduit

Whilst noting the above observations for flow in the freeboard of cylindrical and rectangular fluidised beds, it is still worth considering what the fully-developed steady gas flow in a very tall freeboard section will be, in conditions of low particle concentration (*e.g.* when $U \sim U_{mf}$). Square-sided fluidised beds were used in this research for reasons of optical convenience, but this choice of geometry means that the standard laminar and turbulent velocity profiles for flow in a circular pipe are not applicable.

In this experimental work, $\rho_f = 1.2\,\mathrm{kg\,m^{-3}}$ and $\mu = 1.8 \times 10^{-5}\,\mathrm{N\,s\,m^{-2}}$ for air at room temperature and pressure, whilst D (a characteristic length) will be 0.30 m for the larger fluidised bed. This means that a superficial air velocity of $U = U_{mf} = 0.10\,\mathrm{m\,s^{-1}}$ gives a Reynolds number of about 2000. The superficial velocities used were about $0.07\,\mathrm{m\,s^{-1}}$ to $0.20\,\mathrm{m\,s^{-1}}$, so laminar flow would be expected to develop at the top of a very tall freeboard above the larger fluidised

bed if the particle concentration in the freeboard is low and if relatively low air flowrates are used, whereas the flow at the higher gas flowrates will probably have made the transition to turbulent flow. Meanwhile, the limiting velocity profile in the smaller bed ($D = 0.076\,\mathrm{m}$) will be laminar at mild fluidisation conditions; a superficial velocity as high as $0.4\,\mathrm{m\,s^{-1}}$ ($U > 5\,U_{mf}$) gives a Reynolds number of only 2000. Thus, the fully-developed laminar velocity profile in a square duct is potentially of interest when considering all work carried out in the smaller fluidised bed and the lower-velocity work carried out using the larger bed. On the other hand, large-scale industrial fluidised beds might well exhibit turbulent flow.

Knudsen & Katz (1958) quoted work by Timoshenko & Goodier in 1951, in which the torsion in rectangular beams was solved. Timoshenko and Goodier found a series solution for the torsion profile in a rectangular beam of dimensions $a \times b$, with perpendicular co-ordinates x and z aligned with the edges of the beam and intersecting at its axis, where $x = z = 0$. In the geometry quoted here, the edges are given by the lines $x = \pm b/2$ and $z = \pm a/2$.

Laminar flow in a vertical rectangular duct is an analogous problem, with boundary conditions given by similar differential equations. The differential equation for the vertical fluid velocity u_y (parallel to the axis of the duct) is:

$$\frac{\partial^2 u_y}{\partial x^2} + \frac{\partial^2 u_y}{\partial z^2} = -k_1^2, \tag{2.12}$$

for some real k_1. The boundary conditions are that $u_y = 0$ at each wall and that $\partial u_y/\partial x = \partial u_y/\partial z = 0$ when $x = z = 0$ (*i.e.* no velocity gradient across the axis).

When the solution is translated into the fluid mechanical situation, the velocities in the vertical y-direction are given by:

$$u_y = k \sum_{n=1,3,5,\ldots} \frac{(-1)^{(n-1)/2}}{n^3} \left(1 - \frac{\cosh(n\pi x/a)}{\cosh(n\pi b/2a)}\right) \cos\left(\frac{n\pi z}{a}\right),$$

where k is a constant which is proportional to the pressure gradient along the pipe ($\partial P/\partial y$) and inversely proportional to the fluid viscosity.

It can be noted that the (1−cosh) term gives $u_y = 0$ at $x = \pm b/2$, whilst the cos term gives $u_y = 0$ at $z = \pm a/2$. There is, therefore, no slip at any wall. Further, the first derivatives will be proportional to either $\sinh(n\pi x/a)$ or $\sin(n\pi z/a)$, so $\partial u_y/\partial x = 0$ along $x = 0$ and $\partial u_y/\partial z = 0$ along $z = 0$: no velocity gradients are encountered when moving across the x-axis or the z-axis. Thus, the boundary conditions appear to be satisfied.

Arbitrarily choosing $a = b = 100$ units to model the square geometry used in this research, this result is plotted in Figure 2.8 with the velocity scaled so that its maximum value is 1.000 units s^{-1} along the bed axis. In the left-hand half of this Figure, contours of equal vertical velocity within a horizontal cross-section through the flow are shown, for vertical velocities of 0.05, 0.15, 0.25, ..., 0.95 units s^{-1}. Selected velocity profiles in the x-direction, along various constant values of z, are shown in the right-hand half of this Figure. For example, the continuous curve represents the velocity profile obtained by varying x when $z = 0$.

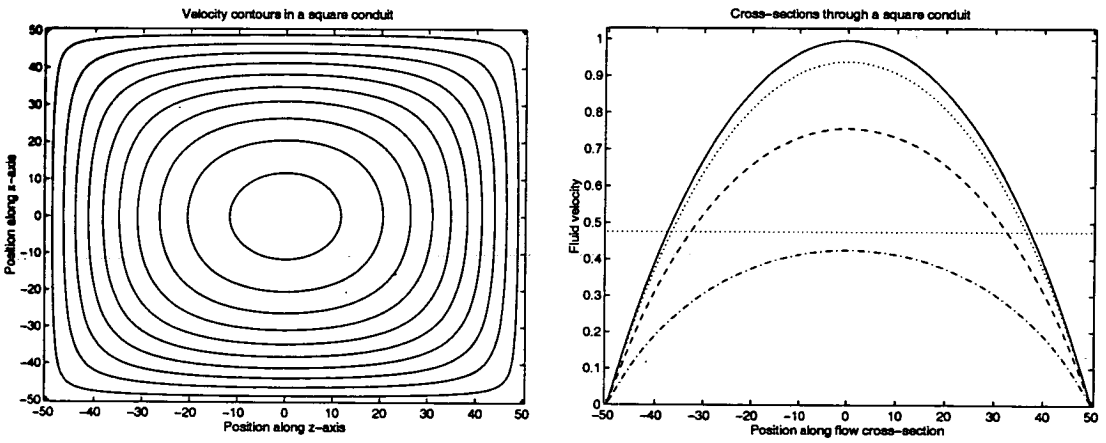


Figure 2.8: The fully-developed laminar velocity profile for flow in a square duct with a vertical axis, calculated using the result quoted by Knudsen & Katz (1958). Left: Contours of equal vertical velocity, for points in a 100×100 horizontal cross-section through the duct. Right: The velocity profiles along various lines in the left-hand figure, namely $z = \pm 38$ (dash-dot curve), $z = \pm 25$ (dashed curve), $z = \pm 12$ (dotted curve) and $z = 0$ (continuous curve).

On this scale, taking the first 25 terms of the series expansion and averaging across the whole of the flow area reveals that the 'superficial' mean y-velocity is $\overline{u_y} = 0.47692$ units s^{-1}, *i.e.* 47.69% of the axial velocity (Figure 2.8 right, dotted horizontal line). This was calculated by evaluating u_y at 1 unit intervals across the

100×100 section, summing u_y and then dividing by the area, *i.e.* by evaluating a numerical approximation to $a^2\,\overline{u_y} = \int u_y\,dA$ (for area elements dA).

An examination of the profile along $z = 0$ shows that the arithmetic mean velocity along a section through a square bed, running parallel to one pair of walls and passing through the bed axis, is 0.68653 units s^{-1}. This is obtained by evaluating $\int u_y\,dx/a$ along the path $z = 0$. The difference between this value and the superficial velocity is simply due to the different method of averaging.

The significance of this is that the 'mean' vertical velocity obtained by PIV in most of this work should be about 44.0% greater than the superficial velocity, because test-sections through the centre of the bed include the largest velocities found in the flow and ignore the corners, where the velocities are significantly lower. Put another way, an averaging method using data from PIV takes into account the frequency with which each vertical velocity is found along only one particular path across the test section, whereas the superficial velocity takes into account the frequencies across the whole of the section. Similarly, the magnitude of the greatest vertical velocity in each picture should be slightly more than double the superficial velocity (1/0.47692=2.0968 if laminar flow applies). This factor of 2.0968 is actually quite close to the factor of exactly 2 which is obtained when the same equations are solved for a cylindrical system.

In the PIV experiments in this work, vertical sections through the centre of a square-sided duct were photographed. Assuming that fully-developed laminar flow is found, a superficial velocity of 0.0700 m s^{-1} should give a 'mean' velocity along the test section of 0.1008 m s^{-1} and a maximum velocity of 0.1468 m s^{-1}.

Dryden *et al.* (1956) presented a similar solution. Equation 2.12 can be solved by assuming that the velocity is the sum of two functions $f_1(x)$ and $f_2(x, z)$:

$$
\begin{aligned}
u_y(x,z) &= f_1(x) + f_2(x,z), \\
\text{where} \qquad \frac{d^2 f_1}{dx^2} &= -k_1^2 \\
\text{and} \quad \frac{\partial^2 f_2}{dx^2} + \frac{\partial^2 f_2}{dz^2} &= 0.
\end{aligned}
$$

The function $f_1(x)$ is found to be a quadratic, by integrating the simple Ordinary Differential Equation which it must satisfy. The function $f_2(x,z)$ is solved using a separable variables method. The two functions are added together, then the boundary conditions are applied. The solution is:

$$
\frac{u_y}{k} = \left(1 - \frac{4x^2}{b^2}\right) - \frac{32}{\pi^3} \sum_{n=1,3,5,\ldots} \frac{(-1)^{(n-1)/2}}{n^3} \frac{\cosh(n\pi x/a)}{\cosh(n\pi b/2a)} \cos\left(\frac{n\pi z}{a}\right).
$$

Setting $a = b$ in order to obtain a square-sided duct and choosing k so that the maximum velocity is $u_y = 1.000$ units s^{-1} at $x = z = 0$ when the first 25 terms are taken, it is found that the superficial velocity is 0.47693 units s^{-1} (averaging over the whole of the cross-sectional area) and that the centreline profile has a mean velocity of 0.68654 units s^{-1} (averaging along either $x = 0$ or $z = 0$). Thus, a superficial mean velocity of $0.0700\,\mathrm{m\,s^{-1}}$ gives a central-section 'mean' velocity of $0.1008\,\mathrm{m\,s^{-1}}$ and a maximum vertical velocity of $0.1468\,\mathrm{m\,s^{-1}}$.

This agrees with the results from the equation given by Knudsen & Katz (1958) for the number of significant figures quoted. Indeed, if both solutions are correct, they should turn out to be identical if the summation involves very many terms (rather than just the values $n = 1, 3$ or 5, as listed by Knudsen & Katz). Both series converge rapidly over the first few terms, but less quickly later on. Considering only the first three terms gives an error of only 1%.

Further, the result for a square duct should be independent of the choice of axes, *i.e.* the dependence of u_y on x should be the same as its dependence on z in the special case where $a = b$. Whilst proving this looks to be difficult, the similarities can be demonstrated numerically by plotting the two equations.

Chapter 3

Stream functions for bubbles in gas-fluidised beds

3.1 Preamble

In this chapter, the simplest potential flow analyses of a fluidisation bubble will be considered. According to the 'ghost bubble' theory of Pemberton & Davidson (1984), a vortex ring circulation develops in and around each rising bubble in a fluidised bed and persists into the freeboard region after the bubble passes through the bed surface. Potential flow models have been developed which include growth terms, so as to model the increases in size observed in experimental fluidised beds and to allow better predictions of the bubble shape to be made.

3.1.1 Frames of reference

One problem encountered in this flow measurement work is that the potential flow solutions for rising bubbles are worked out in a bubble reference frame, *i.e.* a frame of reference which rises at the same speed as the bubble, thus making the bubble appear stationary. In contrast, a PIV photograph uses a fixed camera, making the solid fluidised bed appear stationary but allowing bubbles to move.

The velocities observed by PIV, therefore, can only be expected to resemble the vortex ring velocity profile predicted by Pemberton & Davidson if PIV vector maps of the freeboard region are obtained in the usual way and are then either translated into a ghost bubble reference frame (by subtracting the ghost bubble rise velocity from the experimental velocities) or else compared to translations of ghost bubble flow into the laboratory frame (obtained by adding the ghost bubble rise velocity to the theoretical ghost bubble flow pattern).

The required ghost bubble rise velocity in the freeboard region will probably be between zero and the expected bubble rise velocity at the top of the fluidised bed, as ghost bubbles are expected to slow down during their ascent through the freeboard region (due to the entrainment of slow-moving material with which momentum needs to be shared). As this ghost bubble rise velocity is somewhat uncertain, a number of different frames of reference may need to be considered if the second approach is adopted, each based on a particular vertical ghost bubble velocity (measured relative to the laboratory frame).

3.1.2 Co-ordinate system

An inviscid, irrotational flow in two-dimensions can be described by a stream function (ψ), a potential function (ϕ) or velocity components in two orthogonal directions. These descriptions are equivalent, *i.e.* given one of the three, the other two can be determined.

For flow past a fluidisation bubble, it is convenient to use polar co-ordinates with the angle $\theta = 0$ in the vertical direction and the origin ($r = 0$) at the bubble centre. The angle is chosen to lie in the range $-\pi < \theta \leq \pi$. In polar co-ordinates, the velocity terms at a point (r, θ) are v_r radially outwards and v_θ in the anticlockwise tangential direction, as shown in the left-hand side of Figure 3.1. The points indicated in the right-hand side of this figure have co-ordinates $A = (r, \theta)$, $B = (r, \theta + \delta\theta)$ and $C = (r + \delta r, \theta + \delta\theta)$).

The Cartesian equivalent of this system has the x-axis vertical, the y-axis horizontal and the origin at the bubble centre.

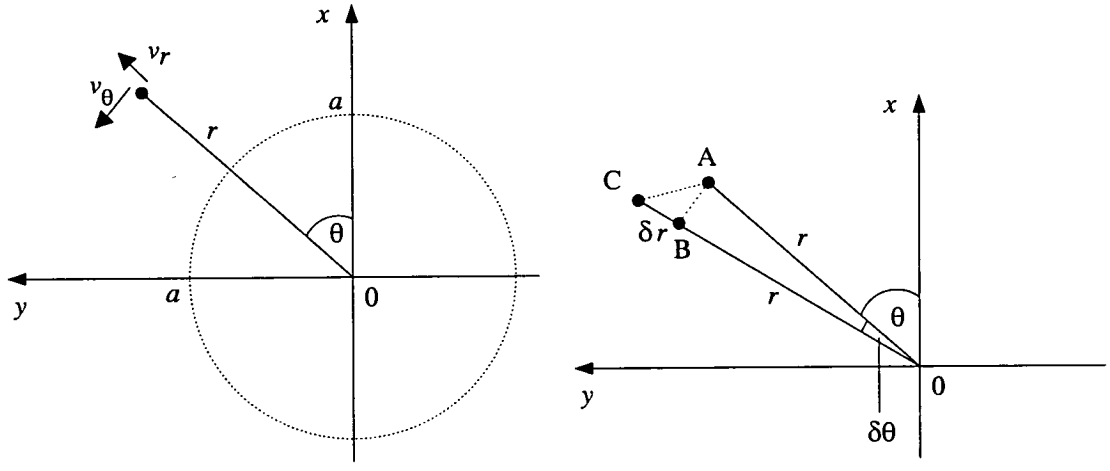


Figure 3.1: The co-ordinate system around a bubble of radius a, showing the direction of the velocity components v_r and v_θ (left); three closely-spaced points A, B and C (right).

3.1.3 Stream functions and velocities

The stream function is defined by a volume balance for incompressible flow. The medium is porous, with a constant mean void fraction ε. The stream function at a point A in a given situation is a measure of the flow rate across a line between the origin and the point A. Contours of stream function are parallel to the direction of the flow velocity vectors, as no flow can cross these streamlines.

In two dimensions, the sign convention of Davidson (Various, 1961) will be followed. This takes the stream function ψ to be positive remote from the bubble in the $0 < \theta < \pi$ quadrants, where $y > 0$.

A function $\psi = f(r, \theta)$ can be differentiated to yield:

$$\delta\psi = \frac{\partial\psi}{\partial r}\delta r + \frac{\partial\psi}{\partial\theta}\delta\theta. \tag{3.1}$$

The volume balance from A to C via point B (Figure 3.1, right) is used with the definition of ψ to give:

$$\begin{aligned}
\psi(C) - \psi(A) &= (\psi(C) - \psi(B)) - (\psi(A) - \psi(B)) \\
&= \text{flow out through } BC + \text{flow out through } AB \\
\Rightarrow \delta\psi &= [\delta r.\varepsilon](v_\theta) + [r\delta\theta.\varepsilon](-v_r), \qquad (3.2)
\end{aligned}$$

assuming that the void fraction ε is constant in the region ABC. The velocities v_r and v_θ are interstitial local gas velocities.

Comparing the terms in δr and $\delta\theta$ in equations 3.1 and 3.2:

$$v_r = -\frac{1}{r\varepsilon}\frac{\partial\psi}{\partial\theta} \qquad (3.3)$$

$$\text{and } v_\theta = \frac{1}{\varepsilon}\frac{\partial\psi}{\partial r}. \qquad (3.4)$$

Usually $\varepsilon = 1$, so the inclusion of the void fraction has no noticeable effect. If the gas flows through regions of varying voidage, as in a fluidised bed, its significance becomes clear when streamlines cross the bubble boundary and enter a region of significantly different voidage. If defined in this manner, the stream function can be made continuous even when it crosses the bubble boundary.

The equivalent expressions for three-dimensional axi-symmetric flow are:

$$v_r = -\frac{1}{\varepsilon\, r^2 \sin\theta}\frac{\partial\psi}{\partial\theta} \qquad (3.5)$$

$$\text{and } v_\theta = \frac{1}{\varepsilon\, r \sin\theta}\frac{\partial\psi}{\partial r}. \qquad (3.6)$$

3.2 Stream functions for fluidisation bubbles

For a minimum fluidising velocity U_{mf} and a mean void fraction ε in the fluidising medium, the average interstitial velocity for gas flowing between particles in the dense phase is given by:

$$U_o = \frac{U_{mf}}{\varepsilon}. \tag{3.7}$$

According to the two-phase theory of fluidisation, $\varepsilon = \varepsilon_{mf}$ throughout the dense phase of a bubbling fluidised bed, if $U > U_{mf}$. If so, U_o must be constant and any gas flow in excess of $A\,U_{mf}$ must rise through the bed as bubbles. For a Geldart Class A material, U_{mf} should be replaced by U_{mb}.

3.2.1 Two-dimensional bubbles, constant volume

Davidson (Various, 1961) derived stream and potential functions for the flow of fluidising gas around a cylindrical void of radius a and of constant volume, rising at a velocity U_b (in the laboratory frame) in a 2-D fluidised bed. Working in a reference frame moving with the bubble, Davidson modelled the particles as moving in streamline flow around the apparently-stationary void and then calculated interstitial gas velocities using Darcy's Law. Here, Davidson's stream function has been multiplied by ε to make it consistent with the modified definition and subscripts have been added to emphasize that this applies for flow outside the bubble only. Later, a subscript 'i' will be used to denote a property which only applies inside the bubble.

Davidson's stream function ψ_o outside the bubble (*i.e.* where $r > a$) was:

$$\Rightarrow \psi_o = \varepsilon\,(U_b - U_o)\left(1 - \frac{A_2^2}{r^2}\right) r\sin\theta, \tag{3.8}$$

where the constant length A_2, given by

$$A_2 = a\left(\frac{U_b + U_o}{U_b - U_o}\right)^{1/2} \tag{3.9}$$

is the radius of a cylindrical 'limit of penetration' when $U_b > U_o$.

After differentiation using equations 3.3 and 3.4, this gives rise to velocities:

$$v_{r,o} = -(U_b - U_o)\left(1 - \frac{A_2^2}{r^2}\right)\cos\theta \tag{3.10}$$

$$\text{and } v_{\theta,o} = (U_b - U_o)\left(1 + \frac{A_2^2}{r^2}\right)\sin\theta. \tag{3.11}$$

Whereas the particles have zero radial velocity at the bubble surface, the gas has $v_{r,o} = 0$ at $r = A_2$. If $U_b > U_o$, as is the case with Geldart B materials, this circle is concentric with the bubble boundary ($r = a$) but has a larger diameter. Gas in regions where $r > A_2$ flows by just as it would flow around an imaginary cylindrical object of radius A_2, *i.e.* no gas flows into or out of the zone bounded by the limit of penetration ($\psi_o = 0$ along $r = A_2$).

At the bubble boundary, setting $r = a$ in equations 3.8, 3.10 and 3.11 shows that:

$$\psi_o = -2\,U_{mf}\,a\,\sin\theta,$$

$$v_{r,o} = 2\,U_o\,\cos\theta \tag{3.12}$$

$$\text{and } v_{\theta,o} = 2\,U_b\,\sin\theta. \tag{3.13}$$

Davidson's analysis assumed that the gas flow was incompressible and that the void fraction was constant. Collins (1989) investigated the results of allowing ε to vary as a function of r, but found that the differences were small.

Pyle & Rose (1965) extended Davidson's work by assuming that streamline flow continues into the bubble. They used the following boundary conditions:

1. the tangential velocity of the gas at the bubble surface is the same as the local particle velocity, for gas just outside or just inside the bubble, *i.e.*

$$v_{\theta,i} = v_{\theta,o} \tag{3.14}$$

2. continuity is satisfied for the gas flow across any small element of the bubble surface. There is a sudden decrease in the magnitude of v_r on moving from

a region of low voidage ($\varepsilon \sim 0.5$) into a particle-free bubble, where the void fraction $\varepsilon_b = 1$. Radial velocity is therefore not continuous across the bubble surface:

$$v_{r,i} = \varepsilon\, v_{r,o}. \tag{3.15}$$

They then obtained:

$$\psi_i = 2\, U_{mf} \left[(1+\beta)\,\frac{r}{a} - (2+\beta)\right] r \sin\theta, \tag{3.16}$$

where

$$\beta = \frac{U_b}{U_{mf}} = \frac{U_b}{\varepsilon\, U_o}. \tag{3.17}$$

No voidage term appears in the expression for ψ_i, because the void fraction ε_b inside the bubble has been taken to be one. Evaluated inside the bubble, $U_{mf}/\varepsilon_b = U_{mf}$, not U_o.

Differentiating:

$$v_{r,i} = -2\, U_{mf} \left[(1+\beta)\,\frac{r}{a} - (2+\beta)\right] \cos\theta \tag{3.18}$$

$$\text{and} \quad v_{\theta,i} = 2\, U_{mf} \left[2\,(1+\beta)\,\frac{r}{a} - (2+\beta)\right] \sin\theta. \tag{3.19}$$

On again setting $r = a$, it is found that:

$$\begin{aligned} \psi_i &= -2\, U_{mf}\, a \sin\theta = \psi_o, \\ v_{r,i} &= 2\, U_{mf} \cos\theta = \varepsilon\, v_{r,o} \\ \text{and} \quad v_{\theta,i} &= 2\, U_b \sin\theta = v_{\theta,o}. \end{aligned}$$

It can be seen that the stream functions used by Pyle & Rose do indeed give rise to velocities which satisfy their boundary conditions at the bubble surface, *cf*

equations 3.14 and 3.15. Further, the introduction of the void fraction ε into the definition of stream function ensures that $\psi_i = \psi_o$ at the bubble boundary.

A plot of velocity vectors and stream function ψ against position in or around the bubble is shown in Figure 3.2. The bubble is centred on the point $(0, 0)$ in the Cartesian co-ordinates shown and has a radius of 10 units. The following velocities were used: $U_b = 2$, $U_o = 1$ and $U_{mf} = 0.5\,\text{units}\,\text{s}^{-1}$, in the arbitrary scale used here.

The contours of $\psi = 0$ should run down the line $y = 0$ and around the limit of penetration, a cylinder of radius 17.3 units concentric with the bubble. There is, however, some noise due to the vagaries of the contour-plotting method.

Figure 3.2 shows that gas within the zone bounded by the limit of penetration is trapped within a 'cloud' of radius A_2 and does not mix with other fluidisation gas. The gas within the cloud recirculates along closed streamlines, all of which pass through the bubble ($r = a$). However, not all of these closed streamlines pass through the denser gas-particle phase, so some of the fluidisation gas cannot ever leave the bubble.

Hence, all of the gas trapped within the cloud spends at least some of its time in the bubble, with some gas spending all of its time in the bubble. If a chemical reaction can take place which involves both the gas and the particles, then the gas which happens to be inside each bubble at any instant cannot react at that particular moment in time. Furthermore, some of this gas will completely bypass the particle phase and will not be able to react at all during the whole of the time the bubble spends within the bed. The gas in the freeboard region will be a mixture of interstitial gas (well reacted), bubble/cloud gas (partly reacted) and bubble-only gas (unreacted).

The flow both inside and outside the bubble is highly symmetric, *e.g.* the stream functions ψ_o and ψ_i are both symmetric about the horizontal y-axis ($x = 0$; *N.B.* $\sin\theta = \sin(\pi - \theta)$) and, albeit with a change of sign, about the vertical x-

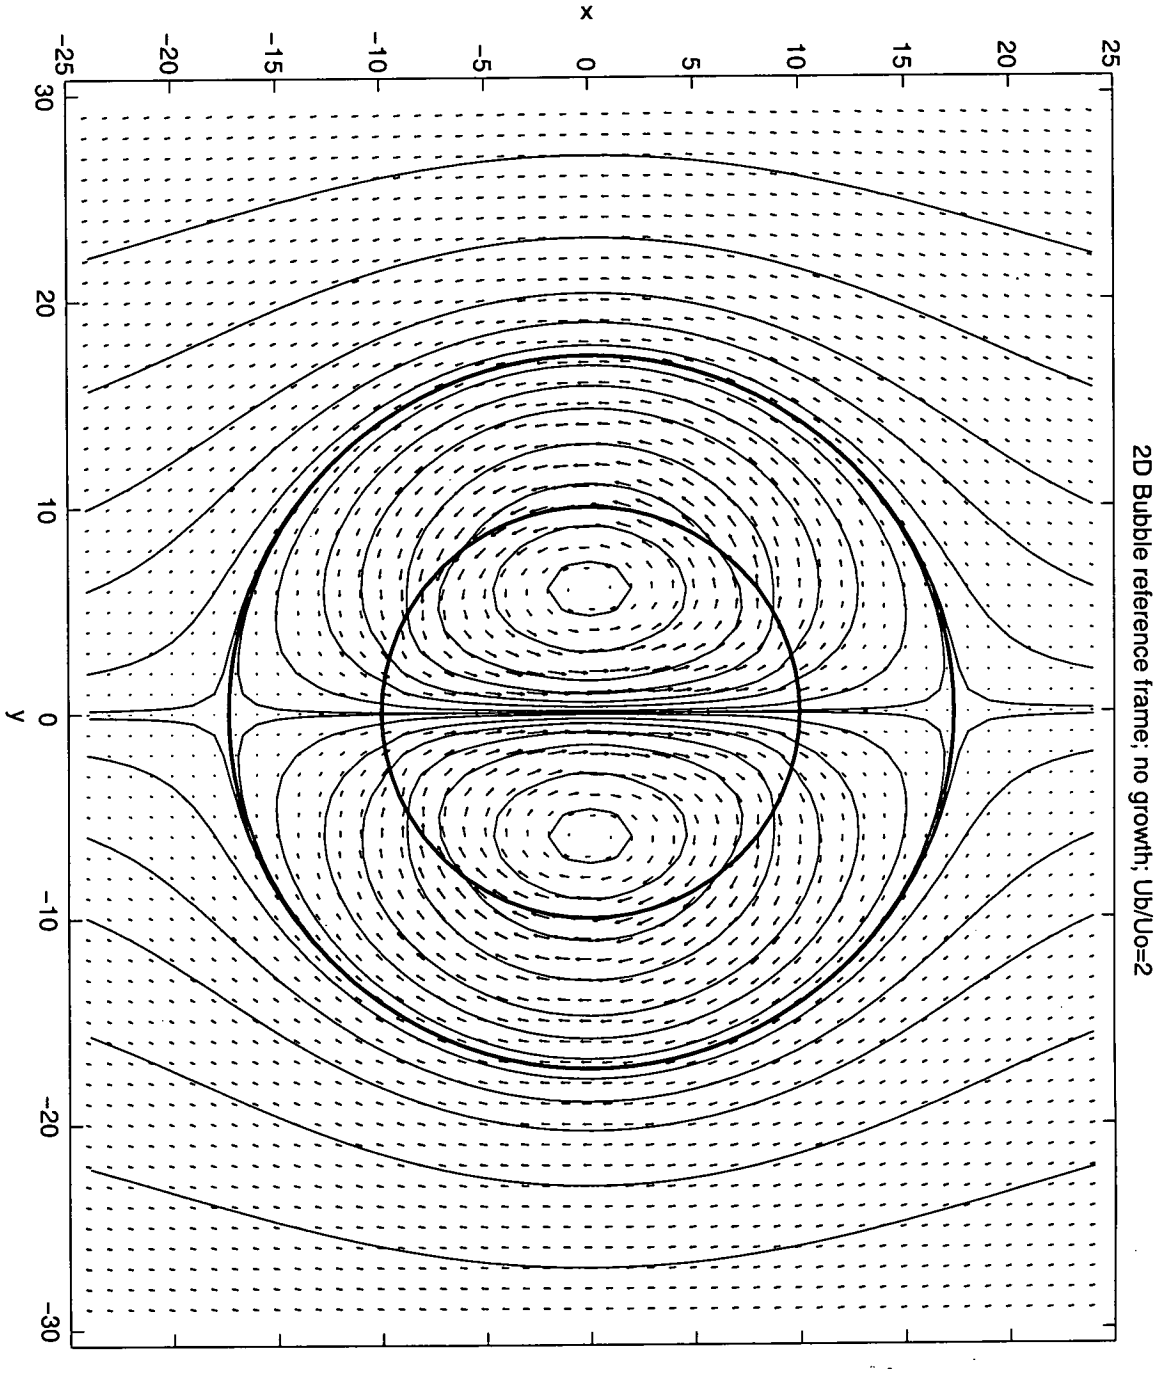


Figure 3.2: Velocity vectors and streamlines for a constant-volume 2D bubble

axis ($y = 0$; *N.B.* $\sin(-\theta) = -\sin\theta$). These symmetries mean that, in this case, velocities and streamlines need only be plotted in one quadrant in order to see every feature of the gas profile. The flows in the other quadrants will be simple reflections in the x- and/or y-axes of the flow in this quadrant.

Figure 3.3 shows plots of the vertical component of gas velocity from various different horizontal sections through or around the rising cylindrical bubble shown in Figure 3.2, along lines given by x = constant.

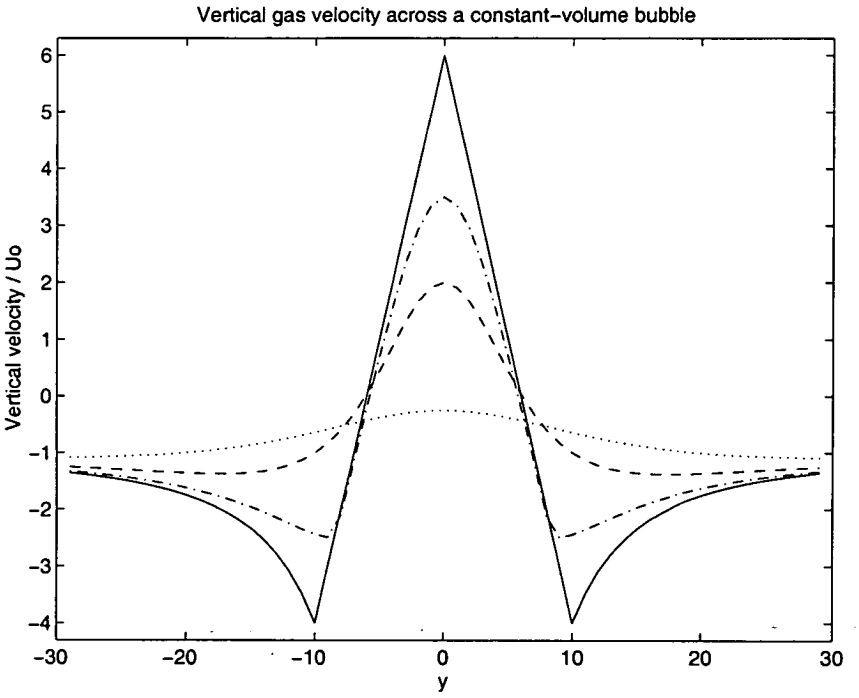


Figure 3.3: The vertical velocity along horizontal sections through a 2-D bubble. These sections run through the bubble centre ($x = 0$, solid curve), half a diameter above or below the bubble centre ($x = \pm 5$, dash-dot curve), just missing the top or the bottom of the bubble ($x = \pm 10$, dashed curve) or two bubble diameters above/below the origin ($x = \pm 20$, dotted curve).

Sample plots of stream function contours for various other bubble rise velocities in this two-dimensional case are shown in Figure 3.4. These can be compared to Figure 3.2, where U_b/U_o was 2. In each graph, no velocity vectors are plotted and streamlines are shown using dotted lines. The positions of the bubble boundary and the limit of penetration (where appropriate) are also plotted, using continuous curves. The particular contours plotted were selected in order to illustrate features of the flow; the contours were not taken at equal intervals of stream function.

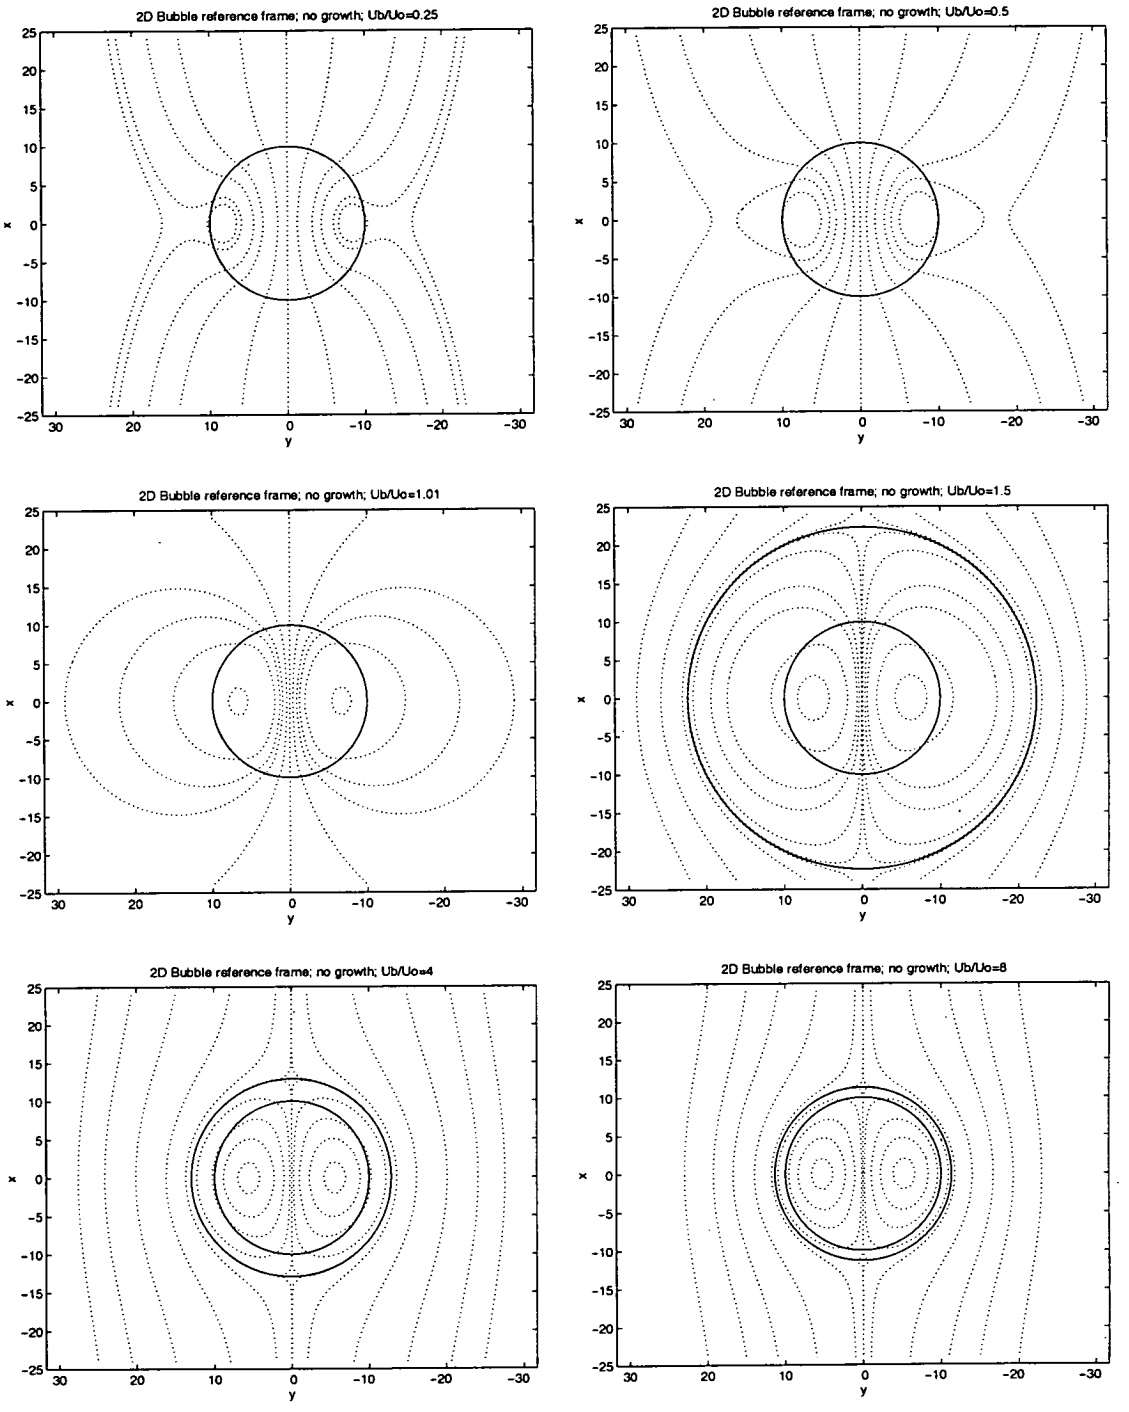


Figure 3.4: Contours of stream function (dotted curves) around a two-dimensional bubble (inner continuous circle), plotted for various bubble rise velocities: from left to right and from top to bottom, $U_b/U_o = 0.25$, 0.5, 1.01, 1.5, 4 or 8 units s^{-1}, with $U_o = 1$ unit/s and $U_{mf} = 0.5$ units s^{-1} in all cases.

The first two plots in Figure 3.4 show the case of 'slow' bubbles, which rise less quickly than the interstitial gas. Only the bubble boundary is shown here, as the limit of penetration (given by equation 3.9) is complex if $U_b/U_o < 1$. These bubbles contribute to the gas flow by acting as void regions, which offer a low-resistance route up the bed. This behaviour is typical of Geldart D materials.

When U_b/U_o is very slightly greater than one, the limit of penetration has a very large radius (equation 3.9). This is illustrated in Figure 3.4 for the case $U_b/U_o = 1.01$.

The 'fast' bubbles in Figure 3.4 rise at a higher speed than the interstitial gas ($U_b/U_o > 1$). Their limit of penetration is real and has a greater diameter than the bubble. Faster bubbles have smaller limits of penetration, so a greater proportion of the gas trapped within this cloud region spends all of its time within the bubble, never coming into contact with the particles. Faster bubbles also cause less deflection to streamlines passing a given horizontal distance away from the bubble centre (*e.g.* 1.5 bubble diameters).

3.2.2 Three-dimensional bubbles, no growth

Davidson (Various, 1961) presented the result outside a constant-volume spherical bubble in a three-dimensional fluidised bed, namely:

$$\psi_o = \varepsilon\,(U_b - U_o)\left(1 - \frac{A_3^3}{r^3}\right)\frac{r^2\sin^2\theta}{2}, \tag{3.20}$$

which has been modified here to include the void fraction. When $U_b > U_o$, the limit of penetration describes a sphere, centred on the origin, with radius A_3:

$$A_3 = a\left(\frac{U_b + 2\,U_o}{U_b - U_o}\right)^{1/3}. \tag{3.21}$$

The velocities, obtained by differentiation and then equations 3.5 and 3.6, are:

$$v_{r,o} = -(U_b - U_o)\left(1 - \frac{A_3^3}{r^3}\right)\cos\theta \tag{3.22}$$

$$\text{and } v_{\theta,o} = (U_b - U_o)\left(1 + \frac{A_3^3}{2\,r^3}\right)\sin\theta. \tag{3.23}$$

Using the same boundary conditions as before, Pyle & Rose again derived a stream function for the inside of the bubble:

$$\psi_i = \frac{3\,U_{mf}}{4}\left[(2+\beta)\,\frac{r^2}{a^2} - (4+\beta)\right] r^2 \sin^2\theta, \tag{3.24}$$

which again does not appear to depend on the void fraction because $\varepsilon_b = 1$. The parameter β is the same as that used earlier (equation 3.17).

The velocities inside the bubble are:

$$v_{r,i} = -\frac{3\,U_{mf}}{2}\left[(2+\beta)\,\frac{r^2}{a^2} - (4+\beta)\right]\cos\theta \tag{3.25}$$

$$\text{and } v_{\theta,i} = \frac{3\,U_{mf}}{2}\left[2\,(2+\beta)\,\frac{r^2}{a^2} - (4+\beta)\right]\sin\theta. \tag{3.26}$$

At the bubble boundary, $r = a$, it can be shown that $\psi_o = \psi_i$ and that the velocity boundary conditions are satisfied (equations 3.14 and 3.15).

A plot of these velocity vectors and stream functions is shown in Figure 3.5, for the case $U_b/U_o = 2$.

In addition, stream function contours are plotted in Figure 3.6, for various bubble speeds. These graphs are quite similar to those obtained by solving the two-dimensional case.

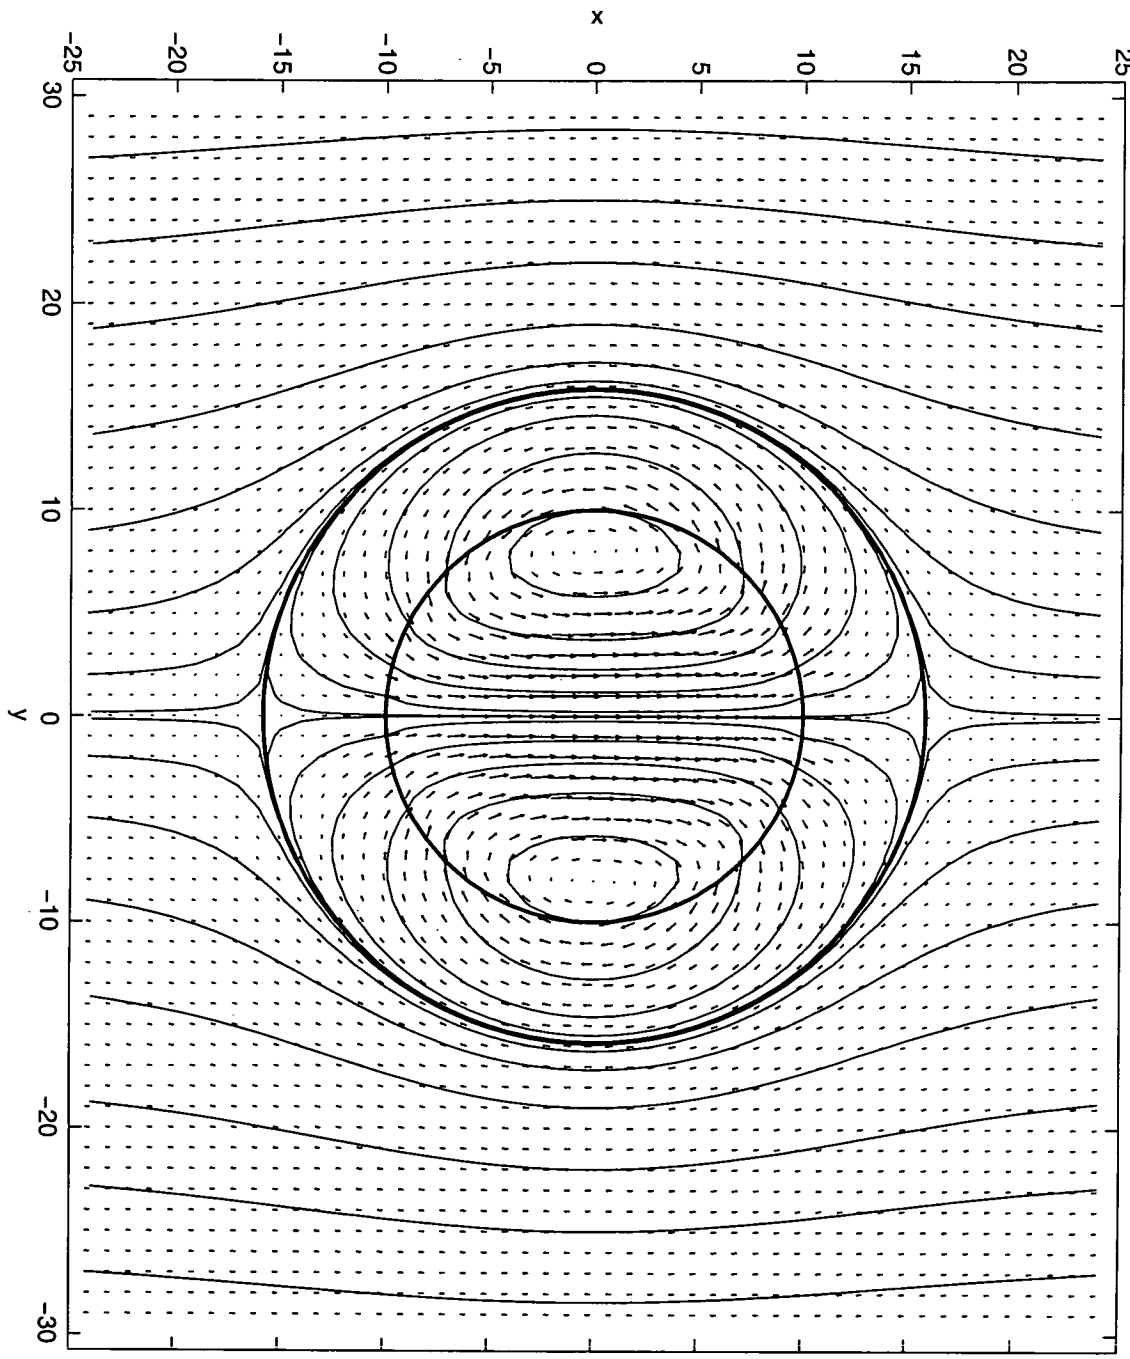


Figure 3.5: Velocity vectors and streamlines for a constant-volume 3D bubble

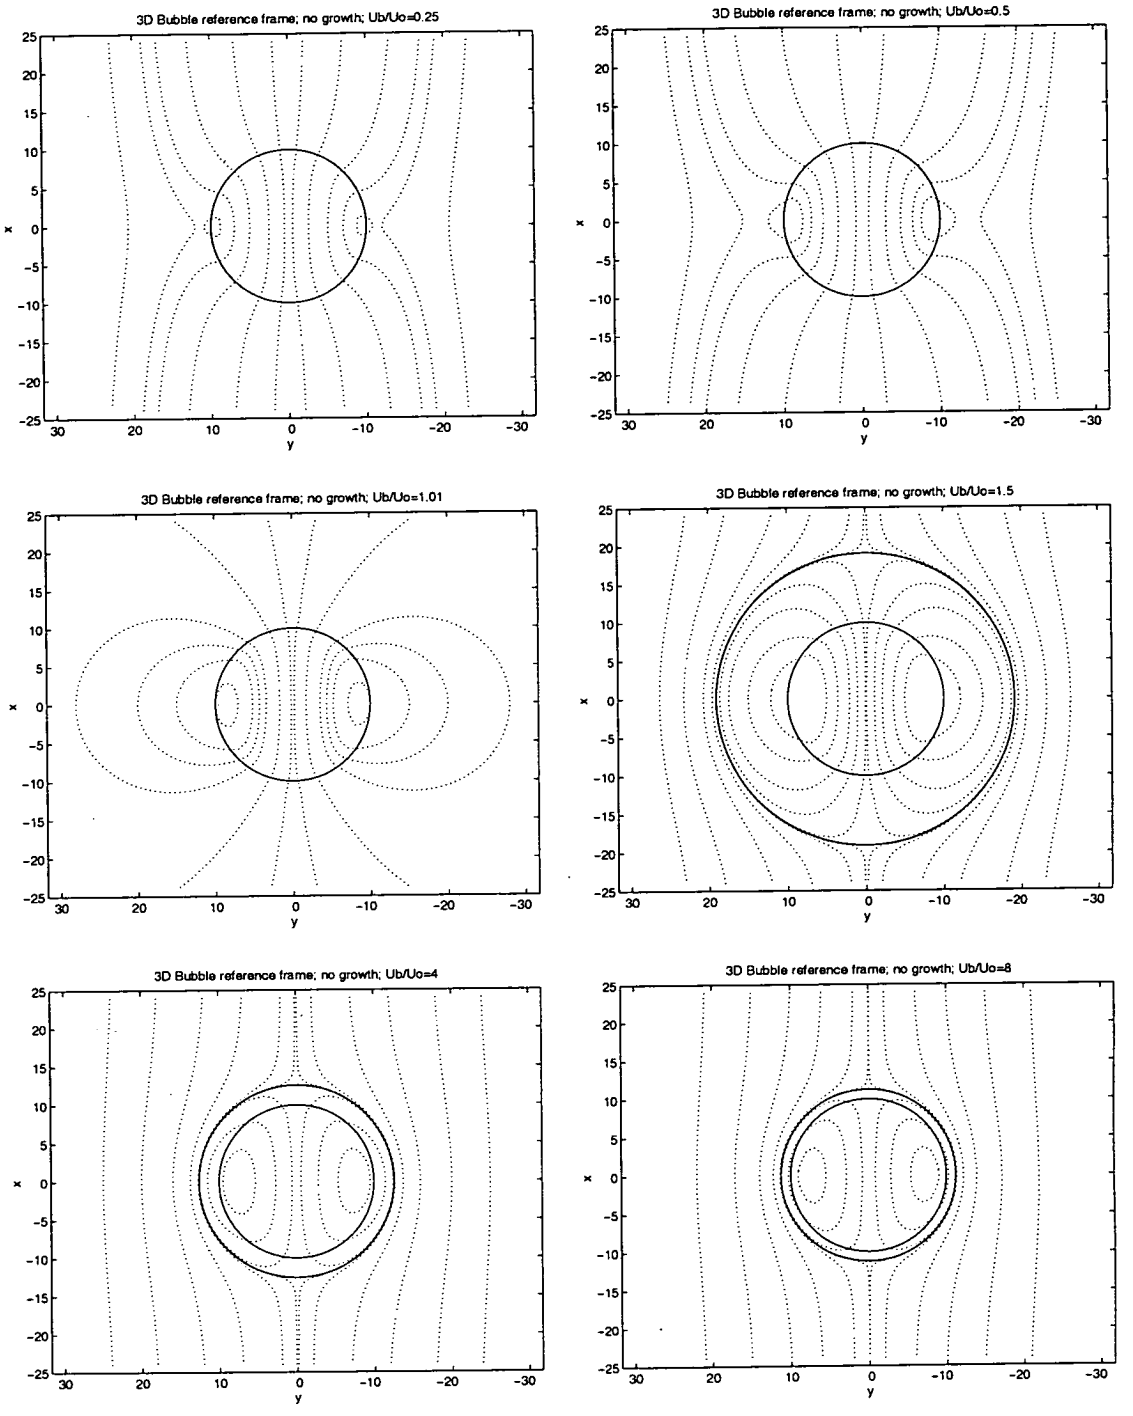


Figure 3.6: Contours of stream function (dotted curves) around a three-dimensional bubble (inner continuous circle), plotted for various bubble velocities. From left to right and from top to bottom, U_b/U_o = 0.25, 0.5, 1.01, 1.5, 4 or 8 units s^{-1}. $U_o = 1$ units s^{-1} and $U_{mf} = 0.5$ units s^{-1} in all cases.

3.2.3 Growing two-dimensional bubbles

Collins (1982) extended the above analysis in order to describe the flow outside a growing cylindrical bubble. His approach will be described here, then an attempt will be made to extend Collins' work to the inside of the bubble.

Collins added a potential flow source term to Davidson's velocity potential function for the particle phase, to cause the particles to move away from the bubble centre with extra velocities given by:

$$\begin{aligned} \Delta v_{r,p} &= a\,\dot{a}/r \\ \text{and } \Delta v_{\theta,p} &= 0. \end{aligned}$$

In Davidson's model, particles do not flow into the bubble and so have zero radial velocity at $r = a$. Collins' source term gives all of the particles at $r = a$ a net outward velocity of $\dot{a} = da/dt$, which is the rate of change of the bubble radius. This uniform particle outflow allows the particle-free bubble region to grow in a uniform manner. It might be argued that a different particle outflow expression could be used, *e.g.* one including a dependence on θ. However, this sort of uniform source/sink term is often used in potential flow calculations and is a good point from which to start.

The corresponding rate of gas flow towards the bubble can be found by a volume balance through a circular element concentric with the bubble and of radius r (the surface AA' in Figure 3.7):

$$\begin{aligned} \text{Volumetric flowrate of solids out} &= [r\,\delta\theta]\,(1-\varepsilon)\,\Delta v_{r,p}|_{r=r} \\ \text{Volumetric flowrate of gas in} &= [r\,\delta\theta]\,\varepsilon\left(-\Delta v_{r,o}|_{r=r}\right) \\ \Rightarrow \Delta v_{r,o} &= -\frac{(1-\varepsilon)}{\varepsilon}\,\dot{a}\,\frac{a}{r} \\ &= -\gamma\,U_o\,\frac{a}{r}, \end{aligned}$$

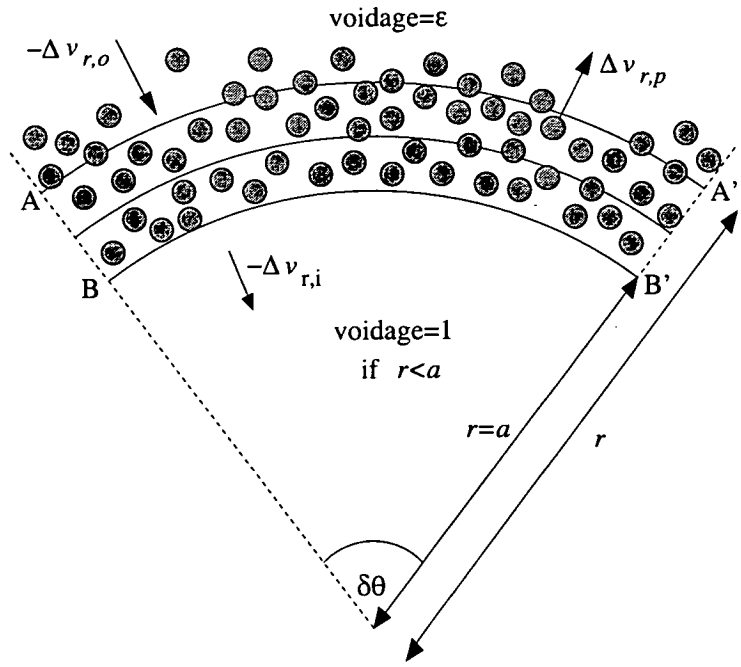


Figure 3.7: Volume balance over an element AA' outside the bubble, which is bounded by the surface BB'.

where γ is a dimensionless bubble growth parameter, given by:

$$\gamma = \frac{(1-\varepsilon)\,\dot{a}}{\varepsilon\, U_o}. \tag{3.27}$$

Taking the tangential velocity arising from this growth term to be zero, the corresponding stream function can be found by integrating the velocity terms using equations 3.3 and 3.4:

$$\Delta\psi_o = \varepsilon\, U_o\, a\, \gamma\, \theta = U_{mf}\, a\, \gamma\, \theta. \tag{3.28}$$

Thus, overall:

$$\begin{aligned} \psi_o &= \varepsilon\,(U_b - U_o)\left(1 - \frac{A_2^2}{r^2}\right) r\sin\theta + U_{mf}\, a\, \gamma\, \theta, \\ \text{so that } v_{r,o} &= -(U_b - U_o)\left(1 - \frac{A_2^2}{r^2}\right)\cos\theta - \frac{U_o\, a\, \gamma}{r} \\ \text{and } v_{\theta,o} &= (U_b - U_o)\left(1 + \frac{A_2^2}{r^2}\right)\sin\theta. \end{aligned} \tag{3.29}$$

Collins' work only considered the outer problem. One way to extend his work, in order to describe the flow inside the bubble, is to apply the boundary conditions used by Pyle & Rose at $r = a$ and to attempt to match the additional stream function term outside the bubble. A separable-variable expression can be used for the flow inside the bubble:

$$\Delta\psi_i = U_{mf}\, a\, \gamma\, \theta f(r), \tag{3.30}$$

for some function $f(r)$ of r only. The velocities inside the bubble will be those due to Pyle & Rose plus extra terms $\Delta v_{r,i}$ and $\Delta v_{\theta,i}$, derived from $\Delta\psi_i$.

The aim of this extension to Collins' work is to allow velocities and stream functions to be plotted inside the bubble as well as outside it, as was done for a non-growing bubble in Figures 3.2 and 3.5. When $\gamma = 0$, the results presented earlier will become special cases of more general results.

Comparing velocities at $r = a$, first tangentially (equation 3.14):

$$\begin{aligned} \Delta v_{\theta,i}\big|_{r=a} &= \Delta v_{\theta,o}\big|_{r=a} \\ \Rightarrow U_{mf}\, a\, \gamma\, \theta \left.\frac{df(r)}{dr}\right|_{r=a} &= 0 \quad \forall\theta \\ \Rightarrow \left.\frac{df(r)}{dr}\right|_{r=a} &= 0 \end{aligned} \tag{3.31}$$

and then radially (equation 3.15):

$$\begin{aligned} \Delta v_{r,i}\big|_{r=a} &= \varepsilon\, \Delta v_{r,o}\big|_{r=a} \\ \Rightarrow -\frac{1}{r\,\varepsilon_b} \left.\frac{\partial \Delta\psi_i}{\partial\theta}\right|_{r=a} &= -\varepsilon\, \gamma\, U_o \\ \Rightarrow \frac{1}{a} U_{mf}\, a\, \gamma\, f(a) &= \gamma\, U_{mf} \\ \Rightarrow f(a) &= 1 \end{aligned} \tag{3.32}$$

If the radial velocities inside and outside the bubble satisfy $\Delta v_{r,i} = \varepsilon\, \Delta v_{r,o}$ at the bubble boundary, then $f(a) = 1$. This also means that the stream functions must be continuous here too (equations 3.28 and 3.30).

An additional physical constraint on the choice of function $f(r)$ is that the radial velocity should remain finite as $r \to 0$:

$$\begin{aligned} \Delta v_{r,i}|_{r=0} &= -\frac{1}{r\,\varepsilon_b}\frac{\partial \Delta\psi_i}{\partial\theta}\bigg|_{r=0} \\ &= U_{mf}\, a\, \gamma\, \frac{f(r)}{r}\bigg|_{r=0} \\ \Rightarrow \lim_{r\to 0}\frac{f(r)}{r} &\quad \text{is} \quad \text{finite} \qquad (3.33) \\ \text{and} \quad \lim_{r\to 0} f(r) &= 0. \qquad (3.34) \end{aligned}$$

For this reason, the trivial solution $f(r) = 1$ for all r will not do, as it would generate infinite velocities as $r \to 0$. The constraints on $f(r)$ are that:

- $f(0) = 0$ (equation 3.34);
- $f(r)$ is proportional to r (or a higher power) at $r \sim 0$ (equation 3.33);
- $f(r)$ has a turning point (most likely to be a maximum) at the point $r = a$ (equation 3.31).
- $f(a) = 1$ (equation 3.32);

Some plausible functions for $f(r)$ are a quadratic in r, a sine term or a first order Bessel function of the first kind. Each function needs to be scaled:

$$\begin{aligned} f(r) &= 1.719\, J_1(1.8412\, r/a) \qquad (3.35) \\ \text{or } f(r) &= \sin(\pi\, r/2\, a) \qquad (3.36) \\ \text{or } f(r) &= \frac{r(2\,a - r)}{a^2} = \frac{r}{a}\left(2 - \frac{r}{a}\right). \qquad (3.37) \end{aligned}$$

The Bessel function suggested itself because it arose in the theory of de Kock & Judd (1965a), where gas velocities inside a two-dimensional void within a porous mass were obtained by assuming that vorticity generated at the void boundary remained constant along streamlines within the void. The resultant J_1 term was then found by de Kock & Judd to run from $(0,0)$ up to its first maximum ($J_1(1.8412) = 0.5819$), by applying their boundary conditions. (It is not physically realistic to expect the function $f(r)$ to have any maximum or minimum turning points for $0 < r < a$.)

These functions are plotted in the left-hand half of Figure 3.8, where the dotted and continuous lines are the sine and quadratic expressions respectively. Some values of the Bessel function are plotted, using '+' signs. It is apparent that the Bessel function (equation 3.35) and the sine function (equation 3.36) are almost identical after this rescaling, for the range in question. In the right-hand half of this figure, equations 3.35 and 3.36 have been evaluated for various values of r/a, then the quotient of the corresponding results has been calculated. It can be seen that the Bessel function is never more than 0.75% greater than the sine function. The greatest absolute difference was about 3.3×10^{-3}, found at $r/a \approx 0.45$. Of course, the J_1 Bessel function and the sine function are quite different if r/a is outside the range used here (*e.g.* $r/a = 2$ in equations 3.35 and 3.36).

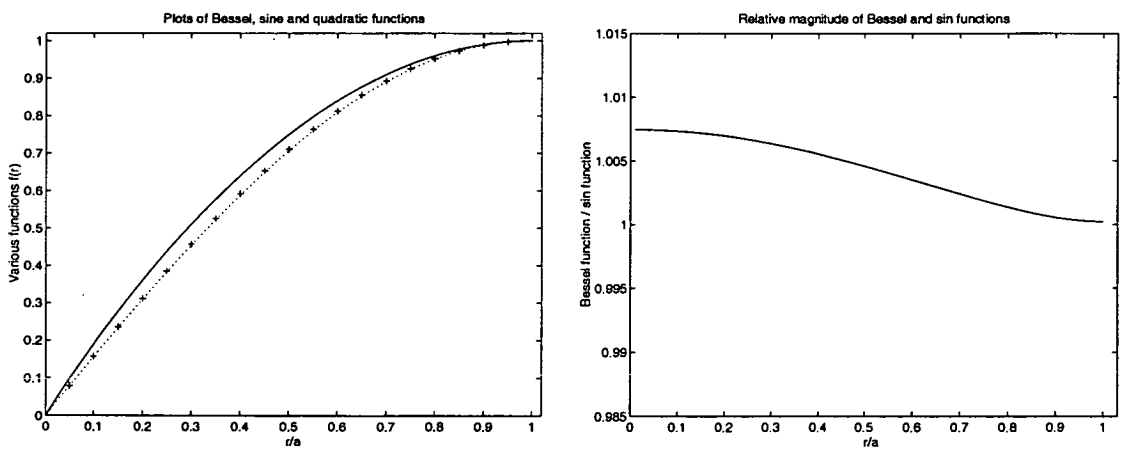


Figure 3.8: Left: the various functions $f(r)$: a Bessel function ('+' signs, equation 3.35), a sine function (dotted line, equation 3.36) and a quadratic (continuous line, equation 3.37); Right: the ratio of the Bessel function to the sine function.

For convenience, the sine expression (equation 3.36) was chosen, as it very closely matches the Bessel function alternative (equation 3.35) over the given range of values but is considerably easier to evaluate and differentiate. Thus:

$$\Delta\psi_i = U_{mf}\, a\, \gamma\, \theta \sin\left(\frac{\pi\, r}{2\, a}\right) \tag{3.38}$$

$$\Rightarrow \Delta v_{r,i} = -U_{mf}\, a\, \gamma\, \frac{\sin(\pi r/2\, a)}{r} \tag{3.39}$$

$$\text{and } \Delta v_{\theta,i} = \frac{\pi}{2}\, U_{mf}\, \gamma\, \theta\, \cos\left(\frac{\pi\, r}{2\, a}\right) \tag{3.40}$$

Slightly disturbing is the fact that the tangential velocity does not equal zero at $\theta = \pm\pi$, for points within the bubble (equation 3.40). This analysis suggests that $\Delta v_{\theta,i}$ at (r, π) will be equal and opposite to its value at $(r, -\pi)$, even though these two points are identical. Further, there is a step change in the magnitude of $\Delta v_{\theta,i}$ when moving down the x-axis, from $\Delta v_{\theta,i} = 0$ at small positive x ($\theta = 0$) to $\Delta v_{\theta,i} = \pi^2\, U_{mf}\, \gamma/2$ at small negative x ($\theta = \pi$). The magnitude of these velocities is always less than a tenth of the magnitude of the vertical velocity, which is dominated at small r by the Pyle & Rose terms.

The sine function used here (equation 3.36) is proportional to r as $r \to 0$. When calculating the radial velocity, an expression involving $\sin(\pi r/2a)/r$ is obtained (equation 3.39) which has a finite but non-zero limit as $r \to 0$. Thus the radial velocity corresponding to the growth stream function is not zero at the bubble centre. This means that there is a net gas flow into any circle centred on the origin, regardless of its radius. A circle centred on the origin becomes indistinguishable from the origin as its radius $r \to 0$. This solution has succeeded in avoiding infinite radial velocities at $r \to 0$ by reducing the effect of the sink term to a certain extent, but has not solved the problem of fluidisation gas disappearing into oblivion at the origin (*i.e.* the sink term still has some effect).

A different function $f(r)$ might give improved results inside the bubble. It can be noted that, if the constraints on $f(r)$ in equations 3.31–3.34 are met by one type of function, then the same constraints will still be met if that function is

raised to some power n, such that $n > 1$. For example, if $f(r)$ is proportional to r^1 or a higher power as $r \to 0$, then $[f(r)]^n$ will also be proportional to r^1 or a higher power, and, if the point $(a, 1)$ is a turning point of $f(r)$, then it will also be a turning point of $[f(r)]^n$. This seems to be an easy way of generating any number of other candidate functions.

Reconsidering the volume balance in Figure 3.7, it can be noted that the volume balance on the additional gas and particle flows (following the addition of Collins' source term for the particles) should only apply outside the bubble, because there are no particles within the bubble. Thus, it could be argued that the extra radial gas velocity term within the bubble should simply be zero (now that there are no particles against which a volume balance can be performed) and that Pyle & Rose's velocity boundary condition (equation 3.15) should not be used to match the radial velocities which arise from Collins' growth terms.

The way to modify the function $f(r)$, in order to take these radial velocity considerations into account, is to raise it to a large positive power n. Starting from $f(r) = \sin(\pi\, r/2\, a) \leq 1$ for all r, any new function $f_1(r) = [\sin(\pi\, r/2\, a)]^n$ will be closer to zero than $f(r)$, for all r in the range $0 < r < a$ (if $n > 1$). Both functions pass through the points $(0, 0)$ and $(a, 1)$.

The new function $f_1(r)$ is plotted in Figure 3.9, for various values of n. As $n \to \infty$, it can be seen that $f_1(r)$ becomes a step-function which is 1 if $r = a$ and zero otherwise. This is good news as far as the radial velocity and stream function are concerned, as the stream function can remain continuous across the bubble boundary whilst the radial velocity falls rapidly to zero (thus removing the effect of the sink function which pulls gas in towards the bubble centre).

On the other hand, equation 3.4 says that the additional tangential velocity inside the bubble is given by $\Delta v_{\theta,i} = \partial \Delta \psi_i / \partial r$. Consider a radius drawn within the bubble, for example at an angle $\theta = 2\pi/3$. If the stream function is continuous across the bubble boundary, then $\Delta \psi_o(a, \theta)$ units of stream function need to be dropped over a distance a, from the bubble centre to the bubble boundary. At

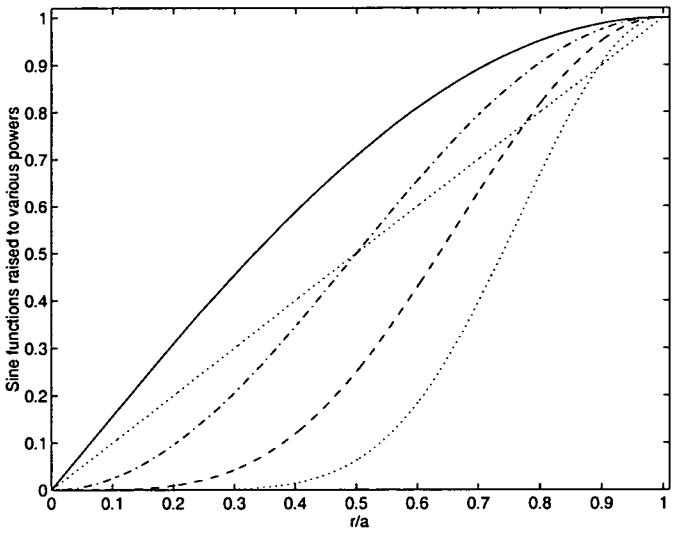


Figure 3.9: The modified sine function $y = f_1(r) = [\sin(\pi\, r/2\, a)]^n$, for various values of n. From top left to bottom right, $n = 1$ (continuous curve), 2 (dash-dot curve), 4 (dashed curve) or 8 (dotted curve). The dotted line plotted is $y = x$.

any point along this radius, the extra tangential velocity due to the growth term is simply the radial gradient of $\Delta\psi_i$ along this path, which is a property of $f(r)$.

With the basic sine expression used earlier (equation 3.36), the change in stream function along each radial line was slow but steady, giving rise to moderate gradients and hence moderate tangential velocities. Using, say, the alternative function $f_1(r) = [\sin(\pi r/2a)]^8$ inside the bubble, $f_1(r)$ is almost zero for $0 < r < a/2$, so the curve must now rise more steeply in the region where $a/2 < r < a$ if it is to reach the required value of 1 at $r = a$ (Figure 3.9). As n increases, the position of the steepest part of $f_1(r)$ moves closer to $r = a$ and the magnitude of its slope increases; hence the maximum gradient of $\Delta\psi_i$ increases and larger tangential velocities are created.

The compromise $n = 2$ can be tried. This allows the additional radial velocity to reach zero at the bubble centre without letting the corresponding tangential velocity get too large. In fact, the greatest slope of a $\sin^2\theta$ function (at $\theta = \pi/4$) has exactly the same gradient as the greatest slope of a $\sin\theta$ function, which is found at $\theta = 0$ (Figure 3.9). Hence, the two functions will generate additional tangential velocities which have the same greatest magnitude, but the $\sin^2$ function will be better behaved as $r \to 0$.

$$\Delta\psi_i = U_{mf}\, a\, \gamma\, \theta \left[\sin\left(\frac{\pi\, r}{2\, a}\right)\right]^2 \tag{3.41}$$

means that the additional velocity terms inside the bubble are:

$$\begin{aligned}
\Rightarrow \Delta v_{r,i} &= -\frac{1}{r\,\varepsilon_b}\frac{\partial \Delta\psi_i}{\partial\theta} \\
&= -\frac{U_{mf}\, a\, \gamma}{r}\left[\sin\left(\frac{\pi\, r}{2\, a}\right)\right]^2 \\
\text{and } \Delta v_{\theta,i} &= \frac{1}{\varepsilon_b}\frac{\partial \Delta\psi_i}{\partial r} \qquad (3.42) \\
&= U_{mf}\, a\, \gamma\, \theta \left[2\sin\left(\frac{\pi\, r}{2\, a}\right)\cos\left(\frac{\pi\, r}{2\, a}\right)\right]\frac{\pi}{2\, a} \\
&= \frac{\pi}{2}\, U_{mf}\, \gamma\, \theta \sin\left(\frac{\pi\, r}{a}\right) \qquad (3.43)
\end{aligned}$$

Contours of the stream function given by equation 3.41 are plotted in Figure 3.10. The stream functions are continuous over the bubble boundary.

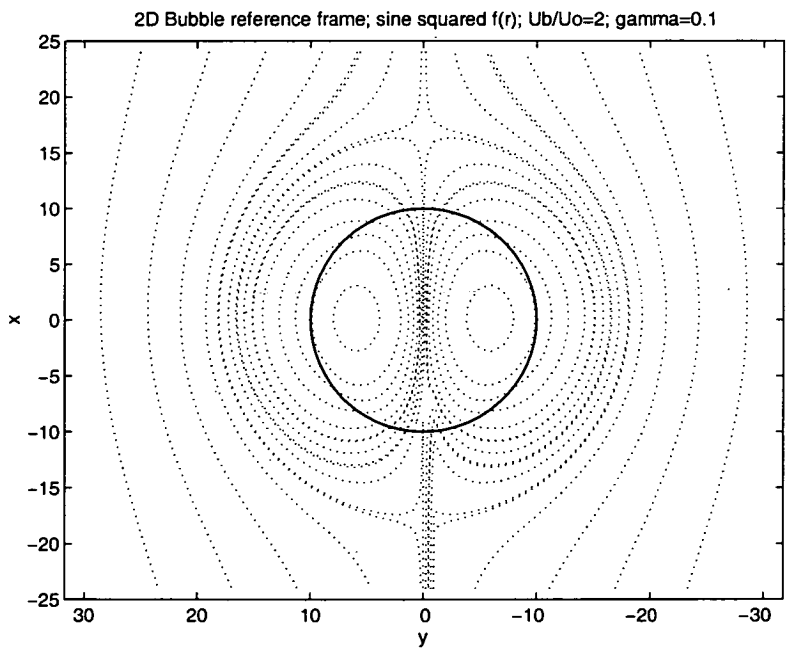


Figure 3.10: Streamlines for a growing 2D bubble, using a sine-squared function $f(r)$ inside the bubble and $\gamma = 0.1$.

A key feature of this plot is the existence of two stagnation points at different values of ψ (*cf* the no-growth case of Figure 3.2, when both stagnation points lay on $\psi = 0$). Gas flowing down towards the bubble flows around the side of the bubble and is drawn in through the lower half of the bubble to be used for growth,

if $-U_{mf}\, a\, \gamma\, \pi < \psi_o < U_{mf}\, a\, \gamma\, \pi$. In the Cartesian co-ordinates of Figure 3.10, this region is given by $-6 < y < 6$ at $x = 24$. Streamlines outside this central zone are deflected around the bubble and settle down closer to the x-axis.

This means that the stream function contours given by $\psi_o = U_{mf}\, a\, \gamma\, \pi$ and $\psi_o = -U_{mf}\, a\, \gamma\, \pi$ converge towards the lower half of the $y = 0$ axis. This is not a result of the particular choice of $f(r)$ for the growth flow inside the bubble, but instead concerns the nature of Collins' growth stream function outside the bubble. Here, θ was chosen to lie in the range $-\pi < \theta \leq \pi$. The points $(20, 3.14)$ and $(20, -3.14)$ (in polar co-ordinates) are in fact very close together (they are only about 0.06 units apart) even though there appears to be a large difference in θ between them. If θ was redefined so that $0 \leq \theta < 2\pi$, then these points would be represented as $(20, 3.14)$ and $(20, 3.143)$ which looks much better, but the corresponding points $(20, 0.00153)$ and $(20, 6.281)$ would now straddle a discontinuity in θ. The discontinuity would simply have been chased away to a different location, but not removed.

Similarly, the concept of giving this planet a series of time zones, which depend on a longitudinal angle of rotation away from the Greenwich meridian, necessarily involves the creation of an International Date Line. On either side of this date line, people are either 12 hours ahead of GMT or 12 hours behind. The $0°$ meridian could instead have been chosen to pass through New York or Tokyo, but, in either case, the date line would not disappear but would simply move to a new position.

In this situation, Collins' two-dimensional sink term $\Delta\psi_o$ is directly proportional to θ. This angle θ has a discontinuity across the lower half of the x-axis, so $\Delta\psi_o$ must have a discontinuity there too. In Figure 3.10, there is therefore a discontinuity in ψ_o about the negative x-axis, moving from 2π at $(-25, 0.001)$ in the Cartesian (x, y) co-ordinates shown in this figure to -2π at $(-25, -0.001)$. Matlab interpolated between these values, which makes the streamlines in the lower right-hand half of the picture veer down towards the negative x-axis instead of continuing in towards the origin. The left-hand half this figure seems correct, so only these contours will be shown from now on.

Overall velocity vectors and stream functions, for the flow both inside and outside a growing two-dimensional bubble, are plotted in Figures 3.11 and 3.12, for $\gamma = 0.1$ and $\gamma = 0.2$.

3.2.4 Growing three-dimensional bubbles

Collins (1983) later added sink and source terms to the gas and particle phases in a three-dimensional fluidised bed, in order to model the flow outside a growing spherical bubble. His extra stream function term for the gas in this case was:

$$\Delta\psi_o = 0.5\,\varepsilon\,U_o\,\gamma\,a^2\,\sin^2\theta.\sec^2(\theta/2),$$

which can be simplified by breaking down the $\sin^2\theta$ term:

$$\begin{aligned}\Delta\psi_o &= 0.5\,U_{mf}\,\gamma\,a^2\,\frac{[2\,\sin(\theta/2)\,\cos(\theta/2)]^2}{\cos^2(\theta/2)} \\ &= 2\,U_{mf}.\gamma\,a^2\,\sin^2(\theta/2) = U_{mf}\,\gamma\,a^2\,(1-\cos\theta). \qquad (3.44)\end{aligned}$$

By equations 3.5 and 3.6, the associated velocity terms were:

$$\begin{aligned}\Delta v_{r,o} &= -U_o\,\gamma\,\frac{a^2}{r^2} \\ \text{and}\quad \Delta v_{\theta,o} &= 0.\end{aligned}$$

The analogous attempt to extend this model to the inside of the bubble involves trying a solution of the form:

$$\Delta\psi_i = U_{mf}\,\gamma\,a^2(1-\cos\theta)\,g(r), \qquad (3.45)$$

where $g(r)$ is a function of r only.

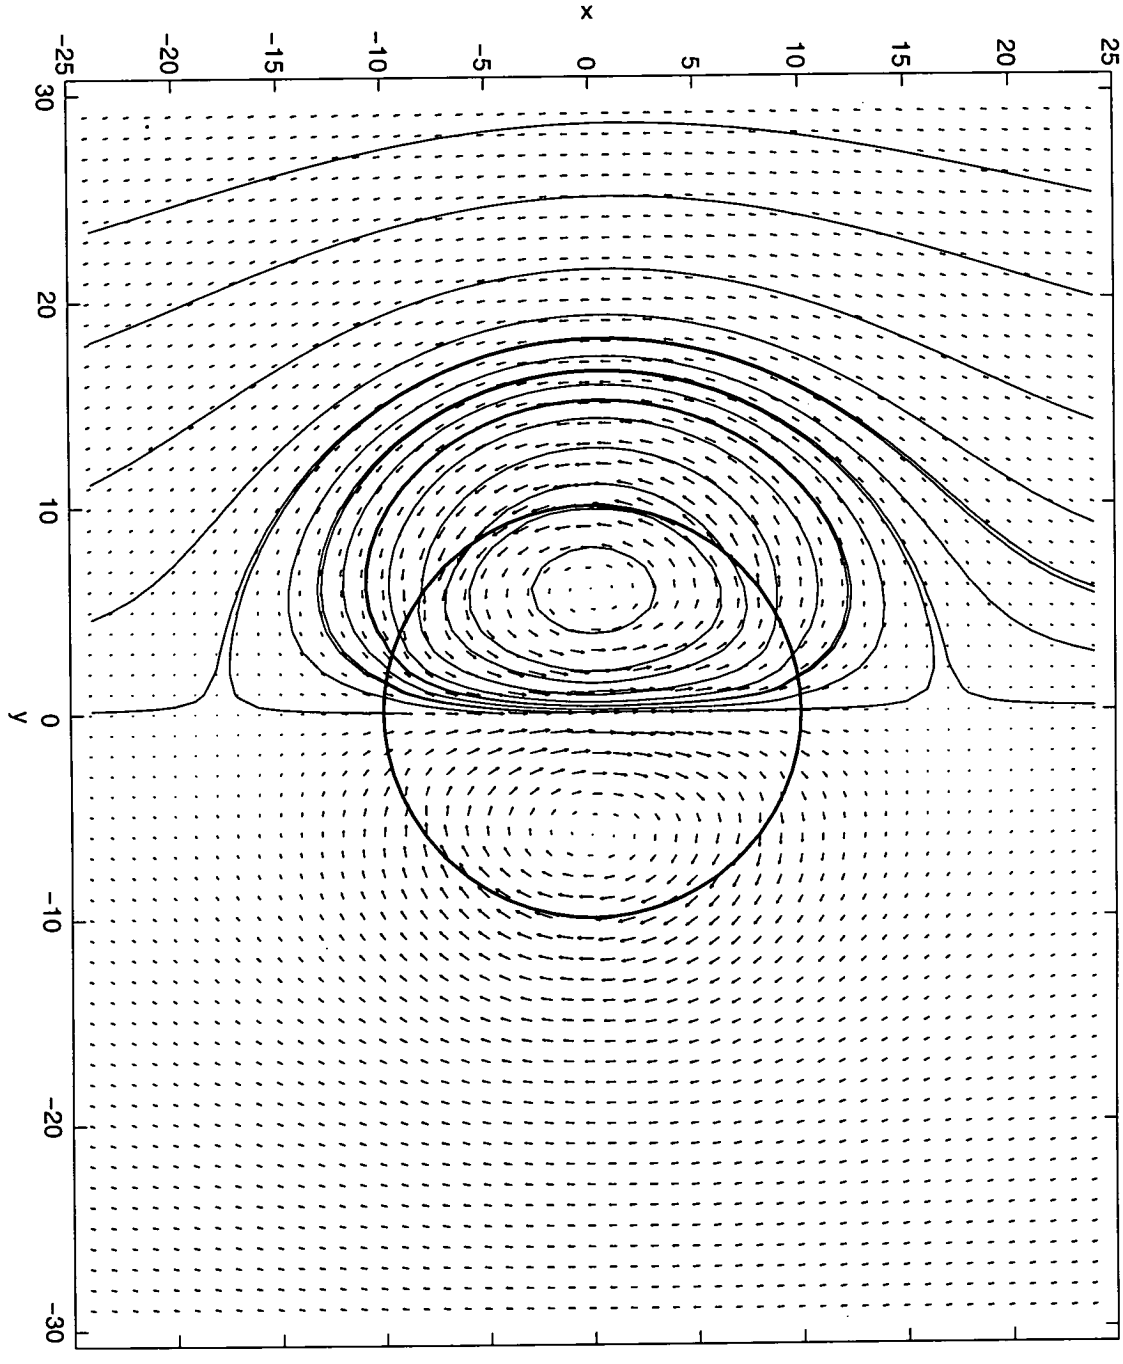


Figure 3.11: Velocity vectors and streamlines for a growing 2D bubble, using a sine-squared function inside the bubble and $\gamma = 0.1$.

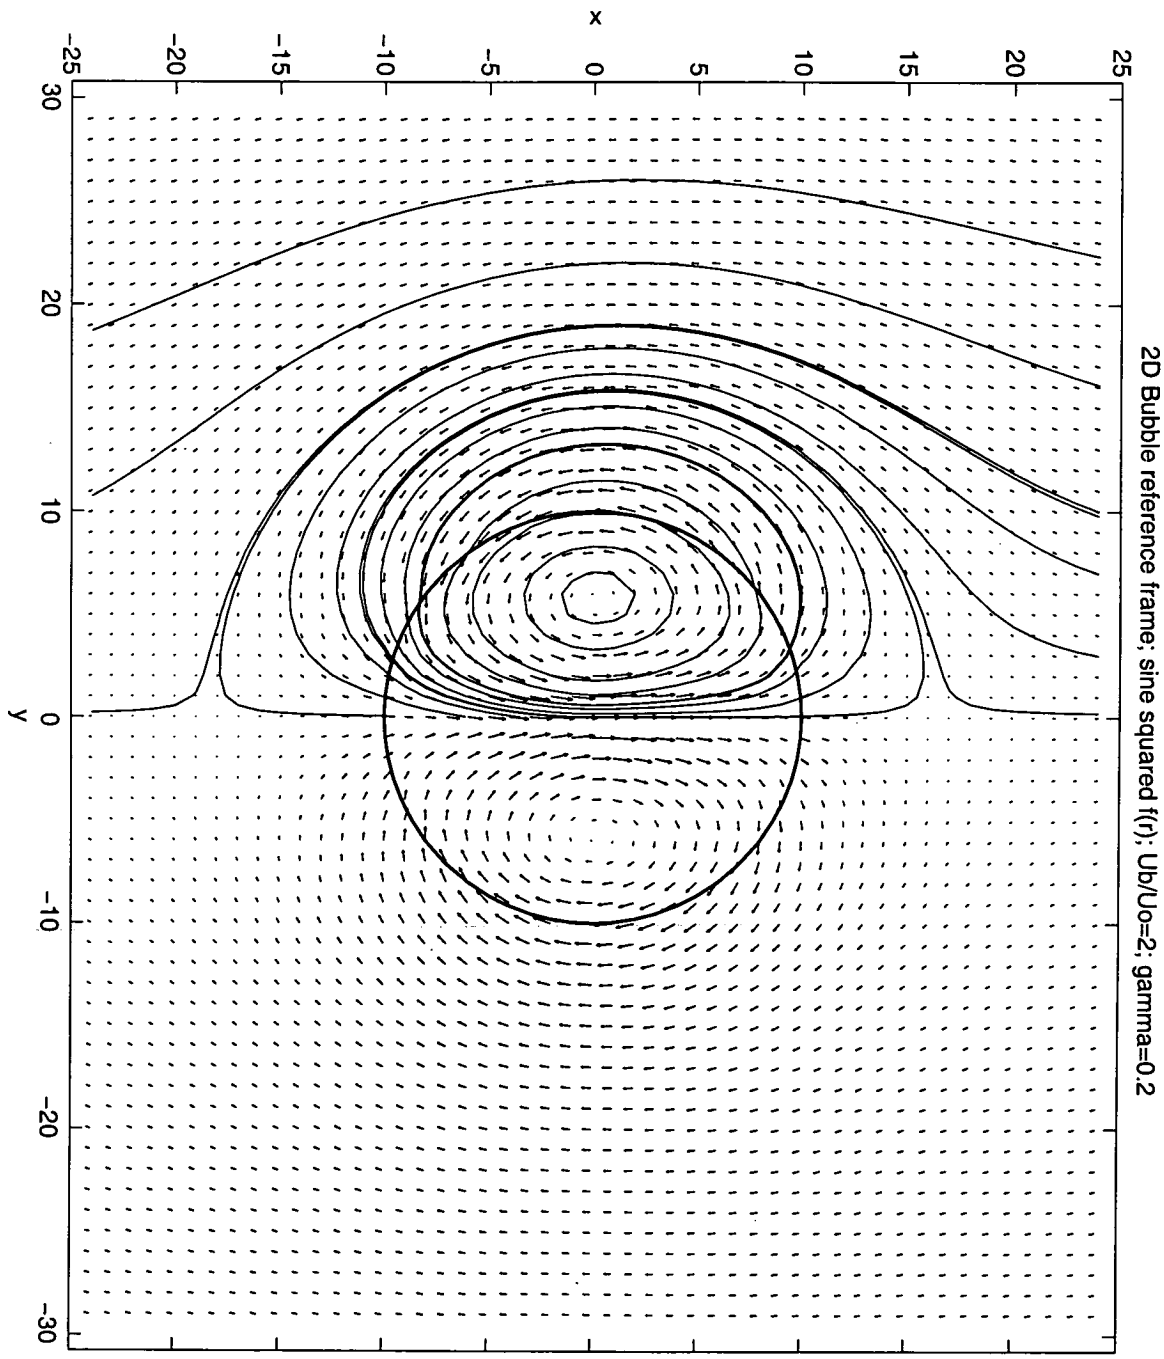


Figure 3.12: Velocity vectors and streamlines for a growing 2D bubble, using a sine-squared function inside the bubble and $\gamma = 0.2$.

Applying the Pyle & Rose boundary conditions for the new velocity terms and requiring the radial velocity to remain finite as $r \to 0$, gives a similar set of constraints to those in the two-dimensional case, above. The new requirements are that $g(r) \propto r^2$ (or a higher power) for small r (hence $g(r) = 0$ at $r = 0$) and that $g(r) = 1$ is a maximum turning point, reached at $r = a$. Apart from the fact that the minimum possible order of $g(r)$ when $r \approx 0$ has increased by one power of r, this is the same as in the two-dimensional case.

Suitable functions of minimum order are a Bessel function of the first kind $[g(r) = 2.055\, J_2(3.0542\, r/a)]$, a sine-squared function $[g(r) = \sin^2(\pi r/2a) = 0.5\,(1 - \cos(\pi r/a))]$ or a cubic polynomial $[g(r) = 3\,(r/a)^2 - 2\,(r/a)^3]$. Plots of these functions show that they are almost identical over the required range, $0 < r < a$. The trigonometric term was preferred for its ease of evaluation and differentiation. As in the two-dimensional case, in order to ensure that $\Delta v_{r,i} \to 0$ as $r \to 0$, it was then decided to use a $\sin^3$ expression instead:

$$\Delta\psi_i = U_{mf}\,\gamma\,a^2\,(1 - \cos\theta)\left[\sin\left(\frac{\pi\, r}{2\, a}\right)\right]^3 \tag{3.46}$$

Thus, applying equation 3.5:

$$\begin{aligned}
\Delta v_{r,i} &= -\frac{1}{\varepsilon_b\, r^2\, \sin\theta}\,\frac{\partial\Delta\psi_i}{\partial\theta} \\
&= -\frac{U_{mf}\,\gamma\,a^2}{r^2\,\sin\theta}\,\sin\theta\left[\sin\left(\frac{\pi\, r}{2\, a}\right)\right]^3 \\
\Rightarrow \Delta v_{r,i} &= -U_{mf}\,\gamma\,a^2\,\frac{\sin^3(\pi r/2a)}{r^2}
\end{aligned} \tag{3.47}$$

and similarly, using equation 3.6:

$$\begin{aligned}
\Delta v_{\theta,i} &= \frac{1}{\varepsilon_b\, r\,\sin\theta}\,\frac{\partial\Delta\psi_i}{\partial r} \\
&= \frac{1}{r\,\sin\theta}\,U_{mf}\,\gamma\,a^2\,(1 - \cos\theta).3\left[\sin\left(\frac{\pi\, r}{2\, a}\right)\right]^2\cos\left(\frac{\pi\, r}{2\, a}\right)\frac{\pi}{2\, a} \\
&= \frac{3\,\pi\,U_{mf}\,\gamma\,a}{2\, r}\,\frac{2\,\sin^2(\theta/2)}{2\,\sin(\theta/2)\,\cos(\theta/2)}\,\sin^2\left(\frac{\pi\, r}{2\, a}\right)\cos\left(\frac{\pi\, r}{2\, a}\right)
\end{aligned}$$

$$\Rightarrow \Delta v_{\theta,i} = \frac{3\pi U_{mf}\gamma a}{2r}\tan\left(\frac{\theta}{2}\right)\sin^2\left(\frac{\pi r}{2a}\right)\cos\left(\frac{\pi r}{2a}\right) \quad (3.48)$$

The velocity vectors and stream functions, both inside and outside this growing three-dimensional bubble, are plotted in Figures 3.13–3.15, for $\gamma = 0.1$, 0.2 or 0.4.

One drawback of this solution is that the tangential velocity depends on $\tan(\theta/2)$ and hence becomes infinite at $\theta = \pm\pi$. In fact, this must be the case if the function g in equation 3.45 depends on r only. If g is a function of θ only, as is the case for $\Delta\psi_o$ (equation 3.44), then $\partial\Delta\psi_i/\partial r = 0$ and so the tangential velocity is zero. An alternative function $g(r,\theta)$ could perhaps be sought instead.

3.3 Ghost bubble vector maps in the laboratory frame

For the case of a three-dimensional fluidisation bubble rising without growth in a fluidised bed, the air velocity vector map and stream functions have already been given (Figure 3.5). These quantities were worked out in the bubble reference frame, but can only be recorded on film with this appearance if the camera rises up in the laboratory at the same speed as the bubble.

Figure 3.16 shows the appearance of such a bubble in the laboratory frame, *i.e.* as photographed by a stationary camera. A vertical velocity U_b has been added to all velocity vectors given in Figure 3.5, to account for this change of reference frame. Velocity vectors have been plotted, but only for every other row or column. Stream functions are not given here, as they can only be defined in the bubble frame. A short time 'after' this plot was made, the bubble will have risen a short distance up the x-axis, so the velocities at each point will change.

It can be seen that the velocities after the change of reference frame appear rather different to those in the bubble frame. There are large vertical velocities at the centre of the bubble and small recirculation regions centred on (or close to) the points $(x = 0, y = \pm 10)$. Imagining a set of path lines at this instant in

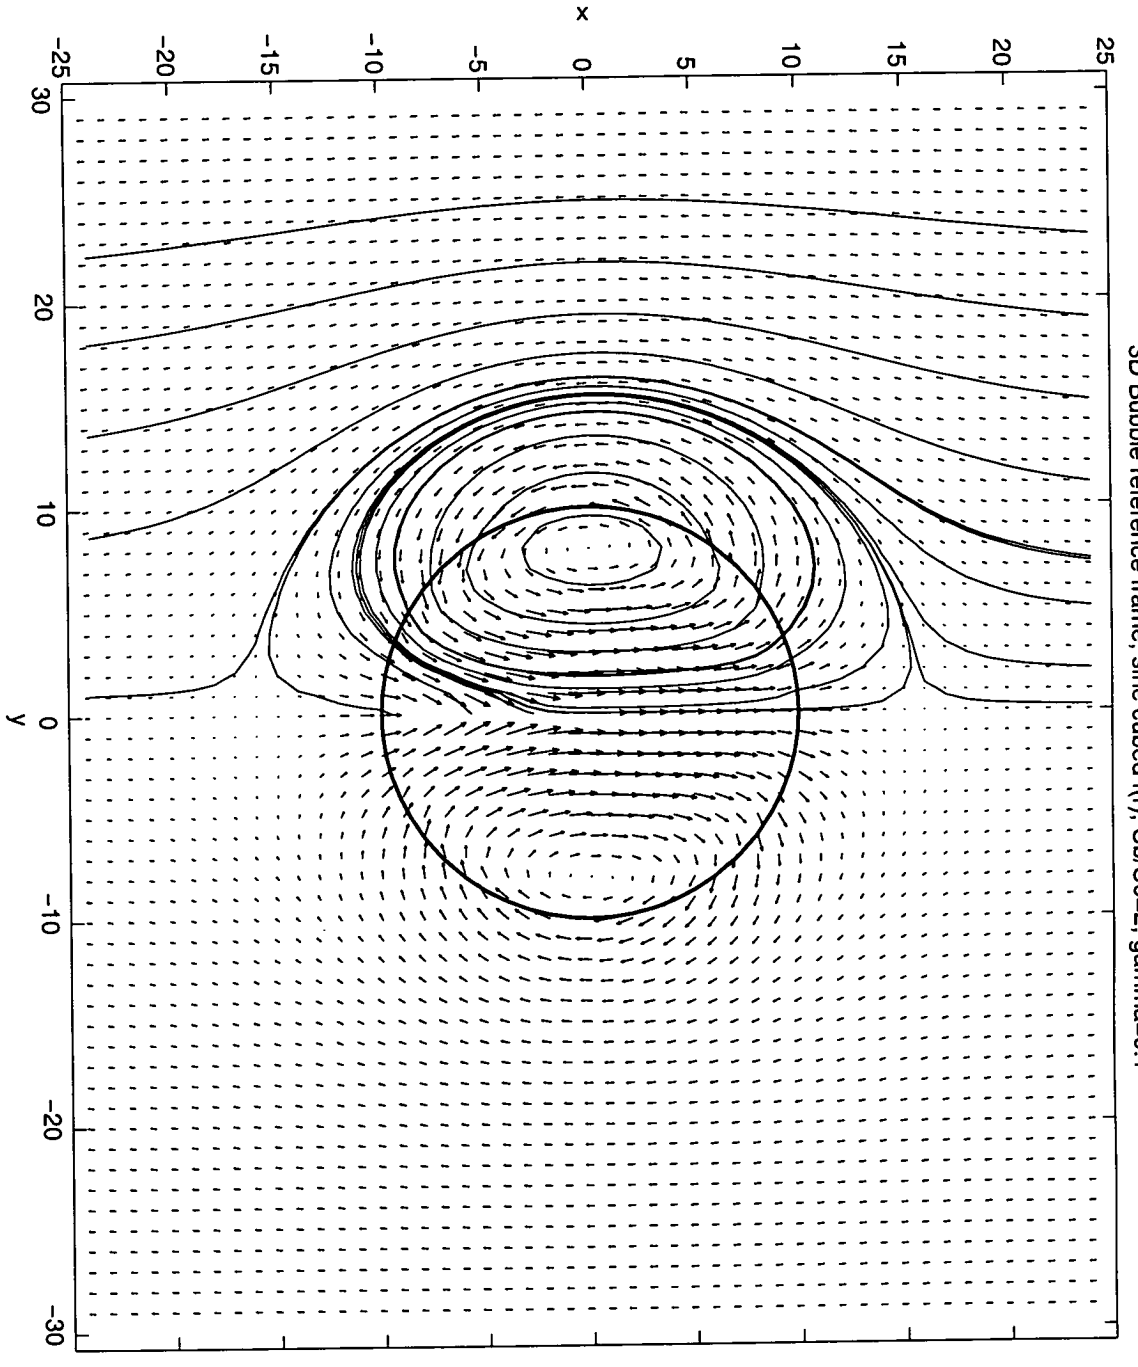


Figure 3.13: Velocity vectors and streamlines for a growing 3D bubble ($\gamma = 0.1$)

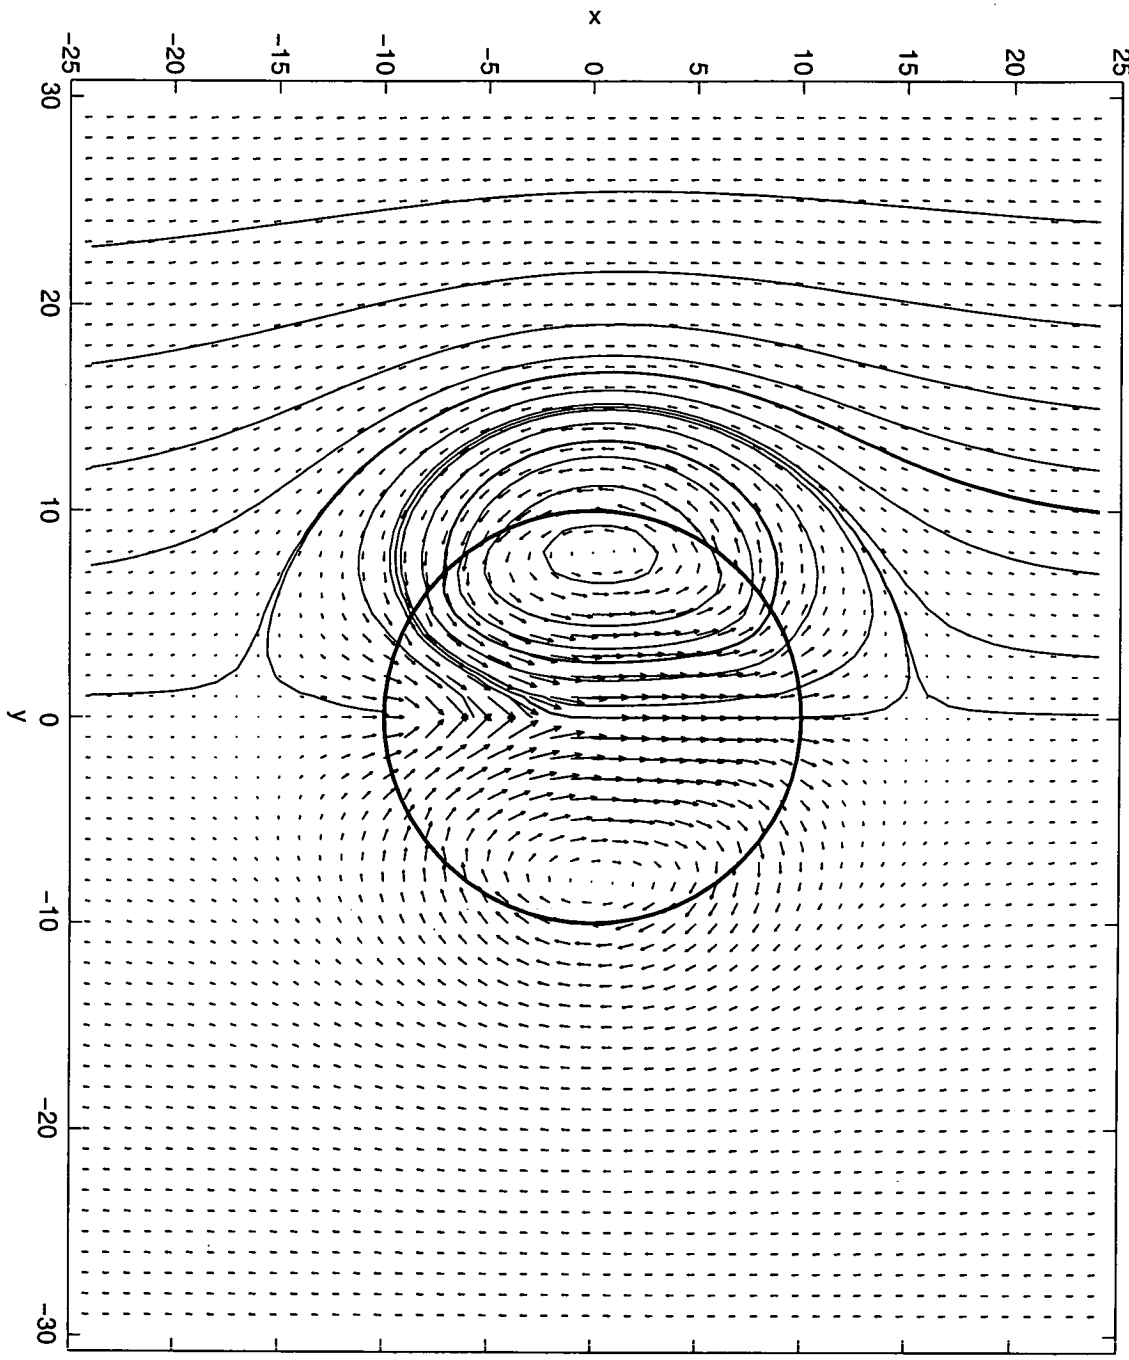


Figure 3.14: Velocity vectors and streamlines for a growing 3D bubble ($\gamma = 0.2$)

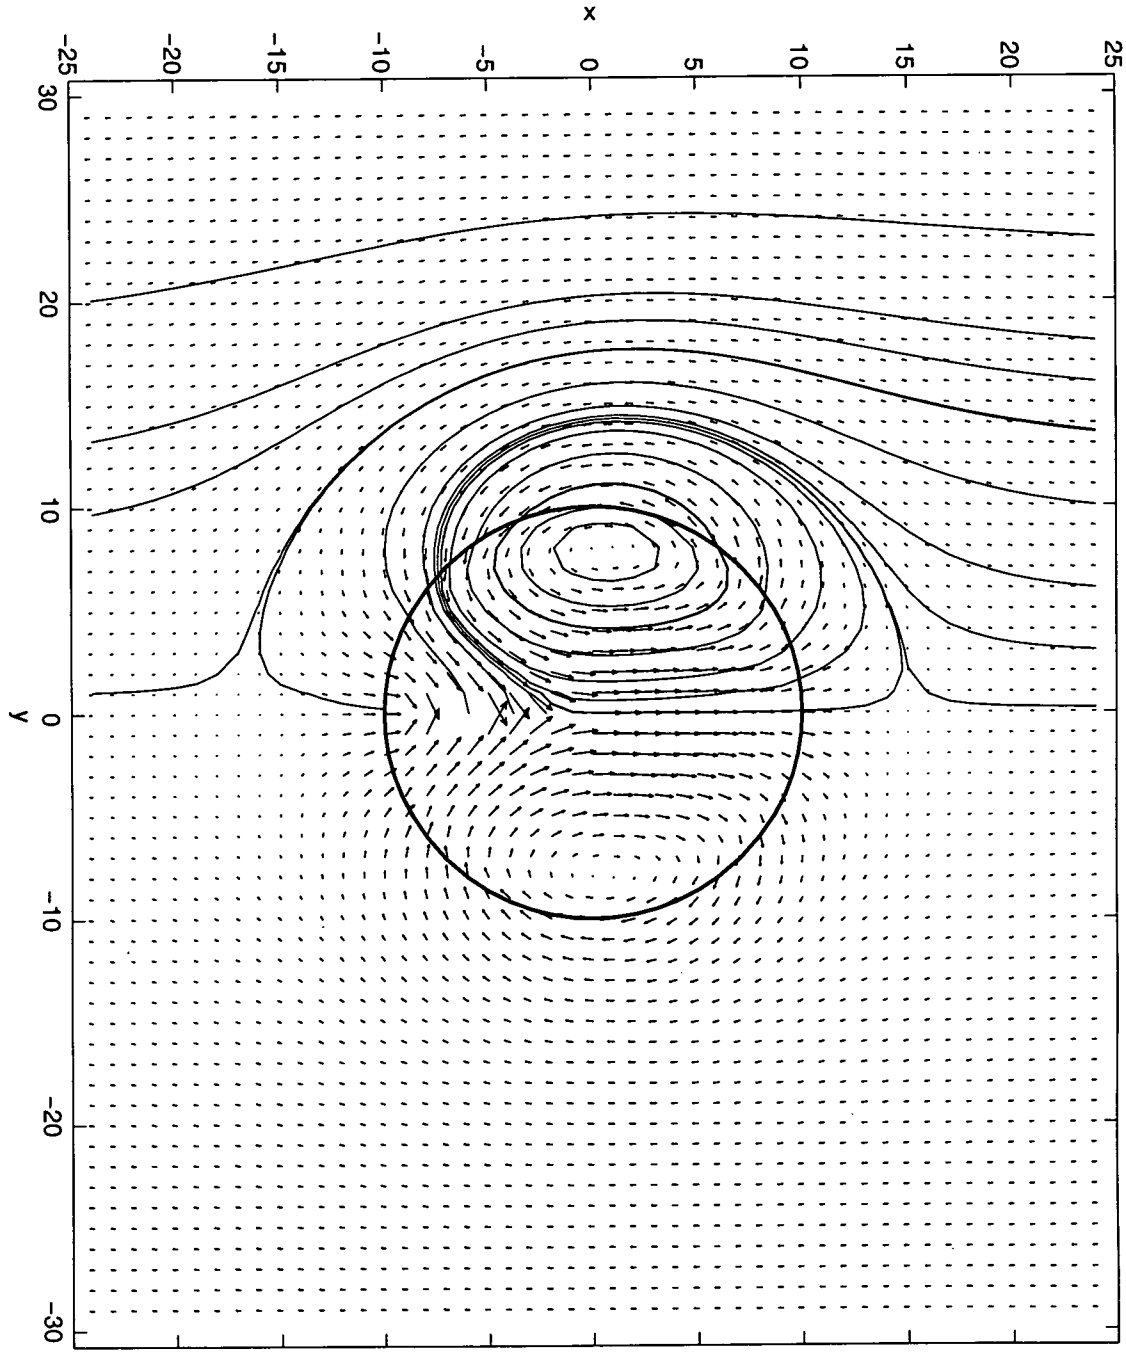


Figure 3.15: Velocity vectors and streamlines for a growing 3D bubble ($\gamma = 0.4$)

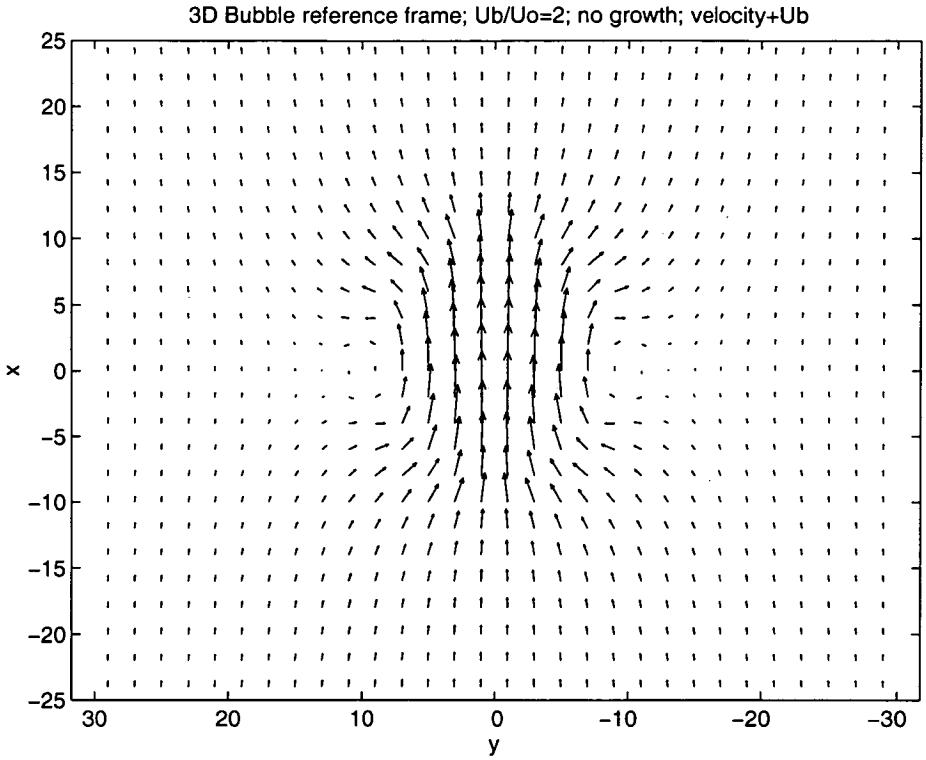


Figure 3.16: Velocity vectors for a 3D fluidisation bubble rising without growth ($\gamma = 0$), after a vertical velocity U_b has been added to all vectors.

time, defined as lines across which no fluid was flowing at this particular moment, the appearance is similar to the stream lines in the bubble reference frame for the case $U_b/U_o = 0.5$ (Figure 3.6).

The change of reference frame made by adding U_b is the biggest velocity shift which needs to be considered. This change indicates the appearance of a fluidisation bubble in a fluidised bed or (making the probably unlikely assumption that the bubble is not affected by breaking through the bed surface and entering the freeboard) at the bottom of the freeboard region.

Once the bubble has entered the freeboard region, it is expected to rise up the freeboard as a 'ghost bubble'. During its rise, other gaseous material from the freeboard is entrained into the bubble, resulting in the bubble growing and slowing down. Plots of the velocity vectors from Figure 3.5 are given in Figure 3.17, after the addition of smaller reference frame shift velocities of $0.75\,U_b$ (top) or $0.5\,U_b$ (bottom). These velocity shifts model the effect of ghost bubbles slowing down.

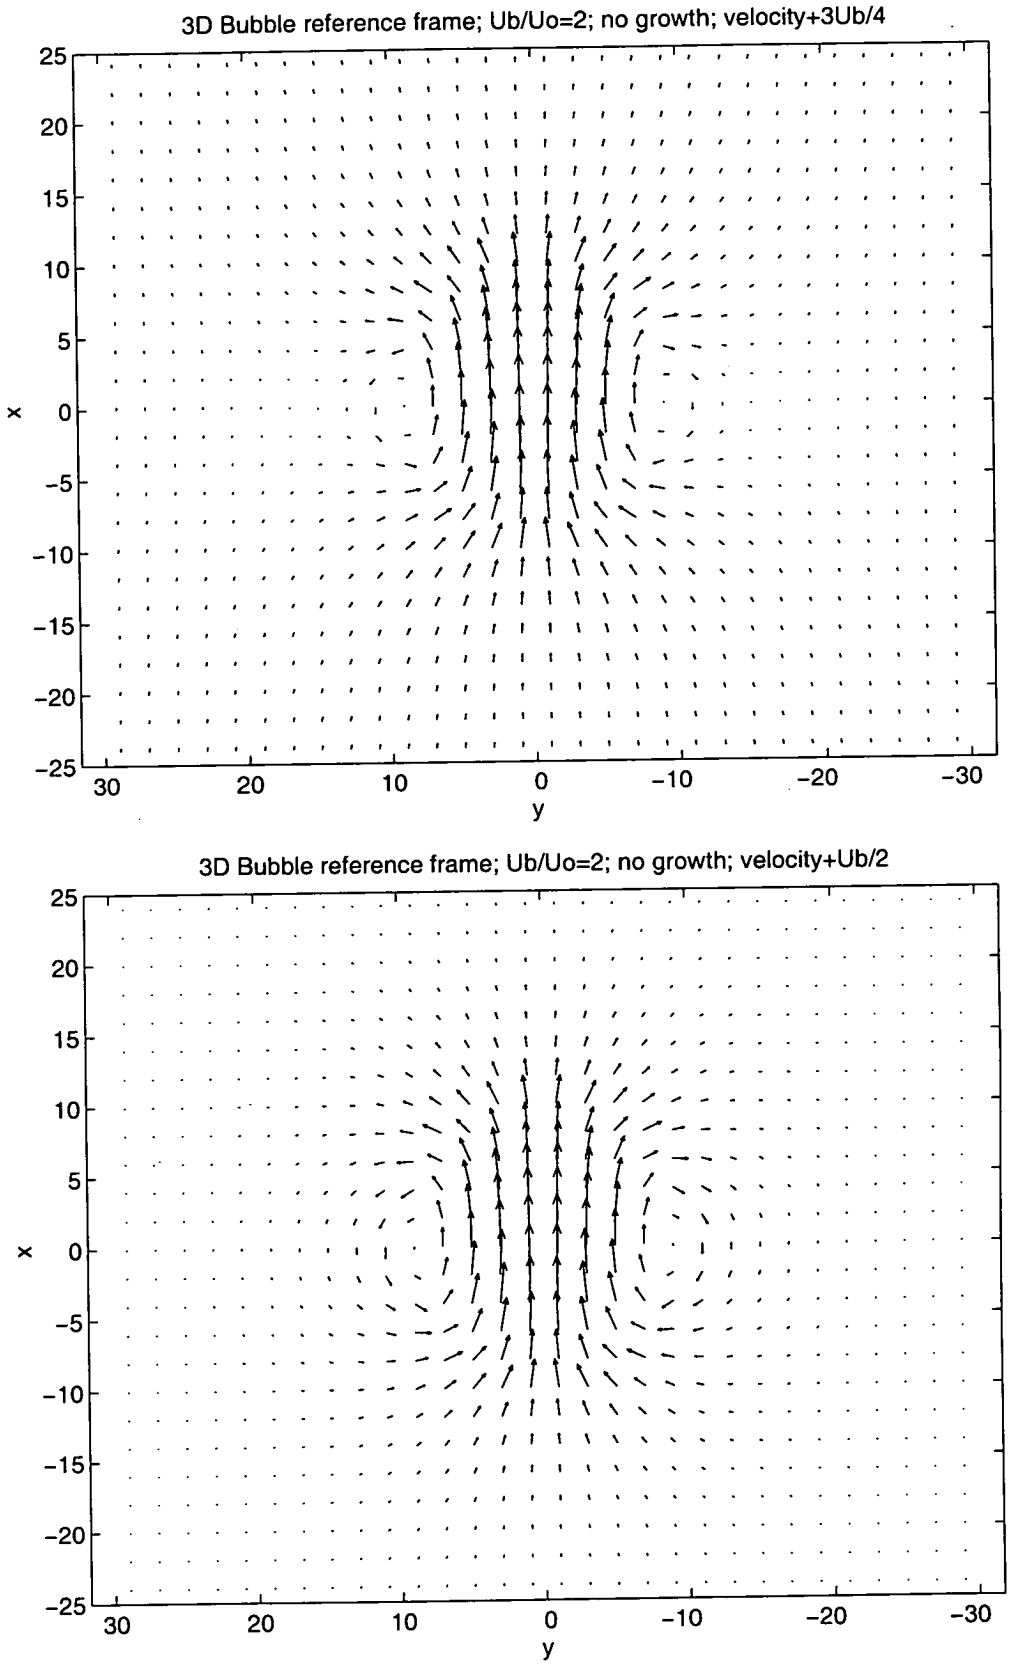


Figure 3.17: Velocity vectors for a 3D bubble rising without growth, with a vertical velocity of $0.75\,U_b$ (top) or $0.50\,U_b$ (bottom) added to all vectors.

In the limit, ghost bubbles might grow so large as to become almost stationary, in which case the bubble reference frame will be almost the same as the laboratory frame. If this happens, the velocities will resemble those given earlier in Figure 3.5.

Considering Figures 3.16 and 3.17 in combination with Figure 3.5, it can be seen that the paired, counter-rotating vortices, centred on or about the points $(x = 0, y = \pm 10)$, are more apparent when the reference frame shift velocity is low (*i.e.* when the new reference frame is similar to the bubble reference frame, which is most likely to be the case if the bubbles move slowly at large freeboard heights). Nevertheless, the vortex ring pair should still be visible even when the difference between reference frames is greatest, *i.e.* at low freeboard heights.

As well as the vortex ring, something else to look out for in PIV work is the small region at the centre of each ghost bubble, which contains vertical velocities which are very much larger than those in the surroundings. This feature can be seen in all reference frames. In contrast, a jet of rapidly-rising gas in the freeboard will be expected to occupy a much greater height of the freeboard space and might not involve the formation of any recirculation regions.

3.4 Alternative stream functions

Some alternative stream functions were deduced by de Kock & Judd (1965a), at about the same time as the work of Pyle & Rose (1965). However, de Kock & Judd considered a stationary void in a fluidised bed, rather than a rising bubble. Their boundary conditions were the same velocity conditions as those used by Pyle & Rose, *i.e.* equations 3.14 and 3.15. The velocities used were those from Davidson's theory, given earlier in equations 3.12 and 3.13. In particular, it can be noted that $v_\theta = 0$ at the edge of a stationary void, because $U_b = 0$ (equation 3.13).

There is no vorticity outside the void, where Laplace's equation is obeyed. When gas enters the void, de Kock & Judd (1965a) recognised that a vorticity component must be created.

1. First, they constrained the vorticity inside the void to be constant along streamlines. This is equivalent to assuming that the flow is inviscid. By matching velocities at the edge of the void, they found:

$$\psi_i = -3.437\, a\, U\, J_1\,(1.8412\, r/a)\, \sin\theta \tag{3.49}$$

2. By assuming that the vorticity inside the void region was proportional to the distance from the axis (*i.e.* $r\,\sin\theta$) only, they concluded that:

$$\psi_i = -U\left(3\,r - \frac{r^3}{a^2}\right)\sin\theta \tag{3.50}$$

This solution assumes that viscous flow applies in the horizontal direction, because streamlines are not always parallel to the axis. No viscosity applies to fluid elements moving vertically, however.

3. In a later reply in the same journal, following some highly critical comments by Collins (1965), de Kock & Judd (1965b) added that the vorticity inside the void could also be proportional to $\sin\theta$ only. This assumes that viscous forces apply tangentially but not radially. The result was a special case of Pyle & Rose's (1965) more general solution for a rising bubble, which can be found by substituting $\beta = 0$ into equation 3.16.

The assumption that viscous dissipation is anisotropic (solutions 2 and 3) does not seem physically realistic. The first solution (equation 3.49) is therefore the most convincing, giving inviscid (but rotational) flow.

It can be noted that none of these equations for the flow inside a two-dimensional bubble (or void) satisfies Laplace's equation. This would require:

$$\nabla^2\psi = \frac{\partial^2\psi}{\partial r^2} + \frac{1}{r}\,\frac{\partial\psi}{\partial r} + \frac{1}{r^2}\,\frac{\partial^2\psi}{\partial\theta^2} = 0.$$

Azevedo & Pereira (1991) extended de Kock & Judd's analysis to cover rising two-dimensional bubbles. They assumed that the vorticity inside the bubble was proportional to ψ_i (*i.e.* inviscid flow) but did not require $U_b = 0$. They found:

$$\psi_i = C\, U_o\, J_1(k\, r)\, \sin\theta. \tag{3.51}$$

The constants C and k can be determined using equations 3.12–3.15. For k:

$$\frac{(k\,a)\;J_0(k\,a)}{J_1(k\,a)} = 1 - \frac{U_b}{U_{mf}} = 1 - \beta. \tag{3.52}$$

Kreysig (1983) invites the reader to prove that $\int x^n\; J_{n-1}(x) = x^n\; J_n(x) + c$, for some n. Choosing $n = 1$ then differentiating both sides, it can be shown that:

$$\frac{x\;J_0(x)}{J_1(x)} \;=\; 1 + \frac{x}{J_1(x)}\;\frac{dJ_1(x)}{dx}. \tag{3.53}$$

In the earlier work of de Kock & Judd (1965a), $U_b = 0$ hence $\beta = 0$ also. Therefore, in their version of equation 3.52, the right-hand-side was simply equal to 1. Equation 3.53 now reveals that this requires $[k\,a/J_1(k\,a)] \times [dJ_1(k\,a)/d(k\,a)] = 0$, *i.e.* the first derivative term must be equal to zero. The first turning point of $J_1(k\,a)$ satisfying $k\,a > 0$ was considered to be the only physically-realistic choice ($J_1(1.8412) = 0.5819$, a maximum). Hence they settled on equation 3.49.

Azevedo & Pereira did not assume $U_b = 0$, so equation 3.52 must be solved iteratively. However, their Bessel function term will clearly *not* normally run from $r = 0$ up to a maximum turning-point at $r = a$ (*i.e.* $k \neq 1.8412/a$), because $U_b \neq 0$ in general. Their other constant was given by: $C = 2\,a\,\varepsilon/J_1(ka)$.

Inside a three-dimensional void, de Kock & Judd (1965a) realised that an equivalent to their first two-dimensional solution (equation 3.49) was not possible. Their assumptions, which led to equation 3.50 in two dimensions, now gave rise to an equivalent of Pyle & Rose's (1965) solution for a three-dimensional bubble. To obtain the expression for a void, $\beta = 0$ can be substituted into equation 3.24.

Davidson & Harrison (1963) showed that the flowrate of fluidising gas into a three-dimensional bubble is $3\,U_{mf}\,\pi\,a^2$. Taking $3\,U_{mf}$ to be a typical velocity to use in a Reynolds number based on the bubble diameter and the properties of air, values of Re $\sim$ 1000 can be found, using the bubbling conditions found in the larger fluidised bed used in later experiments ($U_{mf} = 0.07\,\mathrm{m\,s^{-1}}$ and $D_b = 0.07\,\mathrm{m}$). The

potential flow assumption that $\mu = 0$ is most valid at large Reynolds numbers. With Re = 1000, it is not obvious that inertial forces will actually dominate viscous forces, but, equally, the assumption is not obviously wrong.

3.5 Experiments on fluidisation bubbles

Bubble/dense phase clouds were photographed by Wace & Burnett (1961), after injecting NO_2 gas into an incipiently-fluidised two-dimensional bed and photographing the bubbles using a cine camera. It was found that the shapes of the bubble and the cloud were closer to that of a section through an upside-down apple, rather than the cylinders predicted by Davidson (Various, 1961). Other workers, *e.g.* Rowe & Partridge (1963) and Reuter (1966) also observed deviations from a cylindrical shape in the wake region below rising bubbles. The shapes predicted by Collins (1982, 1983) for growing bubbles seem better.

A careful analysis of such photographs allowed Rowe *et al.* (1964) to make measurements of the 'diameters' of various bubbles and their associated clouds. They found that the experimental cloud diameter was usually somewhat lower than that predicted by Davidson's model (equation 3.9). Better agreement was found when alternative flow models developed by Jackson (1963b) and Murray (1965) were considered. In these models, gas momentum was taken into account. Clift *et al* (1972) state that Jackson's and Murray's approaches are actually equivalent, although differences in notation obscure the fact.

On the other hand, Littman & Homolka (1973) found that the experimentally-measured pressure field around a bubble in a two-dimensional bed is Laplacian and well described by Davidson's model; the models of Jackson and Murray fared less well under this test, predicting dynamic pressures which were much too low. Stewart (1968) found that the pressures predicted by Jackson's model were much closer to data obtained by Reuter when a different equation for the bubble rise velocity was used.

Rowe & Partridge (1963) obtained reasonably good agreement between streaklines obtained by time-stepping the gas velocities around a rising Davidson bubble and experimental observations (based on the injection of NO_2 gas to the side of an injected air bubble in a two-dimensional bed).

Whilst Davidson's model is relatively simple, it does appear to describe the gas flow around a fluidisation bubble reasonably well. Alternative models have been developed by Jackson and Murray, but their refinements did not solve all the problems. Pyle & Rose (1965) extended Davidson's stream functions and velocities into the bubble. Collins (1982, 1983) modelled bubble growth by adding a source term to Davidson's stream function outside the bubble, but it is not clear how this model can be extended into the bubble. Despite the problems, it is hoped that these models can give at least a rough indication of the expected velocity fields around ghost bubbles in the freeboard region above a fluidised bed.

Chapter 4

Methods of flow measurement

4.1 General comments

4.1.1 The aims

Most fluids are colourless and transparent, the most common examples being water and air. As one constituent element of a homogeneous fluid is identical to its neighbours, it is not possible to pick out an element in order to follow its progress and track its position as a function of time.

A numerical record of a flow might consist of a three-dimensional grid of closely-spaced points, each one recording, at one instant in time, the local fluid velocity. (Other fluid properties, such as temperature, pressure and chemical composition, could also be useful if chemical reactions are taking place.) Comparing nearby points would allow secondary quantities such as vorticity to be determined, whilst a similar record a short time later would allow the development of the flow with time to be evaluated. A record of this type would be easy to compare with a Computational Fluid Dynamics (CFD) simulation. But how can this be obtained from a vessel of fluid?

4.1.2 Point measurement techniques

There are a number of well-established techniques available for measuring the velocity at a point in a flow. These techniques tend to yield data for only one point in the flow at a time; if information is required at other points, then the probes need to be moved to different locations.

A Pitot-static tube measures the relative pressure at a point in the flow; because the tube is a dead-end leading only to a manometer, the flow immediately in front of it is static. The difference in velocities creates a pressure difference relative to other points in the same cross-section of the flow, which is governed by Bernoulli's Law. This pressure difference can be related to velocity.

A Hot Wire Anemometer inserts a conductive probe into the flow. At equilibrium, the heat generated within the probe from the resistance encountered by a constant electric current will balance the heat lost to the flow. The equilibrium temperature of the probe can be measured by means of the potential difference across it, and this temperature depends on the fluid flow past the probe.

Laser-Doppler Anemometry (LDA) is a well-established technique. Whilst it does not rely upon anything as invasive as inserting a wire or a tube into the flow, it does nevertheless require the flow to be seeded with small particles. For one-dimensional LDA, a laser beam is split into two derivative beams, which are directed at a point in the flow. When these coherent beams intersect, an interference pattern is set up. As a seeding particle moves through this interference pattern, it scatters light which can be collected by photomultiplying optics. The scattered light is relatively bright when the centre of a seeding particle lies within a zone of constructive interference, and is relatively dim when the LDA beams interfere destructively. The distance between zones of constructive interference is a property of the laser beams, and the time taken for the particle to move between these zones can be obtained from the signal obtained by the photomultiplier. A more advanced version of this technique allows two orthogonal velocity components to be measured.

Although these techniques have proved to be useful in many flow situations, they do have certain limitations. In a simple example such as laminar flow around a cylinder, consistent and reproducible readings can be taken by moving the probe point across the flow and upstream or downstream. A technique such as LDA will yield a time-history of the flow velocity at each point, which can be compared with the time-history at other points.

The flow above a bubbling fluidised surface is a different matter. Whilst a lot of useful information was obtained by Ramírez de Santiago (1990) using LDA measurements of the gas velocity above a bubbling liquid, his discovery that gas velocities lay in the range $u_g \pm 12\,u_g$ tells its own story. Comparing two points at the same height above the bubbling surface can be difficult. For example, if the time-histories are significantly different, is this just because the velocity sampling caught different numbers of bubbles at each point? If the time-histories are identical, are the velocities in phase? If the two points are very close together then they should be, but there could be a constant phase difference.

4.1.3 Flow visualisation in liquids

The first step towards whole-field velocity measurement is flow visualisation. This has been achieved in liquids by many methods (Werle, 1989). One way is to add streaks of a second, visible fluid to the fluid under investigation. Examples include coloured inks, fluorescent dyes, chemicals (such as potassium permanganate) and emulsions (*e.g.* milk). These fluids can be selected so as to match the physical properties of the test medium. A classic example of this type is Reynolds' paper from 1883, concerning laminar and turbulent flow regimes; extracts from this paper have been reproduced by Rott (1990) and in *Proc Royal Soc*, Oct 1995.

It is possible to introduce dye by injection upstream, through a row of injection nozzles. Combining such a line of sources with a clear overall direction of fluid flow produces a plane of visualised flow, which can be photographed by a camera placed perpendicular to the flow. If the flow is steady-state, the visualised streaks will

follow the local direction of the flow, revealing many streamlines simultaneously; if the flow is not too turbulent and diffusion is not excessive, the streaks will remain sharp. Van Dyke (1982) gives a number of examples of flow past an immersed object using dye streaks.

Another class of visualisation method in liquids is to add tracer particles, which may be solid (glass spheres, aluminium particles) or gaseous (hydrogen bubbles, generated by electrolysis in water). Where such particles are used, they will generally be distributed throughout the fluid under investigation.

In contrast to the case of hydrogen bubbles emanating from a single, straight wire, there can be problems due to this dispersion: a region of particular interest needs to be illuminated brightly, so that particles within this particular region are much brighter than their neighbours. When a photograph is taken, only those tracer particles which lie within the illuminated zone will be bright enough to be recorded. A proviso, however, is that the total particle concentration must not be so large as to form an opaque cloud, obscuring the illumination plane.

Particle-laden flows can be illuminated by a spotlight combined with optics in order to produce a thin focused sheet, or by a laser. Lasers produce a narrow, intense beam of light with little divergence. A laser beam is coherent, *i.e.* all the waves within it are in phase, but this feature is not usually exploited.

4.1.4 Constraints on the choice of seeding

The Stokes Law terminal velocity (u_t) for the descent under a gravitational acceleration (g) of a small particle of diameter d_p and density ρ_s, in a fluid of density ρ_f and viscosity μ, is given by equation 2.6: $u_t = (\rho_s - \rho_f)\, g\, {d_p}^2 / 18\, \mu$ (as stated in Chapter 2). From this, it can be seen that u_t can be minimized by matching the particle density to that of the fluid and/or by using small particles. The terminal velocity represents the velocity with which a particle within the flow tends to fall out of suspension, due to gravity.

Another factor to consider is the reaction time of a seeding particle in the particular fluid flow, which can be defined as the time it takes to accelerate from rest up to 95% of the local fluid velocity. If this time is too long, then although there may be minimal velocity error when measuring a constant-velocity flow, the particle will not be able to follow the velocity changes encountered in a real flow. For example, when an air stream flows past a car body, the air flow is deflected up by the windscreen and passes over the roof. The velocity of a fluid element changes as it approaches the car (even when the velocity at each point in space is independent of time) and there are also random, short-term fluctuations due to local turbulence. A seeding particle should be small enough to follow every detail of the fluid flow and should not crash into the windscreen, like an airborne insect.

Rudinger (1980) analysed the one-dimensional dynamics of a small isolated particle, released from rest into a gas flow and experiencing a Stokes Law drag force. For the case of a uniform gas velocity u_0, the particle velocity u_p is given, as a function of the time t after the release, by

$$u_p = u_0 \, [1 - e^{-t/\tau}],$$

where the constant $\tau = \rho_s \, d_p^2 / 18 \, \mu$ is known as the velocity relaxation time of the motion. Similarly, when a particle is ejected at $u_p = u_0$ into a stagnant medium, its velocity will be given by $u_p = u_0 \, e^{-t/\tau}$.

Putting some numbers into the above equation, it is found that the particle velocity will rise from zero to about $0.95 \, u_0$ when the time $t = 3 \, \tau$. This velocity is fairly close to u_0, so it seems that $3 \, \tau$ is a useful time to bear in mind in the context of applying a given seeding material to a particular flow. When it comes to following turbulent fluctuations in a gas velocity, a seeding particle will not be starting from rest but from a velocity fairly close to the mean gas flow. In this time $3 \, \tau$, the *difference* in velocities will once again decrease by 95%, although the absolute error will be smaller to start off with.

4.1.5 Particulate seeding in air

Various neutrally-buoyant seeding particles are available for the study of water flow; examples include conifer pollen and hollow glass spheres. However, when investigating the flow of typical gases, neutral-buoyancy is harder to achieve. Because of the very low density of gases, most candidate seeding particles, if solid or liquid, will have a density many hundred times greater than that of the gas. Alternative methods involving hydrogen-filled soap bubbles are not always suitable.

One way around the problem is to make the particles small. Consider two groups of seeding particle, where particles in group A are half the size of particles in group B. For particles in group A, the Stokes Law terminal velocity and the velocity relaxation time are one quarter of those from group B, which means that they will follow the fluid flow more reliably. However, the amount of light scattered by an illuminated particle is also proportional to d_p^2, if the seeding particles are much larger than the wavelength of the illuminating light (Adrian, 1991). Particles in group B will therefore appear to be about four times times as bright as those in group A, given identical conditions of illumination.

As a result, a compromise size must be sought, which is small enough to follow local variations in fluid velocity yet large enough to be photographed successfully with the available methods of illumination.

4.2 Light sheets

4.2.1 The expanded beam method

A typical experimental arrangement for flow visualisation using small seeding particles is shown in Figure 4.1. A cylindrical lens expands the beam from a laser into a thin, diverging triangular sheet. This two-dimensional light sheet passes through a converging lens to collimate it and is then directed into the test section

of the flow, where it illuminates a narrow plane (typically about 1 mm or so thick). Usually the illuminated plane is vertical or horizontal.

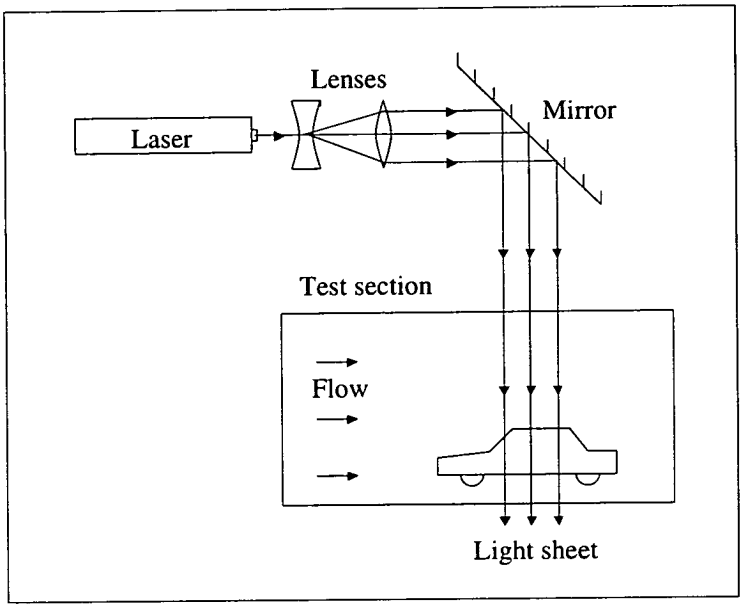


Figure 4.1: A light sheet from a cylindrical lens.

Comparing particles within the light sheet to those outside it, the former will be brighter and, if a photographic recording is made, much more accurately focused (due to the limited depth of field). A photograph will record only the particles within the thin, well-lit illumination plane.

Photography along the normal to the illumination plane records a streak of images of each particle, revealing information about velocity (streak length in the given exposure time) and the shape of the streamlines, provided that the seeding particles are small enough to follow the flow reliably.

Using a light sheet can create shadows of a test object (*e.g.* flow around an opaque cylinder). These problems can sometimes be avoided if appropriate care is taken when designing the experiment.

A problem with expanding a laser beam in this manner is that the intensity distribution across a laser beam is usually given by a two-dimensional Gaussian function. When expanded into a light sheet, this distribution function is stretched in one dimension. The central region of the resulting sheet is significantly brighter

than its edges, which complicates subsequent photographic processes. Further, the focusing method of Figure 4.1 means that the resultant light sheet is limited in size by the diameter of the converging lens: either a narrow light sheet is obtained, or a large and expensive lens has to be used.

An alternative method of collimating a sheet of light diverging from a point is to replace the convex lens with a concave mirror. This will bring light from a distant object into focus at its focal point, or, conversely, convert light diverging from a source placed at its focal point into a parallel sheet. This avoids the problems associated with large lenses, should a sheet wider than (say) 1 m be required.

4.2.2 A scanning-beam light sheet

In the previous subsection, a laser beam was expanded using a lens, *i.e.* by refraction. Another way to obtain a divergent light source is to direct the laser beam at a multi-faceted rotating mirror. The combination of a rotating mirror with a large converging mirror produces an arrangement illustrated in Figure 4.2.

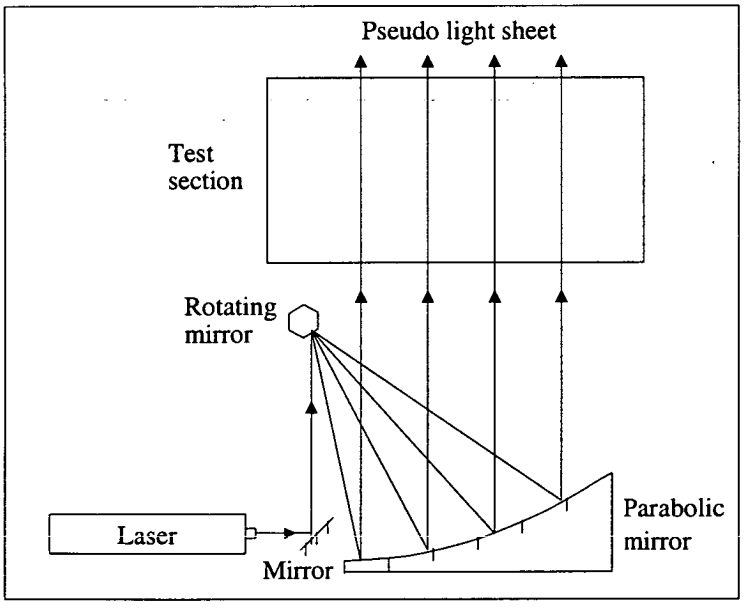


Figure 4.2: A light sheet from a combination of mirrors.

In this case, only a pseudo-light-sheet is obtained, because the 'sheet' is in fact made up of repeated scans by a single beam. Illumination is not simultaneous

across the sheet and the sheet is inherently pulsatile. If these features are not desirable, then the time difference between pulses can be made small by the use of a rapidly-rotating mirror with many sides. For example, an eighteen-sided mirror spinning at 250 Hz produces 4500 scans per second and a time interval between successive scans of only 0.22 ms. This time is much too short for a human eye to detect any flicker and there is no theoretical reason why even shorter scan times cannot be obtained. Compared to an expanded light sheet from a cylindrical lens, scanning-beam light sheets have a more uniform effective intensity profile.

4.2.3 The geometry of a scanning beam system

Using a scanning beam system, the light sheet emerging from the parabolic mirror is not uniformly bright, because the speed at which the laser beam traverses the 'sheet' is not uniform. Gray *et al.* (1991) discuss a diagram similar to Figure 4.3, for a parabolic mirror of length L in the x-direction. The converging mirror's focal point is at $(0, L/2)$ in the Cartesian co-ordinate system used here, and this is where the rotating mirror is placed.

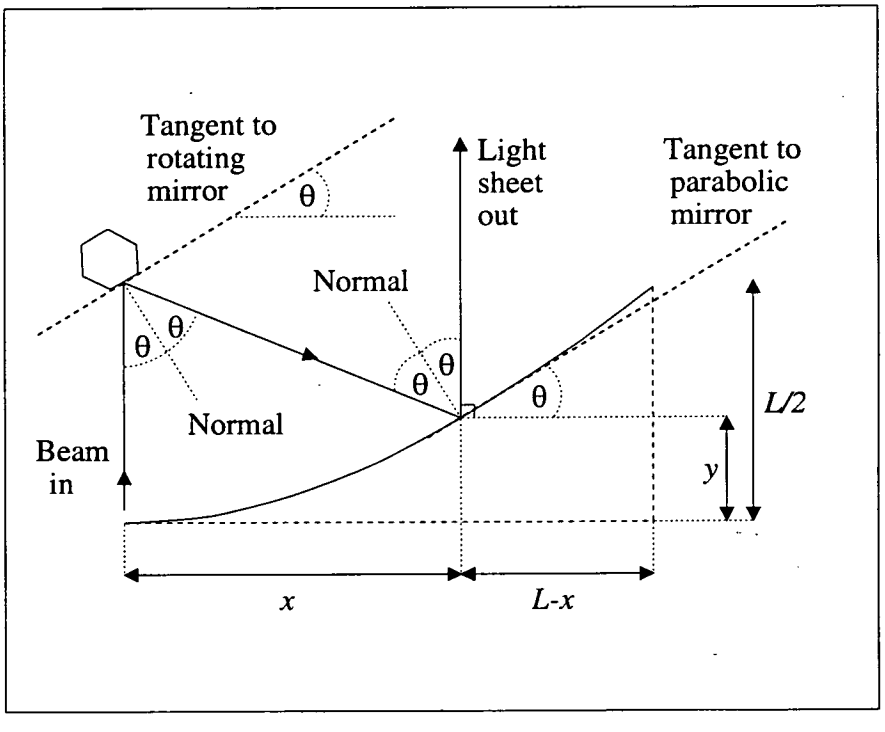


Figure 4.3: Analysis of the scanning beam method.

The z-direction, not indicated on Figure 4.3, will be perpendicular to the x- and y-directions, *i.e.* going into and coming out of the plane of the paper.

The path of the laser beam at any instant in time is cylindrical, assuming that its divergence is low. A cross-section through the beam leaving Figure 4.3 lies in the xz–plane and the beam propagates in the y-direction. At any y along the beam path, the beam's profile is a two-dimensional Gaussian function of x and z.

Gray *et al.* work in terms of the total angle the beam is deflected through at each mirror, which on Figure 4.3 represents 2θ. However, a geometric approach based on half of this angle yields a simpler form of their result, more easily.

If the emerging light sheet is to be collimated for all θ, then the tangent to the rotating mirror must always be parallel to the tangent to the converging mirror. The angles marked in Figure 4.3 must, therefore, all be equal to θ, including the angle between the x-axis and the tangent to the collimating mirror. This condition can be used to determine the shape of the collimating mirror, which is given by $y = x^2/2L$; the mirror surface needs to be parabolic.

The slope of the tangent to any point on a curve is equal to the tan of the angle between the tangent and the x-axis. Here:

$$\frac{dy}{dx} = \frac{d}{dx}(x^2/2\,L) = \frac{x}{L} = \tan\theta$$

The above expression is a simplification of the relation between x and $2\,\theta$ given by Gray *et al.* Clearly the function $x(\theta)$ is not strictly linear because of the tan term, although this non-linearity is less significant for 'small angles' where $\tan\theta \approx \theta$. The equation can be differentiated:

$$\frac{dx}{dt} = \frac{dx}{d\theta}\frac{d\theta}{dt} = \frac{d}{d\theta}(L.\tan\theta).\frac{d\theta}{dt} = L.\sec^2\theta.\frac{d\theta}{dt}$$

4.2.4 The intensity profile in the x-direction

When considering the profile of the laser beam in the x-direction, a point to note is that the scans of the laser beam are also in the x-direction. The light energy reaching each particle during each scan involves an averaging effect. As the beam approaches the particle, points on the particle surface will be, in sequence: not illuminated by the approaching beam, illuminated dimly by the leading edge of the beam, illuminated brightly by the central region of the beam, illuminated dimly by the trailing edge of the beam and then no longer illuminated.

If the seeding particles are significantly smaller than the width of the laser beam (*e.g.* 10 μm seeding particles illuminated by a 1 mm–wide beam), the light energy delivered to the particle is the product of the laser power (averaged in the x-direction) and the time of illumination. The illumination time is approximately the time taken by the scanning-beam system to move the beam a distance equal to its width. The energy delivered to each particle is therefore inversely proportional to the beam's scanning velocity in the x-direction, *i.e.* dx/dt:

$$\text{Energy} \propto 1/\frac{dx}{dt} \propto \frac{1}{\sec^2\theta.\frac{d\theta}{dt}} = \cos^2\theta/\frac{d\theta}{dt}$$

Usually the speed of the rotating mirror is kept constant, in which case the term $d\theta/dt$ will be constant whilst the laser beam strikes a given facet (*i.e.* when $0° < \theta < 45°$ for an octagonal mirror). Whenever the laser beam starts to strike a different facet, there is a step-change in θ, as this angle suddenly returns to $0°$. During the constant-rate periods, $d\theta/dt$ is proportional to the rotational frequency of the rotating mirror and to the number of facets of this mirror.

For a rotating mirror with eight faces, the angle between faces is $360°/8 = 45°$; this doubles (because of the optical lever rule) to produce a maximum beam deflection of $2\,\theta = 90°$. As $\cos^2(0°) = 1$ and $\cos^2(45°) = 0.5$, the beam's progress in the x-direction will be twice as fast at $\theta = 45°$ as it will when $\theta = 0°$. Thus regions near $x = 0$ ($\theta = 0°$) will receive twice as much light as those near $x = L$ ($\theta = 45°$). This difference in illumination is equivalent to one aperture stop on a

camera, which is significant but usually acceptable.

For convenience, the rotating mirror and the parabolic mirror can be clamped in place relative to one another and mounted in a box. For a system using an octagonal rotating mirror, the dimensions of a scanning-beam box which produces a 1 m light 'sheet' are about 1 m × 0.5 m (*cf* Figure 4.3 with $L = 1$ m).

4.2.5 The intensity profile in the y-direction

Following Figure 4.3, the light sheet from a scanning-beam system will have a certain intensity as it passes the level of the rotating mirror (the line given by $y = L/2$ in this figure). As discussed above, this intensity is a function of x, the distance across the light sheet.

As a laser beam propagates forwards (in the y-direction of Figure 4.3), the amount of energy it carries will decrease as it moves away from its source. This attenuation is due to light scattering as the beam travels through the various media it encounters, *e.g.* dusty air, glass surfaces and seeding particles.

Within a seeded test section, the intensity of the beam is given by Beer's Law:

$$I = I_0 \, e^{-k\,C\,L},$$

where the beam's intensity has the value I_0 at a reference point and I at another point a distance L away ('downstream'). C is the concentration of scattering particles in the test section and k is a constant for the given beam-particle system.

If a flow is being measured by a method involving seeding particles, then there must be some attenuation of image intensity if the particles are to be photographed successfully from a direction perpendicular to the plane of the light sheet; after all, the purpose of the seeding particles is to scatter light towards the camera. Beer's Law gives a warning that the seeding concentration within a test section should

not be made too large, as this would cause the intensity of illumination to vary significantly within the flow field and could lead to images of seeding particles in the dimmer regions not being detected.

4.2.6 The intensity profile in the z-direction

As described above when referring to the profile in the x-direction, the intensity profile of a laser beam travelling perpendicular to the xz-plane is a Gaussian function of x and z. The dependence of this on x is reduced by the averaging effect of having the beam scan in the x-direction.

There is no such averaging effect on the profile in the z-direction. This will be similar to a Gaussian function, *i.e.* the intensity is quite a strong function of z. Particles in the centre of the light sheet will be significantly brighter than those near the edges. Tracks made by particles with significant velocity components out of the illuminated plane (*i.e.* in the z-direction) will not be uniformly bright.

4.2.7 The number of rotating mirror facets

If a more evenly-illuminated light 'sheet' is required than that produced by an eight-sided rotating mirror, then this mirror could be replaced by a twelve-sided unit, to smooth out the intensity profile in the x-direction. Here, the range of 2θ is from 0 to 60° and the maximum scanning velocity (at $x = 0.577\,L$) is 1.33 times faster than that at $x = 0$.

A drawback is that, because the parabolic mirror is now being used only in the range $0 < x < 0.577\,L$, the value of L has to be made 1.73 times larger in order to provide the same width of light sheet. Although the required parabolic mirror still needs to extend only 1 m in the x-direction of Figure 4.3 in order to provide a 1 m-wide light sheet, the rotating mirror now has to be 1.73 times further away from the parabolic mirror, in order to provide the same expansion. The cross-section of such a scanning box would need to be 1 m × 0.87 m.

Table 1 lists the characteristics of scanning-beam boxes based on various rotating mirrors. It can be seen that eight-sided rotating mirror systems offer a good compromise between size of apparatus and uniformity of intensity (in the x-direction of Figure 4.3).

No. sides	max θ	Range of brightness	Cross-section ($x \times y$)
6	60°	0.25–1	1 m × 0.29 m
8	45°	0.50–1	1 m × 0.50 m
12	30°	0.75–1	1 m × 0.87 m
18	20°	0.88–1	1 m × 1.37 m
36	10°	0.97–1	1 m × 2.84 m

Table 4.1: The properties of various scanning beam boxes, each providing a 1 m wide light sheet

4.2.8 The brightness of scanning-beam and expanded light sheets

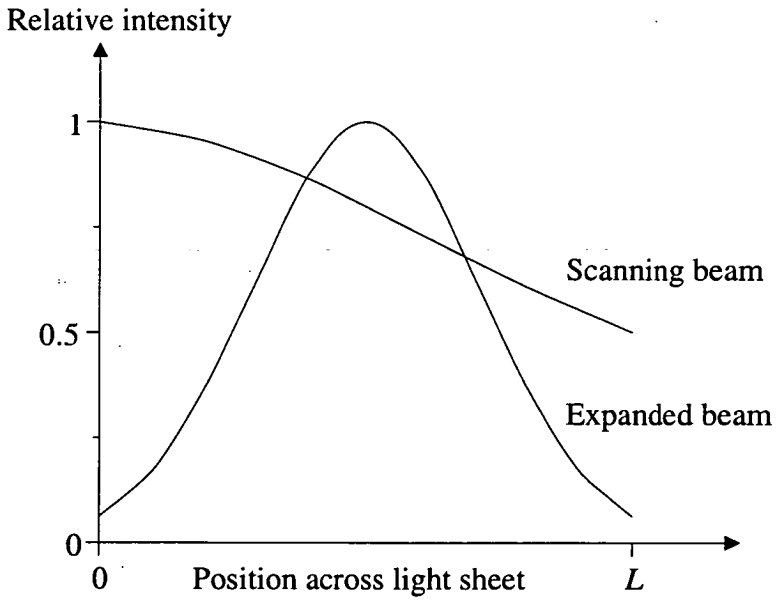


Figure 4.4: The brightness of a light sheet produced by the scanning beam and expanded beam methods.

Figure 4.4 shows the relative intensity of a light 'sheet' produced using a scanning-beam box fitted with an octagonal rotating mirror, as a function of position along the 'sheet' (the variable x of Figure 4.3). Also shown is a possible profile for an expanded beam, which is Gaussian.

It can be seen that the scanning beam method produces a more uniform light 'sheet' than the expanded beam method, assuming that it is necessary to use most of the power from the laser in each case.

If a more uniform 'sheet' from a scanning beam system is required, a rotating mirror with more facets could be used (as discussed above). This would still use all of the light from the laser but would occupy more space in the laboratory.

Under the expanded beam method, the light sheet has a Gaussian intensity profile, stretching (in theory) from $-\infty$ to ∞ (although the size of the optical components will limit this in practice). When sizing something with an intensity distribution function, it is necessary to talk about a length which includes 90% of the light energy, or a length containing only light above a certain brightness threshold. In each case, a certain amount of light in each tail of the intensity function has to be neglected, which means that not all of the light from the laser is being used to form the light sheet. For the expanded beam profile shown in Figure 4.4, almost all of the available light is being used, but this means that the light intensity at each edge of the light sheet is only 6% of the maximum.

If laser power is not limited, then a more uniform light sheet of given width could be obtained by expanding the laser beam to a greater extent, so as to ignore a larger fraction of each low-intensity 'tail'. Of course, this requires more power to be delivered by the laser if the brightness of the light sheet is to be kept constant.

4.3 Streak photography

If a two-dimensional flow is suitably seeded and illuminated, then those seeding particles present in the flow and lying within the light sheet will scatter this incident light in all directions. All that needs to be done is to observe what happens or photograph the proceedings; the flow has been visualised.

If the illumination is achieved using a continuous illumination method, *e.g.* an expanded laser beam or a scanning beam system operating at a high scanning frequency, then each photograph will record the passage of many particles during the time of exposure.

Streak photography can be used in different modes, according to the seeding density in the flow. If the seeding density is high, then a short exposure can be taken, perhaps lasting a fraction of a second. Each of the many particles will produce a short streak, the length of which will be proportional to the particle's velocity. If the seeding concentration is low, then a longer exposure can be made, perhaps lasting several seconds or even minutes; in this case, particles can create streaks running from one edge of the photograph to another and the streaks from different particles will not necessarily be simultaneous.

Before moving on to a discussion of PIV, it is helpful to think about factors which can cause problems with streak photography, because these can also influence PIV.

The method of illumination involving a thin sheet of light is convenient to use in conjunction with photography, because a camera has a finite depth of field (often only a few mm under typical PIV conditions). Particles outwith the light sheet are not illuminated and their images are out of focus. However, this thinness can cause problems in flows where the velocities have significant components in all three dimensions. Particles which lie within the light sheet and move parallel to it can be photographed over the whole of their path. On the other hand, particles with significant velocity components perpendicular to the light sheet may move into or out of the light sheet, and thus may only be illuminated during a portion of their journey. Therefore, out-of-plane motion gives rise to artificially-shortened streaks: even if the length of each streak can be measured very accurately, it is not a reliable measure of velocity because the time taken to make the streak is not certain. All that is known is that this time must be less than or equal to the time taken to make the photographic exposure.

This problem of 'signal dropout' is compounded by the Gaussian intensity profile across the light sheet: a particle only has to move a short distance across the beam for its brightness to drop substantially. Clearly it is better to use short exposure times if the flow has significant out-of-plane components.

On the theme of out-of-plane motion, streak photography does not usually measure the velocity component perpendicular to the light sheet. All that can be determined is the projection of the local fluid velocity onto the illumination plane.

Spinning non-spherical particles may appear to have a fluctuating intensity, even if they move parallel to the light sheet. This is because a stationary aspherical particle will have a direction of maximal light scatter. If the particle is spinning, then the intensity observed by a camera will vary as a function of time. It is easy to imagine that such a particle might not photograph well at every stage of its journey, giving rise to a bright-dim-bright-dim streak. In extreme cases, the flow record might appear to contain many short streaks, rather than one long one.

Finally, the streak photography method described above records the path of a seeding particle between a start point and an end point. No information is stored *by the technique itself* as to the order in which the points were actually visited. Often this does not matter (*e.g.* for laminar flow in a pipe, all velocities will be in the same direction as the bulk flow), but it is a fact that light sheet techniques have an inherent directional ambiguity unless special steps are taken.

Whilst a short-time streak photograph might not yield as accurate a local velocity measurement as is obtained using a point-measurement technique such as LDA, any velocity results which are obtained can be put into context by comparing them to velocities at neighbouring points. Structures in the flow, *e.g.* vortices, should be easy to spot. A whole-field technique such as this contains a lot of qualitative information about the flow at one instant in time, which is particularly useful if the flow is unsteady. An LDA experiment provides a series of velocity measurements at different times for one particular point in space, whereas streak photography and PIV provide velocity measurements at many points in space, for one particular moment in time.

4.4 Particle Image Velocimetry

PIV is a fairly natural development of streak photography for well-seeded flows. Its aim is to obtain accurate information about flow velocities across the whole of a flow field, without suffering from the relatively-large experimental errors obtained using streak methods.

4.4.1 The advantages of a pulsing light sheet

The main problem with streak photography is the apparent shortening of displacement vectors due to out of plane motion: the actual time taken to make each streak is not known. Compared to streak photography, a PIV record is obtained using a pulsing light sheet. The records obtained resemble streaks, but have the advantage that the number of images recorded can be used to determine the exact period of illumination. It is helpful to set the pulsing rate so that each particle moves at least a distance equal to its diameter during the time between exposures, so that images do not overlap and can easily be counted.



(a) Continuous streak record



(b) Particle-image record

Figure 4.5: A comparison between a short streak and a PIV record.

Figure 4.5 illustrates this. It shows how the number of distinct particle images can be used to determine the time between the first and the last recorded image of a particle which has a significant out-of-plane velocity component. The time taken to record the continuous streak record in Figure 4.5(a) is uncertain; if the

camera took an exposure lasting one sixtieth of a second, then all that is known is that this time was less than about 17 ms. In Figure 4.5(b), however, it can be seen that the particle-image method recorded nine distinct images; if the stroboscopic illumination rate was 1000 pulses per second, then the particle track must have been made over a period of 8 ms (N objects have $N-1$ gaps between neighbours).

With a particle-image record, it can even be possible to estimate the out-of-plane displacement component, in some cases. If the particle in Figure 4.5(b) started outside the light sheet, moved across the light sheet and then exited, it will have left behind various clues. The particle will have been illuminated on only some of the 16 or 17 stroboscopic flashes during the exposure. Also, the first and the last images might be dimmer than the other images, as they would have been taken on the fringes of the light sheet (where the intensity is lower). In Figure 4.5(b), for example, the out-of-plane velocity component can be estimated to be the thickness of the light sheet divided by 8 ms. However, it is also possible that a particle with a curved trajectory moved into and out of the light sheet via the same side, without passing through the centre of the light sheet. Halloin & Jottrand (1995) obtained three-dimensional velocity measurements using this method, but added long pulses of light at the start and end of each stroboscopic sequence; the third velocity component was only calculated for tracks without streaks.

Further, if there are velocity gradients within the flow, these can be exploited by a particle-image method. Using a particle-image method, the displacement between successive images will be seen to vary when each particle track is examined in detail. A displacement gradient can be calculated and then used to determine the fluid displacement with more accuracy (Hind & Christy, 1996). In the worst case, displacement gradients can at least be noted.

4.4.2 PIV illumination methods

Three main methods of illumination can be used in PIV:

1. If a beam expansion method is to be used in combination with a Continuous Wave (CW) laser, then some form of optical chopper must be used in order to obtain pulses of light. This could be a rotating disk, a section of which has been cut away in order to allow the transmission of the light sheet for only a part of each revolution.

   A disadvantage of this method is that some of the light from the laser is already being discarded for the sake of a reasonably uniformly-lit light sheet (due to the Gaussian intensity profile in a cross-section of the expanded beam) and now much of what remains is discarded by the optical chopper. This technique is not suitable when light energy is at a premium, *e.g.* for the investigation of high-velocity gas flows.

2. Because a pulsing sheet of light is essential to PIV, the scanning beam method is highly suitable because it creates a relatively-uniform pulsing pseudo light 'sheet' from a relatively-convenient CW laser. Unlike method 1, this technique uses all of the light energy generated by the laser (typically between 10 mW (Helium-Neon lasers) and 15 W (Argon Ion lasers)) because the laser is always directed at some part of the flow. The rate at which the light 'sheet' pulses can easily be controlled, by adjusting the speed of the rotating mirror. A drawback is that the time for which each particle is illuminated is inversely proportional to the rotating mirror speed, so increasing this at constant laser power means that the energy given to each particle is reduced. CW lasers may not be suitable for high-velocity flows.

3. For high-velocity flows, a pulsed laser source might be required. These lasers deliver a constant amount of energy in each pulse. The time between pulses is controlled by the laser's electronics, so no rotating mirror is needed. The beam is expanded through a cylindrical lens and then directed at the flow. Pulsed laser systems can produce very intense illumination for short periods of time (*i.e.* anything obtained from such a laser is already pulsatile) so no optical chopper is needed. Drawbacks are that pulsed lasers are harder to align than CW lasers and are more hazardous.

4.4.3 Mean particle displacements

In practice, the PIV technique, as used in this research, is only slightly different from the method described above. The modifications made are a result of the analysis method which is used.

In keeping with the discussion above, it is possible to calculate fluid speeds by determining, for each particle track, the total distance moved and the actual time taken to make the track. The fluid speed can then be found by evaluating the quotient of these quantities (assuming that the seeding particles follow the fluid flow). To obtain the fluid *velocity*, it would also be necessary to measure the angle between the particle track and a reference direction, *e.g.* the horizontal. Alternatively, the total particle *displacement* could be determined for each track, as this is a vector quantity and therefore includes directional information. This method gives rise to the technique of particle tracking, which is most applicable at low seeding concentrations.

Instead, PIV usually calculates velocity from the *average* particle displacement during the time interval between successive pulses of light, for each of a series of small *interrogation areas*, arranged in a grid across the flow record. This yields an array of average-velocity vectors, covering the area photographed. The method of averaging particle displacements is useful, as it takes into account the effect of small velocity gradients across each interrogation region. Further, the averaging process avoids the need to worry about exactly how many images of each particle have been recorded (provided, of course, that there are at least two such images). There are a number of automated methods available for calculating this mean particle displacement, which will be considered in the following subsections.

Usually a PIV record contains between two and six images of each particle, with a seeding concentration high enough to ensure that each small interrogation region contains images from at least four different particles. A small piece of a typical PIV record might resemble Figure 4.6.

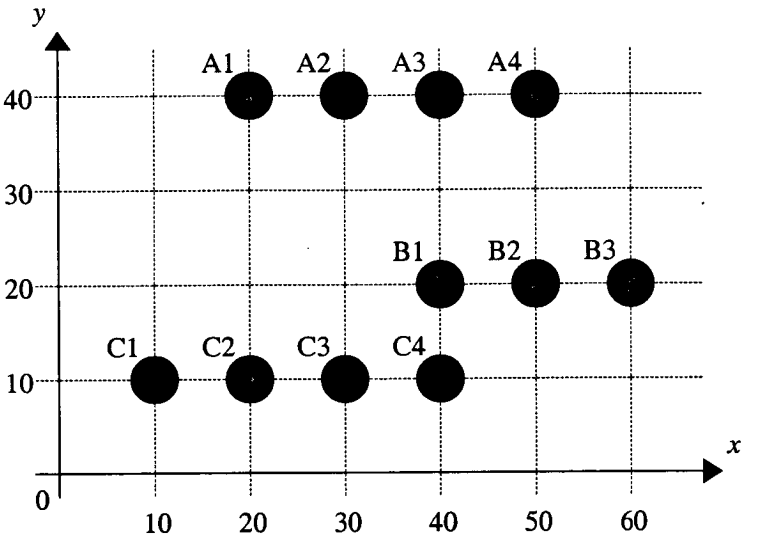


Figure 4.6: A typical PIV record.

Two of the particles have been caught by four illumination pulses. A third particle has only been illuminated three times, which could be because it had a noticeable out-of-plane velocity component and started off near the edge of the light sheet, or because the fourth image lay outside this interrogation area. The other particles, with the same velocity, may have started off in the centre of the light sheet or near the leading edge of the interrogation area, thus having further to travel before escaping.

However: on a photographic negative, particle images appear dark and the background appears light. An enlargement print made from film will be a negative of the photographic negative, *i.e.* it will show well-lit particles against a dark background, just like the photographed object. It is perhaps worth noting that, for convenience, Figure 4.6 is shown in 'inverse-video' mode, whereby bright objects appear dark and dark objects appear bright. This Figure resembles what might be seen on a photographic negative.

4.4.4 Young's fringes techniques

In the classical Young's fringes experiment, two pin-holes are made, close together, on an opaque surface. Both pin-holes are illuminated by light from the same

source, which must be coherent and of only one wavelength. A screen (*e.g.* a plane white surface) is placed beyond the pin-holes and is illuminated by light passing through the pin-holes. As far as points on the screen are concerned, each pin-hole acts as a distinct source of light. There is a small difference in the path length from each pin-hole to any point on the screen surface. If this difference is an integer number of light wavelengths, the light beams from each pin-hole always arrive in phase and add constructively, whereas if there is a 180° phase difference, the beams always add destructively and cancel each other out. The result is that an interference pattern is seen on the screen, consisting of a series of light and dark bands. The spacing between adjacent bands is inversely proportional to the distance between the pin-holes. The orientation of the fringes is perpendicular to the displacement vector from one pin-hole to the other.

PIV data can be stored on photographic film, each particle image appearing as a small dot on the photographic negative. In the Young's fringes method of analysis, small portions of the photographic negative are probed by a beam of coherent light (typically from a low-power He-Ne laser). In each interrogation area, each particle image acts as a source of transmitted light. A pattern is produced by the interference of light from each particle image with that from every other particle image. The most commonly-occurring image separation will produce the dominant fringe pattern. If the fringe separation and orientation of the dominant fringe pattern can be measured, then the displacement of particle images between successive illuminations can be calculated.

In a Young's fringes PIV analysis system, a photographic negative is mounted on a translation stage, which is used to move the negative between interrogations. A laser beam passes through a spatial filter and an aperture, before encountering the negative. Some of this light is diffracted around the particle images, whilst other rays pass through the clear regions of the negative, without deflection. The light is then focused by the lens of a small CCD camera. The bright central region of the image is removed; this is a direct image of the undiffracted laser beam and contains no information about the flow. The remainder of the image, hopefully containing a number of Young's fringes, is digitised by the CCD at an appropriate

resolution. A Fast Fourier Transform (FFT) is performed on the digitised data. The peak value of this function lies at a point in the image plane which corresponds to the image displacement.

An advantage of optical interrogation is that the PIV recordings are probed directly, avoiding errors due to digitisation. The Young's fringes pattern is created virtually instantly. Typically 3–15 fringes are recorded by the CCD camera. According to sampling theory, the images of the fringes need only be recorded at a resolution of about 32×32 pixels, so as to sample at about twice the frequency of the smallest fringe spacing. This means that a low-resolution CCD camera will suffice. Higher resolution than this would yield a small increase in accuracy at the expense of a much longer FFT computation. FFT algorithms require the number of data points to be a power of two, so the next step up would be to 64×64 resolution (which would take 3–4 times longer to compute).

A further refinement is to use a Spatial Light Modulator (SLM) to produce an optical image of the autocorrelation plane. This can be recorded using a CCD array, allowing the particle displacement to be determined without calculating any FFTs (Jakobsen *et al.*, 1993).

Disadvantages with these methods are that the translation steps are slow. If the time interval between successive interrogations is 1 second, then a vector map with 4000 vectors will take over an hour to be processed. There are limits on how fast negative translation steps can be made.

4.4.5 Digital processing

In recent years, computing hardware has become considerably faster, meaning that the time saved by carrying out some FFT calculations optically is less significant. It is already possible to carry out a digital PIV analysis in 5–10 minutes and there is every prospect of computational speeds continuing to improve. Even if direct number-crunching speeds do begin to level off, it is possible to implement parallel-processing algorithms on multi-processor machines.

Whilst wet photography gives very good image resolution at a low price, digital CCD cameras are beginning to catch up. If it becomes possible to record by digital means as good an image as can be obtained using wet photography, then the time required to develop negatives in a darkroom can be saved. This amounts to about 40 minutes of time, by the time one has prepared solutions of chemicals at the required temperature, developed, rinsed, fixed and then washed the film, before hanging it up to dry. In contrast, a digital photograph can be taken and stored on a hard disk drive in a matter of seconds. In the future, most PIV recordings will probably be made digitally and optical analysis methods will become redundant. Ideally, it will be possible to obtain results whilst still in the laboratory, allowing experimental variables such as the time interval between illuminations and the seeding concentration to be adjusted on the spot.

The analysis method used here involves a digitised flow record. This was obtained by photographing each PIV record onto rolls of film and then digitising each negative using a scanner. Each exposure was made onto an area of 100 ASA film, which measured approximately 36 mm×24 mm.

The digital record is a computer file, consisting of about 3500 × 2400 data points. The photographic negative is therefore digitised at a resolution of about 100 pixels/mm, which is close to the grain size of the negative. Each entry in this array stores the grey-scale brightness of a very small region of the photographic negative, as a whole number in the range 0–255. Low grey-scale values mean that the negative was light due to no particle image occurring in that region, whereas high values mean that the negative was dark there (due to an image of a brightly-lit particle).

During the analysis, the digital record is processed as a series of small interrogation areas. In early work, these measured 64 × 64 pixels, with the centre of each interrogation region spaced 42 pixels away from the centre of its nearest neighbours. This method means that there is a certain amount of overlap between adjacent interrogation areas. These settings were found to work well in practice.

One factor which was only appreciated late on in this work is that the interrogation area spacing should be chosen carefully for each run. For example, a 3500×2400 bitmap recording, analysed using 64×64 interrogation areas spaced 42 pixels apart, gives a vector array of 82×56 vectors. However, the 34 columns of pixels nearest the right-hand edge of the bitmap (and the 26 rows nearest the bottom of the recording) will not be analysed, as a total of 42 pixels are required before a new column (or row) of vectors is started. The vector map will contain limited information about the velocities in the extreme right (or at the very bottom) of the flow.

If instead an interrogation area spacing of 44 pixels is chosen, the analysed area will measure 3496×2396 pixels, covering almost the whole of the recording (and producing 79×54 vectors). Alternatively, the computer could be allowed to vary some of the interrogation area spacings, *e.g.* it could choose a combination of 41 and 42 pixel separations which allows every pixel of the bitmap to be analysed.

4.4.6 Autocorrelation

An analogy to the process of digital autocorrelation is to imagine superimposing two copies of one interrogation region (such as that of Figure 4.6) and then moving one copy relative to the other, using only geometrical translations, *i.e.* avoiding rotation. (A couple of examples of this appear in the next few Figures, *e.g.* Figures 4.8 and 4.9.)

For each possible displacement between the two copies, the area of overlap of particle images can be determined. When this is plotted as a function of the autocorrelation displacement $(\Delta x, \Delta y)$ between the two copies, local maxima will be found in a number of places. The magnitude and location of these maxima can be related to the displacements of particle images in the interrogation area.

The digital analysis process involves calculating the autocorrelation function of small interrogation areas within a PIV recording. Let each such region be a

64×64 array of pixels. Each point within the array has Cartesian co-ordinates (x, y) and the image brightness at each point is stored as $I(x, y)$.

Autocorrelation $A(\Delta x, \Delta y)$ can be thought of as the multiplication of the grey-scale value at a point (x, y) by that at the point $(x + \Delta x, y + \Delta y)$, summed over all points in the array:

$$A(\Delta x, \Delta y) = \frac{1}{64^2} \sum_{x=1}^{64} \sum_{y=1}^{64} I(x, y)\, I(x + \Delta x, y + \Delta y) \tag{4.1}$$

However, something needs to be done if the co-ordinates $(x + \Delta x, y + \Delta y)$ lie outside the range of the interrogation array (*e.g.* $x + \Delta x = 100$ but the array is only 64×64). The intensity of these points can simply be taken to be zero; when the autocorrelation function is calculated using FFTs, this is equivalent to embedding the data in a 128×128 array of zeros ('zero-padding'). Otherwise, the values at $x + \Delta x = 100 - 64 = 36$ can be used instead ('no padding'). This method allows faster computations to be made but can introduce inaccuracies.

The method of zero-padding is *biased*, in that relatively-small displacements produce more non-zero terms in the summation than larger ones and hence yield a larger autocorrelation value. For example, with $\Delta y = 0$ in what follows: if $\Delta x = 14$ then the multiplications produce non-zero products for $1 \leq x \leq 50$ ($50 \times 64 = 3200$ terms); points such as $(51, 1)$ do not contribute to the summation, as they are mapped beyond the range of the array to points of zero brightness. The larger displacement $\Delta x = 39$ produces non-zero terms only if $1 \leq x \leq 25$ (*i.e.* $25 \times 64 = 1600$ terms, half as many as for the displacement $(14, 0)$). Consequently, a pattern in the interrogation region which produces correlations every 14 pixels might be given twice the significance of a pattern which correlates every 39 pixels, all things being equal. An alternative to zero-padding is to pad the data with an average intensity value.

These summations are evaluated for all possible displacements in the range satisfying $0 \leq \Delta x < 64$ and $-64 < \Delta y < 64$, giving rise to an autocorrelation function $A(\Delta x, \Delta y)$. Only two quadrants of this function need to be evaluated, because autocorrelation is an even function: $A(\Delta x, \Delta y) = A(-\Delta x, -\Delta y)$.

In a digital PIV record, the brightest points represent images of a seeding particle in the flow at one instant in time. These regions have the highest grey-scale values and should be a few pixels in image diameter (d_τ). Consider an idealised case such as that shown in Figure 4.6, where the background intensity in the digitised image has a grey-scale value of 1, the particle images are circular and the brightness of all pixels which make up a particle image is 255.

The autocorrelation function has a maximum value at $\Delta x = \Delta y = 0$, which represents the correlation of the interrogation region with itself. Here, the sum of grey-scale multiplications is greatest, because all the largest grey-scale values are multiplied together. For example, the summation will include many terms of magnitude 255×255 (*e.g.* all points of brightness 255 will self-correlate in this way). The magnitude of $A(0,0)$ will depend on the number of particle images captured in the interrogation region; each image will produce a strong self-correlation.

Approximating, the value of the autocorrelation function $A(\Delta x, \Delta y)$ depends on the number of pixels of brightness 255 which are correlated with other pixels of the same brightness, under the given displacement. This number depends on the area of overlap between particle images, which is greatest for self-correlation.

Suppose that no particle images have a centre-to-centre separation of less than two particle diameters ($2d_\tau$) from each other (as was the case in Figure 4.6). A very short displacement away from self-correlation, *e.g.* at $(\Delta x = 1, \Delta y = 0)$, the autocorrelation will be reduced, because some of the brightest points will now be correlated with dimmer background points, giving rise to products of magnitude 255×1. This will happen first to points close to the edge of a particle image, where there are large gradients of grey-scale brightness (Figure 4.7(a)). As the magnitude of the autocorrelation displacement $(\Delta x, \Delta y)$ tends towards d_τ, the autocorrelation function will decrease further still, because there will be fewer and fewer of the 255×255 terms which dominate the summation.

Figure 4.7(b) shows a much reduced overlap. Eventually a point is reached where there is no image overlap (Figure 4.7(c)) and the autocorrelation function

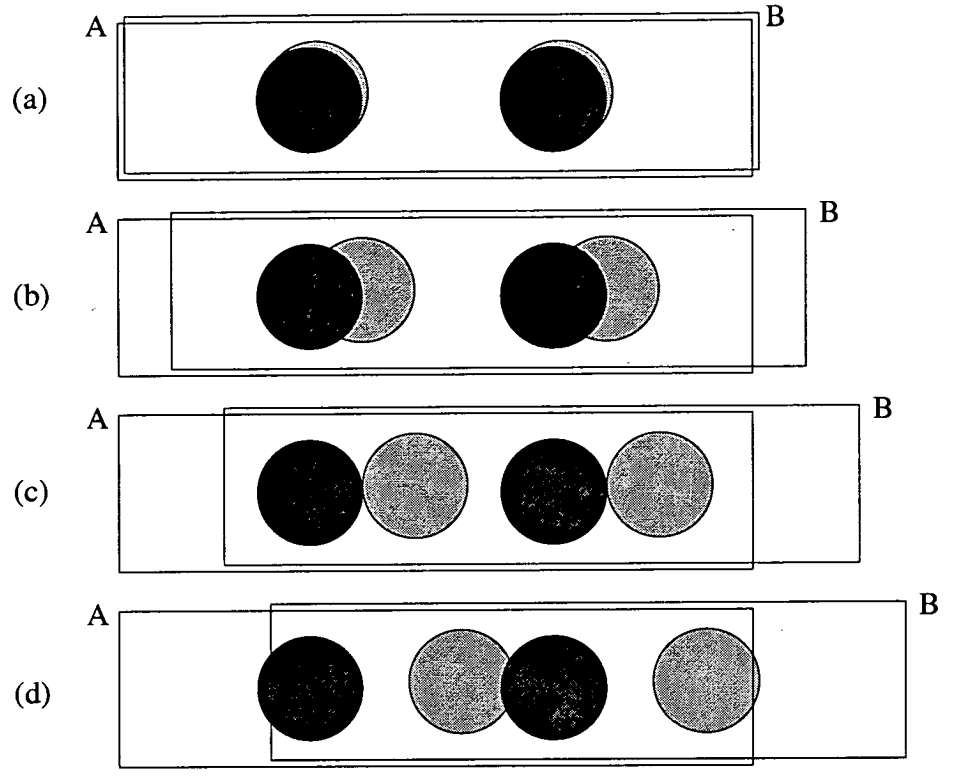


Figure 4.7: How the extent of particle image overlap changes as a function of increasing autocorrelation displacement.

reaches a minimum value. Increasing the displacement to values larger than d_τ can give rise to new overlaps, between different particle images (Figure 4.7(d)); this causes the autocorrelation function to increase again.

The area of particle image overlap is greatest for self-correlation and there are likely to be autocorrelation minima in regions where there is no overlap. For example, if no two particle images have centres within two diameters of each other, then $A(\Delta x, \Delta y)$ will be very much lower than $A(0,0)$ at values satisfying $\Delta x^2 + \Delta y^2 = d_\tau^2$ (Figures 4.7(a) and (c)).

The exact area of overlap of two circles of radius r and separation $0 < h < 2r$ can be calculated using circle geometry. This area consists of two back-to-back segments of a circle (*cf* Figure 4.7(b)). Hind & Christy (1992) stated that:

$$\text{Area}(h) = 2\left[r^2 \cos^{-1}\left(\frac{h}{2r}\right) - \frac{h}{2}\sqrt{r^2 - \left(\frac{h}{2}\right)^2}\right].$$

If the above assumptions of particle separation and image intensity hold, then the separation $h = \sqrt{\Delta x^2 + \Delta y^2}$ and the autocorrelation $A(\Delta x, \Delta y) \propto \text{Area}(h)$ as given above, for displacements such that $0 < h < 2r$.

If particle images are closer than this, then additional overlaps need to be considered. In terms of the paper/transparency analogy, suppose that the lighter image of Figure 4.7(b) is on the transparency. If the separation of particle images is between d_τ and $2d_\tau$, then the overlap between the transparency image and that on the paper will still diminish between Figures 4.7(a) and (b). In some directions, however, the image separation shown in Figure 4.7(b) may be sufficient to cause an overlap to start with a different particle image, along the lines of Figure 4.7(d). Only if no particles have a separation of less than $2d_\tau$ will the idealised autocorrelation fall to the greatest possible extent.

4.4.7 Maxima of the autocorrelation function

Considering Figure 4.6 again: this PIV record contains 1 track of three particles and 2 tracks which are four particles long. This means that the self-correlation $A(0,0)$ is based on 11 particle image overlaps (Figure 4.8).

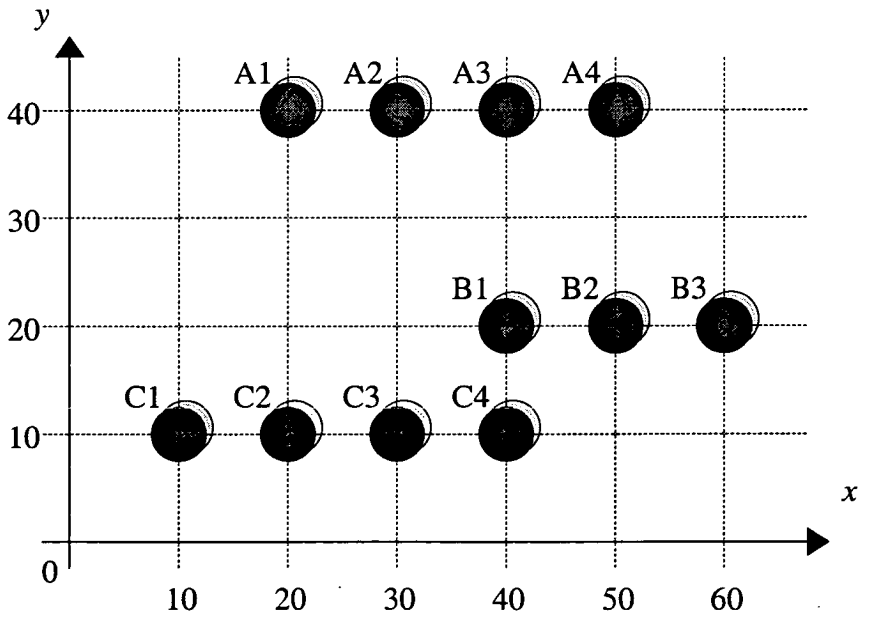


Figure 4.8: The self-correlation of Figure 4.6.

It can be seen that there are many different translations $(\Delta x, \Delta y)$ available which map one track of particle images onto another track of images. For example, a displacement $(10, 30)$ will map image C1 onto image A1, C2 onto A2 *etc.*, providing four particle image overlaps.

This displacement from track C onto track A could produce more image overlaps than this in theory, as the same displacement might also happen to map a third track of images onto a fourth. This, however, is unlikely to happen if the seeding particles are randomly distributed throughout the flow and are not attracted towards one another, *e.g.* by electrostatic forces.

It can also be seen that the equal and opposite displacement $(-10, -30)$ maps A1 onto C1, A2 onto C2, *etc*, achieving an equivalent result: $A(10, 30) = A(-10, -30)$. In general, if all correlations $A(\Delta x, \Delta y)$ satisfying $\Delta x \geq 0$ are evaluated, then the correlations involving $\Delta x < 0$ need not be calculated.

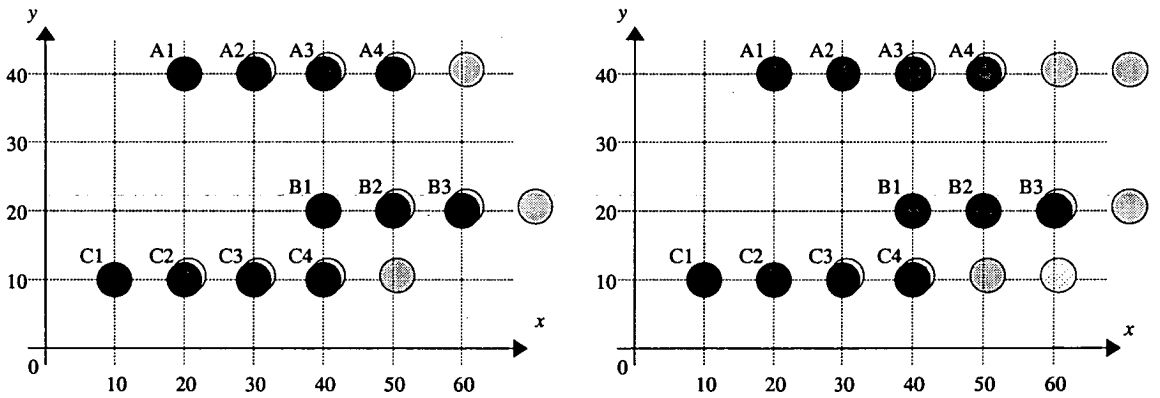
The autocorrelation $A(10, 30)$ will be a local maximum, with a magnitude based on four image overlaps (36% of $A(0, 0)$). There will be a number of other local maxima of similar size, based on mappings between different particle tracks, *e.g.* the displacements $(20, 30)$ and $(20, -20)$ will both produce three image overlaps.

For a flow to be suitable for PIV analysis, the time of the photographic exposure should be short compared to the time-scale over which the flow velocities change at each point in space. In other words, each particle within the flow undergoes a reasonably constant displacement between successive illuminations. Further, PIV works best if there are only small local velocity gradients, so that all the particles within each interrogation region have very similar velocities.

Aside from the case of self-correlation, the largest correlation peak will occur at a displacement (X, Y) which maps the first image of any given track of particle images onto the second image of that same track. Because this particle's velocity is constant, this displacement will also map the second image of this same particle track onto the third, the third onto the fourth, *etc.* Further, if the velocity gradients are low, then this displacement will not only map image A1 onto A2,

image A2 onto A3 *etc.*, but also produce corresponding particle image overlaps from every other track within the interrogation region. If these assumptions are not quite met, then the optimum displacement (X, Y) will produce some partial overlaps; there will still be an autocorrelation peak here, but it will be somewhat 'squashed', having a reduced amplitude and a greater width. Hind & Christy (1996) describe an iterative PIV analysis method which can compensate for this: knowledge of the flow field from one analysis iteration allows velocity gradients to be calculated, which are used to deform the data within each interrogation region so as to obtain a sharper velocity peak during the next iteration.

It can be seen that, when applied to Figure 4.6, the displacement $(10, 0)$ will produce two particle image overlaps for tracks containing three images, and three overlaps for tracks of four. This correlation will therefore be based on $3+2+3 = 8$ overlaps (Figure 4.9(a)). For comparison, $A(10, 0)$ can be expected to be about 73% of $A(0, 0)$ and about twice as large as the correlations between different particle image tracks.



(a) the main velocity peak (b) the first harmonic peak

Figure 4.9: Image overlaps at (a) the fundamental velocity peak and (b) the second harmonic peak. It can be seen that these produce eight image overlaps and five image overlaps, respectively.

Further, the related displacement $A(2X, 2Y)$ will map the first image within each track onto the third, the second onto the fourth, and so on. In Figure 4.6, this correlation will be based on $2 + 1 + 2 = 5$ overlaps (Figure 4.9(b)). In this case, $A(20, 0)$ is only about 45% of $A(0, 0)$. Even so, it is still larger than the correlations between different particle tracks.

Similarly, higher order harmonics can be found at $A(3X, 3Y)$, *etc.* The number of harmonics and the relative magnitude of the peaks depend on the number of images of each particle within each track and the displacement between successive images. In particular, only the first displacement peak (the fundamental peak) can be obtained if only two images of each particle are found in each interrogation area, which will happen if the image displacement is large or if only two images of each particle are actually recorded.

The size of the displacement correlation peak also depends on the number of different particle tracks in the interrogation area. For example, suppose that an extra particle track of four images was added to Figure 4.6. This would increase the self-correlation peak to a value based on 15 image overlaps; the velocity peak would now be based on 11 image overlaps, the second harmonic peak on 7 overlaps, *etc.* The correlations between different particle tracks, on the other hand, would probably still have maxima based on 4 image overlaps, if the new particle track was added at a random position. Thus, adding an extra particle track serves to increase the desired 'signal' (by enlarging the velocity peak) whilst being unlikely to worsen the 'noise'; there will be more noise peaks at random locations, but these will have the same magnitude as before.

An analysis of Figure 4.6 has been performed by hand. A two-dimensional plot of the resultant autocorrelation surface is shown in Figure 4.10, along with a section taken through the main displacement peaks, parallel to the Δx-axis. If autocorrelation is performed on discrete data points, the position of each displacement peak can be determined to sub-pixel accuracy, by using interpolation methods to construct a continuous best-fit surface to cover the discrete data points.

In summary, the autocorrelation function should have a strong maximum at the particular displacement (X, Y) which corresponds to the movement of each seeding particle between consecutive illuminations. Further, there will be smaller harmonic maxima at the points $(2X, 2Y)$, $(3X, 3Y)$, *etc.*, if three or more pulses are recorded by the camera. For the case of Figure 4.6, the autocorrelation function $A(\Delta x, \Delta y)$ was shown in Figure 4.10. The mean flow displacement can be

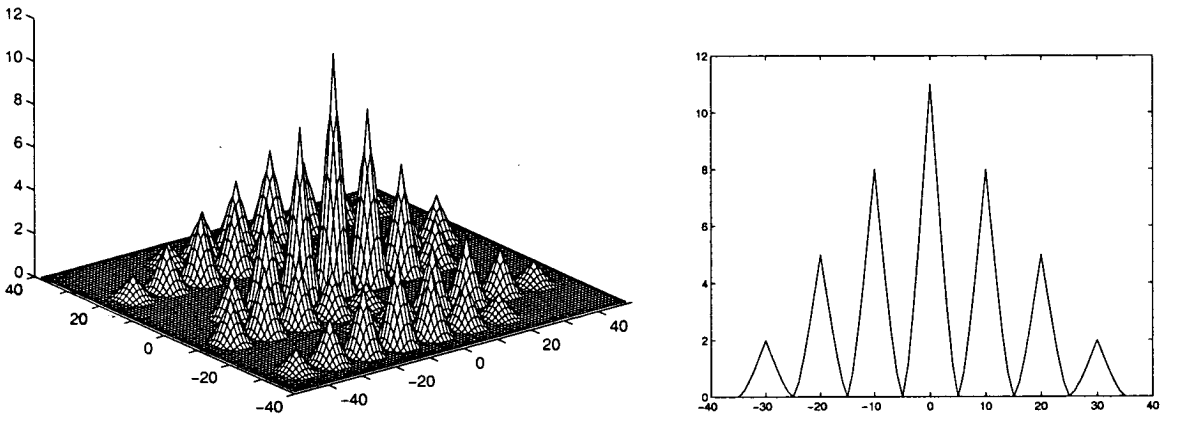


Figure 4.10: The complete autocorrelation surface from the analysis of Figure 4.6 and a cross-section of this surface along $\Delta x = 0$ (from bottom left to top right).

seen to be either $(10, 0)$ or $(-10, 0)$, based on the positions of the highest peaks outside the central self-correlation zone. The two symmetrically-located peaks are a result of the directional ambiguity of PIV: the data recorded in Figure 4.6 does not include any information as to whether the particles were moving to the left or to the right.

The higher harmonic peaks can be used to validate the displacement measurement, if more than two images of each particle appear in each interrogation area. After determining the position of the maximum value of $A(\Delta x, \Delta y)$ (ignoring the peak at $(0, 0)$, which is due to self-correlation), the values of $A(\Delta x, \Delta y)$ can then be examined in the region around the point $(2X, 2Y)$. If a well-defined second harmonic peak is found, then the point (X, Y) can be considered to be valid; if not, $A(\Delta x, \Delta y)$ might perhaps be re-examined in order to search for an alternative peak. However, because of its greater amplitude, the position of the peak at (X, Y) can be determined more accurately than the position of any higher-order harmonic peaks, as the latter are more influenced by background noise. As a result, it is not usually wise to take the position of the displacement peak to be half of the displacement at the second harmonic.

Obviously the position (X, Y) of the autocorrelation peak can be related to the velocity in the real flow. If the real flow has a local velocity $\boldsymbol{v} = (v_x, v_y)$, the

test section is photographed at an image:object magnification M : 1 and the time between successive illuminations is Δt, then:

$$(X, Y) = (M\, v_x\, \Delta t,\ M\, v_y\, \Delta t)$$

4.4.8 Windowing

If 48 × 48 interrogation areas are required, it is necessary to extend these digital samples to the nearest higher power-of-two side length before Fourier transforms can be used. This can be done using partial zero-padding, by adding 16 rows and columns of zeros to the 48 × 48 data from the sample.

An alternative is to use 'windowing', in which the sampled data is scaled before autocorrelation is carried out. For example, 64 × 64 samples of data can be multiplied by a windowing function, which is 1 for the centre-most 48 × 48 pixels but falls to zero near the edges. In this scheme, the autocorrelations will be dominated by the central pixels, but some useful, extra correlations can also be made between particle images near the centre and others near the edges (which would not be included in 48 × 48 interrogation areas). These correlations will be weak, because pixels near the edges have their brightness artificially reduced by the windowing function. Thus the intended correlations on 48 × 48 pixel interrogation areas are made as normal and are boosted by additional (weak) correlations to images just outside the region of interest (which would not be obtained using zero-padding). Compared to results using 64 × 64 interrogation areas without windowing, velocity gradients cause fewer problems and edge effects are reduced (*e.g.* when half-particle images at the edge of an interrogation area produce abnormal correlations with full-particle images).

Various windowing functions can be tried. Square-based functions are an obvious choice, because digital interrogation areas are square. Another possibility is a two-dimensional Gaussian function centred on the 'point' (32.5, 32.5) in a 64 × 64 interrogation area. In this case, autocorrelation analysis will mimic the results

obtained when analysing a PIV negative using an optical method (based on a low-power laser beam, with an axi-symmetric Gaussian intensity profile).

4.5 Setting PIV parameters

There are many experimental parameters which need to be considered if a good PIV photograph is to be taken. For example, the magnification M determines how much of the flow is recorded onto each negative; the light sheet has an intensity controlled by the laser and a pulsing rate controlled by the rotating mirror; a camera has many adjustable parameters, such as exposure duration and aperture (f-number), *etc.*

4.5.1 Laser power

Seeding particles should be illuminated with an intensity which is much brighter than the background against which they are photographed. It is often desirable to use as much laser power as possible, because this gives more freedom when considering other PIV parameters, such as film speed, camera aperture number, *etc.* Further, an increased illumination intensity allows smaller seeding particles to be used (or else the same seeding particles can be used in larger-scale apparatus).

If a CW laser can be used, the laser power should usually be made as high as possible. However, photographic film has a 'gross fog level', which is the maximum brightness which can be recorded. This is due to the nature of the film. There is a finite concentration of silver halide crystals on the surface of the film; once all the silver ions have reacted on exposure to light, further darkening is not possible.

Adrian & Yao (1985) recommend that 45% of each particle image should be saturated, in order to minimize the effect of 'noise': if the illumination is less intense then the autocorrelation peaks become harder to detect; if the illumination

is too intense then the intensity of the particle images is still limited by the gross fog level on the film, whilst the background regions become brighter.

If a CW laser cannot deliver the required light intensity, then a pulsed laser will have to be used (with an expanded-beam method).

4.5.2 The choice of film

Photographic film is usually '35 mm'; this provides an area on the negative measuring 36 mm×24 mm, onto which each photograph is taken. The negative scanner used in this work could only accept this standard size of film. Other film sizes exist but these need less widely-available cameras, *e.g.* Hasselblad cameras take exposures onto an area 50 mm×50 mm.

Black & White Kodak TMAX film was used in this work, in the 35 mm size. Versions of this film have a number of different sensitivities to light, ranging from 'slow' 25 ASA film up to 'fast' 3200 ASA versions. Faster film is more sensitive to the available light but has a larger grain size. In most of this research, 100 ASA film was used, although some photographs were taken to test 400 ASA film.

In the future, it will be more attractive to take PIV photographs using digital CCD cameras, particularly if the subsequent analysis is to be performed digitally. Currently conventional film offers better resolution at a much lower price.

4.5.3 Magnification

The dimensions of the experimental apparatus, the film and the light sheet may suggest convenient values of magnification. For example, two camera arrangements suggest themselves when using a 0.5 m-high light sheet (propagating horizontally) to illuminate the flow in a fluidised bed 0.3 m× 0.3 m in cross-section. In each case, the image:object magnification can be calculated from the ratio of a dimension of the film to the corresponding dimension of the apparatus.

'Portrait' mode involves rotating the camera so that the 36 mm×24 mm film has its longer dimension arranged vertically. If all of the light sheet is to be photographed, then the image:object magnification $M = 36\,\text{mm}/0.5\,\text{m} = 0.072$ (based on vertical distances). Horizontally, the width photographed is $24\,\text{mm}/0.072 = 0.33\,\text{m}$, just 10% larger that the width of the flow section. This arrangement is sensible, because most of the film area is used to record a large area of flow.

'Landscape' mode has the camera arranged so that the film's longer dimension is horizontal. Based on horizontal distances, the magnification $M = 36\,\text{mm}/0.3\,\text{m} = 0.12$ if the full width of the test section is to be photographed. In this arrangement, the vertical distance photographed will be $24\,\text{mm}/0.12 = 0.2\,\text{m}$, *i.e.* only 2/5 of the light sheet will be used. This method records a closer-up view of the object, but is inefficient if the light sheet comes from a 0.5 m scanning-beam box. A 0.2-0.25 m light sheet would waste less energy.

The above arrangements offer a range of magnifications. The landscape mode was selected in most experiments, photographing slightly more than the full width of the illuminated flow field. This gave a relatively high magnification, which had a number of advantages: less illuminating light was required, as the camera could be closer to the test section; larger particle images were recorded, at better resolution; the light sheet had a relatively even intensity in the region of interest.

Photographing slightly more than the full width of the apparatus had an advantage. Light was scattered when the laser beam passed through the two side glass walls of the test section, giving rise to bright lines on the photographic negatives. As the separation of these walls was known, the magnification M could easily be calculated for each picture, by relating this distance in the object plane to the separation recorded in the image plane.

4.5.4 Image size

Given an image:object magnification M, the diameter of particle images on the negative (d_τ) can be calculated using the formula given by Adrian & Yao (1985):

$$d_\tau = \sqrt{(M\,d_p)^2 + (2.44\,(1+M)\,f_\sharp\,\lambda)^2 + d_g^2} \qquad (4.2)$$

Here, d_p is the diameter of the particles being photographed in the object plane (*i.e.* the true size of the seeding particles), $f_\sharp$ is the f-number of the camera (its aperture size), λ is the wavelength of the illuminating light from the laser and d_g is the average grain size of the light-sensitive emulsion on the photographic film. The three terms within the square-root represent, in order, the geometric image size, the diffraction-limited spot size and the grain size of the photographic film.

It can be seen that the image diameter increases if the f-number is increased. However, this may or may not make a significant difference to d_τ, according to the values which the other parameters happen to take. As far as the aperture setting is concerned: a larger f-number means that more laser light is needed (as the camera's aperture is smaller) and the image size will increase; however, the camera's depth of field increases, so it becomes easier to record focused images.

Prasad *et al.* (1992) recommended that $d_\tau \sim 4\,d_g$. This compromise size avoids quantisation errors which would occur if $d_\tau \sim d_g$ (it being harder to locate the centre of the image) whilst limiting random errors and keeping the magnitude of the smallest resolvable velocity relatively low (see later). The optimum image size is investigated in the next chapter, using simulated PIV recordings.

The above formula assumes that the particles do not move significantly during the time they are illuminated. This is usually the case, *e.g.* if the velocities to be measured are about $1\,\mathrm{m\,s^{-1}}$, then a 0.5 m light sheet repeating scans every $5 \times 10^{-3}\,\mathrm{s}$ has an average scanning velocity of $100\,\mathrm{m\,s^{-1}}$. The beam's scanning velocity is therefore much larger than the velocities of the seeding particles.

4.5.5 Illumination time interval

When the time Δt between successive illumination pulses is chosen, it is necessary to bear in mind the way in which the analysis method determines the displacement.

The particle image displacement $\boldsymbol{d}$ on the film gives rise to a peak in the autocorrelation function, which should be sufficiently far away from the large self-correlation peak at $\boldsymbol{d} = (0, 0)$ for there to be no significant interference. Image separations in the range d_τ to $2d_\tau$ cause overlaps to occur between the intended displacement peak and the tail of the self-correlation peak. The addition of these signals causes the measured displacement to be found somewhat closer to the origin than it should be, giving rise to a biased measurement.

Various authors have suggested values of how many particle diameters the smallest displacement on the flow record ($|\boldsymbol{d}|_{min}$) should be. In the formula:

$$|\boldsymbol{d}|_{min} = M\,|\boldsymbol{v}|_{min}\,\Delta t > n.d_\tau,$$

Keane and Adrian (1990, 1991) suggest $n = 1.5$ whereas McCluskey *et al.* (1994) propose, more cautiously, that $n = 2$. The latter value represents the situation shown earlier in Figure 4.7(c): this displacement means that there should be no interference in the velocity peak from the self-correlation peak.

Simulations, presented in the following chapter, suggest that $n = 1.5$ is acceptable if the image diameter is 4 pixels ($|\boldsymbol{d}|_{min} = 5.7$ pixels) whereas the more restrictive $n = 2$ is better if particle images are only 2 pixels in diameter ($|\boldsymbol{d}|_{min} = 3.7$ pixels). In fact, a better correlation of the minimum resolvable displacement will be shown to be:

$$|\boldsymbol{d}|_{min} > 1.1\,d_\tau + 1.6\,\text{pixels}.$$

One thing to note is that these suggestions are guidelines, not hard and fast rules. Measured displacements can be obtained in the range $d_\tau < |\boldsymbol{d}| < 2\,d_\tau$,

although the experimental errors will be larger than usual. These measurements can still be useful in qualitative work, however, because they do indicate regions of low velocity.

Similarly, if the mean particle image displacement is too large compared to the chosen size of interrogation area, then the quality of the autocorrelation peaks is affected in two ways.

First of all, the strength of the autocorrelation function $A(\Delta x, \Delta y)$ decreases at larger displacements, because fewer terms contribute to the summation. This gives a biased autocorrelation, which favours peaks at smaller displacements.

Secondly, as the displacement $\boldsymbol{d}$ increases, fewer images from each particle track can fit into each interrogation area (size: d_I). For example, consider a 64 × 64 interrogation area which covers a number of particle tracks, each six particle images long. If $|\boldsymbol{d}| = 10$ pixels, then a complete particle track can be contained within the interrogation area. If however $|\boldsymbol{d}| = 40$ pixels, then only one or two images from each track can possibly lie within the interrogation region. Larger displacements cause a drop in the number of particle image pairs, which reduces the size of the displacement peaks and so makes them more susceptible to errors due to autocorrelation noise. These errors become apparent as an increase in the spread of valid measurements around the true value or by an alternative peak being found (leading to a wildly inaccurate measurement). The causes of pair loss will be examined in more detail in the next two chapters.

Keane & Adrian (1990, 1991) recommended that:

$$|\boldsymbol{d}|_{max} = M\,|\boldsymbol{v}|_{max}\,\Delta t < 0.25\,d_I.$$

Once again, this is only a guideline: displacements larger than the above *can* be measured accurately. The problem is that the yield of valid vectors decreases as the image separation increases, *e.g.* 99% of the interrogations might well give rise to a valid vector if the images are 20 pixels apart, but this figure might drop to about 80% in regions where images are about 30 pixels apart. The results

presented in the next chapter suggest that reasonable results can be obtained for $|\boldsymbol{d}| < 0.5\, d_I$, suggesting that the upper velocity bound is somewhat larger than the cautious recommendation given by Keane & Adrian (1990, 1991).

The type of expression given by Keane & Adrian (1990, 1991) for the maximum displacement of particle images makes sense for analysis systems which perform optical interrogations of the PIV negative. In this case, the interrogating laser beam has a circular cross-section, so the displacement recommendations should also be axi-symmetric. In a digital analysis using square interrogation areas and a given magnitude of image displacement, the best yield of valid vectors is obtained if the direction of these vectors is parallel to two of the four interrogation area sides. The worst results are obtained if the displacement is diagonal. A recommendation that, say, $|\boldsymbol{d}| < 0.45\, d_I$ for square-sided interrogation areas might be good for parallel displacements but over-optimistic for diagonal displacements.

One absolute limit does arise if autocorrelation analysis takes place with un-padded FFTs. If this is the case, a 64×64 interrogation area will give rise to a 64×64 autocorrelation plane. In order to be detected, a displacement peak must lie within this plane, *i.e.* the two components of the measured displacements (parallel to adjacent sides of the interrogation areas) must both be between –32 and +32 pixels. If the analysis takes place with fully zero-padded FFTs, then the autocorrelation plane is 128×128, the maximum detectable displacement is twice as large and this sort of restriction need not be limiting.

In practice, the magnitude of displacement vectors in a flow can be described by a distribution function.[1] Figure 4.11 contains data from a PIV record obtained during the course of this research. This particular velocity size distribution is approximately Gaussian, but bi- and tri-modal distributions have also been found. This pattern is similar to that found by Khan & Ahmed (1995), who found mono-, bi- and tri-modal velocity histograms in different regions of the flow field under investigation (air flowing past a wing-body joint).

[1] Histograms of this type are useful for vector validation, as they suggest that all vectors above a certain magnitude can be deleted — this is very easy to do.

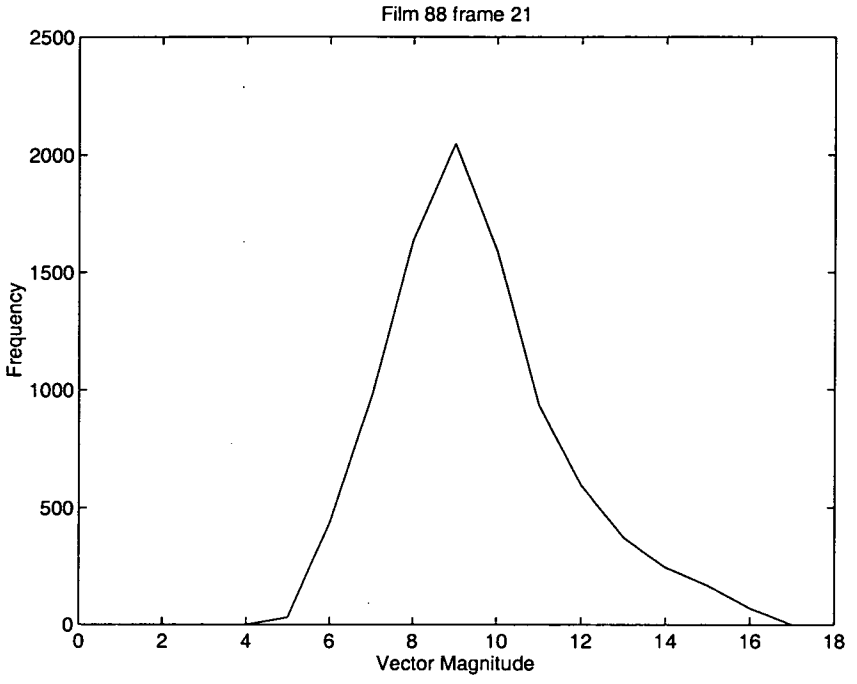


Figure 4.11: The size distribution of displacement magnitudes (in pixels per illumination interval) from a manually-validated PIV record. The array of displacement vectors was 85 × 55, hence a total of 4675 vectors were generated by the analysis. 4521 valid vectors were found (97% of the total).

Once a half-decent PIV picture has been taken, the subsequent analysis can suggest improvements to the technique, to be implemented in later experiments. In this case, the idea of using a longer time interval between PIV illuminations is attractive, in order to stretch out the displacement distribution function. This would make the largest displacements lie closer to the theoretical maximum and would also provide more accuracy at lower velocities. Thinking in terms of a displacement distribution function, an experimenter might be prepared to accept some inaccuracies (or reduced yields) in the fastest 1% of all measurements by increasing Δt, if this means that the remaining vectors can be determined with greater accuracy. However, a longer Δt might require a longer camera exposure time, making the flow more likely to change during the course of the exposure.

For comparison, artificial vector maps have been created for laminar flow and turbulent flow in a circular pipe. The laminar velocity profile is quadratic, whereas the turbulent case involves a 1/7th power law (Figure 4.12). In practice, random fluctuations will cause the turbulent flow result to become more Gaussian.

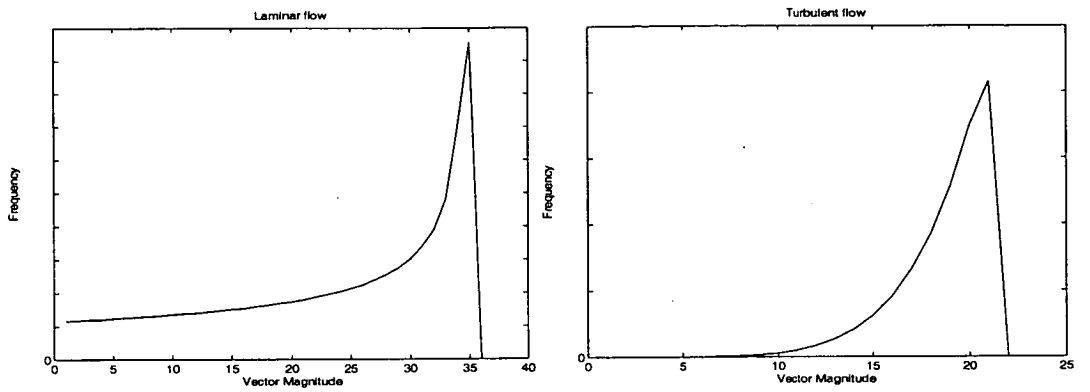


Figure 4.12: The size distribution of displacement magnitudes expected in a PIV measurement of laminar and turbulent flow. The mean flow displacement is 17.5 units in magnitude, in both cases.

4.5.6 The exposure time

For a PIV system, the illumination time is determined by the above constraints. The camera exposure time only affects the number of images of each particle on the film. Generally, about 2–6 images should be recorded.

There is a factor of two between the lower and upper ends of this range, because the exposure time settings on most cameras vary in factors of two. The exposure times most likely to be used in PIV work are 1/125, 1/60, 1/30 and 1/15 of a second. An element of choice may still be present. For example, suppose that the recommended illumination time interval is 10.0 ms. A good PIV photograph will use the 1/30 s camera setting (giving about 3 or 4 images of each particle, and using a relatively high concentration of seeding particles) or possibly 1/15 s (6 or 7 images, using a lower seeding concentration) if the flow does not change too much as a function of time. Shorter exposure times such as 1/60 s would definitely not be suitable, because some regions of the flow would produce fewer than two images of each particle during the $\sim$ 17 ms exposure, and so could not possibly yield a flow velocity; the actual exposure time must be at least twice the illumination time interval. An eighth of a second is probably too long an exposure.

An advantage of recording at least three images is that the digital autocorrelation analysis is better than in the minimum case of two images. For example, 9

double-pulsed images in a given interrogation area mean that the self-correlation peak is based on 18 particle images and the displacement peak is based on 9 image pairs. If instead six different seeding particles are each imaged three times, the self-correlation peak will still be based on 18 images but the displacement peak will now be based on 12 image pairs. In the latter case, the main displacement peak is 33% larger than in the double-pulsed example and so should be less vulnerable to random errors; also, a harmonic peak will be present, which might possibly be used to validate the result. These considerations are particularly important at high image displacements, as will be shown in the next chapter.

As the photographic exposure duration is made longer and longer, however, the assumption that a PIV photograph is an instantaneous record of the flow may become more suspect. In some cases the flow is time-invariant and so this does not matter, *e.g.* in laminar flow around a cylinder or through an orifice, where velocity vectors are independent of time at each point in the flow. In other cases, *e.g.* where the flow velocities display random fluctuations due to turbulence, it can be important to keep the recording time short. The general guideline seems to be that no more than six images should be recorded of each particle.

One point to note is that the exposure duration has no effect on the brightness of each individual image, assuming that the seeding particles scatter incident light (as opposed to emitting light through fluorescence). In normal photography, the intensity increases as the exposure time increases. In PIV, however, each particle is illuminated at each position for a time which is much shorter than the duration of the exposure. If illumination comes from a pulsed laser which produces two pulses only, then the only requirement is for the exposure to be long enough to capture both images. If illumination comes from a continually-pulsing source (*e.g.* a scanning beam box), then negatives taken at longer exposure times *will* appear darker because they contain *more* particle images, but the individual images will still have the same brightness.

The number of images taken cannot usually be predicted exactly. In the example above, the illumination time interval was exactly 10.0 ms and the best PIV

photograph used the 1/30 s camera setting. This exposure time is not guaranteed, however, so the exposure may turn out to last 33 ms plus or minus several milliseconds. During this time, the PIV system will go through somewhat more than three complete scan cycles. For some time $t = 0$ ms at the start of the exposure, some regions of the flow may be illuminated at $t = 1, 11, 21$ and 31 ms (*i.e.* 4 times) whereas others may be illuminated at $t = 8, 18$ and 28 ms (*i.e.* only 3 times). If the maximum and minimum displacement recommendations are treated only as guidelines, the possibility of changing the illumination time interval might be considered. At one point in this work, it was possible to calculate that the actual camera exposure time was about 31 ms. As a result, the illumination time was set to 10.3 ms, so as to record exactly three images of every particle.

4.5.7 Seeding concentration

Keane & Adrian (1990, 1991) recommended that at least 15 particle image pairs should be recorded in each interrogation area. As discussed earlier, the number of image pairs found by autocorrelation analysis depends on the number of particles in each interrogation region and the number of images of each particle. If there are n particles within the illumination plane and m illuminations of each particle, then about $n \times (m - 1)$ particle image pairs will be found (ignoring the complications caused by particles moving into or out of the interrogation area).

If m is relatively large, then fewer seeding particles need to be present. Having a relatively low seeding concentration can be an advantage, because problems of light losses due to scattering by seeding particles (*e.g.* attenuation governed by Beer's Law and fogging) are minimized.

Further, a basic principle of seeding is that the addition of a small concentration of fine particles will not affect the flow being measured; obviously this is more likely to be the case at lower seeding concentrations.

However, the distribution of seeding particles within the flow will follow a Poisson probability distribution (Westerweel, 1994). The average seeding concen-

tration should be somewhat greater than the absolute minimum, so that there is a high probability that the local concentration is adequate at all points in the flow.

4.5.8 The dynamic range of a PIV system

As recommended by Prasad *et al.* (1992), each particle image should be about 4 pixels in diameter. The minimum image displacement, according to Keane & Adrian (1990, 1991), should be about 1.5 times greater than this, *i.e.* 6 pixels. If interrogation areas measure no more than 64 × 64 pixels because of velocity gradients, then the maximum resolvable displacement could be about 30 pixels.

This means that the magnitude of the largest resolvable displacement ($|\boldsymbol{d}|_{max}$) will be about five times the magnitude of the lowest ($|\boldsymbol{d}|_{min}$). This ratio, known as the *dynamic range*, is likely to be constant for a given PIV system.

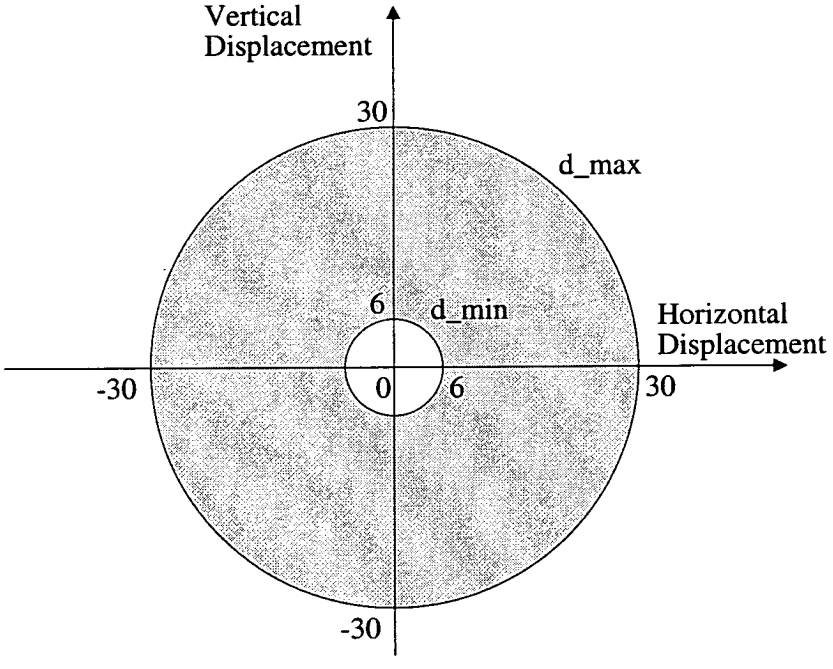


Figure 4.13: The maximum and minimum displacements of a particle image: displacements which lie within the shaded region should yield a valid measurement.

This can be represented diagrammatically (Figure 4.13). Valid velocity measurements lie inside the maximum displacement circle (radius 30 pixels) but outside the minimum displacement circle (radius 6 pixels). This region is shaded.

4.6 Image-shifting

4.6.1 Problems of vector alignment and dynamic range

Given a magnification M and a PIV illumination time interval Δt, the locus of flow displacements can be obtained (*e.g.* by means of a preliminary set of experiments). Suppose that one such test used an illumination time $\Delta t = 10\,\text{ms}$ and yielded displacement vectors which lay within a circle of radius 10 pixels, centred on the point $\boldsymbol{d} = (10, 15)$ — Figure 4.14, dark circle A. Suppose that a different flow, also using $\Delta t = 10\,\text{ms}$, gave rise to another circular locus, of radius 12 pixels and centred at $(10, -10)$ — Figure 4.14, dark circle B.

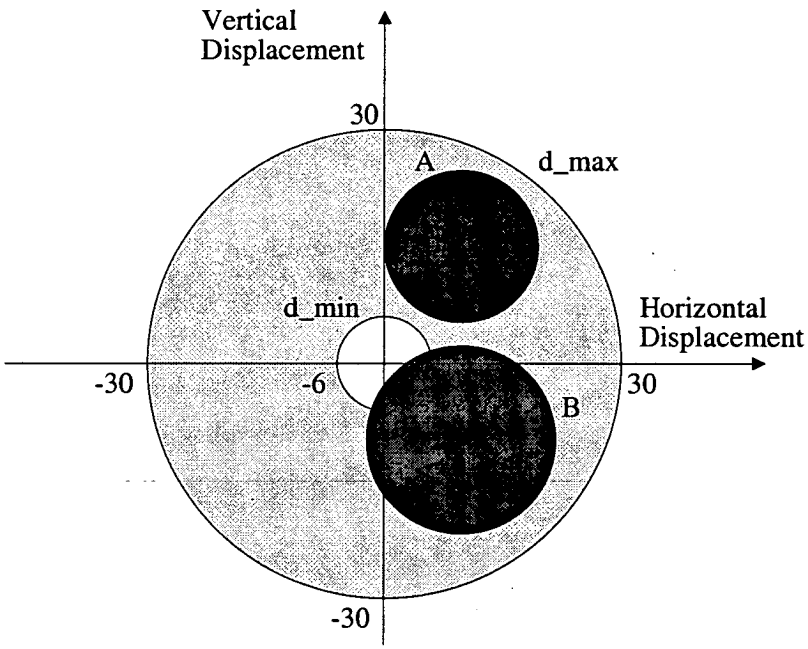


Figure 4.14: The locus of image displacements from two hypothetical experiments. Circle A is well placed, so a good yield of accurate displacement vectors should be obtained. The position of Circle B suggests that there may be regions of low flow velocity, which might not yield satisfactory measurements.

For circle A, everything is fine. The ratio of the largest image displacement to the smallest is not too large, so the circle sits happily in the acceptable region between $|\boldsymbol{d}|_{max}$ and $|\boldsymbol{d}|_{min}$. As the circle lies within one quadrant of the displacement diagram shown in Figure 4.14, a coherent picture should emerge if the displacement vectors produced by the analysis are all made to point upwards or to the right (an easy re-alignment).

A very minor improvement might be made if the illumination time interval was increased to about 10.5 ms. This would move the centre of circle A to the point $(10.5, 15.75)$ and would increase its radius to 10.5 pixels. The circle would still lie just inside the outer $|\boldsymbol{d}|_{max} = 30$ pixels circle, so all the vectors should still be resolvable. The increased illumination time interval should yield a little extra accuracy at the more commonly-occurring displacements ($|\boldsymbol{d}| \sim 20$ pixels). Either way, the directional ambiguity of PIV is not a major problem in this situation.

In the case of circle B, however, there is a problem in that some of this circle overlaps the inner circle ($|\boldsymbol{d}|_{min} = 6$ pixels). This means that some displacement vectors will be too small to be detected by PIV analysis methods. Increasing the illumination time interval might be considered, but this can only be increased by about 20% (*i.e.* to 12 ms) before problems will occur at the highest displacements.

Another issue raised by circle B is that whilst most of the vectors are downward and to the right, some vectors are upward and others point to the left. Bulk alignment to make all vectors point downwards (or to the right) will leave some vectors incorrectly aligned. These vectors will have to be re-aligned during the vector validation process, either individually or by zooming in on a suitable region of the flow so as to align only the vectors within that region. This realignment might be fairly straightforward, but it could well be a matter of guess work.

The location of circle B could be worse. If the circle had included the point $(0, 0)$, then no possible adjustments to the illumination time interval could allow the slowest velocity vectors to be resolved. Furthermore, such a flow might include regions where the true direction of the flow is difficult to determine, *e.g.* if the flow was highly turbulent. Even if the slowest displacements could be resolved, it might prove difficult to align these vectors with the rest of the flow field. Image-shifting attempts to solve both of these problems.

4.6.2 The use of image-shifting to enhance PIV

The basic idea of image-shifting is to superimpose a known velocity onto the velocities within the flow field, so as:

1. to make all of the resultant vectors lie within only one or two quadrants of the displacement graph, and
2. to reduce the dynamic range of the flow displacements measured by PIV.

For example, circle A in Figure 4.15 occupies all four quadrants of the displacement graph and has displacements ranging in magnitude from 0–30 pixels (an infinitely large ratio). This circle is centred on the point $\boldsymbol{d} = (-9, 0)$ and has a radius of 21 pixels; the horizontal component of displacement ranges from -30 to $+12$ pixels. Suppose that these displacements were obtained using an illumination time interval $\Delta t = 10\,\text{ms}$.

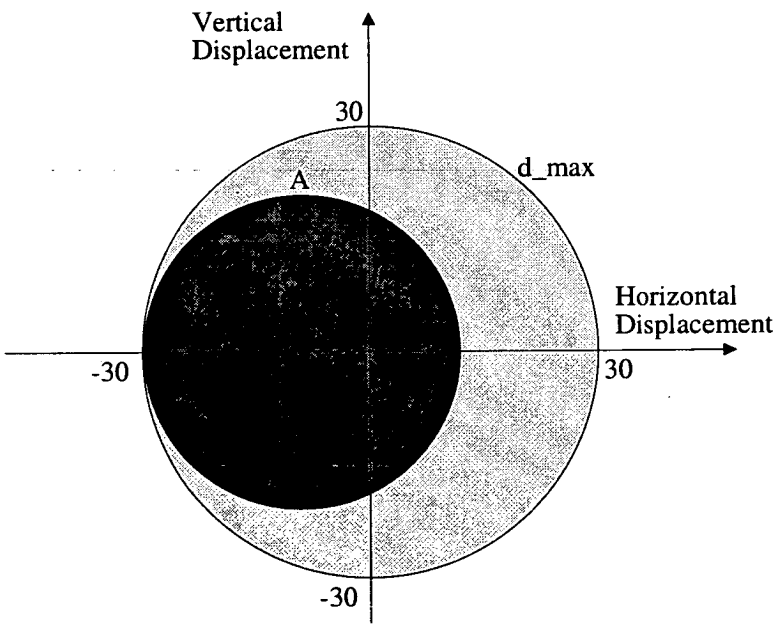


Figure 4.15: The locus of image displacements obtained from a hypothetical flow system, photographed directly with an illumination time interval of 10 ms.

A significant fraction of the displacements in this flow field are too small to be resolved. Orienting the vectors which are resolved could well be difficult.

Suppose that this circle was somehow moved 24 pixels to the left. It would now occupy only two quadrants of the displacement graph and the magnitude of the resultant displacements would be 12–54 pixels. This is shown in Figure 4.16(a).

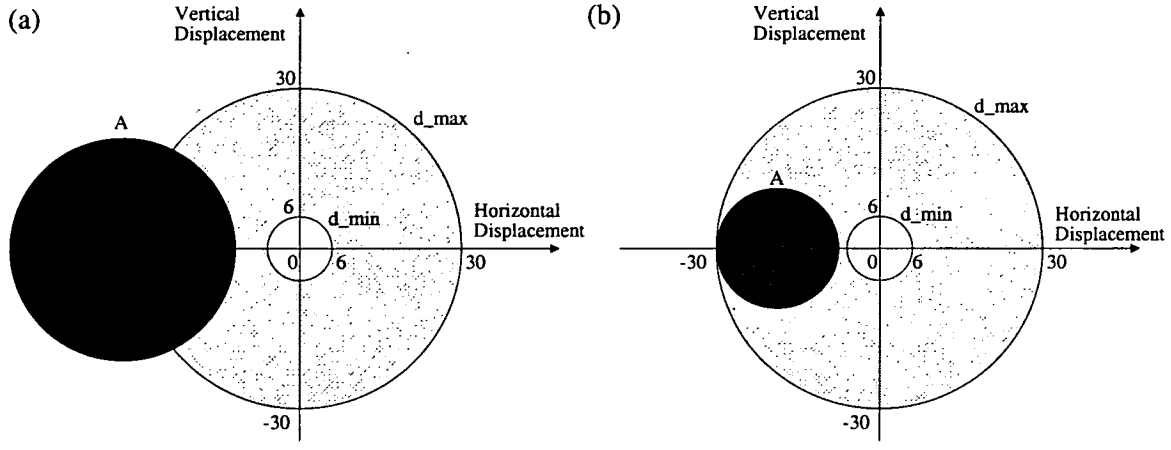


Figure 4.16: The locus of particle image displacements obtained from a hypothetical flow system: (a) after the superposition of an image-shifting velocity such that the added displacement $|\boldsymbol{d}_s| = 24$ pixels to the left on the original basis ($\Delta t = 10\,\text{ms}$); (b) $|\boldsymbol{d}_s| = 24$ pixels to the left on the original basis, but with the illumination time interval now reduced to $\Delta t = 5.56\,\text{ms}$.

A final step is to choose a more appropriate illumination time interval, in order to bring the displacements back into the acceptable, lightly-shaded region. Based on the largest displacements, the time interval should be less than $10\,\text{ms} \times 30/54 = 5.56\,\text{ms}$ (Figure 4.16(b)); based on the smallest displacements, Δt should be at least $10\,\text{ms} \times 6/12 = 5\,\text{ms}$. An illumination time interval between these values would yield a good range of displacements, which could be analysed successfully by this PIV system.

One could go further and search for the particular image-shifting displacement $\boldsymbol{d}_s = (d_{sx}, d_{sy})$ and illumination time interval Δt which allows both conditions to be met simultaneously:

$$\Delta t = 10\,\text{ms} \times \frac{30}{30 + d_{sx}} = 10\,\text{ms} \times \frac{6}{d_{sx} - 12}$$

Values of $d_{sx} = 22.5$ pixels and $\Delta t = 5.71\,\text{ms}$ mean that the rescaled circle would have displacements of 6–30 pixels, thus touching the circles of maximum and minimum permissible displacement. However, an exact solution to this sort

of problem could be difficult to obtain, *e.g.* if the locus of flow displacements is not circular or if its centre does not lie on one of the displacement axes.

After performing an autocorrelation analysis of a photograph taken using image-shifting, every vector can be oriented correctly by simply making all vectors point in the direction of the added image-shifting displacement (in this case, to the left). After this, subtracting the shift displacement from the correctly-aligned vectors yields the true displacements.

In practice, it may be convenient to perform some simple vector validation steps before subtracting the image-shifting displacement. For example, when a horizontal displacement is added, the resulting vectors should all be roughly horizontal. For example, in Figure 4.16(b), it seems fair to say that valid vectors will have an angle $135° < \theta < 225°$ and a magnitude $6 < r < 30$, when expressed in polar co-ordinates. When processing vector maps, it should be easy to delete vectors which do not meet these conditions (*e.g.* spurious near-vertical vectors).

4.6.3 The pros and cons of image-shifting

So, image shifting can be used to avoid problems of directional ambiguity in multi-exposure PIV and can also allow small velocities to be measured.

Image-shifting can also be used to increase the accuracy of PIV measurements. Consider the case of turbulent fluid flowing in a horizontal pipe at $10 \pm 0.1\,\mathrm{m\,s^{-1}}$ and suppose that small-scale turbulence structures are of more interest than the bulk fluid flow. If a straightforward PIV photograph is taken, the mean flow velocity will have to be subtracted from each vector in order to calculate local deviations from this velocity. Even if the errors associated with a PIV measurement are as low as 1% (as is claimed by various authors), the uncertainty in each measured vector is likely to be comparable to the local velocity fluctuation due to turbulence. In a vector displacement diagram such as Figure 4.14, the vectors will lie within a very small circle close to the horizontal displacement axis.

Image-shifting can be used to increase the size of the vector locus, if shifting is performed in the opposite direction to the fluid flow. The magnitude of this shift should be slightly less than the mean flow velocity. For example, an image-shifting velocity of $-9.8\,\mathrm{m\,s^{-1}}$ added to a fluid flow velocity of $10 \pm 0.1\,\mathrm{m\,s^{-1}}$ gives rise to an apparent seeding particle velocity of $0.2 \pm 0.1\,\mathrm{m\,s^{-1}}$. In this case, image-shifting removes most of the bulk flow velocity (which is of less interest) and allows PIV to measure the turbulent fluctuations with more accuracy, after what is left of the mean flow velocity is subtracted.

As is often the case, the technique of image-shifting has drawbacks as well as advantages. When used in the manner illustrated in Figure 4.16, a shifting velocity is incorporated into the PIV record and compensated for when the resultant vectors are analysed. Errors can be introduced by both of these steps.

Considering the case illustrated in Figure 4.16: the largest true displacement was 30 pixels, to the left. If the experimental errors associated with this PIV measurement amount to 1%, then the error in this displacement will be 0.3 pixels. After shifting, this velocity will be increased to 54 pixels (before the illumination time is adjusted). An error of 1% will now be worth 0.54 pixels. Thus, after allowing for image-shifting and the altered illumination time interval, this displacement will emerge as 30 ± 0.54 pixels — a larger error than when the displacement was measured directly.

Other errors are due to the experimental implementation of image-shifting. One method of achieving a velocity shift is to photograph the flow with a rotating camera or to use a stationary camera to photograph the reflection of the flow in a rotating mirror. Landreth & Adrian (1988) used an electro-optical device to perform image-shifting when using double-pulsed PIV.

One of the most commonly-used image-shifting set-ups involves the use of a rotating mirror. In this arrangement, the camera (or other recording device) does not view the flow directly. Instead, the camera axis is aligned with the plane of the light sheet and the camera photographs the reflection of the light sheet in a

mirror, placed at an angle of 45° to the camera axis. Image-shifting is achieved by rotating this mirror at a known angular velocity, so that consecutive images of a stationary seeding particle are photographed in different positions (Figure 4.17).

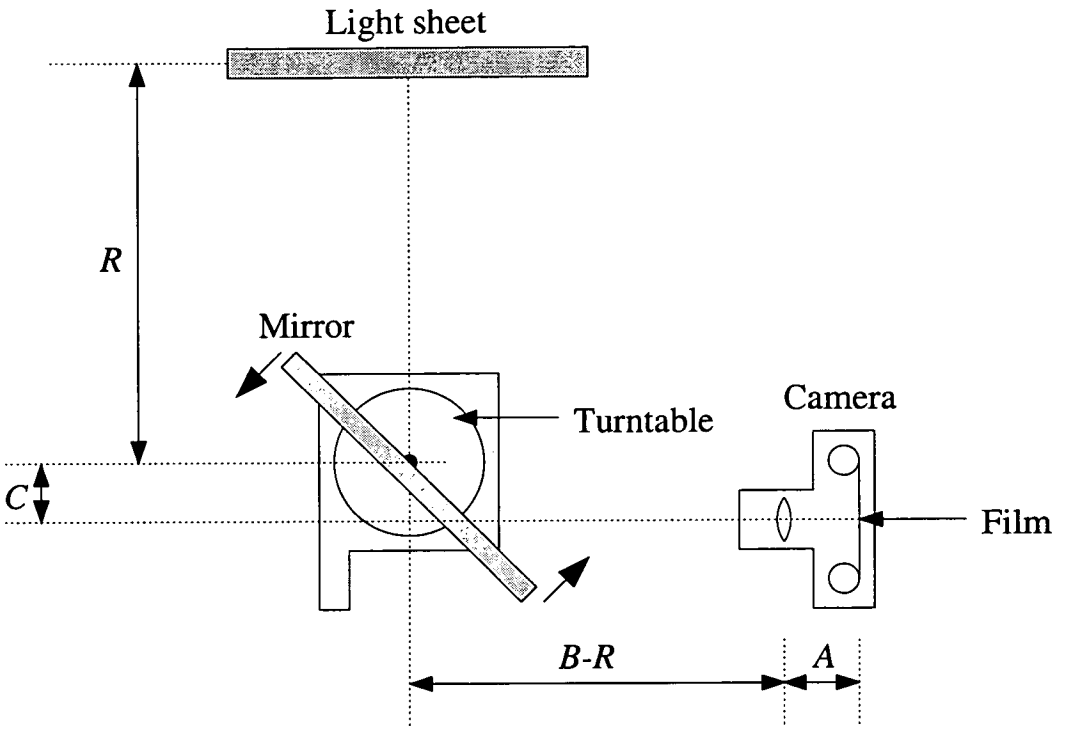


Figure 4.17: A common arrangement for horizontal image-shifting using a rotating mirror, viewed from above. A camera of focal length F views the light sheet (which is vertical and located a total distance B away) via a mirror placed at 45° to its optical axis. Rotating this mirror causes a stationary seeding particle to move to a different position in the image plane (*i.e.* on the photographic film).

Unfortunately, the displacement added by a rotating-mirror image-shifting system is not constant but varies as a function of position within the light sheet. Various authors have derived theoretical expressions for this displacement, *e.g.* Oschwald *et al.* (1995), Zhang & Eisele (1995) and Morrison (1995). The work of Raffel & Kompenhans (1995) is used here to discuss theoretical predictions, as their expressions are easier to use. All of these papers plot results graphically and these graphs reveal a high degree of similarity.

For the usual case of a vertical rotating mirror, the shift displacement $\boldsymbol{d}_s$ (in the image plane) has horizontal and vertical components (d_{sx} and d_{sy}) given by complex functions of various experimental parameters (Raffel & Kompenhans, 1995). These functions include terms involving the sine or the cosine of a certain angle,

equal to twice the angle through which the mirror rotates between consecutive illuminations. If this angle is small, then standard small-angle approximations can be made (*e.g.* $\sin\theta \approx \theta$).

Raffel & Kompenhans applied these small-angle approximations, in order to obtain a simplified relationship for the image-shifting displacement on the film. Taking x and y to be horizontal and vertical axes in the image plane:

$$d_{sx} = \frac{x - M\,R.2\omega\,\Delta t}{(x + C\,M)\,2\,\omega\,\Delta t\;F^{-1}\;(1+M)^{-1} + 1} - x \tag{4.3}$$

$$\text{and } d_{sy} = \frac{y}{(x + C\,M)\,2\,\omega\,\Delta t\;F^{-1}\;(1+M)^{-1} + 1} - y, \tag{4.4}$$

where R is the distance from the mirror's axis of rotation to the light sheet, ω is the angular rotational velocity of the mirror (rad s^{-1}), F is the focal length of the camera lens and C is an offset distance between the optical axis of the camera and the mirror surface's axis of rotation (as shown in Figure 4.17). Also shown in this figure are the distances A (from the camera lens to the film) and B (the effective optical path from the camera lens to the light sheet). As usual, M is the image:object magnification and Δt is the time interval between consecutive pulses of the light sheet.

These expressions can be simplified further. First of all, for the sake of algebraic clarity, let $a = 2\,M\,R\,\omega\,\Delta t$ (metres) and let $b = 2\,\omega\,\Delta t\;F^{-1}\;(1+M)^{-1}$ (metres^{-1}). Then, if the camera points at the axis of rotation of the mirror so that $C = 0$ m:

$$d_{sx} = \frac{x-a}{1+b\,x} - x = \frac{x-a}{1+bx} - x.\frac{1+b\,x}{1+bx} = \frac{-a-b\,x^2}{1+b\,x}$$

$$\text{and } d_{sy} = y\left(\frac{1}{1+b\,x} - 1\right) = y\left(\frac{1-(1+b\,x)}{1+b\,x}\right) = \frac{-b\,x\,y}{1+b\,x}.$$

Further, if $b\,x = 2\,\omega\,\Delta t\;F^{-1}\;(1+M)^{-1}\,x \ll 1$, then the effect of division by $(1+b\,x)$ is negligible:

$$d_{sx} \approx -a - b\,x^2 = -2\,\omega\,\Delta t\left(M\,R + \frac{1}{F\,(1+M)}\,x^2\right) \tag{4.5}$$

$$d_{sy} \approx -b\,x\,y = -\frac{2\,\omega\,\Delta t}{F\,(1+M)}\,x\,y. \tag{4.6}$$

In this work, the following parameters were typical:

- the mirror rotation rate was usually $10^\circ\,\mathrm{s}^{-1}$, so $\omega = \pi/18 \approx 0.17453\,\mathrm{rad\,s}^{-1}$;
- the illumination time interval was usually $\Delta t = 0.0104\,\mathrm{s}$;
- the magnification $M = 0.107$;
- the camera had a 60 mm lens, so $F = 0.06\,\mathrm{m}$;
- the distance from the light sheet to the rotational axis of the mirror gave $R = 0.5\,\mathrm{m}$; and
- the camera pointed directly at the axis of rotation, so $C = 0\,\mathrm{m}$.

First of all, the angle through which the mirror rotated between successive illumination pulses was about 0.1° (or 0.001745 radians). Terms involving the sine and the cosine of twice this angle were used earlier; $\sin 0.00349066^c = 0.00349065$ and $\cos 0.003491^c = 0.999994$, so small-angle approximations seem justified.

Constants a and b were used in the later stages of simplification. Here, $a = 1.936 \times 10^{-4}\,\mathrm{m}$ and $b = 0.05467\,\mathrm{m}^{-1}$. As the camera was used in landscape mode, the longer dimension of the 35 mm film was placed horizontally and so the value of x ranged from $-0.0175\,\mathrm{m}$ to $0.0175\,\mathrm{m}$ (similarly, $-0.012\,\mathrm{m} < y < 0.012\,\mathrm{m}$).

The magnitude of $b\,x$ was therefore less than 10^{-3}, hence the assumption that $b\,x \ll 1$ was found to be plausible but not exact: an error of one part in a thousand (0.1%) will be introduced into the calculated image-shifting displacements at the edges of a vector map if this assumption is made. If the intended image-shifting displacement $d_{sx} = 20$ pixels, then the error introduced by the approximation

will amount to about 0.02 pixels, which might just about be noticeable. The assumption is certainly good enough to help a human to understand the general form of the displacement introduced by image-shifting, but should probably be avoided if precise values are sought; after all, a computer can calculate exact values using equations 4.3 and 4.4 just as easily as it can calculate approximations.

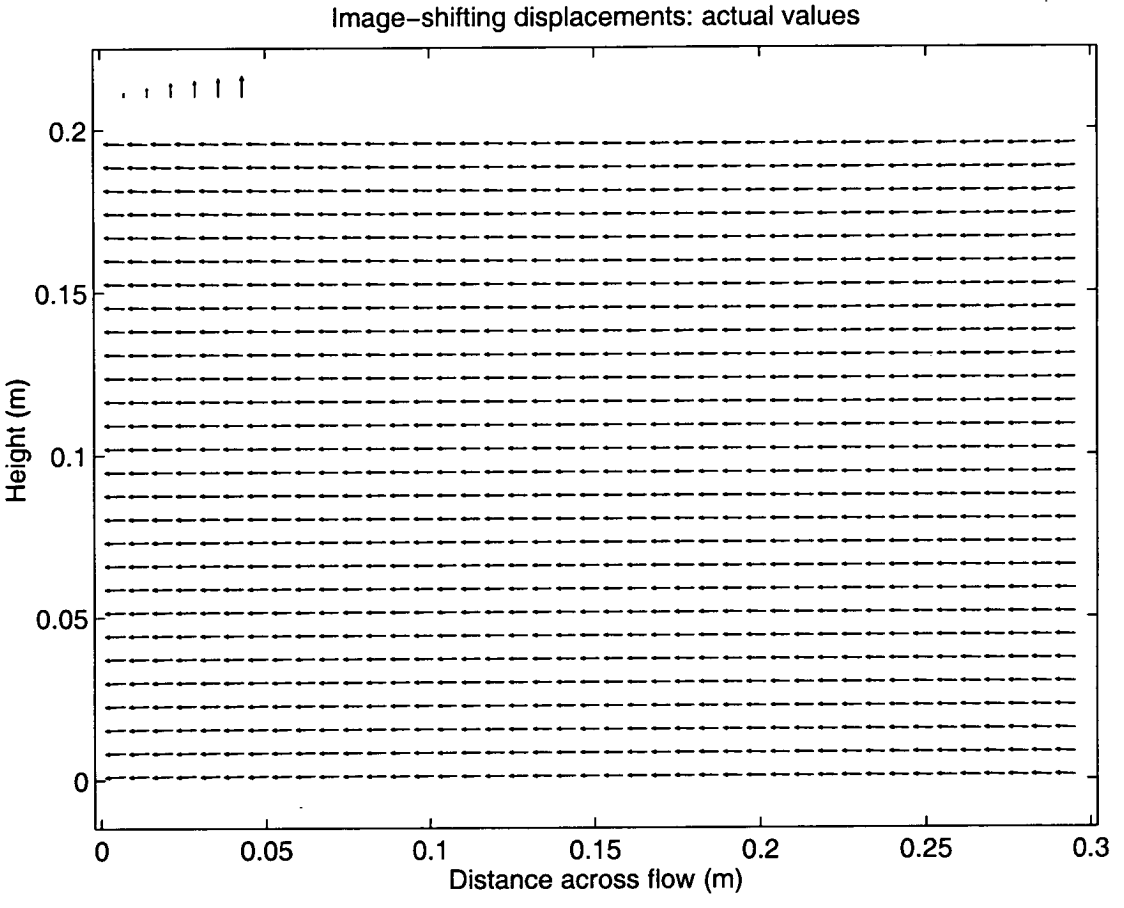


Figure 4.18: A plot of the theoretical image-shifting displacement on the film, using the work of Raffel & Kompenhans (1995). Six vertical marker vectors have been added to the top left-hand corner of this vector map, for reference. Their magnitudes are 50, 100, 150, 175, 200 and 225 μm. The second-largest vector is similar in magnitude to the mean horizontal displacement.

Figure 4.18 is a plot of the displacement predicted by Raffel & Kompenhans (1995), using the experimental parameters listed above. This shows the overall shift velocity, as predicted by equations 4.3 and 4.4. Displacements in the central region of the flow appear to be horizontal but vectors in the corners have vertical components which are just about noticeable. The displacement of the central point is 194 μm; the horizontal displacement of each corner point is 208 μm, to

which a vertical displacement of $\pm 10\,\mu$m must be added. The mean horizontal displacement, found by averaging values across the figure, is $198\,\mu$m.

The deviation of these image-shifting velocities from the central value is shown in Figure 4.19. The maximum horizontal deviation is $14\,\mu$m (7% of $194\,\mu$m) and the maximum vertical deviation is $10\,\mu$m (5%). The same marker vectors have been included as before, but the size of all vectors (*e.g.* in mm on paper per μm of magnitude) has been increased by a factor of about five, so that at least some image-shifting error vectors can be seen.

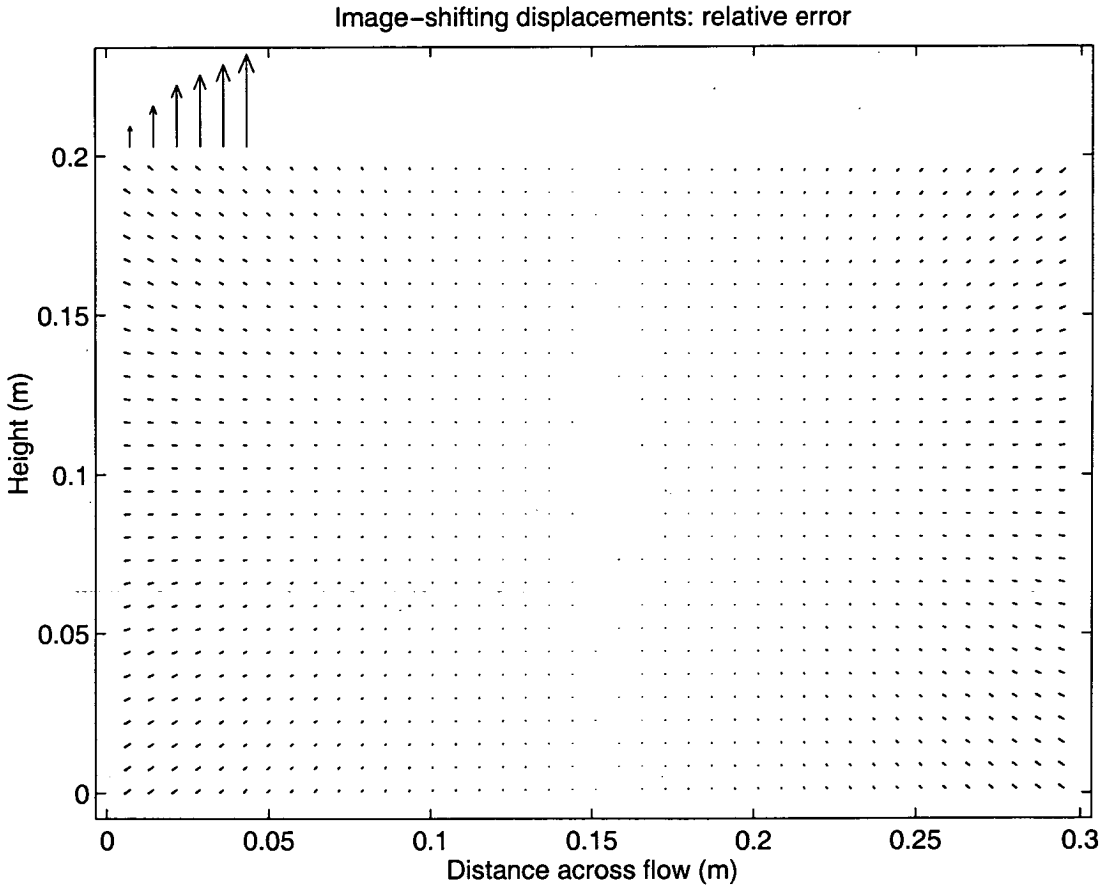


Figure 4.19: The deviation of the theoretical image-shifting displacement on the film, as calculated by Raffel & Kompenhans (1995), from the central value.

4.6.4 Minimising errors

In a rotating-mirror image-shifting system, reasonable approximations to the actual displacements are given by equations 4.5 and 4.6 above. The desired displace-

ment has no vertical component and a horizontal component of $-a$. The actual displacement has a vertical error of $-b\,x\,y$ and a horizontal error of $-b\,x^2$ (assuming that $b\,x \ll 1$). It was also possible to plot graphs of each component against the distance across the flow. These graphs revealed that the simplifications which yielded equations 4.5 and 4.6 above were plausible, but they are not shown here because they are similar to plots given by Oschwald *et al.* (1995).

Both of these errors depend on x, so the errors near the centre of the test section will be lower than those at the edges. For example, the central region of Figure 4.18, where the distance across the flow is between 0.075 m and 0.225 m, has horizontal and vertical errors of up to 4 and 5 μm respectively (less than 2.6%). A tall and thin test section would be geometrically suited to a horizontal image-shift of this type; a wider test section would be less suitable.

It is clear that the error will be reduced at all points (x, y) on the film if the magnitude of b is decreased. One cannot simply set $\omega = 0\,\mathrm{rad\,s^{-1}}$ in order to achieve $b = 0\,\mathrm{m^{-1}}$, however, as ω is also a factor of a. Minimising the ratio b/a allows the error terms to be reduced whilst taking out common factors of a and b:

$$\frac{b}{a} = \frac{2\,\omega\,\Delta t\,F^{-1}\,(1+M)^{-1}}{2\,M\,R\omega\,\Delta t} = \frac{1}{M\,R\,F\,(1+M)}.$$

Increasing the photographic magnification M is not always convenient, *e.g.* it is natural to use the camera to record the full width of the test section onto the full width of the film. M could perhaps be increased by using a larger film size, but this might mean using a different camera. The term $(1 + M)$ is likely to be reasonably constant if $M \approx 0.1$–0.2, *i.e.* for a 'large' test section.

For a given test section and camera lens, the camera needs to be a distance B away from the object, in order to photograph the full width of the test section. This distance is not changed by the introduction of a flat rotating mirror. The distance from the light sheet to the camera, R, can be maximized for a given lens if the camera is placed as close as possible to the rotating mirror. However, care must be taken that no collision will take place if the mirror rotates through 360°.

Errors can be reduced more significantly by using the same camera body with a lens of longer focal length. This increases both F and R, as the camera needs to be further away from the test section if the full width of the test section is still to be recorded. Using the lens formula and noting that the magnification $M = A/B$:

$$\begin{aligned} \frac{1}{F} &= \frac{1}{A} + \frac{1}{B} \\ \Rightarrow \frac{B}{F} &= \frac{B}{A} + 1 = \frac{1}{M} + 1 \\ \Rightarrow B &= F\frac{1+M}{M} \end{aligned} \tag{4.7}$$

For a long-lensed camera placed close to the rotating mirror, $R \approx B$, so:

$$\frac{b}{a} = \frac{1}{M\,F\,(1+M)\,R} \approx \frac{1}{M\,F\,(1+M)}\,\frac{M}{F\,(1+M)} = \frac{1}{F^2\,(1+M)^2}$$

The relative size of the image-shifting errors at each point is proportional to b/a, which should be minimized. If the distance between the camera lens and the rotating mirror is small compared to the total distance between the lens and the test section, then b/a is roughly proportional to F^{-2}. Thus, increasing F is a good way to reduce errors in the image-shifting displacement.

A drawback of a change to a lens with a longer focal length, however, is that such a lens will capture less of the light scattered by each seeding particle, being further away from the light sheet. In order to compensate for this, it may be necessary to use a brighter light sheet or a wider camera aperture setting.

In this work, the seeding particles were illuminated adequately when the laser was delivering the maximum power that could be carried by the fibre-optic cable (12 W), with the camera aperture set to $f4$. It was not certain that good PIV photographs could have been taken if the 60 mm camera lens had been replaced by a 120 mm lens.

Furthermore, doubling the focal length of the camera lens would have approximately doubled the distance B from the light sheet to the camera (equation 4.7). Ideally, this would have allowed R to double as well. In order to maintain the same central image-shifting displacement a, it would then have been necessary to halve the rotational velocity of the rotating mirror. As the experiments were already using the research group's ISS–200 image-shifter at its slowest reliable speed ($\omega = 10^\circ\,\mathrm{s}^{-1}$), however, this would not have been possible. An alternative possibility would have been to keep $\omega = 10^\circ\,\mathrm{s}^{-1}$ and $R = 0.5\,\mathrm{m}$ and to achieve the necessary B by moving the camera away from the rotating mirror; However, the rotating mirror was not large enough to have allowed this and, in any case, keeping R constant reduces the gain in accuracy.

In addition, the scanning-beam box, the image-shifting equipment and the camera were placed on a movable platform during experiments, which allowed PIV photographs to be taken at a variety of heights. This platform was not large enough to allow the increased object:camera distances.

4.7 Other refinements to PIV

As well as image-shifting, a number of other variations on the direct PIV technique have been tried. Some of these are listed in a recent PIV review by Grant (1997).

The idea behind stereoscopic PIV is to photograph a flow from two or more viewing points. Simultaneous PIV recordings are made and analysed, then related to the object plane. The different tracks of the same particle yield different apparent velocities, from which it is possible to calculate a true in-plane velocity *and* an out-of-plane velocity. The different views can be obtained by using two cameras in different positions (*e.g.* offset laterally or placed 'one behind the other' using beam splitters) or a single camera (using a system of mirrors to observe the flow from two different directions, each being photographed on a different half of the film (Arroyo & Greated, 1991)).

Multi-colour PIV is possible, if two or more lasers can be aligned in a single sheet. Busignani *et al.* (1994) took PIV photographs in which the first image of a track was green and the second blue, with these two images connected by a red streak. This solves problems of directional ambiguity (as the order of image capture is known) and also increases resolution at low displacements (if the blue and green images can be separated).

Cross-correlation also has these advantages. It involves recording each particle image in a separate frame, *e.g.* by using a special digital camera or by suitably filtering the red, green and blue levels in a multi-colour digital photograph. Compared to autocorrelation (equation 4.1), cross-correlation involves multiplying the brightness at a point (x, y) on one frame by that at the point $(x+\Delta x, y+\Delta y)$ on a *different* frame, for all points (x, y) in each interrogation area and for all plausible displacements (*e.g.* $-32 < \Delta x < 32$ and $-32 < \Delta y < 32$). Cross-correlation creates no self-correlation peak (so small displacements can be measured) and causes no directional ambiguity (as there is only one signal peak). Velocities are calculated as the object-frame displacement divided by the time between frames.

Chapter 5

Assessing the PIV analysis software

5.1 Background

In PIV, small particles are introduced into a flow and are illuminated by a pulsed light sheet. These seeding particles are photographed by a camera, whose optical axis lies at right angles to the light sheet. Usually between 2 and 6 images of each particle are recorded on each frame.

It is becoming more popular to record and process PIV images by digital means. CCD cameras avoid the need to spend time in a darkroom developing 'wet' photographs but currently have a somewhat poorer resolution. If the cost of CCD devices is prohibitive at present, digital flow recordings can be obtained from photographic negatives using suitable scanners. The speeds of computation available these days allow a 3500 $\times$ 2400 pixel image to be analysed in a few minutes; a few years ago, the same calculations took about an hour. Ever-increasing processing speeds and the prospect of applying parallel-processing techniques on multi-processor machines mean that analysis times will continue to fall.

Digital bitmap files consist of an array of pixels, with each pixel recording the integral of the intensity across a very small region of the flow. This data is divided up into a number of small 'interrogation areas', which are often squares with sides 32 or 64 pixels long. Each such region is analysed by autocorrelation. The position of the second-highest peak in the autocorrelation plane is used to calculate the displacement of particle images between successive illuminations and hence (knowing the illumination time interval) the velocity.

Autocorrelation analysis can be carried out with various degrees of zero-padding applied to the FFTs used for autocorrelation. Analysis without any padding at all is a possibility well worth exploring: a test flow record may be analysed using widely-spaced interrogation areas and unpadded FFTs, to assess the conditions in the illuminated test section whilst the experimenter is still in the laboratory. If conditions are not suitable, then the seeding concentration or the illumination time interval can be altered on the spot.

When setting PIV parameters, Prasad *et al.* (1992) recommended that:

$$\frac{d_\tau}{d_{pix}} \geq 4,$$

where d_{pix} is the pixel size and d_τ is twice that distance away from the centre of an image of a PIV seeding particle at which the recorded image intensity falls by a factor of e^2 (the 'e^{-2} diameter'). In other words, the diameter of each seeding particle image in a digital photograph should be at least four pixels.

For good analysis, successive images of each seeding particle must be distinct. Keane & Adrian (1990, 1991) recommended that:

$$|\boldsymbol{d}|_{min} \approx 1.5\, d_\tau,$$

where $|\boldsymbol{d}|_{min}$ is the magnitude of the smallest displacement $\boldsymbol{d}$ which can be determined accurately. If the image size d_τ increases, then the minimum resolvable displacement is made larger.

Further, the separation of successive images should not be too large. If this is greater than 32 pixels in magnitude and parallel to one of the edges of the interrogation area, and if interrogation areas measure 64 × 64 pixels, then only one or two images of each seeding particle can lie within each interrogation region and the signal peak is weakened. Keane & Adrian (1990, 1991) suggested that:

$$|\boldsymbol{d}|_{max} \approx 0.25\, d_I,$$

where $|\boldsymbol{d}|_{max}$ is the magnitude of the largest displacement which can be resolved reliably and d_I is the characteristic size of the interrogation areas. This suggests that larger displacements can be resolved if the size of each interrogation area is increased. This is not always advisable, however, as larger interrogation areas mean that velocity gradients are more likely to be present.

Taken together: the minimum image size is 4 pixels (Prasad *et al.*, 1992), the minimum image separation is about 6 pixels and the maximum image separation is about 16 pixels (Keane & Adrian, 1990, 1991). This suggests that the 'dynamic range' (*i.e.* the ratio $|\boldsymbol{d}|_{max}/|\boldsymbol{d}|_{min}$) is about 2.6 for a typical PIV system.

Techniques such as image-shifting can solve a number of problems encountered in direct PIV, such as directional ambiguity and difficulties resolving slow-moving regions of the flow. Unfortunately, methods of achieving image-shifting in practice (*e.g.* using a rotating mirror) usually produce a non-uniform image-shift, which needs to be subtracted from the results. Theoretical expressions exist, or the shift can be determined experimentally. However, if the dynamic range of a simple PIV system could be increased from 2.6 to 10, say, there would be less need to rely on image-shifting.

In this work, a number of PIV flow fields have been simulated and analysed using 63 × 63 pixel interrogation areas. The objective was to compare the results to the recommendations of Keane & Adrian (1990, 1991) and to assess the errors introduced by the analysis system.

5.2 Simulating PIV images

A command-line program known as PIVSIM has been developed by Hind (1997), in order to assess the effect of deforming interrogation regions prior to autocorrelation analysis, so as to reduce errors due to velocity gradients (Hind & Christy, 1996). PIVSIM was run under the Unix operating system in this work.

PIVSIM runs with a number of options. These allow the user to specify:

- the size of bitmap to be generated (number of rows and columns of pixels);
- the displacement between successive particle images (pixels);
- the e^{-2} diameter of each particle image (pixels);
- the intensity profile of each image (*e.g.* Gaussian);
- the number of images of each seeding particle;
- the brightness of the centre of each image (from this and the e^{-2} diameter, the average brightness of the surrounding pixels is calculated, subject to a maximum brightness of 255);
- the total number of seeding particles in the bitmap;
- a seeding value used to generate the 'random' location of seeding particles. This allows a deleted bitmap file to be recreated later if required and also allows an alternative bitmap to be created, if a simulation needs to be checked.

In addition, other options allow for the creation of vortices in the bitmap image; however, all particles had the same simulated velocity in this idealised work. No turbulence was added and there were no velocity gradients.

In each simulation, a constant number of images was created of each particle; this number was either 2, 3, 4 or 5. At first, particle images were 4 pixels in diameter. The number of particles present was adjusted whenever the number of

images of each particle was changed, so that about 16 particle image pairs lay within each interrogation area when the displacement was small (Table 5.1).

Images of each particle:	2	3	4	5
Total number of particles:	16,000	8,000	5,400	4000
Total number of images:	32,000	24,000	21,600	20,000
Useful image pairs/interrogation area:	15.9	15.9	16.1	15.9

Table 5.1: The seeding concentrations in each simulation.

The seeding concentration was kept constant over the full range of displacements. Whilst some improvements might have been obtained by introducing more particles into the higher-velocity flows, so as to boost the number of particle image pairs, this was not done here as the aim was to investigate the ability of PIV to resolve a range of velocities, for a given particle concentration.

If one considers the simulated PIV recordings and interrogation regions to be squares with perpendicular sides in the 'horizontal' and 'vertical' directions, then the first set of runs had a vertical displacement of zero pixels and a horizontal displacement of between 5.0 and 45.9 pixels, at intervals of 0.1 pixels. Thus, the characteristics of the PIV system were measured using ten simulated analyses per pixel of horizontal displacement.

5.3 Analysing digital PIV images

Hind has also developed an analysis program called PIVANAL. Command line options allow the user to:

- choose the analysis method (in this work, double-precision Fast Fourier Transforms (FFTs) were used);
- specify an autocorrelation peak location algorithm, *e.g.* Gaussian (as used here), Whittaker's reconstruction, linear, *etc.*;
- enable or disable displacement peak validation (this was disabled);

- specify the size of interrogation areas to be used (here, 63×63 pixels);
- specify the distances between the centres of adjacent interrogation areas, as a number of rows and a number of columns (both of these were 42 pixels);
- specify the degree of padding to be used when calculating the FFTs.

The output of this program is a file which can be read by the mathematical software package MATLAB. A mathematical array is created, in which rows store the results from adjacent rows of interrogation areas and columns represent columns of interrogation areas. The value of each point in this array is a complex number, the real component of which is the horizontal displacement in pixels; the imaginary component stores the vertical displacement, in pixels.

MATLAB supports simple programming, allowing the results to be displayed, edited and manipulated in various ways. For example, one can make all the vectors point in a particular direction, or plot a displacement size distribution function (showing the number of vectors having a certain magnitude of displacement), or one can delete displacement vectors which meet certain criteria.

'Validating' the measured data can also be done by PIVANAL, as mentioned above. For example, applying the 'double' option means that a peak in the autocorrelation plane will only be accepted as valid if the amplitude of this peak is at least 35% of the amplitude of the self-correlation peak and at least 35% larger than the amplitude of the next highest peak in the autocorrelation plane. The 'multi' option is more lenient in terms of amplitudes in the autocorrelation plane (using tolerances of 20% instead of 35%) but also requires a harmonic peak to be found less than two pixels away from the expected position.

The 'multi' option would not always have been suitable here, as harmonic peaks cannot be found within the autocorrelation plane if the image displacement is too large or if only two images of each particle are simulated. It was hoped that the 'double' option could have been used in all cases, but tests suggested that its 35% tolerances were far too strict. The 'none' option was preferred instead.

This method simply takes the highest value in the autocorrelation plane excluding the central self-correlation peak (*i.e.* the second-tallest peak in the autocorrelation plane) as the 'displacement peak', which corresponds to the separation of successive particle images. Validation was carried out as a post-processing step.

The size of the artificial PIV bitmaps and the chosen analysis parameters meant that $47 \times 47 = 2209$ different measurements of the simulated displacement were made from each simulated PIV recording. Using MATLAB, these vectors were neglected if they deviated by more than 1 pixel from the expected displacement, then the mean displacement was calculated for the remaining 'valid' vectors. The process was repeated using the mean displacement from one iteration as the expected displacement in the next cycle, where stricter criteria for validity were applied. Iterations stopped after a number of cycles at a tolerance of 0.33 pixels.

This 'validation' of PIV data measures validity in terms of the self-consistency of the measured vectors. If the displacement peaks are, on average, only slightly larger than the tallest peaks due to random autocorrelation noise, then an incorrect measurement will be obtained whenever this order happens to be reversed. As these noise peaks will probably be much more than 0.33 pixels away from the average position of the 'valid' vectors, this validation method will ignore the resulting spurious results. A poor yield of 'valid' vectors would be expected if the autocorrelation signal:noise ratio was poor, so this would become apparent. The yield of valid vectors should decrease as the image separation increases.

An alternative (but more complicated) method would have involved counting the number of vectors which were less than a certain 'tolerance' distance (*e.g.* 1.0 pixels) away from the simulated displacement and evaluating the mean value of these vectors. An iterative method could then have been used, based on progressive alterations to this tolerance distance for each simulated displacement, in order to search for the particular values which would be needed if some pre-determined fraction of the analysed vectors was to be found in the 'valid' region near each mean value. Each value of tolerance found in this way would then be a measure of the standard deviation in the displacement measurements at each simulated

image separation: in a one-variable situation governed by a normal distribution function, 95% of the measurements will lie within two standard deviations of the mean value. A high tolerance would indicate an unreliable reading, *e.g.* plausible results might be a tolerance of 0.10 pixels at a simulated displacement of 10.3 pixels and a tolerance of 1.35 pixels at a higher simulated displacement of 26.8 pixels. Although this method has its advantages, it was not used in this work.

5.4 Errors in digital PIV analysis when using unpadded FFTs

The first set of results were for recordings with two simulated illumination pulses, which were analysed without applying any zero-padding to the FFTs used for autocorrelation. The particle images had a simulated vertical displacement of zero pixels and a non-zero horizontal displacement. These simulations began at a displacement of 5 pixels. The results are shown in Figure 5.1.

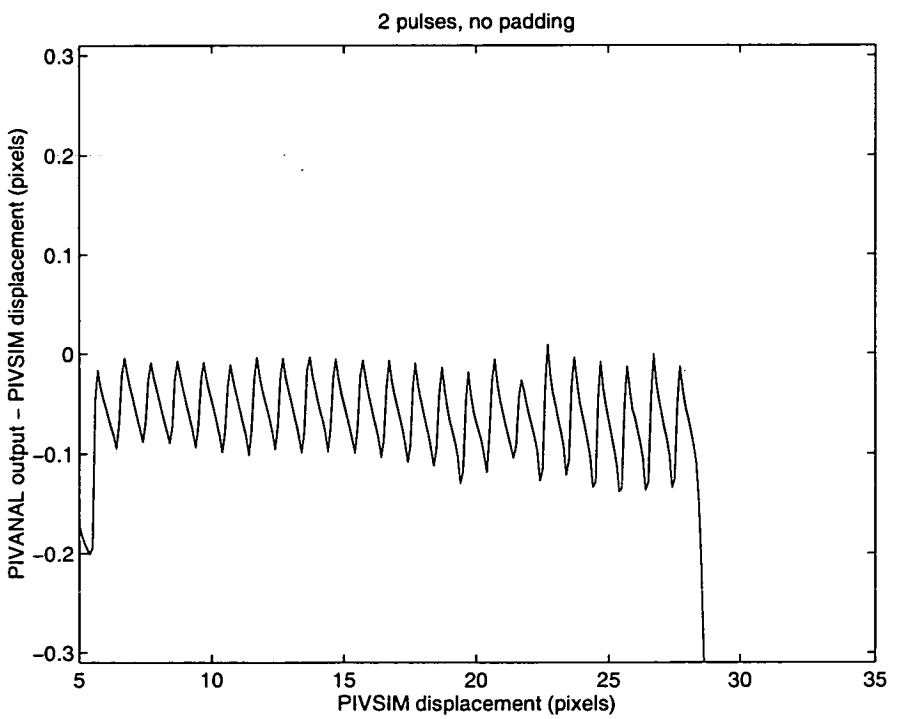


Figure 5.1: The average error in the PIVANAL displacement, calculated without padding and plotted as a function of the simulated displacement, for the case of two images of each particle.

A transient error is apparent at low velocities, beginning at about 4.75 pixels and dying out by the time the displacement reaches about 6 pixels. This error was due to overlaps between the intended velocity peak and the self-correlation peak. The recommendation of Keane & Adrian (1990, 1991) that the minimum displacement should be about 1.5 times the particle diameter seemed to be good.

The analysis software seemed to be unable to detect displacements greater than 29.2 pixels. For example, the yield of valid vectors dropped from 72% at a displacement of 28.3 pixels down to 2% at a displacement of 29.0 pixels. This was due to the peak-searching algorithm.

When analysing a 63×63 pixel image without padding, a 64×64 pixel autocorrelation plane was calculated, which was symmetrical about its centre — the point between pixels $(32, 32)$, $(32, 33)$, $(33, 32)$ and $(33, 33)$. Peaks within this autocorrelation plane were assumed to be due to noise if they failed to meet certain criteria for detectability. One of these criteria required the peak autocorrelation value to be greater than or equal to that of its nearest neighbours (the outer 8 points in a 3×3 array of pixels centred on the peak value) and strictly greater than the autocorrelation values at the 16 points surrounding them in a 5×5 array (Figure 5.2). This meant that peaks whose centres lay in the outermost two pixels of their autocorrelation plane were rejected, simply because they did not have enough nearest neighbours.

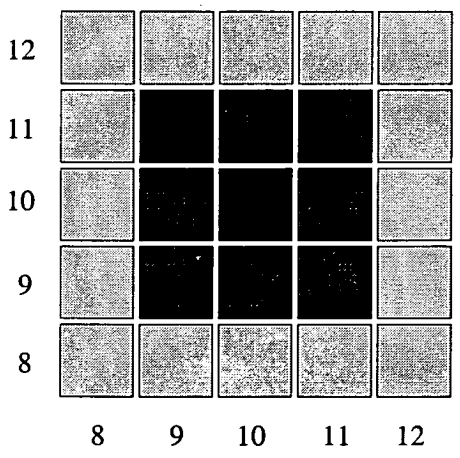


Figure 5.2: An example of the relative intensities of autocorrelation values (shown as degrees of blackness) around a displacement peak at the point $(10, 10)$, which would satisfy the validity criteria imposed by PIVANAL.

The maximum horizontal displacement which could be accepted had an autocorrelation peak at a point 62 pixels across the autocorrelation plane. This would lie 29.5 pixels across from the centre of this plane, in reasonable agreement with the value of horizontal displacement for which the yield was found to be zero.

In Figure 5.1, the average absolute error in the analysed value is approximately -0.05 ± 0.045 pixels. This error can be described as the sum of three terms: a constant offset of -0.05 pixels, a cyclic term with an amplitude of about 0.045 pixels and a period of 1 pixel (the 'variable error') and a lower-amplitude random error. The 'constant offset error' measures the average difference between the measured displacements and the simulated displacements. This is negative, because the measured displacements were almost always smaller than the simulated values (Figure 5.1). The variable error oscillations are due to the peak location method: interpolation is used to estimate the co-ordinates of peaks in the autocorrelation plane to sub-pixel accuracy, but the results can be biased towards integer values.

It should be possible to correct all of the systematic error in each measurement. Considering a typical displacement of about 10 pixels, however, the variable error amounts to less than 0.5% of the measured displacement, so a full correction does not seem worthwhile. It is simpler to correct results obtained using PIVANAL by systematically adding 0.05 pixels to all values and then writing off what remains as an error of about 0.05 pixels, even though the variable error is not entirely random and could therefore be subtracted after calibration tests. (Later work will investigate the image diameter which minimizes the variable error.)

Similar simulations were carried out using three images of each particle. The results are shown in Figure 5.3. Under these conditions, the average error at small displacements was about -0.055 ± 0.035, slightly less than that when two images were simulated. This graph is dominated by the unexpectedly large errors found at horizontal displacements between 19 and 23 pixels in magnitude. The amplitude of these errors is about 0.3–0.4 pixels, which is nearly ten times greater than that of the variable errors found at other displacements.

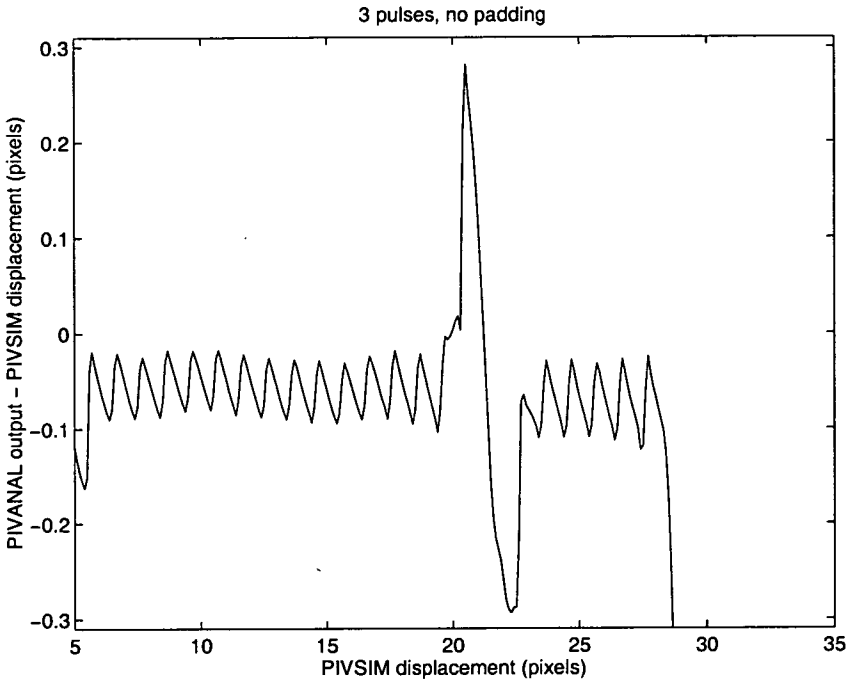


Figure 5.3: The average error in the PIVANAL displacement, calculated without padding for the case of three images of each particle.

These errors arose because the analysis of each interrogation region was carried out using insufficiently-padded FFTs. Here, interrogation areas measuring 63×63 pixels were padded up to 64×64 pixels (as the computation of an FFT is fastest if the data is a power of two in size) and then a 64×64 pixel autocorrelation function was calculated.

When an FFT is calculated, the data being analysed is implicitly assumed to be periodic, with a period which is equal to (or a factor of) the length of the FFT. In one dimension, this is can be thought of as writing the data along the perimeter of a circle, so that the first item of data is adjacent to both the second item and the last item. In two dimensions, one can imagine taking nine copies of the interrogation region and arranging these as 'tiles' in a 3×3 grid (Figure 5.4).

Mathematically, for a 64×64 pixel interrogation area, let the image intensity (brightness) at a point (x, y) be $I(x, y)$. This interrogation area consists of discrete data points and is finite in extent, *i.e.* $I(x, y)$ is defined only for integer values of x and y which satisfy $1 \leq x \leq 64$ and $1 \leq y \leq 64$. As far as the FFT is concerned, however, the data is discrete and infinitely repeating:

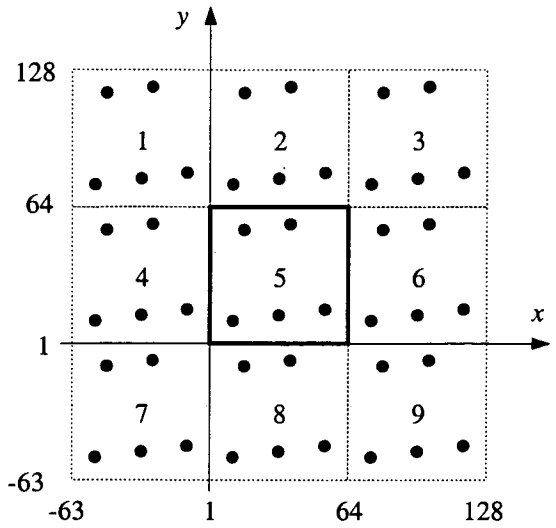


Figure 5.4: Tiling of the data contained in each interrogation area occurs if the data is autocorrelated using Fourier transforms. Copies of the interrogation area in question (numbered 5 and including pixels 1–64 in each direction) are placed around the interrogation area in a 3 × 3 array.

$$I(64\,n + x, 64\,m + y) = I(x, y),$$

for all $n, m \in \mathcal{Z}$, the set of integers. Ultimately, for any autocorrelation displacement $(\Delta x, \Delta y)$, autocorrelation multiplies the brightness of every point (x, y) in the data by the brightness of another point $(x + \Delta x, y + \Delta y)$. These calculations take place without regard to the 'joins' between adjacent tiles.

The situation is illustrated in Figure 5.5, which shows only two tiles along the lines of those shown earlier in Figure 5.2. In this schematic interrogation area, the displacement between particle images is 20 pixels to the right. The first images of tracks A and B lie in the left-hand third of the interrogation area, so three images are obtained. The first image of track C lies in a different interrogation region and only one point is visible. The first image of track D lies in the central third of the region and so two images are found.

An autocorrelation displacement 20 pixels to the right correlates particle images such as A1 and A2 with a later image of the same particle. Images such as A3, B3 and D2, however, lie close enough to the right-hand edge of the interrogation area

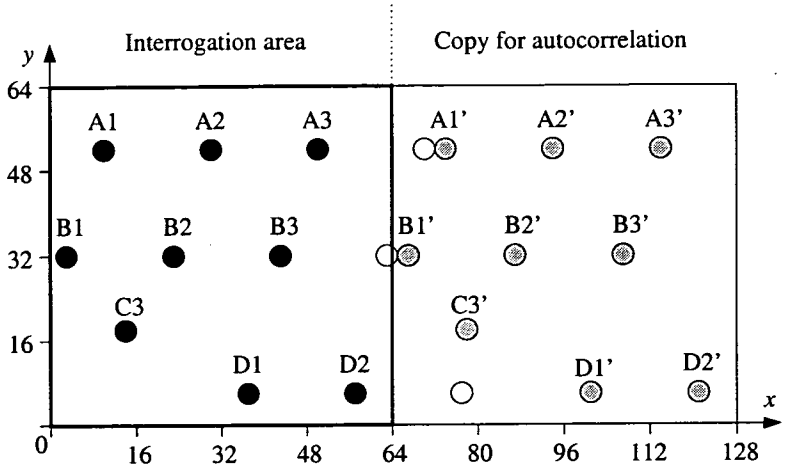


Figure 5.5: A sample interrogation area, where particle images are shown as black circles, being autocorrelated using a 64 × 64 pixel unpadded FFT. Points near the right-hand edge of the interrogation area must be correlated with points from a 'tiled' copy (lightly-shaded images). For this displacement, the particle images A3, B3 and D2 are correlated with the regions shown as unfilled circles.

($x = 64$ pixels) for this displacement to correlate them with regions beyond this edge. In an unpadded FFT, the dataset for autocorrelation is wrapped around onto itself by tiling, as is partly shown in Figure 5.5.

The first image of track A, A1, is centred on $(10, 52)$; A2 is centred on $(30, 52)$ and A3 is at $(50, 52)$. Under an autocorrelation displacement of 20 pixels to the right, A3 correlates with a region centred at $(70, 52)$. Because of FFT wrap-around, a second image of A1 must be considered; this image, A1', is centred at $(10 + 64, 52) = (74, 52)$: right next to the region with which A3 is correlated.

If the artificial image shown in Figure 5.5 is analysed, the intended displacement peak is based on five particle-image correlations, from A1 to A2, A2 to A3, B1 to B2, B2 to B3 and D1 to D2. In this case, however, two extra correlations are just about to have an effect: from A3 to A1' and from B3 to B1'. This wrap-around noise signal will have a peak at the displacement $(24, 0)$, so, in this particular case, there will probably be little effect on the displacement peak at $(20, 0)$.

However, the addition of these spurious wrap-around correlations to the true displacement peak could well be harmful. For example, if the true displacement was 21.0 pixels, then the wrap-around process would introduce a spurious peak at

a displacement of 22.0 pixels. The close proximity of these correlations means that they would overlap significantly, changing the position of the measured peak and thus the measured displacement. Any error introduced by adding them together tends to pull an intended displacement peak at about 20–21 pixels up towards 64/3 pixels. This explains why the measured displacement rises above the expected value, for simulated displacements between 19 and 21 pixels in magnitude.

At a displacement of 64/3 pixels, the error is fairly small. This displacement correlates A1 with A2, A2 with A3 and also A3 with A1'. The wrap-around peak coincides with the main displacement peak, so no error should be found.

At displacements greater than 64/3 pixels, the wrap-around peak lies closer to the origin than the true displacement peak. The addition of these peaks explains the systematic under-read of displacements between 64/3 and (64+4)/3 pixels.

This explanation can be used to make further predictions:

- A further error at about 64/2 = 32 pixels might have been seen, had the analysis system been able to detect displacements that large.
- PIV records obtained using only two illumination pulses should not suffer from wrap-around errors at displacements of 64/3 pixels.
- More than one distortion due to wrap-around noise will occur if more than three images of each particle are simulated. For example, if there are five images of each particle, errors can occur at displacements of about 64/2, 64/3, 64/4 and 64/5 pixels.
- Errors at smaller displacements will be less pronounced. For example, in the three-pulse example illustrated in Figure 5.5, track A gave rise to two correct particle image correlations and one spurious wrap-around overlap when the displacement was about 64/3 pixels. In the case of five particle images and a displacement of around 64/5 pixels, there will be four genuine correlations per spurious wrap-around signal, for those tracks where all five images lie within the interrogation region.

In addition, there will probably be some incomplete particle tracks in each interrogation region (*e.g.* track D of Figure 5.5). These partial tracks contribute to the measurement of the correct displacement but are most unlikely to introduce any wrap-around errors.

- For near-horizontal displacements greater than 32 pixels in magnitude, only one or two images of each particle can fit into each interrogation region. At these large displacements, the wrap-around peak in an unpadded analysis lies closer to the origin than the true displacement peak and has the same amplitude: for every valid correlation from some image A1 to a later image A2, there will be a corresponding spurious correlation from A2 to A1'. If the true displacement from A1 to A2 is $\boldsymbol{d} = (d_x, d_y)$ for some $d_x > 32 > d_y$, the peak detected by autocorrelation analysis might well be that corresponding to the displacement from A2 to A1', *i.e.* $(64, 0) - \boldsymbol{d} = (64 - d_x, -d_y)$.

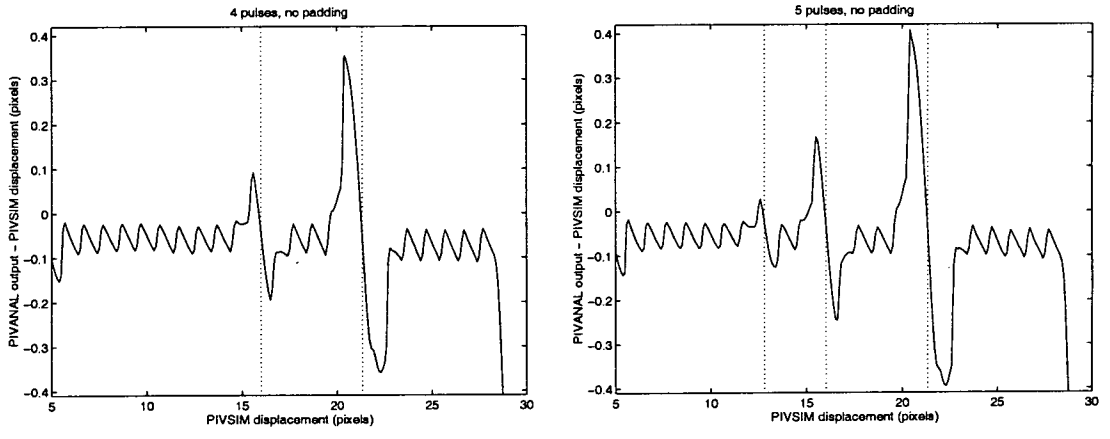


Figure 5.6: The error in the displacement calculated without padding, using four or five images of each particle, plotted against the simulated displacement. The dotted vertical lines mark simulated displacements of 64/5, 64/4 and 64/3 pixels.

Figure 5.6 shows plots of the average error in the measured displacement, for the cases of four and five images of each particle. As before, these results were calculated by PIVANAL, using Gaussian interpolation and with no padding applied to the FFTs. As expected, wrap-around noise caused errors in the determination of horizontal displacements approximately 64/3, 64/4 or 64/5 pixels in magnitude, according to the actual number of images of each particle.

It is also possible to test for wrap-around errors by examining the average measured displacement for simulated displacements larger than 35 pixels. In these cases, the peak due to wrap-around lies within the detectable region of the 64×64 pixel autocorrelation plane, whereas the true peak lies outside the detectable region. Figure 5.7 shows, for the case of two images of each particle, a plot of the average horizontal displacement found by PIVANAL against the horizontal displacement simulated by PIVSIM. The search for a zone of valid vectors was modified, so that the MATLAB code averaged the displacements of vectors a certain tolerance away from 64 pixels minus the simulated displacement, if the latter was greater than 32 pixels. It can be seen that the data obtained fits this model. Results from 3-, 4- and 5-pulse simulations showed similar behaviour.

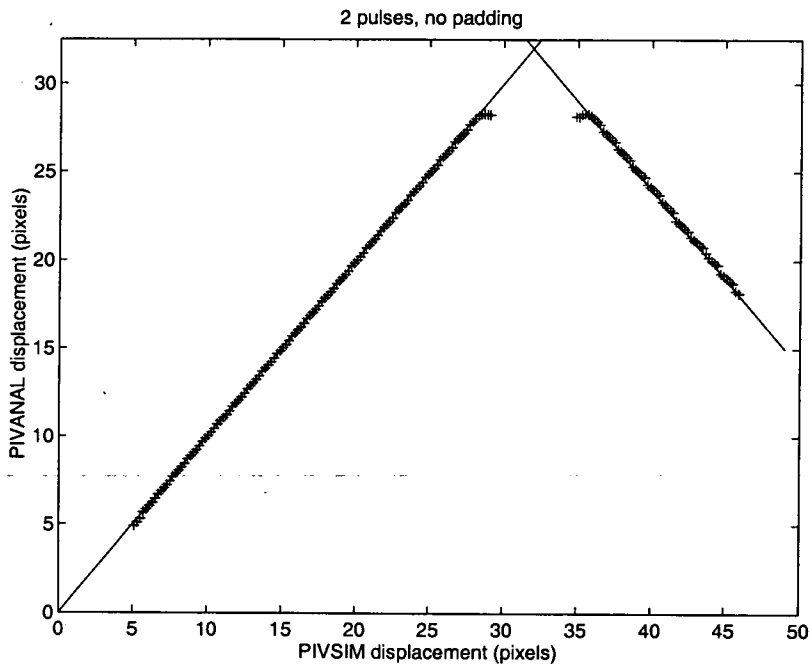


Figure 5.7: A plot of the most common output displacement calculated by PIVANAL against simulated displacement. The continuous lines are plots of ordinate = abscissa and ordinate = 64 − abscissa, shown for comparison.

5.5 Results with padded FFTs

For autocorrelation analysis with full zero-padding, the 63×63 pixel interrogation areas were, in effect, placed in the centre of a 128×128 grid of black pixels (zero brightness). The FFTs used in the autocorrelation analysis were now 128×128

pixels in size. These larger FFTs took about 3–4 times as long to compute, but did not suffer from the effects of wrap-around noise and allowed horizontal displacements of up to 62 pixels to be detected.

Simulated bitmaps with two, three, four or five illumination pulses were created as above, then analysed with a padding ratio of 1.0. As Figure 5.8 shows:

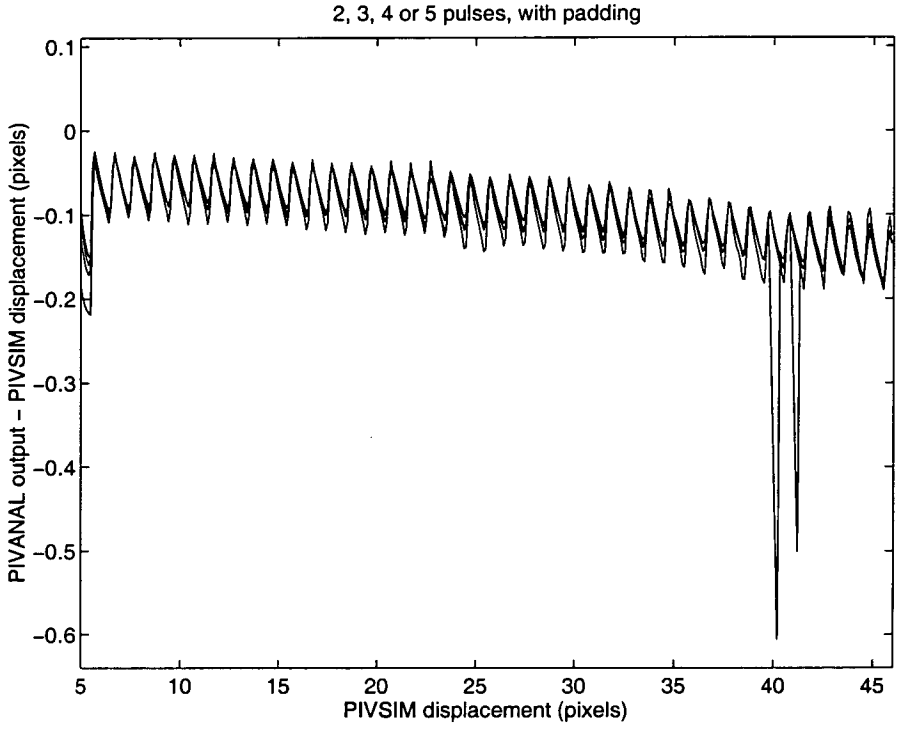


Figure 5.8: The average error when using fully-padded FFTs as a function of the simulated horizontal displacement, for the cases of two, three, four and five particle images.

- in terms of average accuracy, there seems to be very little difference between simulations using different numbers of particle images. One thing that was apparent was that the magnitude of the variable ('bias') error was 25% larger if two images were recorded of each particle, as opposed to five;
- the calculated displacements were systematically about 0.06 pixels less than the simulated displacements, with a variable error of about ± 0.035 pixels; these values were similar to those obtained using with unpadded FFTs;
- there were no discernible errors due to wrap-around noise at displacements of about 64/2, 64/3, 64/4 or 64/5 pixels;

- there was a slight increase in the average under-read error, as the simulated displacement increased;
- for some undetermined reason, there was a sudden drop in accuracy at about 41.2 pixels for the four-pulse case and about 40.2 pixels for the five-pulse case. In these regions, the calculated displacement decreased by about 0.4–0.5 pixels. If genuine, these errors may limit the maximum displacement which can be resolved accurately, if the autocorrelation analysis uses fully-padded FFTs.

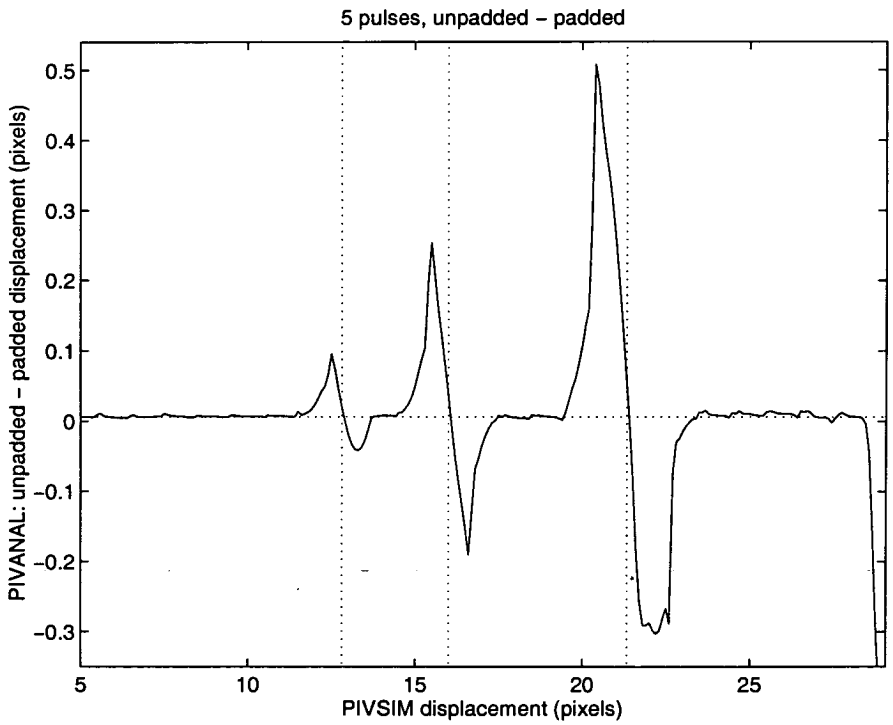


Figure 5.9: The difference in the average horizontal measurements with and without padding for the five-image case, plotted against the simulated horizontal displacement. Data obtained using padded FFTs was subtracted from data obtained without padding. Dotted lines mark values of 64/3, 64/4 and 64/5 pixels.

Because the padded analyses did not show errors due to wrap-around noise, these results could be subtracted from the results obtained without padding, in order to obtain the average signal due to wrap-around noise alone. Such a plot is shown in Figure 5.9, for the case of five images of each particle. It is clear that the wrap-around errors have local maxima and minima on either side of the displacements 64/3, 64/4 and 64/5 pixels. The errors introduced are serious,

e.g. displacements of 20.4 pixels are almost always over-read as 20.9 pixels and displacements of about 22 pixels are under-read as 21.7 pixels.

5.6 The probability of obtaining a valid vector

In the sections above, the accuracy of each PIV measurement was determined by averaging those displacement vectors which lay within a certain tolerance of their mean and so were deemed to be 'valid'. However, the accuracy of the calculated vectors should not be the only consideration when designing a PIV experiment: the probability that each vector will be valid also needs to be taken into account. For example, a good yield of reasonably-accurate vectors might give a better understanding of a flow field than a poor yield of highly-accurate vectors.

In each test above, the number of valid vectors was recorded, along with the mean values which were plotted earlier. This allowed estimates to be made of the workability of each measurement. Figure 5.10 shows the results obtained for autocorrelation analysis using unpadded FFTs. In this Figure, the yield refers to the fraction of vectors lying within 0.33 pixels of the mean displacement calculated by PIVANAL. When analysing with unpadded FFTs, this displacement was close to the simulated horizontal displacement for displacements of less than 29 pixels and was approximately 64 pixels minus the simulated displacement, for displacements greater than 35 pixels (*cf* Figure 5.7). The absence of results for simulated horizontal displacements between 29 and 35 pixels in magnitude is due to the peak detection algorithm in PIVANAL.

Some comments on the data presented in Figure 5.10:

- In the two-pulse case, the yield of valid vectors began to drop off almost immediately. This suggests that the use of two-pulse PIV is limited, even though the accuracy of those vectors which were determined remained good (as shown in Figure 5.1). The average yield dropped to 90% at a displacement of about 15 pixels and to 80% at a displacement of about 23 pixels.

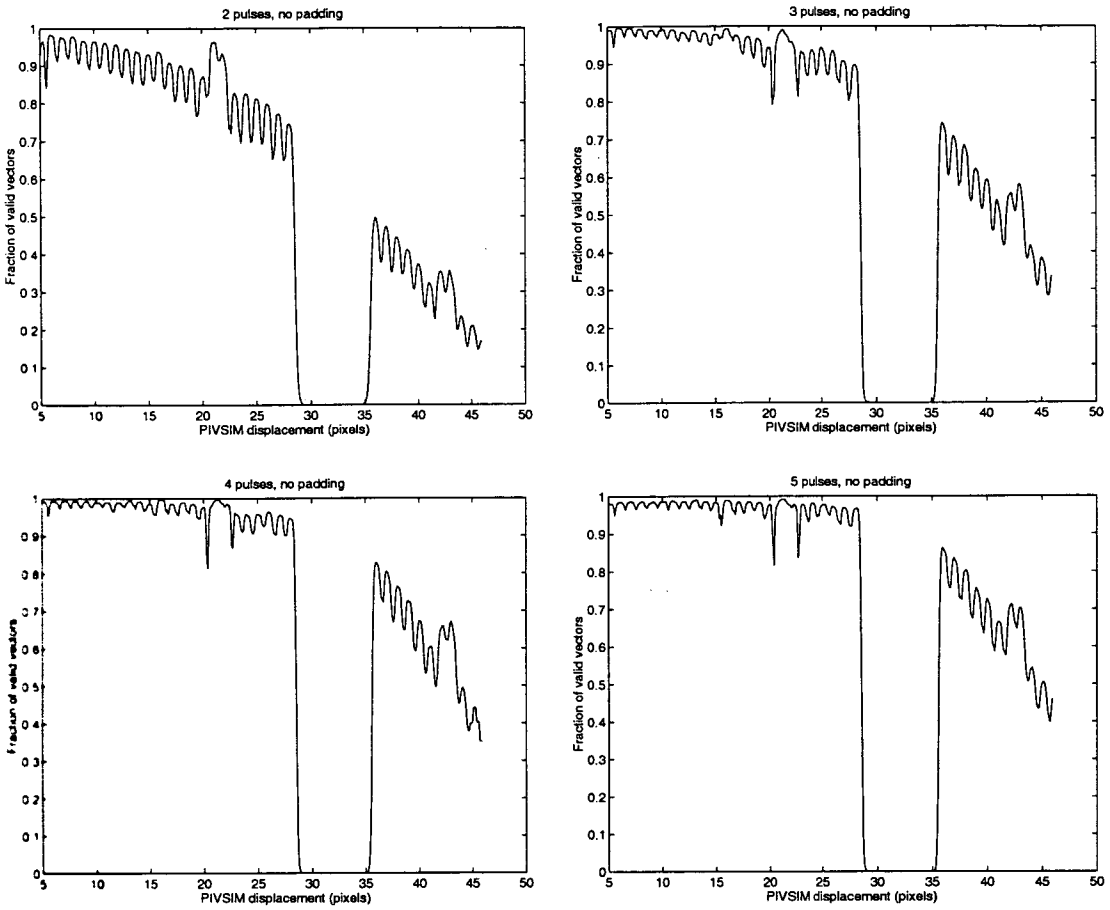


Figure 5.10: The fraction of valid vectors determined by the analysis (without padding) of simulated PIV images, plotted against the simulated displacement. Plots are given for simulations using 2–5 images of each seeding particle.

- The yield improved significantly as the number of illumination pulses increased. The average yield was over 90% for displacements of up to 23 pixels when three images were simulated and almost always above 93% when there were four or more images of each seeding particle. The simulation with five illumination pulses returned the highest yield of consistent 'valid' vectors, but those measurements contained more errors due to wrap-around noise.

- Interestingly, the effects of wrap-around noise are apparent in these graphs of the fraction of valid vectors. There seem to be blips in the yield at both 64/4 and 64/3 pixels in all unpadded simulations, even when only two images were simulated. This was surprising, because wrap-around errors were neither observed in the error plot shown in Figure 5.1 nor expected based on an understanding of the wrap-around phenomenon. Further, blips

in the measured yield seem to occur at $2 \times 64/3 \approx 43$ pixels, regardless of the number of images of each particle.

When padded FFTs were used, the calculated displacement was always close to the simulated displacement. Figure 5.11 shows that there was a significant improvement in the yield of valid vectors if three images of each particle were recorded, rather than just two. Simulations using four or five particle images showed further improvements, which were small but noticeable. In all cases, the yields with zero-padding were better than those without padding (Table 5.2).

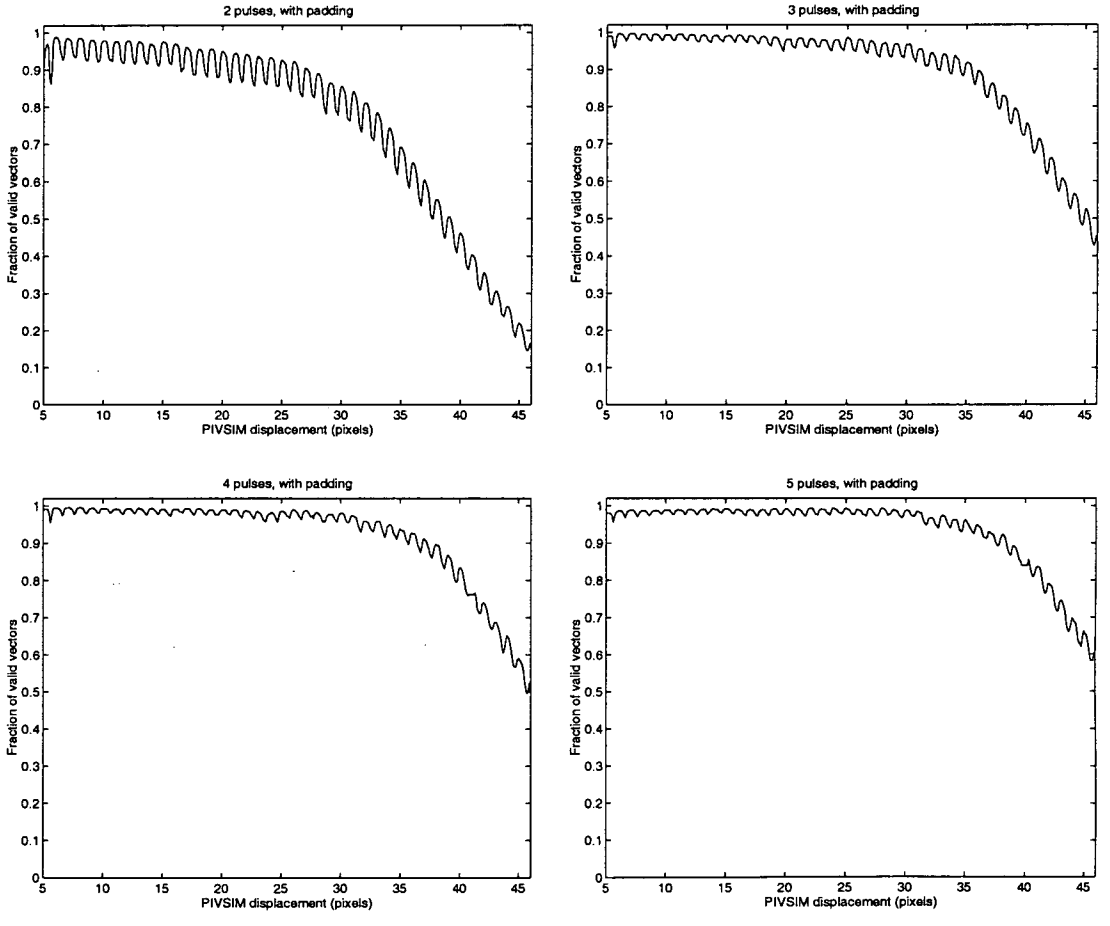


Figure 5.11: The fraction of valid vectors determined by the analysis (with full zero-padding) of simulated PIV images, plotted against the simulated displacement. Plots are given for simulations using 2–5 images of each seeding particle.

	90% yield if $d_x <$ (pixels)	80% yield if $d_x <$ (pixels)
2 pulses	25	31
3 pulses	35	38
4 pulses	36	40
5 pulses	38	41

Table 5.2: The simulated horizontal displacement at which the average yield of valid vectors dropped to 90% and then 80%, when using fully-padded FFTs.

In this idealised situation with no velocity gradients, it appears to be best to design a PIV system to avoid the two-pulse case and to record instead between 3 and 6 images of each particle. Camera shutter times are approximately a factor of two apart (*e.g.* 1/125 s, 1/60 s and 1/30 s) so this can always be achieved. In practice, the flow might well vary as a function of time, so the longer exposure times may not actually yield better results. Recording three images of each particle could be a good compromise, allowing a good yield of valid vectors to be obtained from a relatively-short camera exposure.

In practice, the magnitudes of the measured displacement vectors will follow some sort of distribution function, with a mean value and a variance. It is best to ensure that there is a good yield of vectors at the most commonly-occurring displacement. The 50% of all measurements lying at or below the median value should be satisfactory, provided that the displacement does not become unresolvably small. For displacements greater than the median, the yield will decrease as the displacement is increased (as shown in Figure 5.11). Fortunately, so will the number of measurements required: if only 5% of the vectors to be measured are greater than 38 pixels in magnitude and if 80% of these measurements will be successful anyway, then the overall yield should still be very good.

Based on these criteria, it seems reasonable to extend the dynamic range of a PIV system to horizontal displacements 6–38 pixels in magnitude. If three or more images are recorded of each particle and the digital analysis takes place with fully zero-padded FFTs, then the yields of valid vectors should be very good. .

5.7 Other parameters

In a two-dimensional flow with significant vertical and horizontal displacements, the yield will decrease as either component of velocity increases.

5.7.1 Small vertical displacements

Because wrap-around noise arises when the final image of one track is wrapped-around to the vicinity of an earlier image of the same particle, the errors will be affected by any non-zero vertical component of velocity. If the first image of a three-particle track lies at $(10, 10)$ and the displacement between particle images is $(21, -1)$, then the second image must lie at $(31, 9)$ and the third at $(52, 8)$. The wrap-around peak would be found at a displacement of $(22, 2)$. If the particle diameter is still 4 pixels, then this wrap-around peak will affect the true displacement peak, located at $(21, -1)$. The two peaks will add together, causing the measured displacement to be found somewhere in between, *e.g.* at about $(21.15, -0.80)$. The result of this is that, if a constant vertical displacement of about 0–2 pixels is simulated (instead of the 0.00 pixel vertical displacement used above), then *both* components of the measured displacement can suffer from errors due to wrap-around if the FFTs are not padded.

Figures 5.12 and 5.13 show the effect of various simulated vertical displacements on the PIV analysis, when no padding is applied to the FFTs and when three images of each particle are simulated. The average horizontal error for a vertical displacement of 0 pixels was plotted earlier, in Figure 5.3. Wrap-around errors are still apparent when vertical displacements of up to 1.4 pixels are added to the horizontal displacement, affecting both components of the measurement. Vertical displacements of 1.7 pixels (or more) seem large enough to avoid errors due to wrap-around noise at horizontal displacements close to 64/3 pixels.

The measurement of horizontal displacement becomes increasingly reliable as the vertical displacement is increased. The average error in the measured vertical

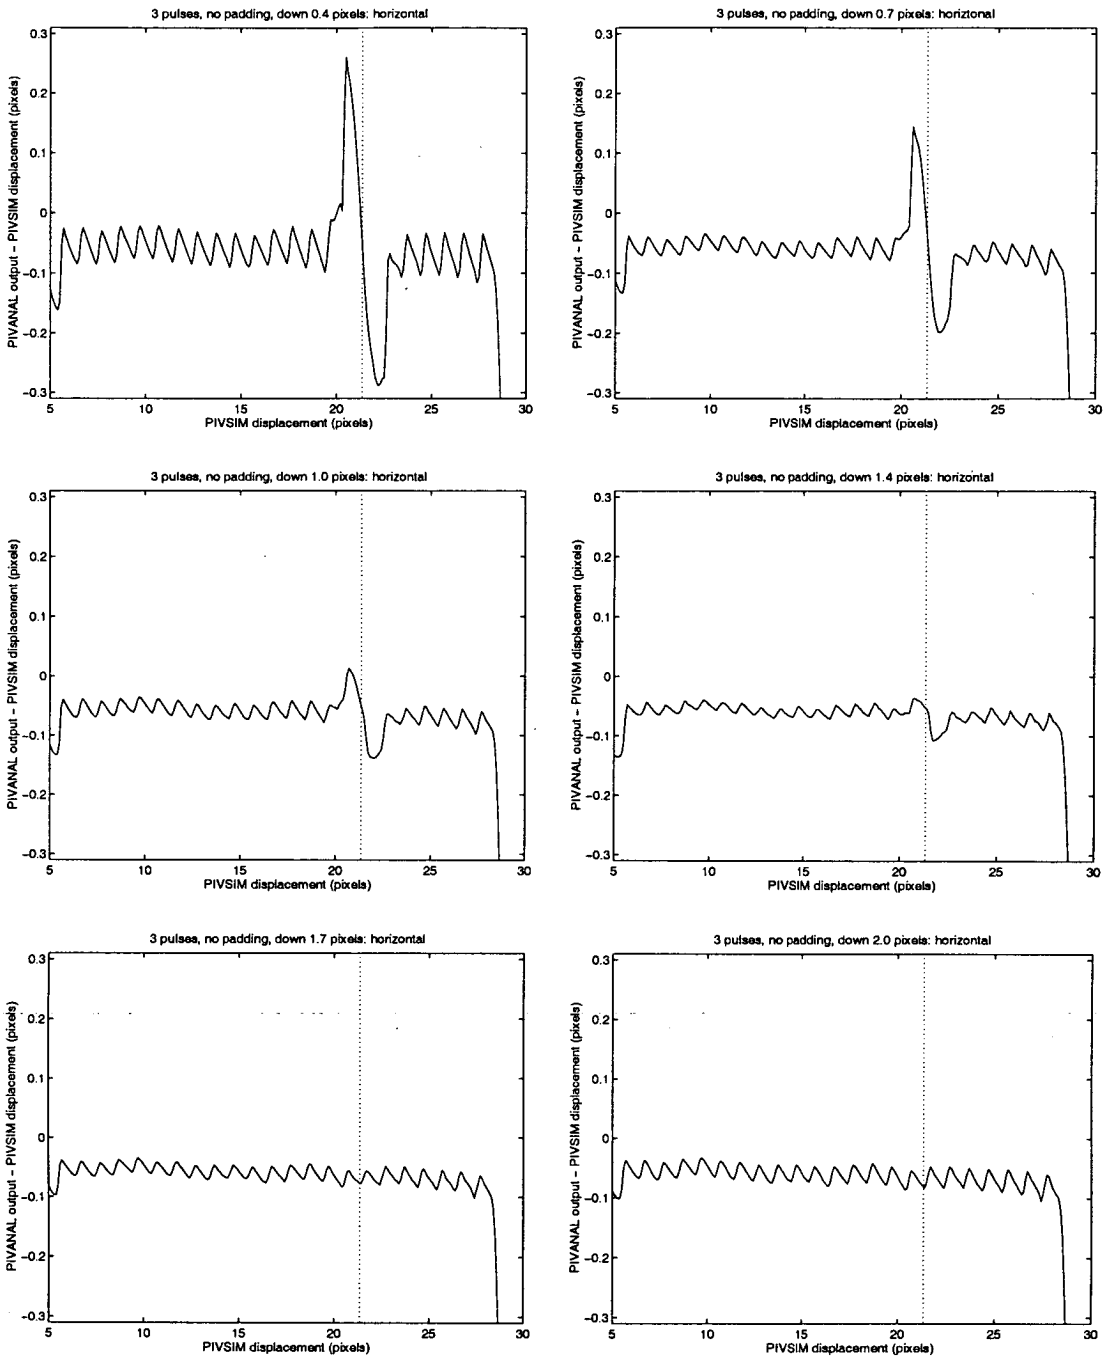


Figure 5.12: The average error in the measured horizontal displacement, plotted against the simulated horizontal displacement. In these simulations, a constant downward displacement of 0.4, 0.7, 1.0, 1.4, 1.7 or 2.0 pixels was added to a range of horizontal displacements. Three images of each seeding particle were simulated and the autocorrelation analysis took place without padding the FFTs. The vertical dotted lines mark a simulated horizontal displacement of 64/3 pixels.

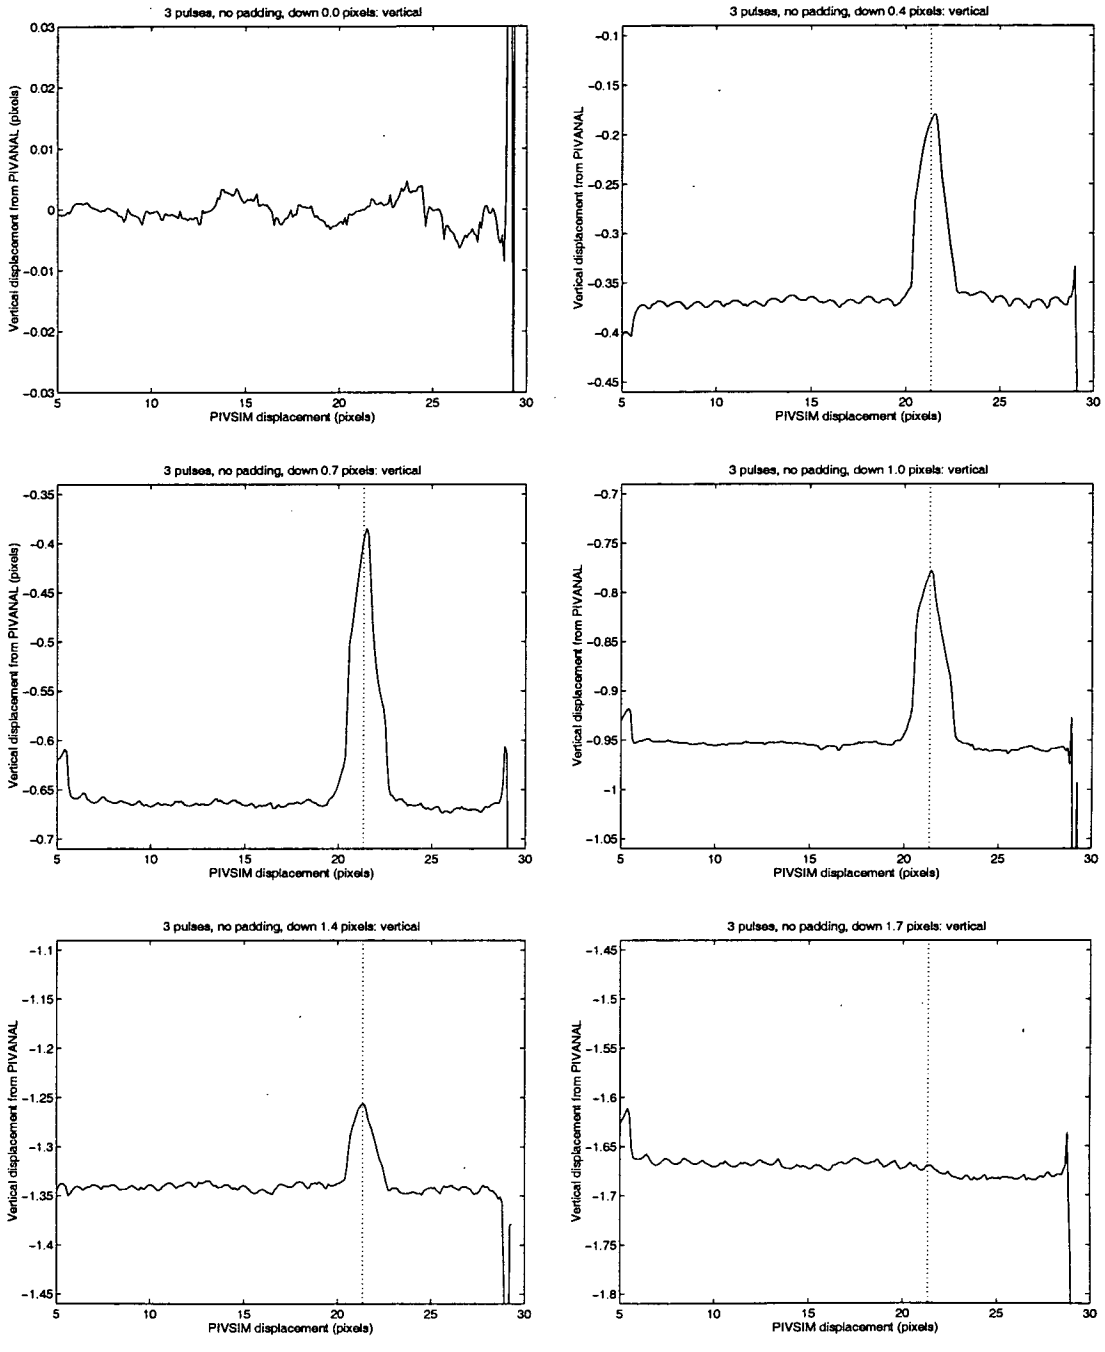


Figure 5.13: The average vertical displacement calculated by PIVANAL, plotted against the simulated horizontal displacement, when a downward displacement of 0.0, 0.4, 0.7, 1.0, 1.4 or 1.7 pixels was added to a range of horizontal displacements. Three images of each seeding particle were created and the autocorrelation analysis used unpadded FFTs. The vertical dotted lines mark a simulated horizontal displacement of 64/3 pixels.

displacement has a maximum value of about 0.28 pixels, when the simulated vertical displacement is about 0.7 pixels and the simulated horizontal displacement is very slightly more than 64/3 pixels. The average errors found in Figures 5.3, 5.12 and 5.13 are summarised in Table 5.3.

Vertical displacement (pixels):	0.0	0.4	0.7	1.0	1.4	1.7
Maximum horizontal error (pixels):	0.30	0.30	0.19	0.06	0.035	0.01
Maximum vertical error (pixels):	0	0.18	0.28	0.18	0.08	0

Table 5.3: Horizontal and vertical errors, at various vertical displacements.

If there are more than three images of each particle, the same vertical displacement of about 1.7 pixels will be required to avoid wrap-around errors at horizontal displacements close to 64/3 pixels, but a lower vertical displacement could suffice elsewhere (*e.g.* about 1 pixel if five images are recorded and $d_x \approx 64/5$ pixels.)

Unidirectional flows such as these *are* likely to be found in practice, *e.g.* flows recorded using image-shifting and turbulent flow in a pipe have a mean velocity and relatively small fluctuations. If the interrogation areas are parallel to a mean displacement which is approximately 64/3 pixels in magnitude, then analysing a three-pulse PIV recording with unpadded FFTs could give very inaccurate results.

5.7.2 Larger vertical displacements

In the next batch of simulations, the vertical component of the simulated displacement was increased in steps of 5.0 pixels. The immediate objective was to determine the yields of valid vectors which could be obtained, if the displacements to be measured were not parallel to two sides of the interrogation areas. After that, a further aim was to use a suitable criterion for successful operation (*e.g.* 'at least 90% of the vectors at each displacement should be measured correctly') to determine the high-displacement limits of this PIV system.

Some of these simulations are shown in Figure 5.14. In each run, a variable horizontal displacement of 5–40 pixels was added to a constant downward displacement of either 0, 5 (not shown), 10, 15 (not shown), 20 or 25 pixels (not

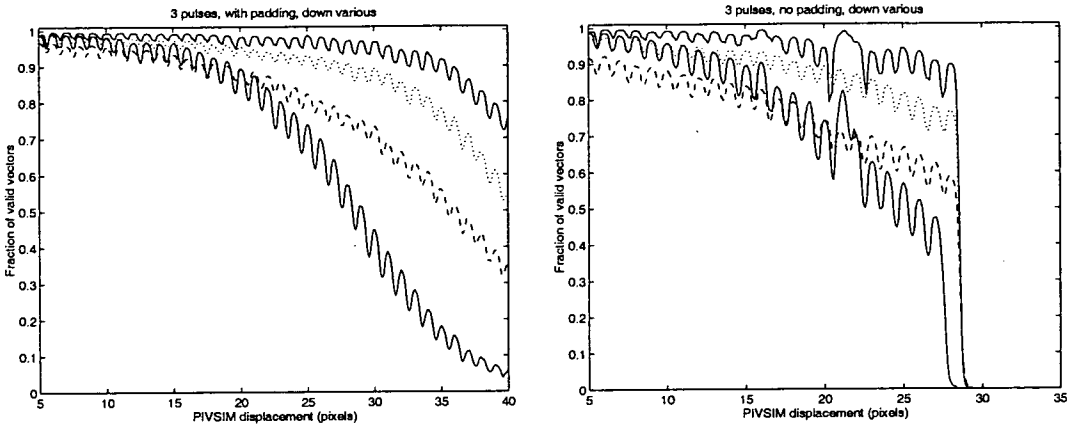


Figure 5.14: The yield of valid vectors obtained by PIVANAL, plotted against the simulated horizontal displacement, when various downward displacements were added to a range of horizontal displacements. Three images of each seeding particle were simulated and the autocorrelation analysis took place with fully-padded (left) or unpadded (right) FFTs. From top to bottom, the curves shown in each graph are for simulated downward displacements of 0 pixels (uppermost solid line), 10 pixels (dotted line) or 20 pixels (dashed line); the lowermost solid curve is for diagonal displacements.

shown). In addition, another simulation was made in which the horizontal and vertical components of displacement had the same magnitude, *i.e.* the particles were moving at 45° to the sides of the interrogation areas used to measure the displacement. In all of these runs, three images of each particle were simulated and the resulting PIV recordings were analysed with either fully-padded or totally unpadded FFTs. The yields could probably have been improved slightly by simulating more particles, or more images of each particle.

The yields obtained using the above conditions are summarised in Table 5.4. In a plot of horizontal displacement against vertical displacement, these results can be used to determine contours of average yield = 0.9 and average yield = 0.8.

Assuming that interrogation areas can be rotated by 90° and/or reflected in vertical or horizontal axes without changing the yield of valid vectors, further data points can be used. For example, the average yield at a displacement $(17, 20)$ was found to be 0.90 when fully-padded FFTs were used; if symmetry applies, then the yields at all displacements matching $(\pm 20, \pm 17)$ or $(\pm 17, \pm 20)$ will also be 0.90. The results appear to indicate that this assumption is reasonable, *e.g.*

Simulated vertical displacement d_y (pixels)	Padded yield > 0.9 if $d_x <$	Padded yield > 0.8 if $d_x <$	Unpadded yield > 0.9 if $d_x <$	Unpadded yield > 0.8 if $d_x <$
0	35	38	24	29
5	31	36	22	29
10	28	34	17	25
15	21	31	12	21
20	17	26	6	16
25	13	21	—	10
d_x	18	22	13	17

Table 5.4: The horizontal displacements which produce given fractional yields of valid vectors, when analysing a range of two-dimensional displacements.

with zero-padding, the yield at a displacement $(20.0, 10.0)$ was 0.954 whereas that at $(10.0, 20.0)$ was 0.955. The data is somewhat approximate, because the yield of valid vectors is not a strictly decreasing function; estimates were made of the simulated horizontal displacement (to the nearest integer) at which the local average value of yield reached 0.9 or 0.8, on graphs similar to those in Figure 5.14.

Plots of yield contours are shown in Figure 5.15. In each case, points at which the yield was either 0.80 or 0.90 are plotted on a displacement diagram, with the two-dimensional displacement between successive images of each particle (d_x, d_y) plotted as a point (*cf* an Argand diagram for complex numbers). Plausible best-fit lines were added, drawn by eye. These best-fit lines ran between a point on an axis (*e.g.* $(35, 0)$ or $(0, 35)$ for a 90% yield and analysis with padding) and a point on the line $d_y = d_x$ (*e.g.* $(19, 19)$ for the above case). An arc of a circle, given by the equation $|\boldsymbol{d}|$ = constant, fitted the yield data poorly, but a single straight line would have been a simpler alternative, offering a reasonable fit (*e.g.* $d_x + d_y = 36$, for a yield of 0.9 with padded FFTs). Dotted lines mark the points which satisfy the guidelines of Keane & Adrian (1990, 1991) for the minimum and maximum displacement of particle images, namely $|\boldsymbol{d}|_{min} = 6$ pixels and $|\boldsymbol{d}|_{max} = 16$ pixels. It appears that the upper limit proposed by Keane & Adrian is unnecessarily conservative, particularly when the FFTs are fully zero-padded.

If the design condition when setting PIV parameters is that 90% of the measured vectors must be 'valid', the maximum acceptable displacements are approx-

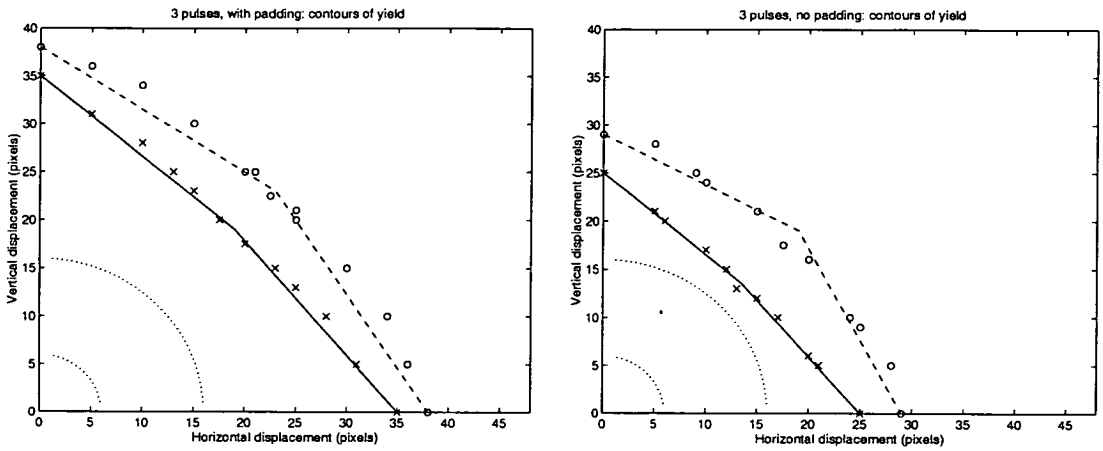


Figure 5.15: For one quadrant of the displacement plane, these graphs show plots of points at which the yield of valid vectors was 80% (circles) or 90% (× signs), for the cases of analysis with (left) or without (right) FFT padding.

imately 40% larger if padding is used. For a given PIV recording and an initial intention to analyse using unpadded FFTs, analysis with padding will allow more of the 'large' vectors to be resolved correctly. Alternatively, given the same flow and an initial choice of unpadded FFTs, the illumination time interval can be made 40% longer if analysis with fully-padded FFTs is used instead; this will decrease the magnitude of the smallest resolvable flow velocity whilst not affecting the yield of larger vectors, allowing more low-velocity vectors to be measured.

Another disadvantage of using unpadded FFTs is that, whilst wrap-around errors for predominantly-horizontal displacements disappear if the vertical displacement exceeds 2 pixels, these wrap-around errors appear again if the vertical displacement increases to about 64/3 pixels. Whilst this might not be considered significant when setting parameters for a three-pulse PIV recording because the yield of vectors at a displacement $(64/3, 64/3)$ would be unacceptably low (about 70%), it could cause problems if four or more images were recorded of each particle. In this case, the wrap-around errors due to near-diagonal displacements would be critical, because the yields of valid vectors at large displacements would be better (so regions close to $(64/3, 64/3)$ would probably produce an acceptable yield) and because wrap-around errors would also occur at displacements close to $(16, 16)$, which would almost certainly be an otherwise-acceptable displacement.

FFT wrap-around causes errors at displacements which do not have to be parallel to the sides of the interrogation areas. Errors can occur at displacements close to $(0, \pm 64/M)$, $(\pm 64/M, 0)$ and $(\pm 64/M, \pm 64/M)$, where M is a whole number between 2 and the number of images of each particle, inclusive. If the interrogation areas are imagined to be 'tiled' in a 3×3 grid, as in Figure 5.4, then these diagonal wrap-around errors occur when images in the central copy of the interrogation area are correlated with those in one of the corner copies (*e.g.* from copy 5 of the interrogation region in Figure 5.4 to copy 3). Wrap-around for near-horizontal displacements occurs between horizontally-adjacent tiles, *i.e.* copies 4, 5 and 6 of the interrogation area shown in Figure 5.4. Wrap-around can also occur for near-vertical displacements.

It appears that digital PIV flow records will give the best results when fully-padded FFTs are used and one side of the square interrogation areas is approximately parallel to the direction of flow. If a flow has local velocities of up to 32 pixels in magnitude pointing in all possible directions (*e.g.* if there is a recirculation zone), then a different PIV analysis scheme might be successful. An initial conventional analysis can be made, in order to ascertain the nature of the flow, then the worst vectors are deleted and replaced with interpolated values. After that, a cross-correlation analysis based on these displacements should produce better yields of valid vectors.

5.7.3 Peak location method

Using a digital PIV analysis method, the data in the autocorrelation plane will be discrete, at intervals of 1 pixel (Figure 5.2). The only way to obtain results to better than the nearest pixel is to use an interpolation method. In the simulations above, the location of each autocorrelation peak was determined to sub-pixel accuracy using Gaussian interpolation. Further tests were carried out using two other methods supported by PIVANAL, namely volume moments and Whittaker's reconstruction (described by Lourenco & Krothapalli, 1995). These tests were limited to horizontal cases (*i.e.* $d_y = 0$ pixels) and were carried out both with and without padding.

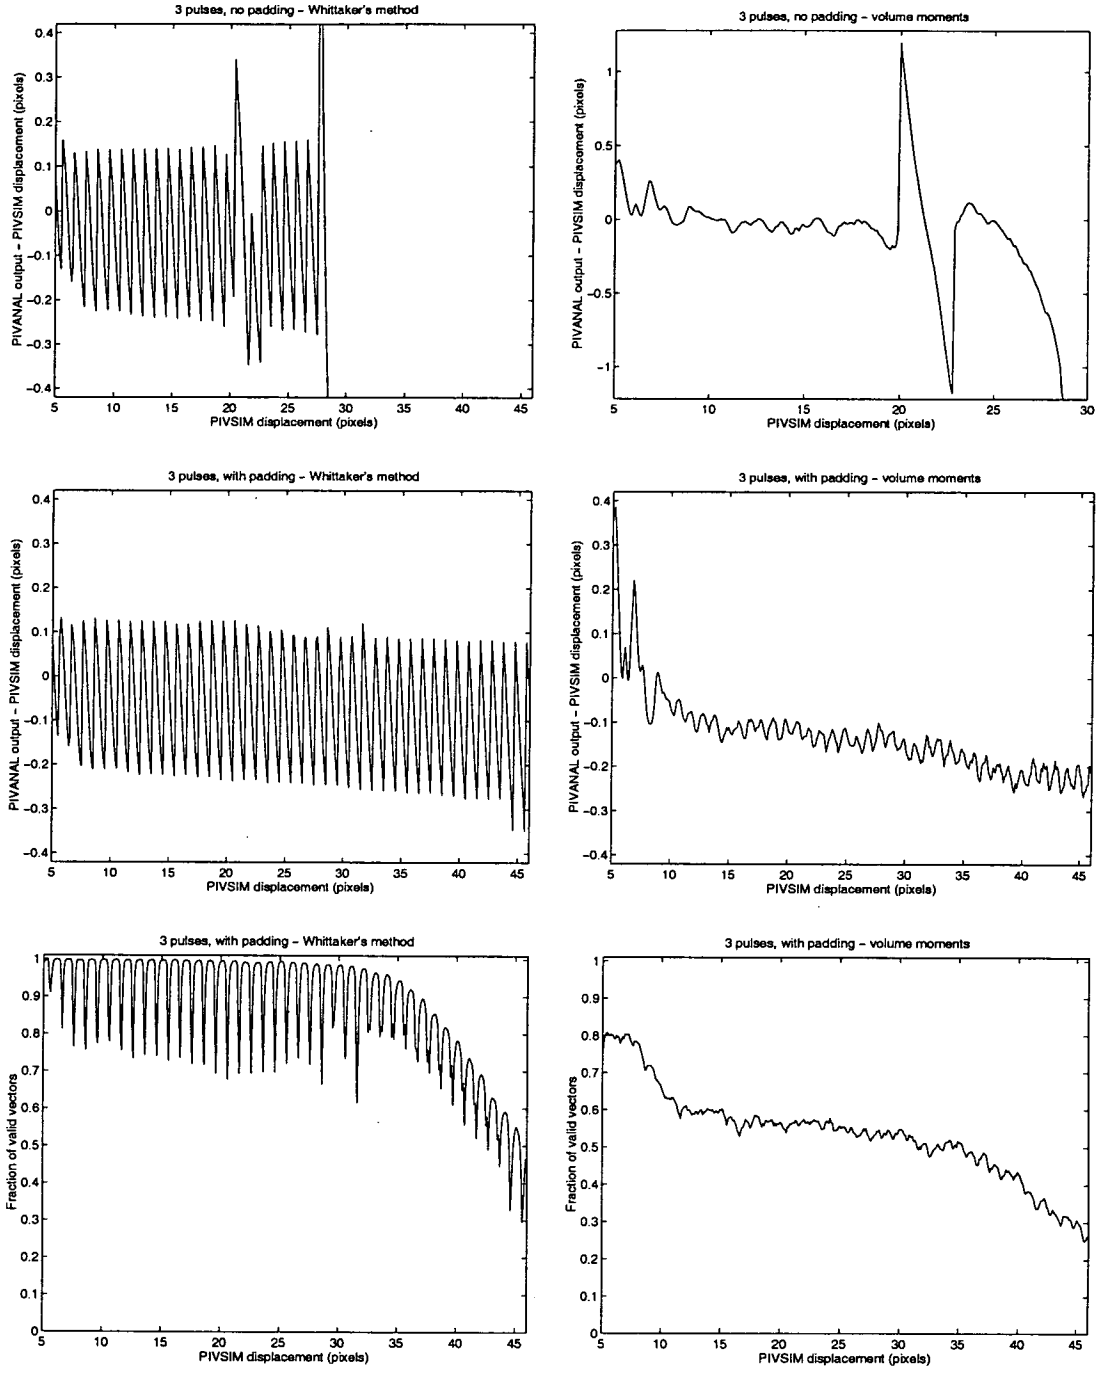


Figure 5.16: Analyses using Whittaker's reconstruction (left) or volume moments (right) to determine the location of the peak value of the autocorrelation function. These analyses were obtained using unpadded FFTs (top row) or fully-padded FFTs (central row). The yields obtained when padding was used are shown in the bottom row.

The results shown in Figure 5.16 reveal that errors due to wrap-around noise were present even when the peak-finding algorithm was changed. This means that the errors found earlier, when padding was not used, cannot simply have been side-effects of Gaussian interpolation. Both of the alternative analysis methods introduce much higher variable (bias) errors than the Gaussian method (*cf* Figures 5.3 and 5.8) and cannot be recommended.

Interpolation using volume moments was clearly the worst choice. The spread of measured displacements was much greater, so the tolerance applied when accepting measured vectors as valid had to be relaxed: vectors were included in the volume-moment averages if their displacements were less than 0.6 pixels from the mean. Even so, the yields were no better than 50–80%. Considering the unpadded case, results to ± 0.1 pixels could be obtained for simulated horizontal displacements in the range 8–19 pixels, but horizontal displacements between 20 and 22 pixels (*i.e.* where wrap-around errors were expected) gave very poor results: all measurements in this range came out at about 64/3 pixels. The average measured displacement was respectable if padded FFTs were used, but the yield of valid vectors was still extremely poor.

The Gaussian method introduced a variable error of about ± 0.04 pixels, whereas Whittaker's reconstruction gave results varying by about ± 0.18 pixels. The maximum error due to wrap-around, for displacements of 20–22 pixels, was approximately ± 0.4 pixels in each case. This was more noticeable when the Gaussian method was used, simply because the 'background' variable error was much smaller. Yields of valid vectors were very good at integer values of the simulated displacement (*e.g.* about 98–100% for $5 < d_x < 30$ pixels) but varied considerably between these values (*e.g.* the yield at a simulated displacement of $(24.6, 0)$ was only 70%). A simple test of Whittaker's reconstruction using only integer values of displacement would show it in a very good light, but these simulations, using displacements at intervals of 0.1 pixels, show that there are hidden problems.

One problem with the higher variable errors found using Whittaker's reconstruction, with or without padding, is that the measured displacements are signi-

ficantly biased towards integer values. In this work, simulations were made at horizontal displacement intervals of 0.1 pixels. If the resulting measurements are plotted in one dimension as points on a number line, unbiased measurements should be evenly spaced, lying close to the simulated values. Figure 5.17 shows one-dimensional plots of the mean horizontal displacement obtained from PIVANAL, when Gaussian interpolation (top) and Whittaker's reconstruction (bottom) were used to estimate the centre of each autocorrelation peak to sub-pixel accuracy.

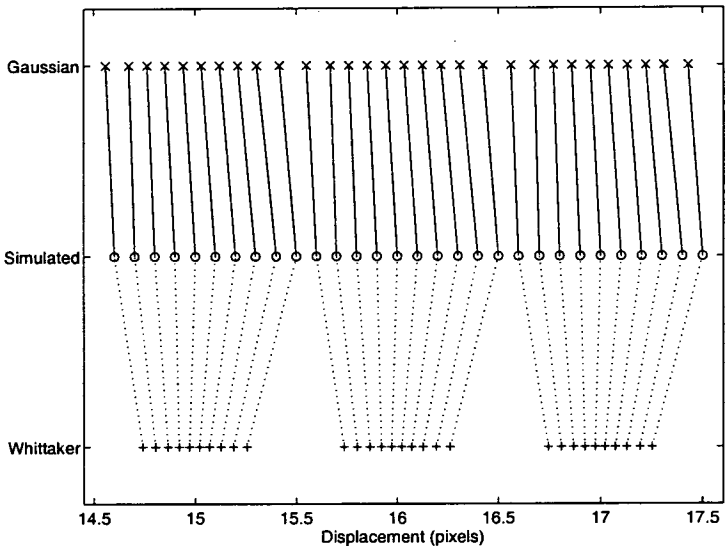


Figure 5.17: Thirty evenly-spaced simulated horizontal displacements in the range 14.6–17.5 pixels are shown in the middle row, using circles to plot these points on a number line. The displacements calculated by PIVANAL from these simulated displacements are also shown, for analysis with unpadded FFTs using Gaussian interpolation ($\times$ characters) or Whittaker's reconstruction ('+' signs).

The horizontal displacements ranged from 14.6 to 17.5 pixels in magnitude and at intervals of 0.1 pixels, so the measured points should be evenly spaced. It can be seen that the distribution of points obtained using Gaussian interpolation was approximately uniform, with only small biases towards integer values such as 16 pixels and away from intermediate values such as 15.5 and 16.5 pixels. When Whittaker's reconstruction was used instead, simulated displacements of 14.6–15.5 pixels (*i.e.* 15.05 ± 0.45 pixels) produced measurements in the range 14.74–15.26 pixels (15.00 ± 0.26 pixels), a considerable biasing towards integer values. A simulated displacement of 15.5 pixels was measured as 15.26 pixels, whereas 15.6 pixels was measured as 15.74 pixels.

More details of this biasing towards integer values can be seen in Figure 5.18. In each scatter graph in this figure, for each of the 2209 displacements calculated by PIVANAL using fully-padded FFTs and the given interpolation method, the horizontal component of the measured displacement was plotted against the corresponding vertical displacement. For example, two of the measurements from a simulated displacement of $(24.5, 0)$ using Gaussian interpolation were $(24.45, 0.15)$ and $(24.43, -0.01)$; these vectors are plotted as points on Cartesian axes.

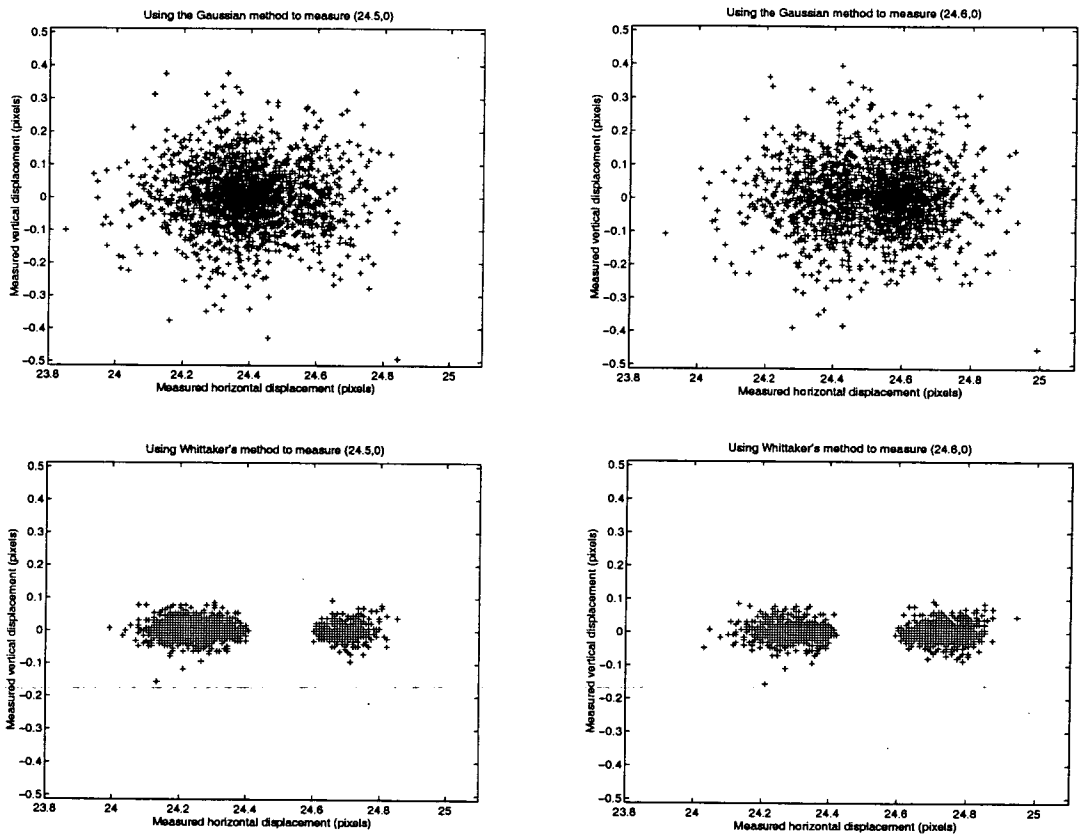


Figure 5.18: Plots of the raw data produced by PIVANAL using fully-padded FFTs, for simulated horizontal displacements of 24.5 pixels (left) or 24.6 pixels (right). Either Gaussian interpolation (top) or Whittaker's reconstruction (bottom) was used, in order to determine the location of peaks to sub-pixel accuracy.

When Gaussian interpolation was used, the mean values for the two measurements were 24.39 and 24.52 pixels, respectively. In each case, over 95% of the measured vectors were valid; in the former case, 95.7% of the vectors were less than 0.33 pixels away from the point $(24.39, 0)$ and it can be seen that most were less than 0.1 pixels away from the mean. These measurements were slightly less than the simulated displacement but showed relatively little bias.

Using Whittaker's reconstruction, the simulated displacements were $(24.5, 0)$ and $(24.6, 0)$ but the average displacements detected by PIVANAL were $(24.26, 0)$ and $(24.70, 0)$. It can be seen in Figure 5.18 that the individual measurements fell into two separate camps, which had radii of about 0.1 pixels and centres close to the points $(24.3, 0)$ and $(24.7, 0)$. In the former case, there were very few vectors within 0.1 pixels of the simulated displacement, $(24.50, 0)$. The main change as the simulated horizontal displacement was increased from 24.5 pixels to 24.6 pixels was that the populations of each camp changed. The region centred on $(24.3, 0)$ was the more populated at the lower displacement (83.9% of the measured vectors) whereas the region at $(24.7, 0)$ was popular at the higher displacement (69.6% of the vectors). Further tests revealed that the two regions were equally populated at a simulated displacement of about $(24.57, 0)$.

Table 5.5 shows the populations of the regions in which the measurement might lie, as the horizontal displacement increases from $(23.2, 0)$ to $(25.1, 0)$. This range covers the rise and fall of the population found in a region based around a measured displacement of $(24.0 \pm 0.3, 0)$. The neighbouring regions, based at $(23.0 \pm 0.3, 0)$ and $(25.0 \pm 0.3, 0)$, show decreasing and increasing populations, respectively. As the image separation increased, the centres of all three regions moved slowly in the direction of increasing horizontal displacement; also, a number of vectors were transferred out of a low-displacement regions and into the adjoining higher-displacement region.

In this work, the first step towards determining the mean was to take the average of all measurements less than 1 pixel away from the expected value. This included all points shown in the plots in Figure 5.18. When Whittaker's reconstruction was used, this mean was biased towards the more populated region. In the final stages, points more than 0.33 pixels away from the previous estimate were ignored and a new mean was calculated, until convergence was achieved. If the initial estimates of the means were good or if the tolerance was high at first then gradually decreased, the final results were likely to include all points in the more populated region and no points from the less populated region, as the centres of these regions were 0.5 pixels apart. Thus, the validation procedure

d_x (pixels)	Number of vectors	at point	Number of vectors	at point	Number of vectors	at point
23.2	2190	(23.1, 0)	1	(23.7, 0)		
23.3	2186	(23.1, 0)	4	(23.7, 0)		
23.4	2156	(23.2, 0)	34	(23.7, 0)		
23.5	1840	(23.2, 0)	350	(23.7, 0)		
23.6	780	(23.3, 0)	1407	(23.8, 0)		
23.7	192	(23.3, 0)	1999	(23.8, 0)		
23.8	53	(23.3, 0)	2140	(23.8, 0)		
23.9	16	(23.3, 0)	2179	(23.9, 0)		
24.0	5	(23.3, 0)	2190	(24.0, 0)		
24.1	1	(23.3, 0)	2194	(24.0, 0)		
24.2	1	(23.3, 0)	2193	(24.0, 0)	0	(24.6, 0)
24.3			2188	(24.1, 0)	3	(24.7, 0)
24.4			2154	(24.2, 0)	35	(24.7, 0)
24.5			1853	(24.2, 0)	337	(24.7, 0)
24.57			1111	(24.2, 0)	1082	(24.7, 0)
24.6			777	(24.3, 0)	1412	(24.7, 0)
24.7			201	(24.3, 0)	1992	(24.8, 0)
24.8			48	(24.3, 0)	2146	(24.8, 0)
24.9			11	(24.3, 0)	2184	(24.9, 0)
25.0			8	(24.3, 0)	2191	(25.0, 0)
25.1			2	(24.3, 0)	2194	(25.0, 0)

Table 5.5: For PIV analyses in which the centre of the autocorrelation peak is calculated to sub-pixel accuracy using Whittaker's reconstruction, this table summarises the regions in which the measured vectors might lie.

selected the most populated region of radius 0.33, ignoring all vectors in a nearby reasonably-populated region. The method did make it apparent that something was amiss, however, because the yields of valid vectors fell sharply.

In similar tests on PIV bitmaps obtained by experiment, Whittaker's method was found to display biasing towards integer values in both the horizontal and the vertical direction. This was found when the horizontal component of the displacement measured by PIVANAL was plotted against the vertical component. Points should have been evenly distributed, but in fact appeared in clusters around integer values, such as $(10, 10)$ and $(10, 11)$. Thus the biasing observed when using Whittaker's reconstruction cannot simply be due to the decision to simulate images with a Gaussian intensity profile, which might have favoured a Gaussian interpolation method. Alternatively, the Gaussian interpolation method

makes certain assumptions about the nature of particle images and the resulting autocorrelation peaks; naturally, simulated images with a Gaussian intensity profile match these assumptions well and give good results, but this behaviour is also shown by PIV images obtained experimentally, so the assumptions seem justified.

The results shown in this section were obtained using autocorrelation analysis with either fully-padded or unpadded FFTs. It appears that padding the FFTs offers few advantages, as far as the bias error is concerned.

5.7.4 Image diameter: no zero-padding

Prasad *et al.* (1992) investigated the effect of image size on PIV measurements. At very small image diameters, the position of each image centre is uncertain. This uncertainty can be reduced if the image diameter (d_τ) is increased.

For example, if $d_\tau \ll 1$ pixel, the location of each particle image can only be determined to the nearest pixel. As the peak corresponding to the image displacement will also occupy only 1 pixel in the autocorrelation plane, the displacement can only be measured to the nearest pixel and measurements of displacement will be 'biased' towards the nearest integer (so bias errors of up to ± 0.5 pixels may be found). If instead $d_\tau = 4$ pixels, as it has been so far in this chapter, the images will cover at least a 4 pixel $\times$ 4 pixel region of the interrogation area, the appropriate peak in the autocorrelation plane will span several pixels and interpolation methods can find the displacement to sub-pixel accuracy.

The other effect of image diameter on accuracy is a random error, due to errors in film grain, digitisation, *etc.* Adrian (1991) suggests that this error is proportional to the image diameter and gets worse for larger images. This suggests that a compromise diameter should be sought, as done by Prasad *et al.* (1992).

Using three images of each seeding particle, simulations have been made and analysed without padding, for horizontal displacements of 1–15 pixels and various

image sizes in the range 0.5 pixels $< d_\tau <$ 7.0 pixels. This models the effect of changing the seeding particles in a flow, whilst leaving unchanged parameters such as their concentration, the number of images of each particle, the displacement between the centres of successive images and the size of the interrogation areas.

Plots such as those shown in Figure 5.19 were obtained. The first plot shown, for $d_\tau = 1.0$ pixels, shows that the bias error can be an absolute error of as much as ± 0.5 pixels if the image diameter is too small. The second plot, for $d_\tau = 1.1$ pixels, shows a significant amount of the transient error found when $0 < |\boldsymbol{d}| < |\boldsymbol{d}|_{min}$: huge errors more than 2 pixels in magnitude were obtained, for some non-zero displacement 'measurements'. The axes in the other plots were chosen so that only a small portion of the transient error is visible.

The bias error is a periodic function of the simulated displacement, measuring the tendency for measurements at all non-integer values of displacement to be biased towards the nearest integer value. It appears that this systematic error oscillates with a period of 1 pixel. Some random noise is also present, however, as fluctuations are seen. The bias error is lowest somewhere around $d_\tau = 2.0$ pixels, where the error appears to be mainly random; the bias error is larger at other values of image diameter, because the obvious periodicity returns.

This observation is contrary to the results of Prasad *et al.* (1992), who suggested that the bias error was a strictly-decreasing function of the image diameter. From their results, random errors were expected to dominate the error signal at larger image diameters, so the expected result at $d_\tau = 3.0$ pixels, say, would be a random signal of higher magnitude. Instead, a periodic signal was found, which was less random than the signal at $d_\tau = 1.2$ pixels (Figure 5.19).

The results from a number of simulations similar to those shown in Figure 5.19 are summarised in the following figures. The variable component of the average error in each measurement is shown in Figure 5.20, whilst the average constant offset error and the minimum resolvable displacement are shown in Figure 5.21.

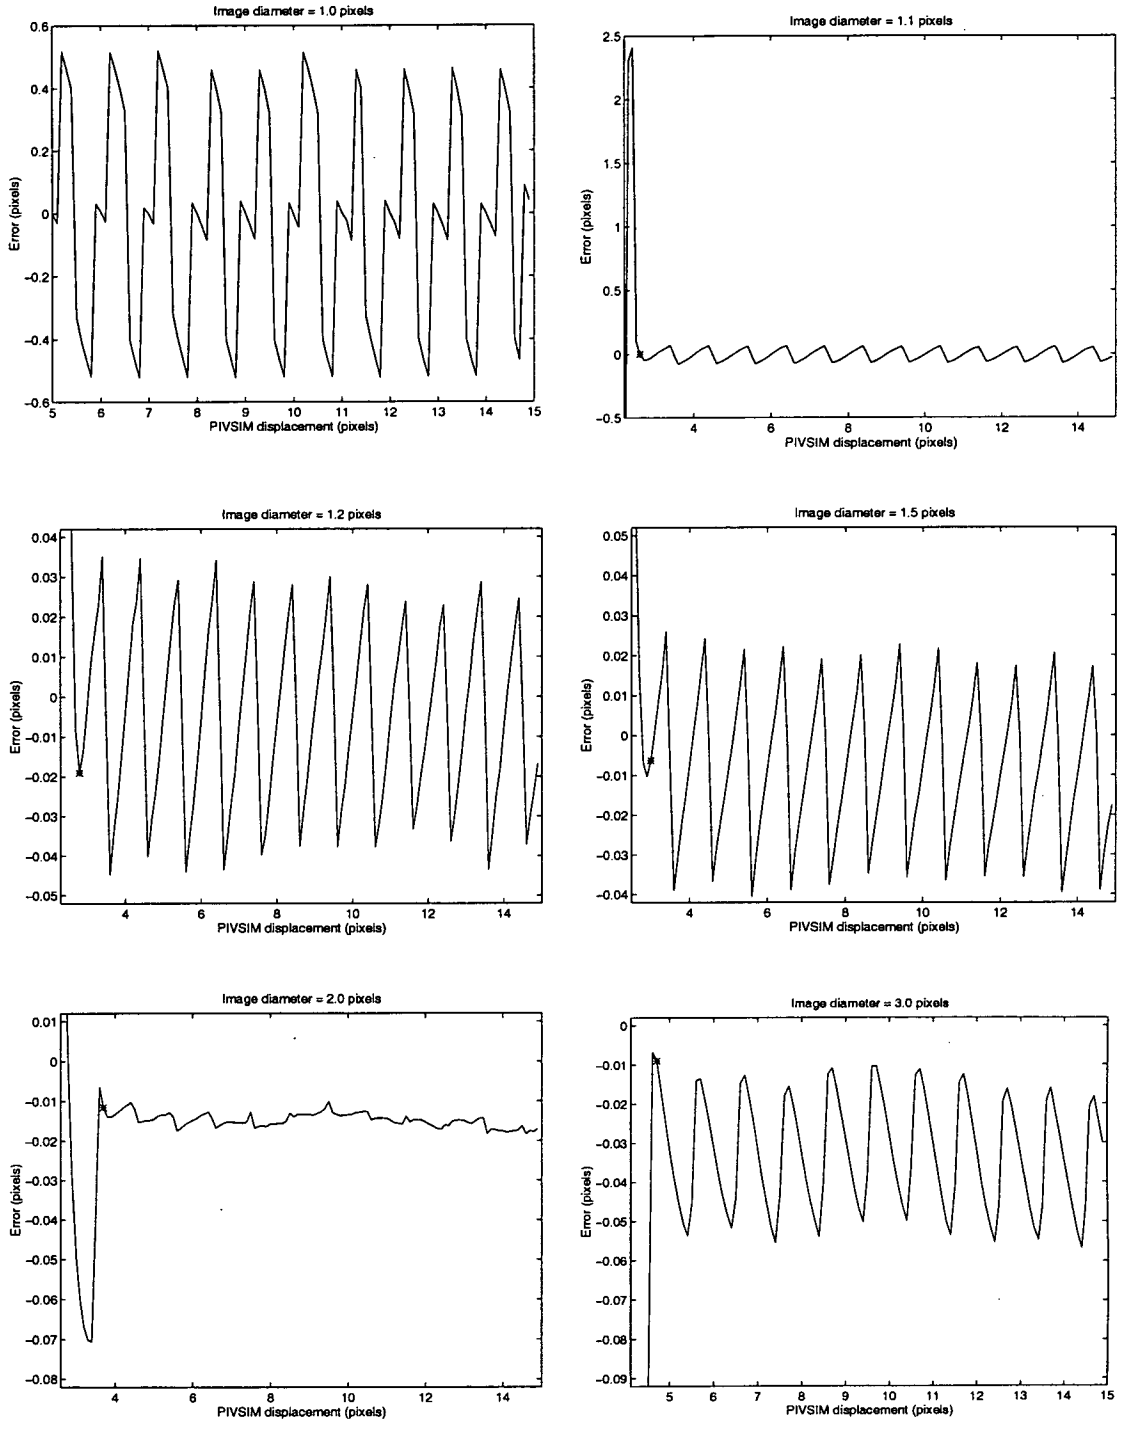


Figure 5.19: The average error in the displacements measured by PIVANAL using unpadded FFTs, plotted against the simulated displacement, for a range of entirely-horizontal displacements. From left to right and top to bottom, the e^{-2} image diameter of each particle image was 1.0, 1.1, 1.2, 1.5, 2.0 or 3.0 pixels. The minimum image separation for accurate results is indicated by an asterisk.

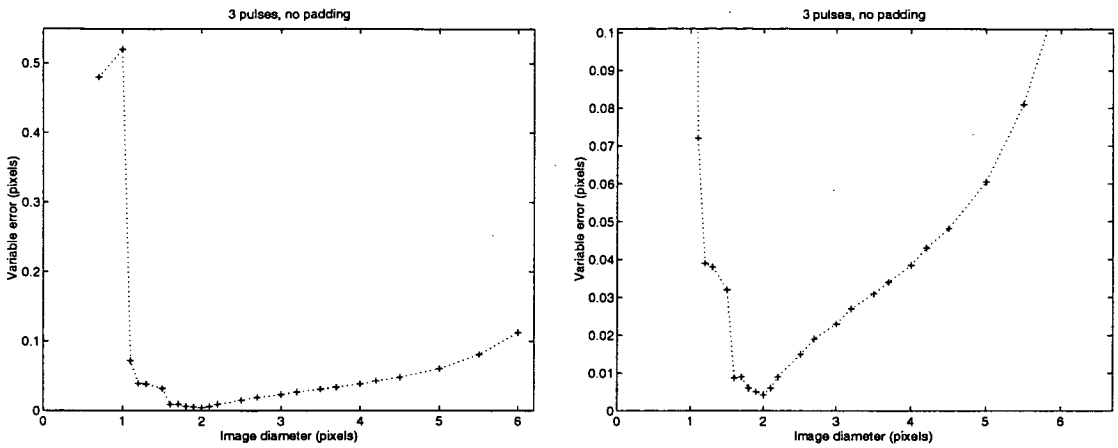


Figure 5.20: The magnitude of the variable error in the displacement measured by PIVANAL (using unpadded FFTs and 63×63 interrogation areas), plotted as a function of the simulated e^{-2} image diameter. A full-scale plot is shown on the left, with a close-up of the lower part of the curve shown on the right.

It was found that the magnitude of the variable error increased markedly at image diameters less than about 1.1 pixels. The variable error was about ± 0.072 pixels at $d_\tau = 1.1$ pixels, but this leapt to ± 0.52 pixels at $d_\tau = 1$ pixel or less.

The magnitude of the variable error was less than 0.01 pixels if 1.6 pixels $< d_\tau <$ 2.2 pixels. At a typical image displacement of 10 pixels, this represents a variable error of only 0.1%; under these conditions, PIV can be an extremely accurate velocity measurement technique. In fact, variable errors as low as this mean that other sources of error are likely to be limiting; examples include errors due to out-of-plane motion or due to seeding particles not following the flow perfectly, as well as the effects of velocity gradients within each interrogation area.

Variable errors of ± 0.01 pixels are very much smaller than the errors due to wrap-around noise found earlier, which were as large as ± 0.4 pixels. It is clear that wrap-around errors must be avoided if a PIV system is to give optimal results.

The ideal image diameter, minimizing the magnitude of the variable error, was found to be very close to 2.0 pixels. At this value, the average error in the measured displacements was -0.014 ± 0.004 pixels, *i.e.* the analysis systematically under-read the displacement by 0.014 pixels and a periodic variable error of ± 0.004 pixels was found on top of this. This variable error is about 0.1% of

the smallest simulated displacement which could be determined accurately at this image diameter (3.7 pixels) and is even less significant at larger displacements.

A couple of features seen in Figure 5.20 differ from the results of Prasad *et al.* (1992), presented in their Figure 13. First of all, the error was found to increase very steeply as d_τ was reduced from 1.1 pixels to 1.0 pixels (almost a step change); Prasad *et al.* obtained a smoother curve.

Further, the variable error found in this work for $d_\tau > 2$ pixels was not proportional to d_τ. A reasonable straight-line fit was obtained for points in this range, but this line passed through the point $(1.78, 0)$ instead of the origin. Its equation was: variable error $= 0.018\, d_\tau - 0.032$ pixels.

A constant offset error was found as well as the above variable error, in that the variable error signal usually oscillated a small distance beneath the zero-error axis. The magnitude of this under-read error appeared to increase according to almost the square of the simulated image diameter, as shown on the left-hand side of Figure 5.21. Ignoring the two points where $d_\tau < 1$ pixel, the data were correlated by the best-fit equation: offset $= -(d_\tau/18.6)^{1.88}$ pixels.

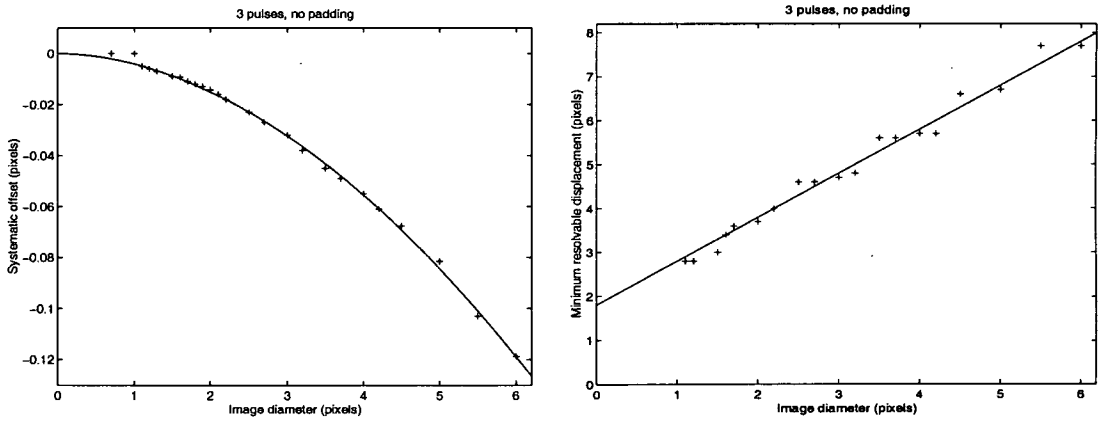


Figure 5.21: For measurements calculated by PIVANAL using unpadded FFTs to analyse simulations in which three images of each particle were simulated: Left: the average systematic error in the PIV measurement, plotted as a function of the simulated e^{-2} image diameter; Right: The minimum resolvable displacement.

For each simulation along the lines of Figure 5.19, the minimum resolvable image displacement was also determined. This was taken to be the smallest value

of displacement which gave a good yield of valid vectors (*i.e.* the yield was similar to that found at a displacement of about 10–15 pixels) and which lay within the periodic region of the plot (*i.e.* the error at this point was similar to that obtained 1 pixel to its right). In other words, there was little change in the performance of the analysis system if $|\boldsymbol{d}| > |\boldsymbol{d}|_{min}$. These points were indicated earlier in Figure 5.19, where they were marked with an asterisk.

The minimum resolvable displacement obtained for each simulated image size is plotted in the right-hand half of Figure 5.21. The relationship $|\boldsymbol{d}|_{min} = 1.5 d_\tau$, as proposed by Keane & Adrian (1990), might not be appropriate for lower-resolution digital images. The model seems intuitively sensible if the images are much larger than the digitisation grain (pixel) size; if images are at least 1.5 image diameters apart and then analysed by autocorrelation, the self-correlation peak decreases sufficiently for its impact on the intended displacement peak to become negligible. A consequence is that the minimum resolvable displacement will be independent of the interrogation area size and the amount of zero-padding applied to the FFTs.

It seems, however, that resolution can be lost if the images are only a few pixels in diameter. For example, suppose that images are 2.0 pixels in diameter. If the displacement is also 2.0 pixels, all that can be seen is a streak: no gaps between successive images will be apparent and it will not be possible to determine the correct displacement using autocorrelation (Figure 5.22, top). If instead the image diameter is 4.0 pixels, then successive images, each a distance 4.0 pixels apart, are likely to show some signs of being distinct (Figure 5.22, bottom). If each image covers a 4×4 region of the digital recording, the central 4 pixels of this region will be somewhat brighter than the outermost 12 pixels and, in addition, the corner pixels of these 4×4 regions will be somewhat dimmer than the 8 other edge pixels. The correlations between successive images will be dominated by the brightest, centremost values, so there is less need for space between adjacent particle images, if these images are large enough to show a significant variation in intensity from their centres to their edges.

A better model for the minimum displacement requires there to be a gap of about two pixels between successive images, *i.e.* $|\boldsymbol{d}|_{min} = d_\tau + 2$ pixels. The

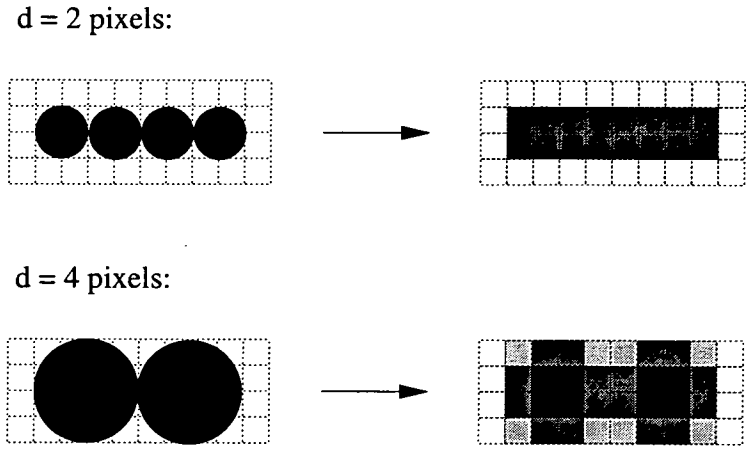


Figure 5.22: The possible appearance of a digital PIV recording, when images are 2.0 pixels or 4.0 pixels in diameter and one image diameter apart.

best-fit line drawn in the right-hand half of Figure 5.21 is for a gap of 1.8 pixels; a gap of 2 pixels seems sufficient to ensure that almost all of the recommended values of $|\boldsymbol{d}|_{min}$ will actually be resolvable. In some cases, *e.g.* at $d_\tau = 4$ pixels, the self-correlation peak and the displacement peak were not completely separate at the minimum resolvable displacement (5.7 pixels), but this did not appear to affect the measurements.

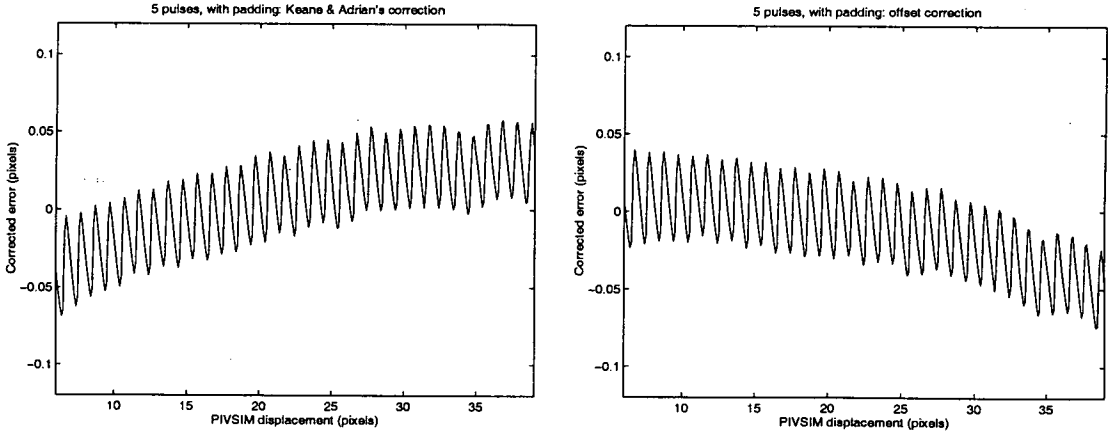


Figure 5.23: The 'corrected' average error in the measured displacement (calculated with fully-padded FFTs for the case of five images of each 4.0 pixel-diameter particle), plotted against the simulated displacement. The measurements were corrected either by multiplying the average measured displacement by $(1+d_\tau^2/d_I^2)$ (left) or by adding 0.07 pixels to each average measured displacement (right).

In their work on optimising PIV parameters, Keane & Adrian (1990) calculated that the measured displacement in an autocorrelation analysis would be underread by a constant factor of $(1 + d_\tau^2/d_I^2)$. This was not found to be the case for

measurements obtained using PIVANAL, where the under-read appeared more likely to be a constant amount, rather than a constant multiplicative factor. Some examples of these correction methods are shown in Figure 5.23.

If the image displacement is less than about 25 pixels, it seems that measurements subjected to an offset error correction have a remaining periodic error only: the corrected results oscillate fairly evenly about the true values within this region. Applying Keane & Adrian's correction factor seems to overcompensate for the slight downward drift which becomes apparent at large displacements in the right-hand half of Figure 5.23. Results corrected in this way still under-read small displacements but now over-read the larger displacements.

5.7.5 Image diameter: with full zero-padding

The corresponding results obtained when full zero-padding was applied to the FFTs were not fundamentally different to those obtained without padding.

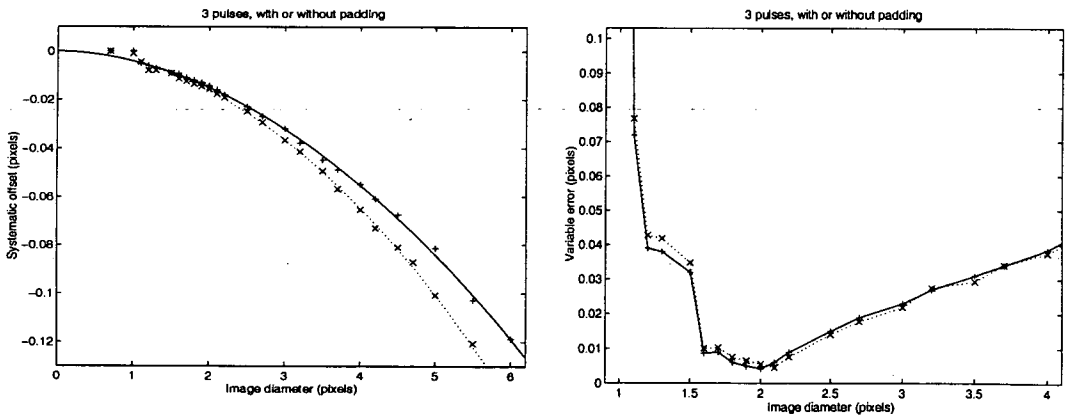


Figure 5.24: For measurements with either no zero-padding ('+' characters; continuous curve) or full zero-padding ('×' characters; dotted curve), using three images of each particle: Left: the average offset error in the PIV measurement, plotted against the simulated e^{-2} image diameter; Right: The variable error.

The constant offset error was somewhat larger when padding was used (see Figure 5.24, left). The plot of variable error against image diameter had a minimum at an image diameter of about 2.1 pixels. For any image diameter in the

region 2 pixels $< d_\tau <$ 6 pixels, the variable error for analysis with padding was marginally lower than that for analysis without padding (Figure 5.24, right).

5.7.6 Magnification/resolution

Prasad *et al.* (1992) found that the variable error was proportional to d_τ. This suggests that there is little point in increasing the resolution of digitisation beyond the value at which the variable error is minimized. If d_τ is increased from 2 pixels to 4 pixels by digitising an existing photographic negative at twice the resolution without altering the seeding particles or the analysis system, the variable errors and the image separations are also doubled; the relative error is unchanged.

On the other hand, if the imaging resolution is doubled, changing the analysis parameters would be advisable. If the image separation doubles, the characteristic length of each interrogation area should also be doubled, so as to allow the larger image-plane displacements to be measured with a comparable yield of valid vectors. This increase in interrogation area size will cause analysis times to increase and may also alter Prasad *et al.*'s constant of proportionality (which they say is a function of the particular analysis method, including the interrogation parameters). Both analysis systems would have to be assessed before such an increase in resolution can be recommended.

For example, the image diameter might be 3.0 pixels when photographed at an image:object magnification of 10,000 pixels per m of object plane, with the sizes of the digital PIV photographs and the image displacements making 63 × 63 pixel interrogation areas the best choice (these conditions were simulated in Figure 5.19). If the whole of this recording is instead photographed at double the magnification, the images of the same seeding particles will become 6.0 pixels in diameter and 127×127 pixel interrogation areas will now be more suitable (to cope with the increased separation of successive images and to avoid the poorer yields of valid vectors which would be found with 63 × 63 pixel interrogation areas). In order to obtain the actual displacements in the object plane, the new measured displacements will need to be divided by a magnification of 20,000 pixels per m.

To assess whether this increase in magnification is beneficial, it is necessary to know whether the variable error after the increase in magnification (*i.e.* that for 6.0 pixel images, analysed using 127×127 pixel interrogation areas and converted into an object-plane displacement by dividing by a magnification of 20,000 pixels per m) is less than it was originally (*i.e.* for 3.0 pixel images, analysed using 63×63 pixel interrogation areas and then divided by 10,000 pixels per m).

The effect of recording PIV images at twice the magnification has therefore been investigated. In effect, similar simulated PIV object planes have been created as were used in Figure 5.19 but these were 'photographed' onto a digital image plane at twice the resolution (*i.e.* with twice as many pixels in the digital image per meter in the object plane). All image plane distances were therefore scaled by a factor of two, *e.g.* the simulated horizontal displacements covered the range from 2 pixels to 29.8 pixels at intervals of 0.2 pixels, image diameters doubled[1] and interrogation areas became 127×127 pixels in size, spaced 84 pixels apart.

Some results are shown in Figure 5.25. In these plots, the image diameter quoted is that after $\times 2$ magnification, *i.e.* the first graph in Figure 5.25 is for particles which appear to be 1.1 pixels in diameter only after the increase in resolution; these particles would have been 0.55 pixels in diameter in terms of Figure 5.19. Further, the simulated and measured displacements have been halved here, to allow direct comparisons to be made with previous results. This adjustment also halves the amplitude of the variable error signal and the constant offset error, whilst doubling the frequency of the bias error oscillations. This correction takes into account the change in the image:object magnification (pixels of digital image per meter of object plane) used to calculate object-plane velocities.

Graphs summarising the results from runs similar to those shown in Figure 5.25 are shown in Figures 5.26 and 5.27. In these figures, the results for analyses without padding using 63×63 pixel interrogation areas are shown using '+' characters, connected by a continuous line or curve. When doubling the resolution,

[1]This doubling of recorded image size will only be obtained in practice if the particle images are neither diffraction-limited nor comparable to the grain size of the film (equation 4.2).

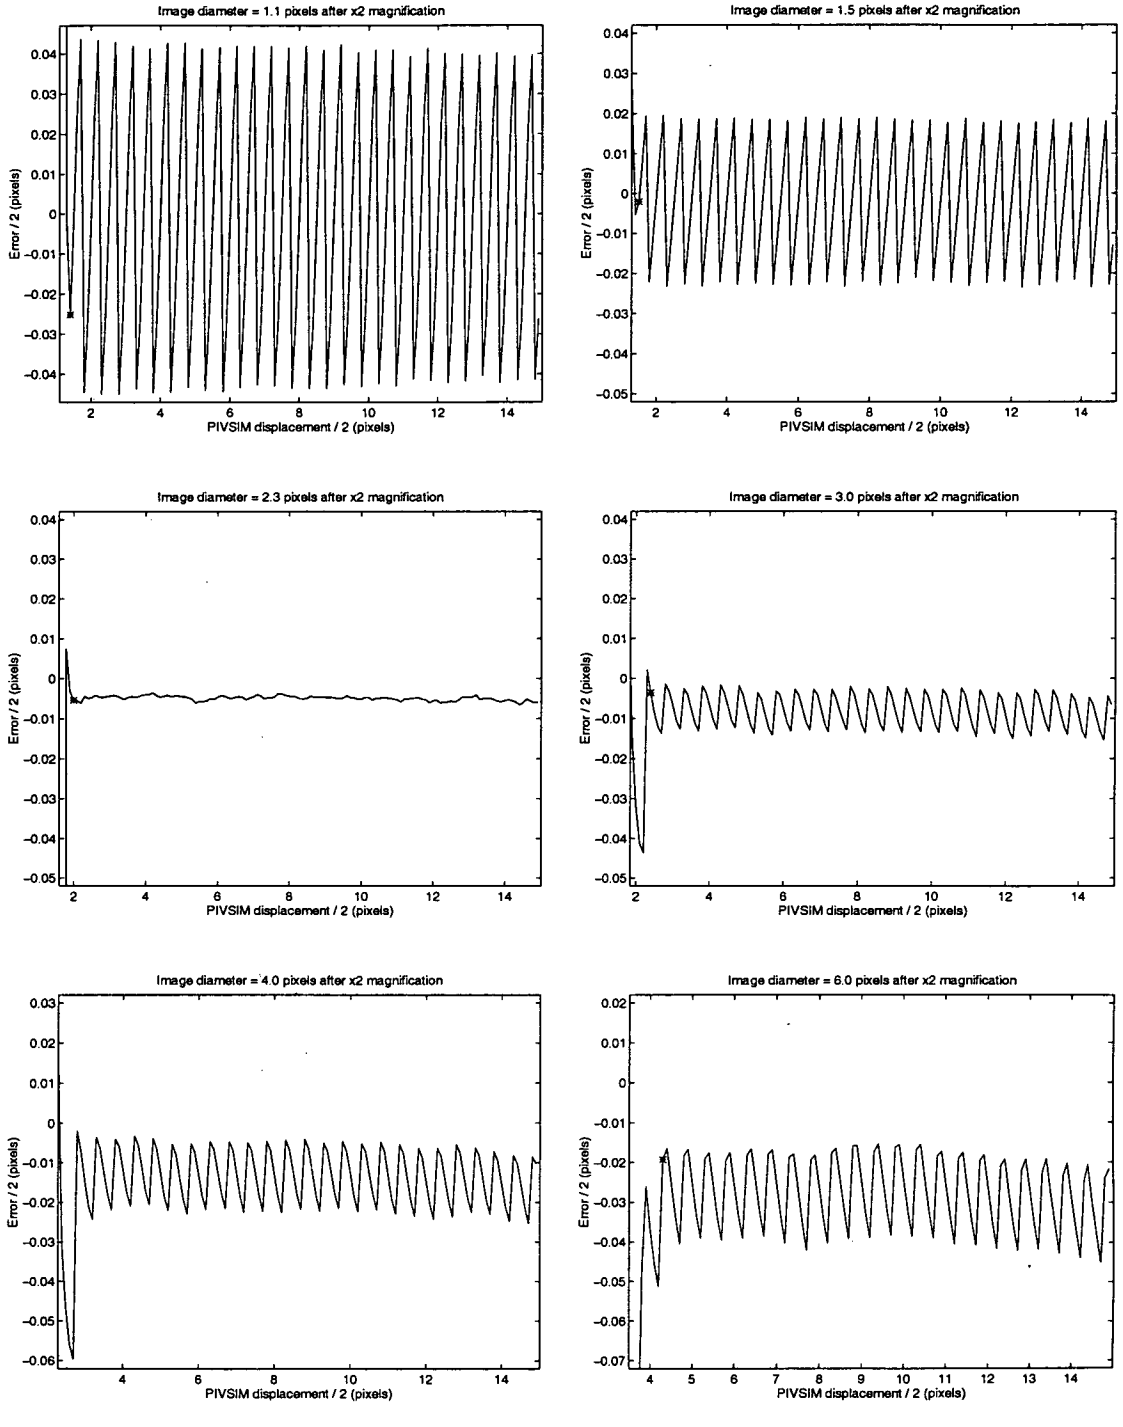


Figure 5.25: Plots of half the average error against half of the simulated horizontal displacement, obtained by simulating PIV images at double the magnification and analysing using 127×127 pixel interrogation areas (without padding the FFTs).

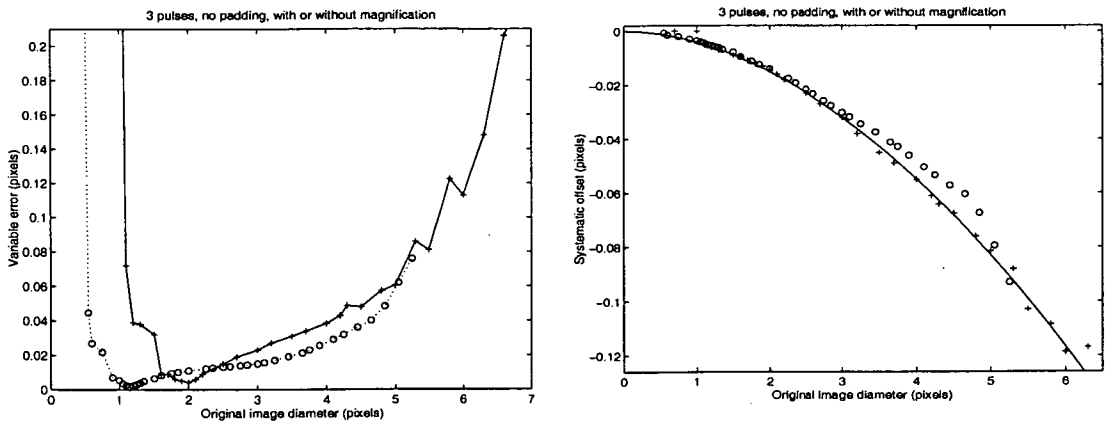


Figure 5.26: For measurements using unpadded FFTs to analyse simulations in which three images of each particle were created at two different magnifications: Left: the variable error in the average measurement, plotted against the simulated diameter at the lower magnification; Right: the average systematic offset error in the PIV measurement, plotted as a function of the simulated e^{-2} image diameter, again at the lower magnification. The points were obtained using 63 × 63 pixel interrogation areas ('+' characters; solid lines) or double the recording resolution then 127 × 127 pixel interrogation areas ('o' characters; dotted lines). The latter results have been halved, due to their doubled magnification.

an image which was originally d_τ in size becomes $2d_\tau$ in simulated size and the analysis takes place using 127 × 127 pixel interrogation areas; the displacements produced by this analysis then need to be halved, to take account of the doubling of the image:object magnification. Points shown by circles in Figure 5.26 are plots, in terms of the *original* image diameter (*i.e.* at the lower magnification) of the variable bias error (left) and the constant offset error (right). In Figure 5.27, the systematic offset error and the minimum resolvable displacement are plotted against the *actual* image diameter, *i.e.* the actual image diameter in the simulated PIV recordings presented to the analysis program.

Comparing the plots of Figures 5.20 and 5.21 to those in Figures 5.26 and 5.27:

- the magnitude of the oscillations was once again lowest when the diameter of the images on the digital PIV recordings fed to the analysis system PIVANAL was about 2.0–2.3 pixels (Figure 5.26, left).

  This optimum image diameter seems to have only a very slight dependence on interrogation area size and the degree of zero-padding. A general rule of

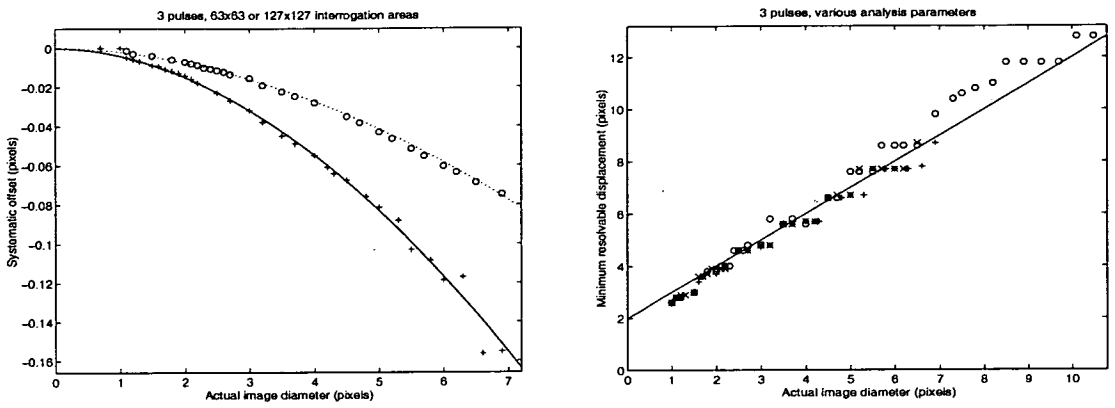


Figure 5.27: Left: the systematic offset error in each measurement, plotted against the actual simulated image diameter; Right: The minimum resolvable image displacement, plotted against the actual simulated image diameter. The analyses used unpadded FFTs to analyse simulations featuring three images of each particle. Either 63 × 63 pixel interrogation areas ('+' characters if zero-padding was not used or '×' characters if it was) or 127 × 127 pixel interrogation areas without zero-padding ('o' characters) were used.

thumb seems to be that the variable error will be lowest if the object:image magnification and the seeding particles are chosen so as to produces images about 2 pixels in diameter. If a first attempt at PIV produces images only 1 pixel in diameter, then using double the recording resolution (and correspondingly-doubled interrogation area dimensions) produces a significant increase in accuracy. If instead the original image diameter is about 4 pixels, then halving the recording resolution will reduce the variable error.

- the magnitude of the variable error oscillations obtained after doubling the resolution was increased if the original image diameter lay between about 1.6 and 2.3 pixels, but was reduced at other image sizes (Figure 5.26, left).

- when padding was not used, images of a given diameter, analysed using 127× 127 pixel interrogation areas, produced only half of the offset error obtained when the same images were analysed using 63 × 63 pixel interrogation areas (Figure 5.27, left). In this graph, the correlations plotted are 1.0× and 0.5× those for images at the lower resolution. The offset error, therefore, appears to be proportional to d_τ^2 and inversely proportional to d_I.

- the constant offset error was almost unaffected by the change in resolution, although some improvement was apparent at large original image diameters

(*e.g.* about 4–5 pixels). For example, images originally 3.1 pixels in diameter and analysed with 63 × 63 interrogation areas became 6.2 pixels in diameter when the resolution was doubled. Doubling d_τ roughly quadrupled the expected offset error, but this was cancelled out by using 127 × 127 interrogation areas (which introduces only half of the offset error found using 63 × 63 interrogation areas, for images of the same diameter) and then dividing by the doubled magnification. For example, the offset error for 3.1 pixel images (analysed with 63 × 63 interrogation areas) was −0.0344 pixels, whereas half of the offset error for 6.2 pixel images (now using 127 × 127 interrogation areas) was about −0.0312 pixels (a slightly reduced error).

Thus, doubling the resolution appears to make little difference to the average amount by which measured object-plane velocities differ from the true values (Figure 5.26, right), but, if anything, errors are marginally reduced.

- the value of the minimum resolvable displacement depended only on the actual image diameter seen by the analysis program: the degree of padding and the size of the interrogation areas had no significant effect (Figure 5.27, right). This suggests that errors in the measurements when $|\boldsymbol{d}| < 2\,d_\tau$ are simply due to overlaps in the autocorrelation plane, between the central self-correlation peak and the intended displacement peak.

  Regression analysis using all available data points gave a best-fit line of $|\boldsymbol{d}|_{min} = 1.10\,d_\tau + 1.56$ pixels. Instead, the line $|\boldsymbol{d}|_{min} = d_\tau + 2$ pixels is plotted in the right-hand side of Figure 5.27, as it is a simpler correlation which is slightly conservative at the recommended values of image diameter (*i.e.* about 2.1 pixels) and only becomes overoptimistic when the image diameter is excessively large (*e.g.* about 8 pixels). The main image separation requirement seems to be that there should be about 2 pixels of clear space between any two consecutive images.

- the value of $|\boldsymbol{d}|_{min}$ was effectively reduced by up to 1 pixel by the increase in resolution; smaller displacements could be resolved and the dynamic range was increased. Thus, doubling the resolution might well be worthwhile, even if the offset error and the bias error are not improved.

For example, a 2.0 pixel image has a minimum image displacement of about 3.7 pixels if good results are to be obtained. If the resolution is doubled, the resultant 4.0 pixel images need to be separated by about 5.6 pixels, which is equivalent to a separation of only 2.8 pixels at the lower recording resolution. Thus, the minimum resolvable displacement is reduced by about 25%, improving the dynamic range of the system.

5.7.7 Interrogation area size

One way to achieve improved yields of accurate vectors in some situations is to use larger interrogation areas to analyse the same bitmap images. It has just been shown that doubling the interrogation area size causes the constant offset error in each reading to halve, so the measured displacements will be more accurate. The variable error will worsen by about 0.01 pixels if the image diameter is between 1.2 and 2.2 pixels, but this should not be too serious.

One attraction of using larger interrogation areas is that, if the image displacements are 'horizontal' and approximately 64/3 pixels to 64/2 pixels in magnitude, wrap-around errors will occur (and yields of valid vectors will be poor) if the usual 63 × 63 pixel interrogation areas are used without zero-padding the FFTs. However, larger 128 × 128 interrogation areas should not give rise to wrap-around errors at these displacements, if only three images of each particle are recorded.

Figure 5.28 shows that this is indeed the case. In the left-hand graph, the difference between the PIVANAL output and the PIVSIM input shows the usual periodic variation, with little change in performance observed at simulated horizontal displacements between about 6 pixels (the minimum resolvable displacement) and about 41 pixels (when wrap-around errors arise, based on a displacement of 128/3 pixels). In contrast, the unpadded results using 63 × 63 pixel interrogation areas were only accurate from 6 to 20 pixels (Figure 5.3). The conclusion is that the maximum analysable image displacement is approximately doubled if the interrogation area side-length is doubled and no zero-padding is used.

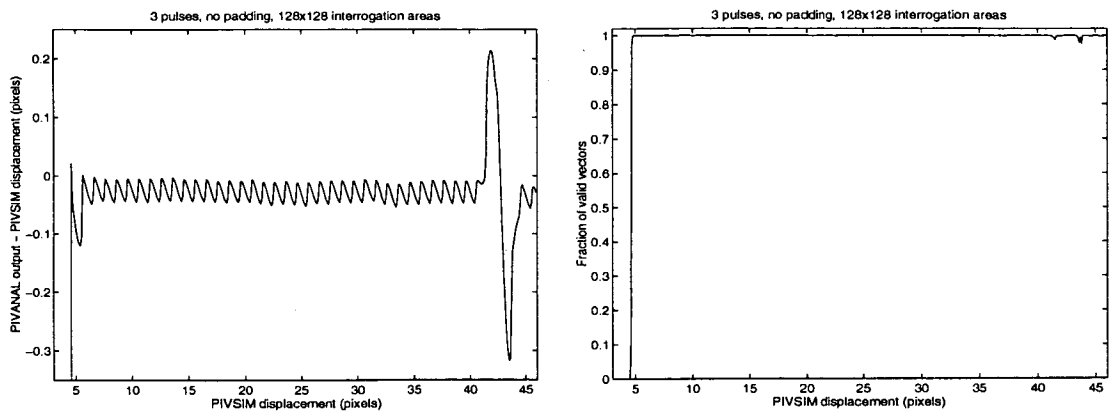


Figure 5.28: The behaviour of the analysis program PIVANAL when 128 × 128 pixel interrogation areas are used to analyse three-image particle tracks, without zero-padding. The simulated image diameter was 4 pixels and the particle concentration was lower than in previous runs. Left: the plot of measurement error against image displacement; Right: the yield of valid vectors.

At the same time, it can also be seen that the yields of valid vectors improved significantly (Figure 5.28, right), being greater than 99.7% for horizontal displacements of about 5–41 pixels. These yields are a considerable improvement over those obtained using 63 × 63 pixel interrogation areas, where the yield fell progressively to about 90% when the image displacement reached about 20 pixels when no zero-padding was used (Figure 5.10) or about 35 pixels when full zero-padding was applied (Figure 5.11). Thus, doubling the side length of the interrogation areas significantly increases the yield of valid vectors.

In this subsection, the simulated particle concentration was halved. The number of simulated tracks was still 8000, but these were placed at random locations within a 2850 × 2850 bitmap image (which had twice the area of the usual 2001 × 2001 bitmaps). The improved yields are probably due to the fact that twice as many particle images were found inside each of the new interrogation areas, as these included four times as many pixels. Thus, the increase in interrogation area size also allows lower seeding particle concentrations to be used.

Similar results have been observed when four images of each particle were simulated, in corresponding conditions. The main difference (compared to the results presented in Figure 5.28) was that two wrap-around peaks were observed, at simulated displacements of about 128/4 and 128/3 pixels. The results were

entirely consistent with the properties of wrap-around errors (as outlined earlier).

Because all of the simulations in this chapter used constant image displacements throughout each simulated bitmap recording, each small interrogation area measured the same image displacement. The use of more than 2000 simulated interrogation areas meant that the output vector maps contained over 2000 independent measurements of each displacement, hopefully providing enough readings for the overall results to have a sound statistical basis. When making recommendations for general PIV use from the observations in this chapter, however, it must be noted that experimental PIV recordings will include velocity variations, both across the flow and within each interrogation area.

Therefore, the general recommendation that larger interrogation areas perform better needs a caveat: larger interrogation areas might not cope well when velocity gradients are present and might not provide enough spatial resolution to reveal small-scale turbulence structures. The effects of velocity gradients have not been considered here.

5.7.8 Seeding particle concentration

Another use of PIVSIM would be to create bitmap images using displacements of 10–20 pixels, for a known (but varied) number of particles of constant diameter. PIVANAL could then measure the displacements, using constant analysis conditions. This would reveal the effect of seeding particle concentration on the yield of valid vectors and on the accuracy of the measured vectors (*e.g.* the constant offset error and the variable error). For example, accuracy might be optimised at a certain seeding concentration, but the yield of valid vectors might be greatest when a different concentration is used. Also, different optimal concentrations will probably be found when using more or fewer images of each particle; these concentrations could be compared to the recommendations of Keane & Adrian (1990, 1991), who advised that there should be 16 image pairs in each interrogation area.

These simulations have not yet been carried out, however, and must instead count as recommendations for future work.

5.8 Other simulations using $d_\tau = 2$ pixels

Replacing the 4-pixel seeding particles (as used earlier in this chapter) with new particles 2 pixels in diameter means that the seeding particles become more likely to follow the flow under investigation. On the other hand, these smaller images will scatter less incident light, so either the laser power will need to be increased or the camera will need to use a wider aperture setting (or a faster film).

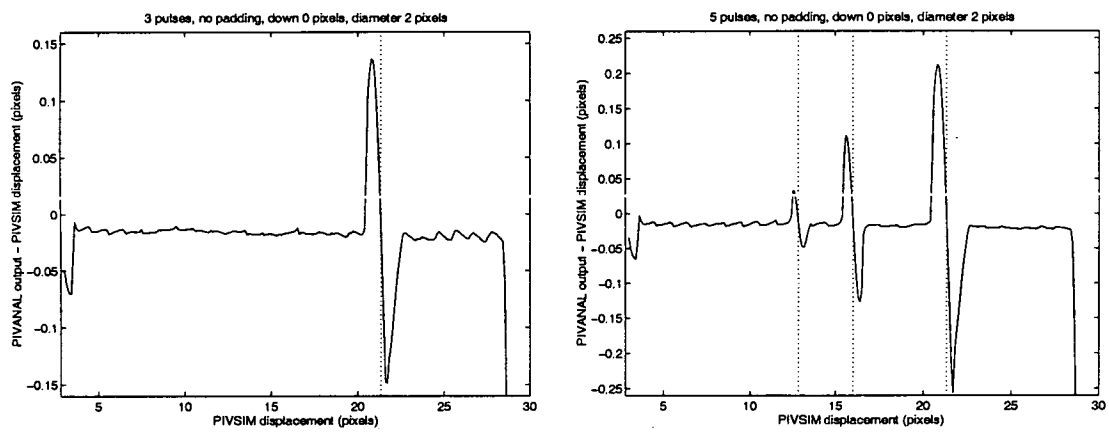


Figure 5.29: The error in the horizontal measurement of displacement, calculated by PIVANAL using unpadded FFTs and $d_\tau = 2$ pixels, plotted against the simulated horizontal displacement, for simulations using three (left) or five (right) images of each particle.

The simulations presented earlier have been repeated with the image diameter set to 2.0 pixels, which appears to be the optimum value for 63 × 63 pixel interrogation areas. Considering simulations based on 3 or 5 images of each particle, it can be seen in Figure 5.29 that the variable errors are much less than those found using $d_\tau = 4$ pixels (Figures 5.3 and 5.6) and the wrap-around errors have about half their previous magnitude.

The yield contours were recalculated for the case of three images of each particle, both with and without padding (Figure 5.30). Higher yields were ob-

tained here with $d_\tau = 2$ pixels, compared to those with double the image size (Figure 5.15). When full zero-padding was applied to the FFTs, the smaller images meant that displacements could be made about 15% larger without causing the yield to drop below 90%. With unpadded FFTs, the improvement was more noticeable; in fact, the relatively-small size of the autocorrelation plane became a limiting condition when padding was not applied.

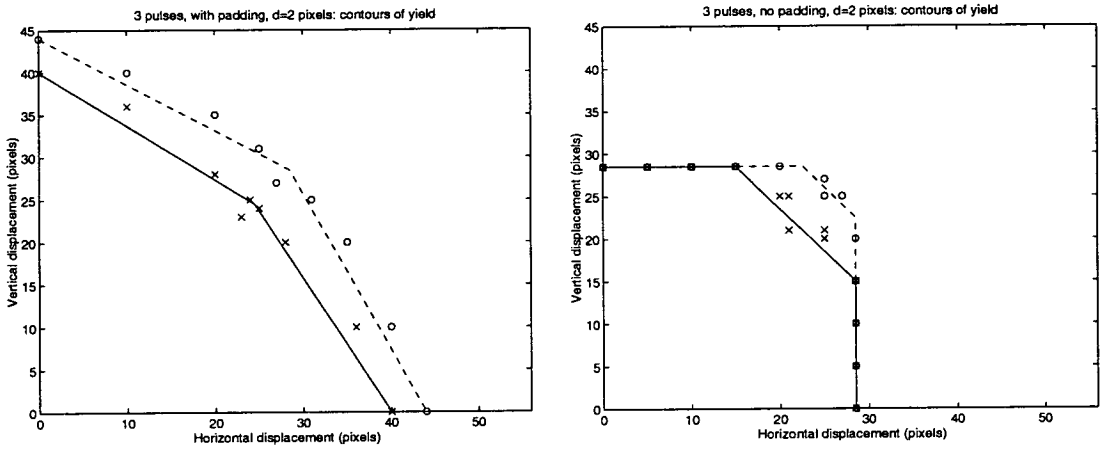


Figure 5.30: Contours of the yield of valid vectors, as measured by PIVANAL using padded (left) or unpadded (right) FFTs and $d_\tau = 2$ pixels, plotted on a graph of the simulated image displacement. The contours are for yields of 90% (solid lines, $\times$ signs) and 80% (dashed lines, 'o' symbols).

5.9 Conclusions

No windowing of the image data is performed by PIVANAL. There were no velocity gradients in these simulations, so the only problems which might have arisen were from partial images at the edges of interrogation areas. The measurements of these uni-directional flows were usually very accurate, so these problems cannot be too serious. Windowing may well be useful when analysing flows with significant velocity gradients, but this has not been tested.

Determining the position of peaks in the autocorrelation plane using Gaussian interpolation gives much more accurate results than the alternatives, namely Whittaker's reconstruction and volume moments. In particular, Whittaker's reconstruction produces results which are heavily biased towards integer values.

For a given number of particle image pairs per interrogation area, higher yields of valid vectors are obtained if the number of images of each particle is increased (reducing the concentration of seeding particles accordingly, so that the number of image-pairs remains constant). The improvement is clear when the number of images of each particle is increased from two to three, but beyond this point there are diminishing returns. Another advantage is that the seeding particles are less likely to distort the flow being measured. A drawback is that longer exposure times mean that the PIV recording is a less instantaneous record of the flow: velocity gradients might be introduced if the flow changes as a function of time. Further, if significant out-of-plane velocity components are present, it might be difficult in practice to record more than 4 or 5 images of each particle.

Analysis with full zero-padding of the FFTs offers a number of advantages, because errors due to wrap-around cannot occur. Given an existing digital PIV recording, the yield of 'valid' vectors is better if padding is applied and the calculated vectors are more accurate (all being very close to the true displacement). Further, the application of full zero-padding doubles the characteristic length of the autocorrelation plane and hence also doubles the magnitude of the largest displacement which can be measured. A drawback is that the analysis times for fully-padded FFTs are about 3–4 times longer.

When padding is not used, vectors must not have a horizontal or vertical component of displacement greater than about 30 pixels in magnitude, because there would otherwise be a risk that a spurious wrap-around signal would be considered to be the displacement peak. For example, a measured displacement of $(31, 2)$ could be due to a true displacement of either $(31, 2)$ or $(33, -2)$ and the surrounding vectors might not make it clear which is correct. In this work, no component of displacement greater than about 29.5 pixels in magnitude could be resolved anyway, due to the autocorrelation peak detection algorithm used by PIVANAL.

Errors in the actual measurement arise when no padding is applied to the FFTs, due to wrap-around noise. Errors of up to 0.4 pixels can be found if the image displacement is nearly parallel to one of the sides of the square interrogation areas

and is approximately $64/M$ pixels in magnitude, where M is any integer between 3 and the number of images of each particle, inclusive. Errors can also occur if the component of displacement in both orthogonal directions is about $\pm 64/M$ pixels.

Better results can be obtained if the image diameter is reduced from 4 pixels to 2 pixels, *e.g.* by using smaller seeding particles. Yields of valid vectors are improved and each measurement is closer to the true value. Similarly, experiments with an image diameter of 1.0 pixels will yield much better results if they can be repeated using double the recording magnification.

For recordings with three images of each particle, regions in which an incorrect measurement is likely to be found are shaded black in Figure 5.31. The regions of satisfactory operation are shown, for analyses with padded or unpadded FFTs and images either 2 pixels or 4 pixels in diameter.

Larger-magnitude displacements can be resolved, if two sides of the square interrogation areas are parallel to the direction of flow. This could well be the case, in practice.

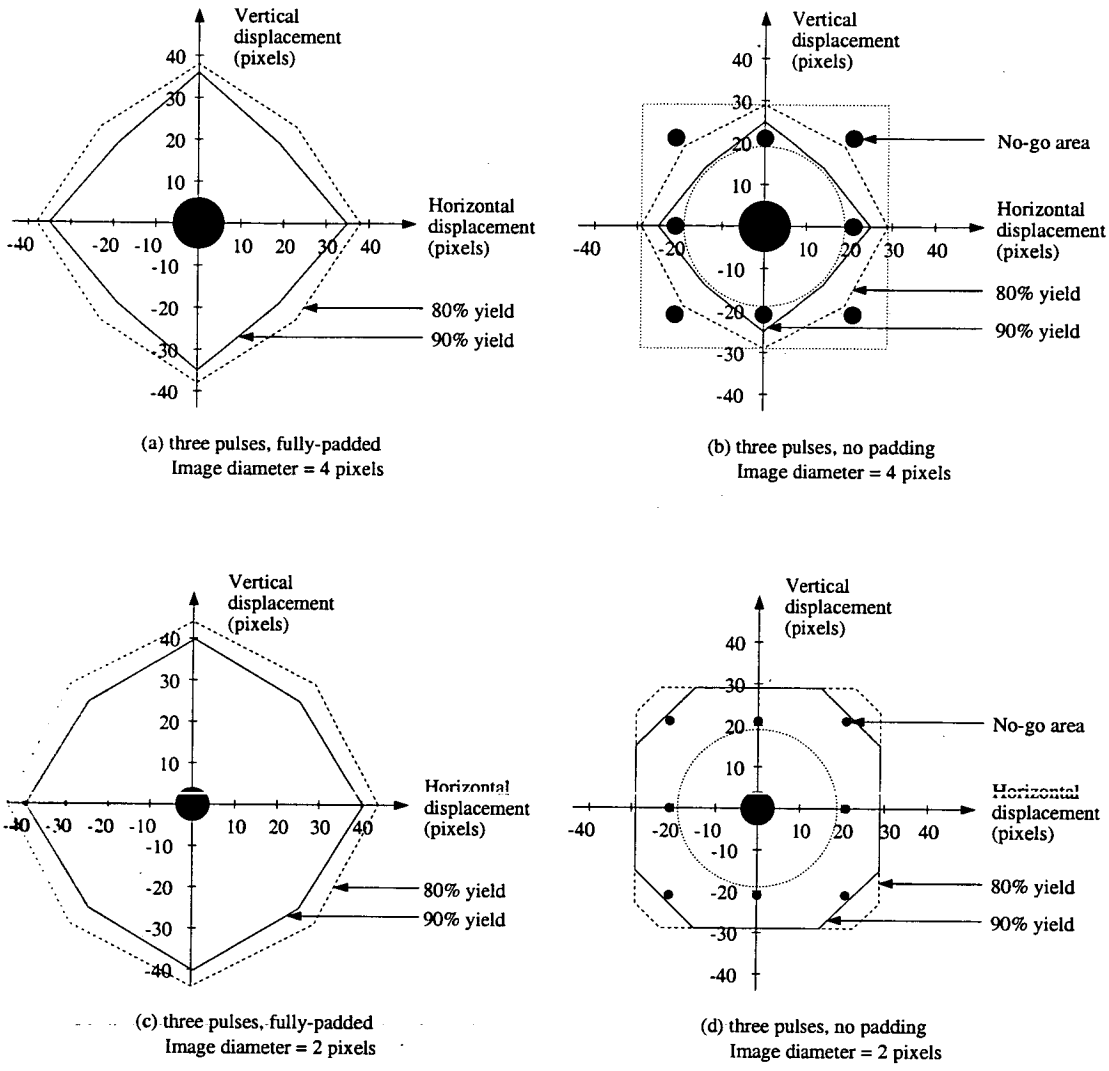


Figure 5.31: If a PIV recording captures three images of each particle, each image being 4 pixels in diameter (top), analysis with fully-padded FFTs (left) gives yields of 90% (solid lines) if the displacement $\boldsymbol{d} = (d_x, d_y)$ is greater than 6 pixels in magnitude and satisfies $d_x + d_y < 36$ pixels or thereabouts. On the other hand, analysis with unpadded FFTs (right) introduces wrap-around errors at many points; these 'no-go' areas are coloured black in these diagrams. The region of safest operation is $6 < |\boldsymbol{d}| < 19$ pixels (dotted circle), which gives yields of over 90% and no wrap-around errors, at the expense of a limited dynamic range. The peak detection algorithm required that $-29.5 < d_x, d_y < 29.5$ pixels (dotted square) for unpadded analyses. Yields improve if smaller seeding particles are simulated (bottom). In both cases, padding the FFTs increases the yields and avoids wrap-around errors, significantly increasing the dynamic range at the expense of requiring more analysis time.

Chapter 6

Loss of pairs and wrap-around errors

As shown in the previous chapter, the yield of valid vectors appears to improve if, at a given displacement, the number of images recorded of each particle is increased. This is due to the greater efficiency at which particle image pairs are created, if more images of each particle are recorded.

For convenience, the following nomenclature will be adopted in this chapter, when considering an 'average' interrogation area:

- images which form an 'image pair' with at least one other image of the same particle will be referred to as 'good' images. These images are separated by the displacement corresponding to the movement of one particle between consecutive illumination pulses and hence contribute usefully towards the determination of this displacement during PIV analysis.

- individual, unpaired images will be referred to as 'bad' images. These images arise when only one image from a particle track is found within the interrogation area, *e.g.* because the other images lie in an adjacent interrogation area or because of out-of-plane movement. During autocorrelation

analysis, no correlations can be made between these images and other images from the same particle track, so they will not help the measurement of the desired displacement. These images can only give rise to spurious particle image correlations, *i.e.* autocorrelation 'noise'.

In the previous chapter, 8,000 tracks were created in each simulated PIV recording, when three images of each particle were simulated. This produced a total of 24,000 particle images in each 2001 × 2001 pixel simulation and an average of 23.8 images per interrogation area. At large displacements, particle tracks were more likely to spread themselves over two or more adjacent interrogation areas, yet the total number of particle images in the whole of each simulated PIV recording (and hence in an 'average' interrogation area) remained constant: for every track where only the last two images lay within a given interrogation area, there was, on average, another track where only the first image lay within this interrogation area. The situation is shown in Figure 6.1, for tracks A and F.

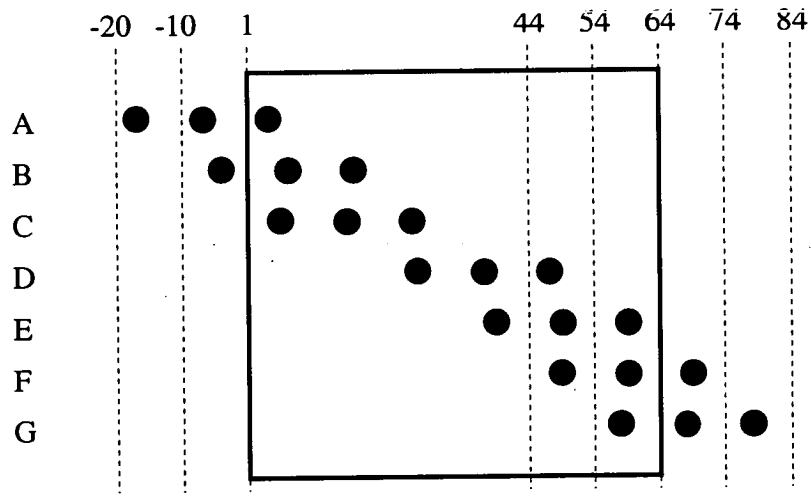


Figure 6.1: For the case of three images of each particle separated by a horizontal displacement 10 pixels to the right, there are a number of ways for one or more images from a particle track to lie within a randomly-placed interrogation area.

The number of particle image pairs which contribute to the measurement of the true displacement is zero for tracks A and G, one for tracks B and F, and two for tracks C–E. The probability of each event is 20/84 for finding either track A or track G, 20/84 for either track B or track F and 44/84 for tracks C–E. Tracks A and G produce only 'bad' images, whereas tracks B–F produce 'good' images.

Given that a particle track produces at least one image in this interrogation area, the expected breakdown of these images is 5/21 individual, unpaired, 'bad' images per $(2 \times 5/21 + 3 \times 11/21)=43/21$ useful, 'good' images, which form an image pair with at least one other image of the same particle. At a seeding concentration of about 24 particle images per interrogation area, about 21 of these images would form part of at least one particle image pair and 3 other unpaired images would also be present, making no useful contribution to the measurement of the simulated displacement. As the number of unpaired images increases, so does the probability that the measurement is inaccurate.

6.1 Useful images from horizontal flow records

In this and the following sections, it will be assumed that the separation of successive images is 'horizontal', *i.e.* parallel to two of the sides of each interrogation area. This is likely to be the case for PIV images recorded using the image-shifting technique, as well as many natural flow situations (*e.g.* flow in a pipe).

In any interrogation area, it is possible that only some pixels from certain particle images will be considered in a 64×64 pixel interrogation area, because the image centres happen to be too close to one of the edges of the interrogation area. In this chapter, however, it will be assumed that all images lie entirely within their given zone. This is most likely to be the case if the particle images are small. Partial images have not been considered here, because they will complicate the analysis.

The probability that several successive particle images all lie inside the same interrogation area decreases as the image separation is increased. Let this displacement be entirely horizontal and d pixels in magnitude, for convenience. With three images of each particle, consider the situation illustrated in Figure 6.2. This is a more general version of Figure 6.1.

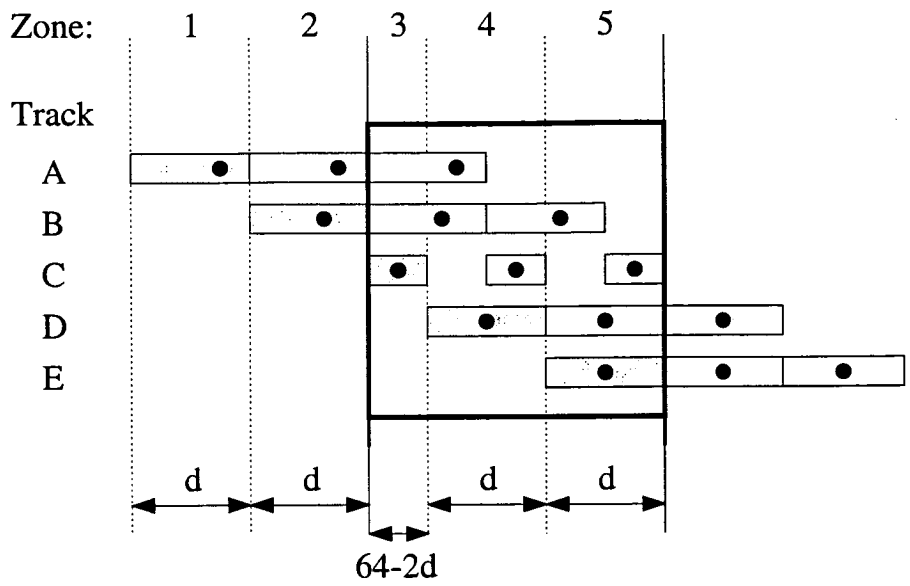


Figure 6.2: For three images of each particle, displacements to the right less than 32 pixels in magnitude and 64 × 64 pixel interrogation areas, consider the position of the first image in a random particle track, relative to the position of an interrogation area containing at least one image of that particle. There are five possibilities to consider, represented by the five lightly-shaded rectangles.

For a horizontal displacement satisfying $d < 32$ pixels, there are five different ways in which at least one image of a three-image track can lie within the interrogation area: the first image of the track can lie in any one of zones 1–5 shown in Figure 6.2. Zones 1 and 2 are each of width d and are located next to one another, to the left of the interrogation area. Zones 4 and 5 are also of width d and are placed so that zone 5 touches the right-hand edge of the interrogation area and the right-hand edge of zone 4. The remainder of the region of interest, namely that part of the interrogation area which does not lie in zones 4 or 5, is zone 3 and its width is $64 - 2d$ pixels. The total width of the region covered by these five zones is $64 + 2d$ pixels. The situation shown in Figure 6.2 is valid only if $d < 32$ pixels: at higher displacements than this, the width of zone 3, $64 - 2d$ pixels, will become negative and a different diagram will need to be drawn.

Given that at least one image of the particle is found inside the interrogation area, the probability that the particle track begins in zone 1, as is the case for track A in Figure 6.2, is $d/(64 + 2d)$. The rectangle representing track A within zone 1

is lightly-shaded, to indicate that it is the locus of the first image of the particle. For track A, two images of the particle lie to the left of the interrogation region and only one image lies inside it; this is true regardless of the exact location of the images within each box. Track E in Zone 5 is similar, but in this case only the first image of the particle lies within the interrogation area. For tracks beginning in zones 2 and 4, two images lie inside the interrogation area: the last two images of track B and the first two images of track D. Finally, as the right-hand edge of zone 3 is a distance $2d$ away from the right-hand edge of the interrogation area, all three images of track C lie within the interrogation area.

Tracks A and E both give rise to only one image inside the interrogation area. Therefore, still assuming that at least one image from a randomly-placed track lies within the interrogation area, the probability that only a single image is found is $2 \times d/(64 + 2d)$. Similarly, the probability of finding exactly two images is also $2 \times d/(64+2d)$ and the probability of finding all three images is $(64-2d)/(64+2d)$.

The ratio of useful, paired ('good') images to lone, unpaired ('bad') images is:

$$\begin{aligned} \text{Good:bad} &= 3 \times \left(\frac{64 - 2\,d}{64 + 2\,d}\right) + 2 \times \left(\frac{2 \times d}{64 + 2\,d}\right) : 1 \times \left(\frac{2 \times d}{64 + 2\,d}\right) \\ &= 64 \times 3 - 6\,d + 4\,d : 2\,d \\ &= 192 - 2\,d : 2\,d. \end{aligned}$$

Thus, if F is the fraction of images which are part of at least one pair, then:

$$F(3 \text{ pulses}, d < 32) = \frac{192 - 2\,d}{192} = 1 - \frac{d}{96}.$$

It can be seen that $F = 1$ at $d = 0$ pixels. F then decreases linearly until $F = 2/3$ at $d = 32$ pixels (the upper limit for results from this method).

For displacements to the right which exceed 32 pixels, see Figure 6.3. Once again, the total width to consider is $64+2d$ and there are five different possibilities. As before, zone 1 lies a distance d to $2d$ to the left of the interrogation area and

zone 5 is the right-hand-most d pixels of the interrogation area. Zone 4 is now what remains of the interrogation area and has width $64 - d$. Zones 2 and 3 cover, between them, the region 0 to d pixels to the left of the interrogation area; these regions are separated by a line drawn $2d$ pixels to the left of the right-hand edge of the interrogation area. The combined width of zones 1, 2 and 3 (and also zones 3, 4 and 5) is $2d$.

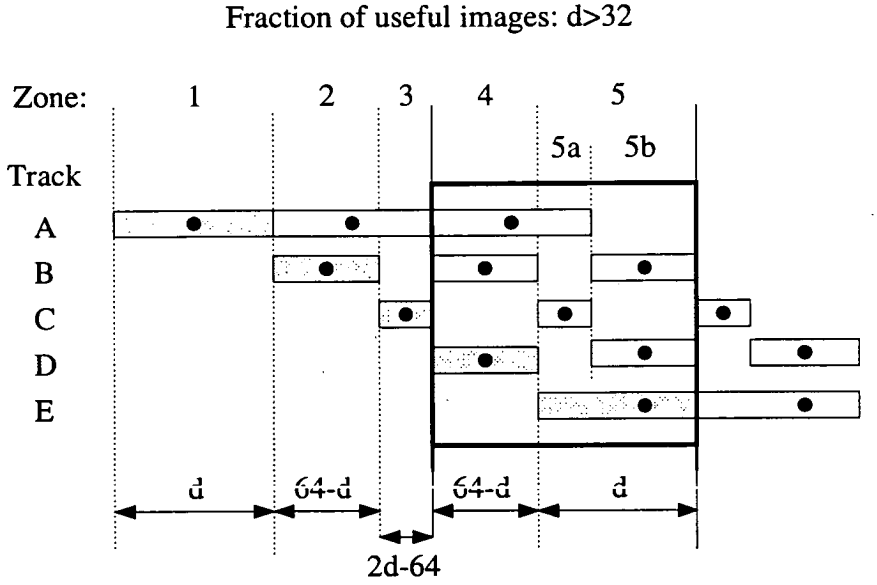


Figure 6.3: For the case of three images of each particle, separated by displacements to the right more than 32 pixels in magnitude, with 64×64 pixel interrogation areas, consider the location of the first image of a particle track (lightly-shaded rectangles), relative to the position of an interrogation area containing at least one image of that particle.

If the first image of a randomly-placed particle lies in zone 1 or zone 5, then only one image of that particle can lie within the interrogation area; the probability of each event is $d/(64 + 2d)$. If the first image of a randomly-placed particle lies in zone 2, then the second and third images of this particle will lie within the interrogation area; the width of zone 2 is $64 - d$ pixels, so the probability of track B being found is $(64 - d)/(64 + 2d)$. If the track begins in zone 3, then the second image, in zone 5a, is too close to the centre of the interrogation region for any further images to be found; the width of zone 3 is $d - (64 - d) = 2d - 64$ pixels, which is positive if $d > 32$ pixels (as was assumed). Track D represents a similar situation to track B and also gives rise to an image pair.

The ratio of paired images to lone, unpaired images if $d > 32$ pixels is:

$$\begin{aligned}\text{Good:bad} &= 2 \times \left(\frac{2 \times (64 - d)}{64 + 2\,d}\right) : 1 \times \left(\frac{2 \times d}{64 + 2\,d}\right) + 1 \times \left(\frac{2\,d - 64}{64 + 2\,d}\right) \\ &= 64 \times 4 - 4\,d : 4\,d - 64 \\ &= 256 - 4\,d : 4\,d - 64.\end{aligned}$$

Thus F, the fraction of paired images, is:

$$F(\text{3 pulses, } d > 32) = \frac{256 - 4\,d}{192} = \frac{64 - d}{48}.$$

It can be noted that $F = 2/3$ at $d = 32$ pixels (which is the same result as was calculated above for the case $d < 32$ pixels) and that F decreases linearly as d increases, reaching zero at $d = 64$ pixels. F is therefore a continuous function which consists of two lines, intersecting at $d = 32$ pixels.

Similar work has been carried out to determine the fraction of images which are part of at least one image pair, as a function of image separation d, for the cases of 2, 4 and 5 images of each particle. When four or more images of each particle were simulated, different sketches along the lines of Figure 6.2 needed to be made for $0 < d < 64/5$ pixels, $64/5 < d < 64/4$ pixels, *etc*, but in each case the same expression was found to apply until d reached 32 pixels. It would be interesting to investigate further why this might be the case.

Number of images of each particle	Fraction paired if $0 \leq d \leq 32$ pixels	Fraction paired if $32 \leq d \leq 64$ pixels
2	$\frac{128-2\,d}{128}$	$\frac{128-2\,d}{128}$
3	$\frac{192-2\,d}{192}$	$\frac{256-4\,d}{192}$
4	$\frac{256-2\,d}{256}$	$\frac{384-2\,d}{256}$
5	$\frac{320-2\,d}{320}$	$\frac{512-2\,d}{320}$
...	...	...
N	$\frac{64\,N-2\,d}{64\,N}$	$\frac{2\,(N-1)\,(64-d)}{64\,N}$

Table 6.1: Calculations of the fraction of paired images.

The results are summarised in Table 6.1 and suggest a general form which can be used if N images are recorded of each particle. The two expressions for a particular number of images of each particle always give the same result at $d = 32$ pixels, so strict inequalities need not be used. Further, if a general interrogation area size d_I is used instead of the 64 pixel value used here, yields can be expressed in terms of a dimensionless image displacement, d/d_I.

Plots of the fraction of paired images are shown in Figure 6.4, as a function of the number of particle images and their separation. The yield is always 1 at $d = 0$ and zero at $d = 64$ pixels. A change in slope can occur at $d = 32$ pixels.

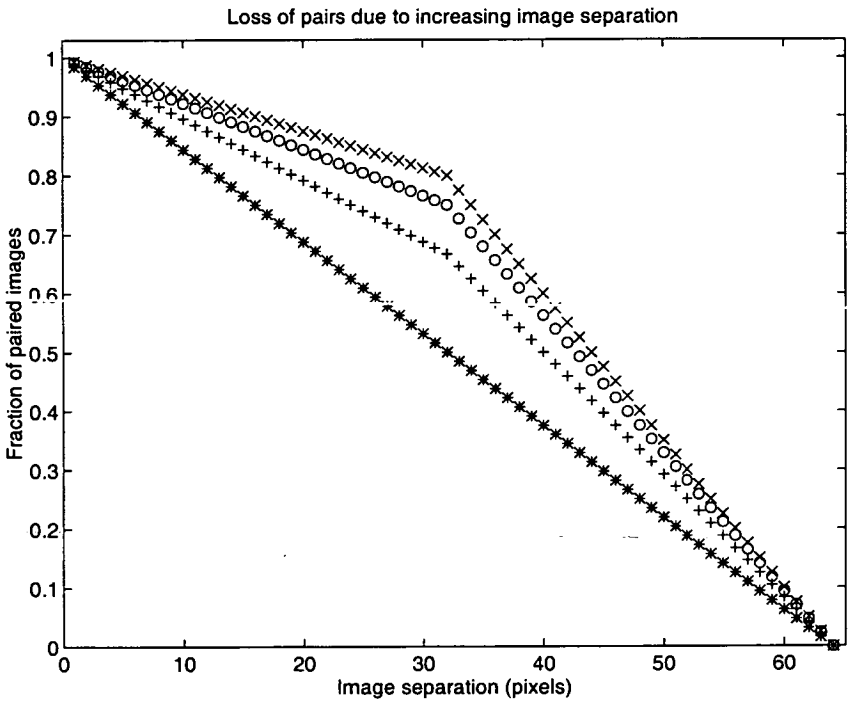


Figure 6.4: Plots of the fraction of useful images within an average interrogation area, as a function of the image separation, for the cases of two ($*$), three ($+$), four ($\circ$) and five ($\times$) images of each particle.

The general form of the expression for the fraction of paired images can easily be derived for the case where $0.5 \leq d/d_I \leq 1$. Following Figure 6.3 above:

If there are N images of each particle, then zone 1 will be a distance $(N-2)d$ to $(N-1)d$ to the left of the interrogation area and the total width to consider will be $d_I + (N-1)d$.

The key region for image pairs is that labelled zone 4 in Figure 6.3. Its width is $(d_I - d)$. There are N ways of selecting one image from the track of N images and placing it in this zone. However, the case where the last image occurs in this zone must be discounted, because that will not actually give rise to an image pair. Therefore there are $(N-1)$ zones of width $(d_I - d)$ in which the particle track can start, if an image pair is to be obtained. The probability of getting a pair of good images is therefore $(N-1)(d_I - d)/(d_I + (N-1)d)$. The probability of getting a lone image is 1 minus this, *i.e.* $(2(N-1)d - d_I(N-2))/(d_I + (N-1)d)$.

Therefore the ratio of good to bad images if $0.5 \leq d/d_I \leq 1$ is:

$$\begin{aligned} \text{Good:bad} &= 2 \times \left(\frac{(N-1)\,(d_I - d)}{(d_I + (N-1)\,d} \right) : 1 \times \left(\frac{2\,(N-1)\,d - d_I\,(N-2)}{d_I + (N-1)\,d} \right) \\ &= 2\,(N-1)\,(d_I - d) : 2\,(N-1)\,d - d_I\,(N-2) \end{aligned}$$

and so the required fraction of paired images, F, is given by:

$$\begin{aligned} F(N \text{ pulses}, d \geq d_I/2) &= \frac{2\,(N-1)\,(d_I - d)}{2\,(N-1)\,(d_I - d) + 2\,(N-1)\,d - d_I\,(N-2)} \\ &= \frac{2\,(N-1)\,(d_I - d)}{d_I\,N} = 2\,\frac{(N-1)}{N}\left(1 - \frac{d}{d_I}\right). \end{aligned}$$

At $d = d_I/2$, the fraction of paired images is $(2(N-1)/N) \times 0.5 = (N-1)/N$. Therefore, if the relationship between the fraction of paired images and d continues to consist of two straight lines either side of $d = d_I/2$, then the line for $0 \leq d \leq d_I/2$ goes through the points $(0, 1)$ and $(d_I/2, (N-1)/N)$. Its equation must be:

$$F(N \text{ pulses}, d \leq d_I/2) = 1 - \frac{1/N}{d_I/2}d = \frac{d_I\,N - 2\,d}{d_I\,N} = \frac{N - 2\,d/d_I}{N}.$$

With the seeding concentrations used in the previous chapter, the number of effective particle image pairs at a given displacement was independent of the number of images of each particle, *i.e.* there were as many paired images in each interrogation area at a displacement $d = 32$ pixels (say) when two images of each particle were simulated, as there were when 3, 4 or 5 images were simulated. In each case, the seeding concentration was chosen with this in mind.

As the image displacement increased, the ratio of useful particle images to unpaired images worsened. The results which would have been obtained with 64×64 pixel interrogation areas are summarised in Table 6.2. It can be seen that there were always 16 paired images at $d = 32$ pixels, with rather more additional unpaired images present when few images of each particle were simulated.

Images of each particle:	2	3	4	5
Per average interrogation area:				
Average number of particle images:	32	24	21.3	20
% useful at $d = 32$ pixels:	50	67	75	80
Useful images:noise images:	16:16	16:8	16:5.3	16:4

Table 6.2: The average ratios of paired images to unpaired images.

For example, when three particle images were simulated a horizontal distance 32 pixels apart, an average interrogation area included about 24 different images. Of these, 16 images formed 8 image pairs which contributed to the measurement of the displacement peak after autocorrelation, whilst the other 8 images lay at random locations, serving no good purpose. There was a finite probability that one of these 8 individual images happened to lie close to one of the 16 paired-up images, or that two of these individual, unpaired images happened to be separated by a displacement close to the true displacement. Any interference caused by this would move the detected peak to a slightly different position in the autocorrelation plane and could lead to the measurement being declared invalid when processing the results. This becomes more likely as the displacement is increased:

- the greater number of unpaired images means that a damaging correlation involving one (or more) of the lone, unpaired images becomes increasingly likely to take place by chance; and

- the displacement peak becomes weaker because, on average, there are fewer useful particle images present in each interrogation area. This means that any spurious correlations have a larger influence on the position of the measured displacement. In addition, the actual concentration of useful images will vary randomly across a PIV recording; some interrogation areas will include fewer image pairs than others and thus will be more prone to error.

Comparing simulations at a given displacement, those with fewer images of each particle have a higher proportion of unpaired, individual images and so there is a greater chance that an uncorrelated image causes interference with the main displacement peak. For example, only half of the two-pulse images in an average interrogation area are useful when the displacement is 32 pixels to the right, whereas 80% of the images are useful if five images of each particle are simulated. This explains why the measured yield at a given displacement is greater if more images are recorded of each particle.

Comparing simulations with a given number of particle images, analysed with or without zero-padding, there are more ways in which errors can occur if padding is not applied to the FFTs. For example, if three images are simulated, each separated by a displacement 18 pixels to the right, then about 20 of the 24 images in an average interrogation area will form at least one useful image pair. Four other images also need to be added, at random locations close to the edges of the interrogation area (in either the left-hand-most 18 pixels or the right-hand-most).

Suppose that one of these individual images, arising from a particle track beginning to the left of the interrogation area, is centred at the point $(8, 5)$ and that no images within the interrogation area overlap. When padding is applied, a spurious correlation can only arise if there is another image close to $(26, 5)$; this image must be the first image of a different particle track, all three images of which will lie within the interrogation area. When no zero-padding is used, this possibility can still cause an error in the measurement and there are also wrap-around problems: if another particle track happens to begin a small distance away from either $(18, 5)$ or $(36, 5)$, or if a lone image happens to lie close to $(54, 5)$, then a spurious wrap-around correlation will be made. These extra ways in which errors can be introduced explain why the yield is better when full zero-padding is used.

Having calculated the average fraction of paired images in each interrogation area, as a function of the number of images of each particle and the image separation, it should be possible to calculate the probability that an average interrogation area gives rise to a correct measurement after autocorrelation analysis.

Other factors need to be considered. For example, random variations in seeding concentration mean that some interrogation areas will include more particle tracks than average. As a result, some interrogation areas might include 24 particle images whilst others might only have 12. If fewer images are present, a spurious reading is more likely.

In addition, not all images are useful. Random variations in initial particle locations mean that some interrogation areas with 24 images will consist of 8 tracks of three images, whilst others might include 8 two-image part-tracks and 8 other, unpaired images. As the image displacement between successive illuminations increases, complete tracks become less likely to fit into an interrogation area, so more 'low quality' unpaired images will be found. As a result, the autocorrelation 'signal' corresponding to the image separation is weakened and it becomes more likely that a spurious reading will be made.

Probability distribution functions now need to be used, to describe the variations of the numbers of paired and unpaired images about their mean values. The effects of unpaired images then need to be considered: these images might all lie in harmless locations, but some of them may well produce spurious correlations which distort the fundamental displacement peak and introduce inaccuracies (especially if unpadded FFTs are used). These probability calculations, however, have not been attempted and are left for future work.

6.2 Second harmonic displacement peaks

This method can also be used to assess the relative magnitude of harmonic peaks in the autocorrelation plane. If the displacement of successive images is $(d, 0)$ and three images are recorded of each particle, then the fundamental displacement peak will be at the points $(\pm d, 0)$ in the autocorrelation plane and second harmonic peaks will be found at $(\pm 2\,d, 0)$.

Figure 6.2 was presented earlier, when discussing the case of three images of each particle and an image separation $d < d_I/2$. Given that at least one image from a particular particle track can be found within a particular interrogation area, the probabilities of finding only one image, exactly two images or all three images were determined to be:

$$\begin{aligned} P(\text{1 image, } i.e.\text{ 0 pairs}) &= \frac{2\,d}{d_I + 2\,d} \\ P(\text{2 images, } i.e.\text{ 1 pair}) &= \frac{2\,d}{d_I + 2\,d} \\ P(\text{3 images, } i.e.\text{ 2 pairs}) &= \frac{d_I - 2\,d}{d_I + 2\,d} \end{aligned}$$

From this, the expected number of particle images from each track can be obtained (if $d < d_I/2$). If this number of images per track is N_{ipt}, then:

$$\begin{aligned} N_{ipt} &= 3 \times \frac{d_I - 2\,d}{d_I + 2\,d} + 2 \times \frac{2\,d}{d_I + 2\,d} + 1 \times \frac{2\,d}{d_I + 2\,d} \\ &= \frac{3\,d_I - 6\,d + 4\,d + 2\,d}{d_I + 2\,d} \\ &= \frac{3\,d_I}{d_I + 2\,d} \text{ images per track} \end{aligned} \tag{6.1}$$

As expected, this expression for N_{ipt} gives 3 images per track as $d \to 0$. Other values are 2 images per track at $d = d_I/4$ and 1.5 images per track at $d = d_I/2$. The same expression was found to apply when $d > d_I/2$, giving rise to the expected result that only 1 image per track can be found when $d = d_I$.

It can be noted that the number of image pairs per track is $N_{ipt} - 1$.

Suppose that the number N_{im} of images in each interrogation area is constant, as was the case in the simulations in the previous chapter. The average number of particle tracks per interrogation area can now be obtained. Let this be N_t:

$$N_t = \frac{N_{im}}{N_{ipt}} = \frac{N_{im} \times (d_I + 2\,d)}{3\,d_I}.$$

The number of fundamental particle image pairs and the number of second harmonic pairs per track can be calculated. If three images of each particle are found then there will be two fundamental image pairs and one second harmonic image pair. If two images are present, there will be one fundamental image pair but no second harmonic pairs. The case of only one image of a given particle is not interesting here, as no image pairs result.

$$\begin{aligned} N_{pairs/track} &= 2 \times \frac{d_I - 2\,d}{d_I + 2\,d} + 1 \times \frac{2\,d}{d_I + 2\,d} + 0 = \frac{2\,d_I - 2\,d}{d_I + 2\,d} \\ \text{and } N_{harmonics/track} &= 1 \times \frac{d_I - 2\,d}{d_I + 2\,d} + 0 + 0 = \frac{d_I - 2\,d}{d_I + 2\,d}. \end{aligned}$$

Simple multiplication of the number of pairs per track and the number of tracks per interrogation area yields the expected number of pairs per interrogation area:

$$\begin{aligned} N_{pairs/IA} &= \frac{2\,d_I - 2\,d}{d_I + 2\,d} \times \frac{N_{im}\,(d_I + 2\,d)}{3\,d_I} = \frac{N_{im}\,(2\,d_I - 2\,d)}{3\,d_I} \\ \text{and } N_{harmonics/IA} &= \frac{d_I - 2\,d}{d_I + 2\,d} \times \frac{N_{im}\,(d_I + 2\,d)}{3\,d_I} = \frac{N_{im}(d_I - 2\,d)}{3\,d_I}. \end{aligned}$$

Thus both the number of fundamental image pairs per interrogation area and the number of second harmonic pairs per interrogation area are decreasing linear functions of the image separation, d. There are $2N_{im}/3$ fundamental image pairs at $d = 0$ and zero at $d = d_I$, whilst there are $N_{im}/3$ second harmonic image pairs at $d = 0$ and none if $d > d_I/2$.

Assuming that a constant number of 24 images are present in each 64×64 pixel interrogation area, these results can be plotted (Figure 6.5). Also shown is the correlation between two different particle tracks, of up to three images each; this will be based on three image pairs when $d < d_I/2$ (*i.e.* 32 pixels) or two image pairs (when $d_I/2 < d < d_I$). It can be seen that the second harmonic peak starts off only 2–3 times larger than the correlations between different particle tracks, which will lie in various random positions. Harmonic peaks are difficult to detect in practice, as they soon disappear into the autocorrelation noise. It is therefore

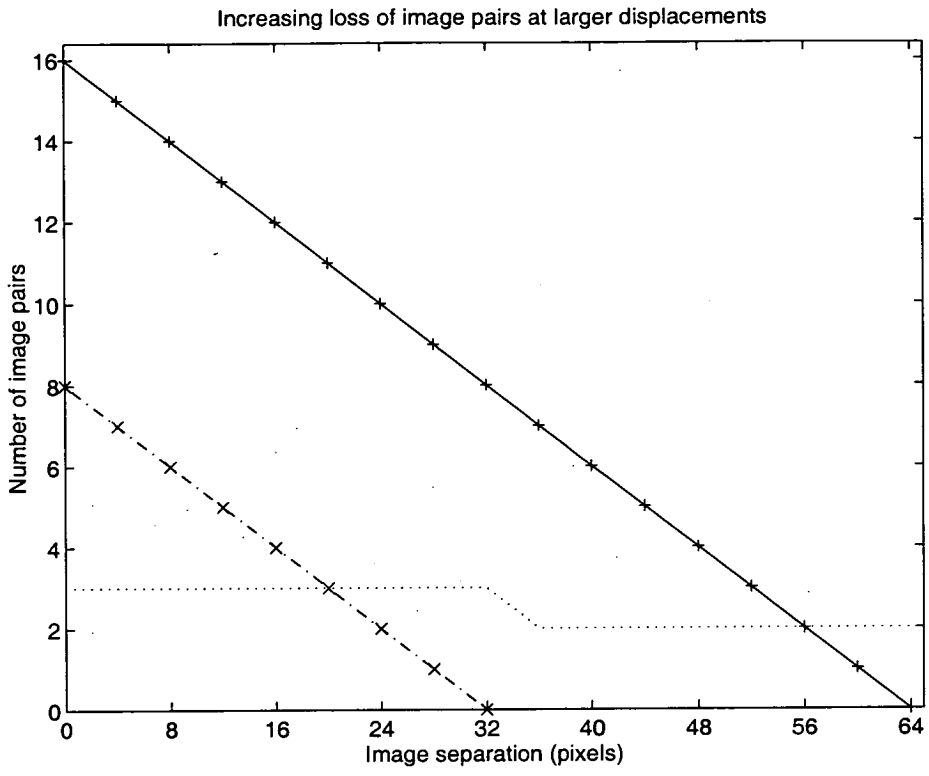


Figure 6.5: For the case of three images of each particle and 24 particle images per 64×64 interrogation area, the number of image pairs (solid line, '+' symbols, images a distance d apart) and the number of harmonic pairs (dash-dot line, '$\times$' symbols, images $2d$ apart) are plotted against the image separation. The dotted line represents correlations between different particle tracks.

best to determine the displacement using the strong fundamental displacement peak, rather than relying in any way on the presence of harmonic peaks.

6.3 The magnitude of errors due to wrap-around

In this section, the ratio of spurious wrap-around correlations to valid particle image correlations will be considered. This will make it possible to determine the relative magnitudes of the corresponding peaks in the autocorrelation plane, where they will be added together. Various shapes of autocorrelation function can be obtained by correlating two images together. By adding two such functions together with suitable weightings, the position of the overall maximum can be estimated. This will be the measured displacement.

6.3.1 The ratio of wrap-around peaks to correct peaks

In Figure 6.6, the case of five images of each particle and an image displacement of $d_I/3$ is shown. After this is examined, the general result can be considered.

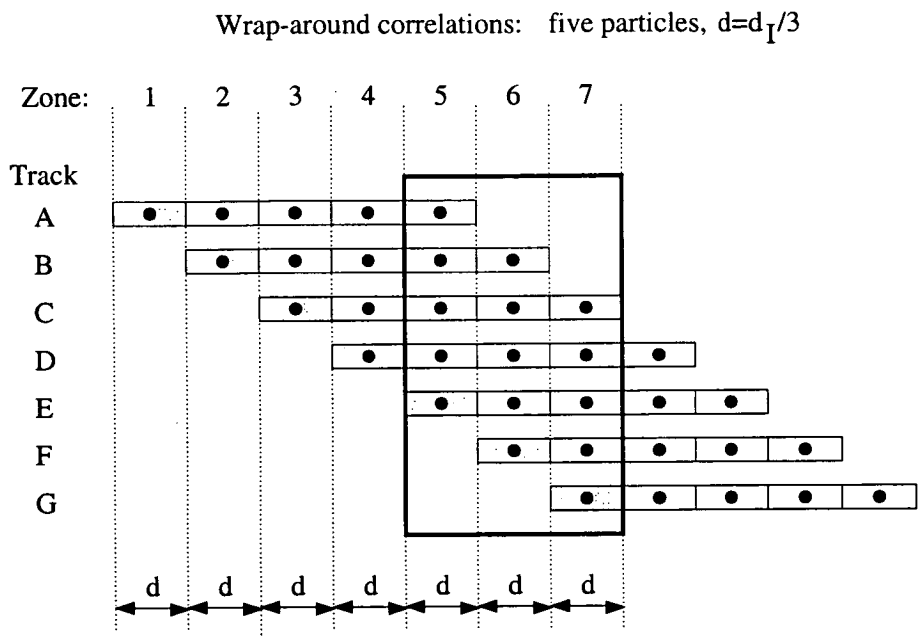


Figure 6.6: For the case of five images of each particle, separated by displacements to the right which are $d_I/3$ in magnitude, with interrogation areas measuring $d_I \times d_I$, consider the location of the first image of a random particle track, relative to the position of an interrogation area containing at least one image of that particle. The various starting locations are represented by lightly-shaded rectangles.

There are various possibilities to consider. Some of these involve starting zones located to the left of the interrogation area. There are four such zones, because at least one image of the particle must lie within the interrogation area. There are a further three starting zones located within the interrogation area itself, as the chosen displacement is $d_I/3$. All seven zones have the same width.

Tracks A and G are those which give rise to only a single, unpaired image, so there are two ways in which this can be obtained. Similarly, there are two ways in which exactly two consecutive images can arise (tracks B and F).

The number of ways in which three images of each particle are obtained can be evaluated by one of two different methods. One way is to say that there are seven zones in total and four zones have already been accounted for, so three remain.

Another way is to count the number of ways of selecting three consecutive images from a sequence of five. As illustrated in Figure 6.6 by tracks C, D and E, there are three ways. This can be explained by the tracks of five particle images giving rise to three images within the interrogation area and two outside it. The possible arrangements for the two *external* images are that there can be zero, one or two such images to the right of the interrogation area (zones 3, 4 and 5 respectively), so there are three possibilities in all.

So, given that at least one image of a particle occurs within this interrogation area, the probability of obtaining a single, unpaired image is 2/7. The probability of finding exactly two images of this particle is also 2/7, giving rise to one pair of 'good' images (which are too far apart to cause spurious wrap-around correlations to be made). The probability of finding a full house of three images is 3/7, in which case there will be two correct particle image pairs and one 'pair' which can cause wrap-around problems (as seen in the previous chapter).

From this track, the expected number of useful particle image pairs is $2 \times (3/7) + 1 \times (2/7) + 0 = 8/7$ pairs per track, whilst the expected number of spurious wrap-around correlations is 3/7 pairs per track. There are therefore going to be $8/3 \approx 2.67$ valid image correlations for every undesirable wrap-around correlation.

Similar results have been derived for other situations, *e.g.* four images of each particle separated by a distance $d_I/4$, as summarised in Table 6.3.

Displacement	Number of images	wrap-around:valid correlations
$d_I/3$	3	1:4
	4	1:3
	5	1:2.67
	6	1:2.5
$d_I/4$	4	1:9
	5	1:6
	6	1:5
$d_I/5$	5	1:16
	6	1:10

Table 6.3: Calculated ratios of wrap-around correlations to valid correlations between paired images.

Further consideration of Figure 6.6 suggests that there is a pattern to these results and that the general situation can be solved. Suppose that there are N images of each particle and that errors occur at an image displacement of d_I/M, for some natural numbers N and M satisfying $N \geq M$.

There will be a total of $M + N - 1$ starting zones to consider: M zones within the interrogation area and $N-1$ to the left of it. Two of these zones will cause only one image of the particle track to be found in the interrogation area (*cf* tracks A and G in Figure 6.6), whilst another two zones will give rise to exactly two images (*cf* tracks B and F). If $M = 3$ then all remaining zones must give rise to a full complement of three images (as was the case in Figure 6.6). If instead $M \geq 4$, however, there can only be two zones which yield tracks containing exactly three images, although other zones will yield four or more.

Particle tracks can be divided into two classes, according to whether or not all M zones within the interrogation area include an image of the particle.

There will, in general, be $2(M - 1)$ zones which give rise to only partially-filled interrogation areas, containing between 1 and $M - 1$ images. Each possible number of images from 1 to $M - 1$ will be associated with exactly two starting zones and the probability of each event will be $2/(N + M - 1)$. If $M = 5$ and $N = 10$, say, the probability of finding only one image will be the same as that of finding exactly two images, or exactly three images, or exactly four. In all of these cases, there will be no danger of encountering a spurious wrap-around correlation. If the actual number of images from a particular track in this interrogation area is n, for some $1 \leq n \leq (M - 1)$, then there will be $n - 1$ valid image *pairs*.

The remaining zones will all contain M images. Only these zones can give rise to wrap-around correlations. The number of these zones can be determined by subtraction to be $(M + N - 1) - 2(M - 1) = (N + 1 - M)$, or counted by realising that the number of images from each track sticking out to the left (or to the right) of the interrogation area might be zero, one, two, or any natural number up to $N - M$. Either way, there are a total of $(N - M) + 1$ possibilities,

hence the probability must be $(N+1-M)/(N+M-1)$. For this full quota of M images from a particular particle in the interrogation area, an autocorrelation analysis using unpadded FFTs will find $M-1$ valid image pairs and one spurious wrap-around correlation.

A well-known mathematical result is that the sum of the first n natural numbers is $n(n+1)/2$, *i.e.* the number of terms multiplied by their average value. This makes it possible to deduce that the ratio of wrap-around to valid correlations is:

$$\begin{aligned}
&\frac{N+1-M}{N+M-1}\times 1 \;:\; \frac{N+1-M}{N+M-1}\times(M-1)+\frac{2}{N+M-1}\times\sum_{n=1}^{M-1}(n-1)\\
&=\;(N+1-M)\;:\;(N+1-M)(M-1)-2(M-1)+2\sum_{n=1}^{M-1}n\\
&=\;(N+1-M)\;:\;(N+1-M-2)(M-1)+2\,\frac{(M-1)M}{2}\\
&=\;(N+1-M)\;:\;(N-1-M)(M-1)+M(M-1)\\
&=\;(N+1-M)\;:\;(N-1)(M-1). \qquad (6.2)
\end{aligned}$$

Thus, the general case of N particle images a distance d_I/M apart has been solved: the ratio of spurious wrap-around correlations to genuine correlations is $(N+1-M):(M-1)(N-1)$. The values in the previous table all fit this expression. This ratio should be made as low as possible, if errors due to wrap-around are to be minimized in an autocorrelation analysis using unpadded FFTs. (Alternatively, displacements could be calculated more carefully, using fully-padded FFTs.)

Having obtained this result, it is possible to estimate the magnitude of the errors which will be found in a PIV analysis using unpadded FFTs, for speed. For example, major problems will occur at displacements of $d_I/2$, where the ratio of wrap-around image pairs to genuine correlations will be 1:1, *i.e.* the fundamental displacement peak may not even be the largest peak in the autocorrelation plane.

To consider an example in detail: an image displacement of $d_I/3$ gives a 1:4 ratio if three images of each particle are recorded ($M=3$, $N=3$). If the displacement is decreased slightly to $(d_I/3-\delta x, 0)$ for some small distance δx, there

will be four valid correlations at this correct position in the autocorrelation plane for every spurious wrap-around correlation at the point $(d_I - 2 \times (d_I/3 - \delta x), 0) = (d_I/3 + 2\delta x, 0)$. For the given number of images of each particle and the given image displacement, calculations can be made which determine not only the position of the displacement peak and the spurious wrap-around peak in the autocorrelation plane, but also their relative magnitudes.

The next step required is to add these correct and incorrect signals together in the autocorrelation plane. The displacement detected by an analysis program will be the location of the global maximum of this sum, which will lie in between the above points if δx is small. To proceed further, however, it is necessary to know the properties of ideal autocorrelation peaks, *i.e.* their size and shape. These properties depend in turn on the particle images in the PIV record, *e.g.* the size and shape of each image, as well as the intensity profile across each image.

Some possible image shapes and intensity profiles will be examined in the next subsection. These images will be autocorrelated, so that the autocorrelation peaks can be assessed and their shapes determined. It will then be possible to add the correct image correlation signal to the spurious wrap-around signal with appropriate weightings, so as to calculate the theoretical position of the global maximum. The errors caused by these wrap-around effects can then be determined and compared to the results from the simulations given in the previous chapter.

6.3.2 Shapes of autocorrelation peaks

Spherical particles will give rise to circular images in the flow record, but aspherical particles will produce non-circular images. The recorded images of spherical particles could have a flat intensity profile or a Gaussian one, according to PIV properties such as the illumination intensity, the reflectivity of the particles and the photochemical response characteristics of the recording medium. For example, a flat intensity profile might arise if all pixels are as bright as can be measured, due to the particles reflecting enough incident light to saturate the film.

A uniform square image with sides 5.0 pixels long can be simulated, by creating a grid of zero intensity then giving selected points a non-zero value. This shape is shown in Figure 6.7 by means of an intensity surface profile (top left) and a plot of intensity contours (top right). This image can be autocorrelated; the autocorrelation surface is shown in the middle left graph, with contours plotted in the middle right. The autocorrelation surface appears to be a cross between a cone and a square-based pyramid. A cross-section through the centre of this shape, taken parallel to one of the sides of the autocorrelation plane, is triangular (bottom left). As a result, a Gaussian fit to the discrete data points is not particularly good (bottom right).

A similar procedure can be carried out with circular images (Figure 6.8). Pixels which were only partly included within a circle 6.0 pixels in diameter were set to intermediate values of intensity, in order to achieve smoother intensity contours. After autocorrelation, the peak shape is found to be similar to a cone, but with a slightly flattened top and a more curved base. Whilst not perfect, a Gaussian fit gives a very reasonable result.

A conical image was created, in which the decrease in intensity on moving away from the centre of the image was proportional to the distance from the apex. After autocorrelation, it became apparent that a Gaussian surface made a very good fit to the autocorrelation function (Figure 6.9). Similar results were observed for a Gaussian image intensity profile (Figure 6.10) and for a flattened Gaussian image intensity profile, in which the effect of film saturation was modelled by subjecting each pixel in a Gaussian profile to a maximum value (Figure 6.11).

So, uniform-intensity images tend to produce autocorrelation peaks with a rather triangular cross-section, regardless of whether the images are square or circular (Figures 6.7 and 6.8). Saturating the recording medium by exposing it to too much light means that the autocorrelation peaks become more conical.

On the other hand, particle images may vary in intensity according to some function of the distance away from the centre of the image, with most (if not all) of the recording medium avoiding saturation. In this case, Gaussian autocorrelation

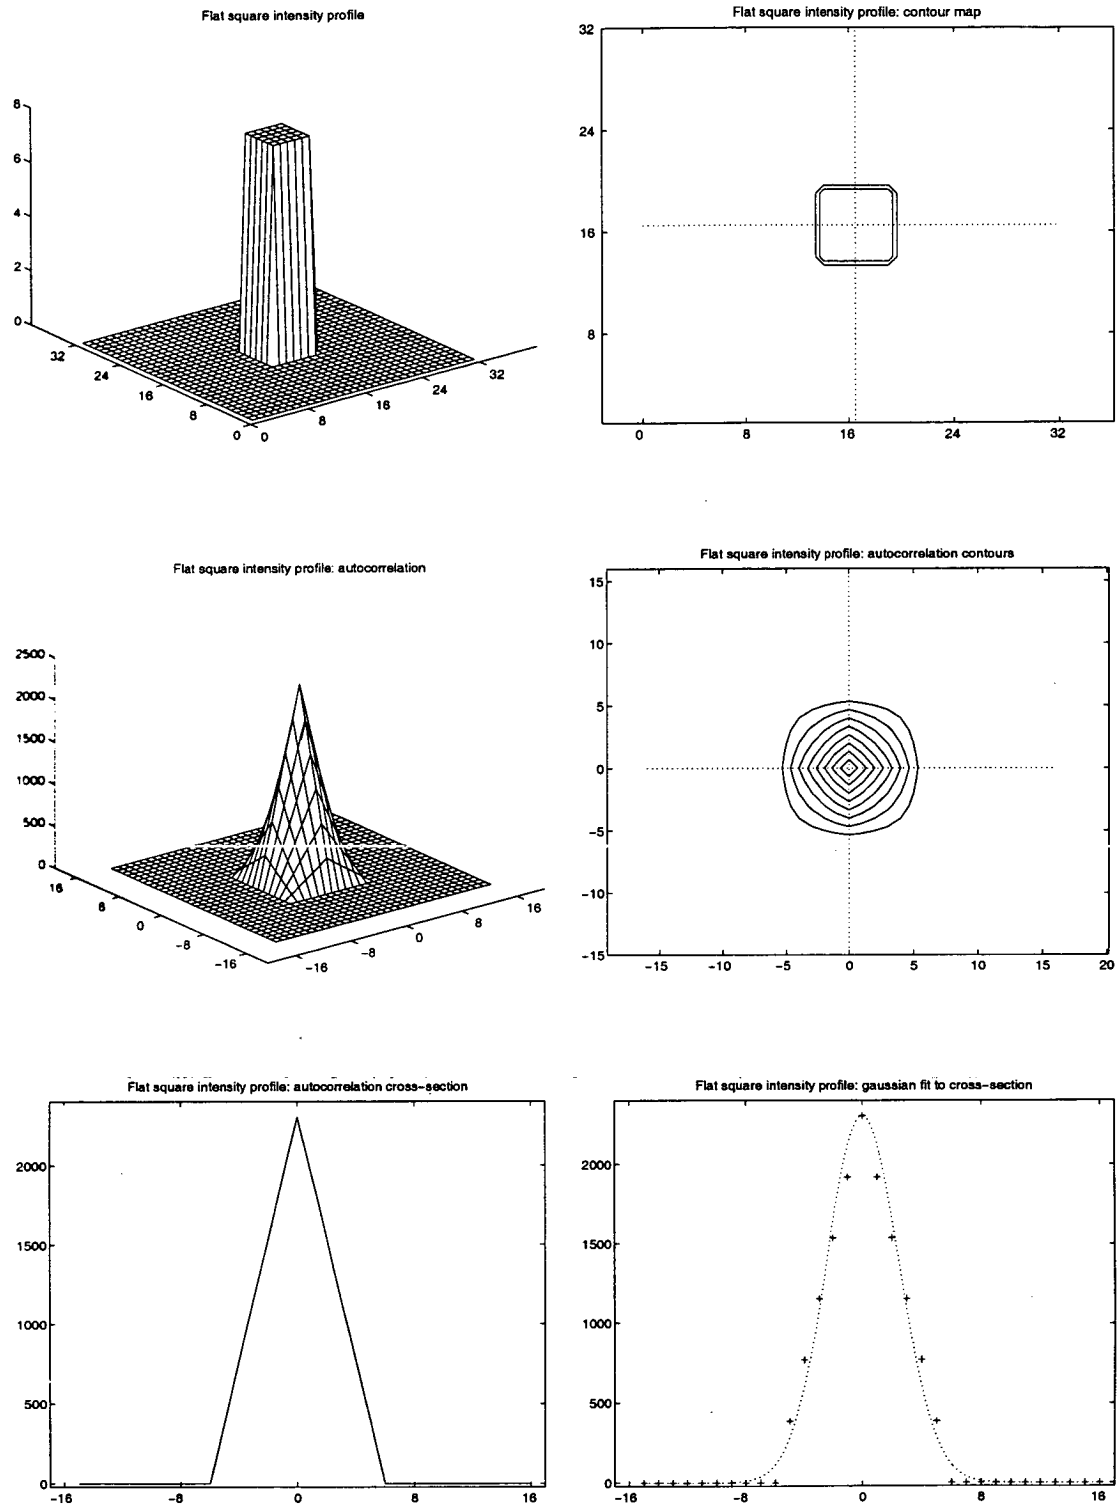


Figure 6.7: Square images of uniform intensity: profile in the image field, contours in the image field, profile of the autocorrelation plane, contours in the autocorrelation plane, cross-section through the autocorrelation plane along $\Delta y = 0$ pixels and a Gaussian fit to discrete data points 1 pixel apart.

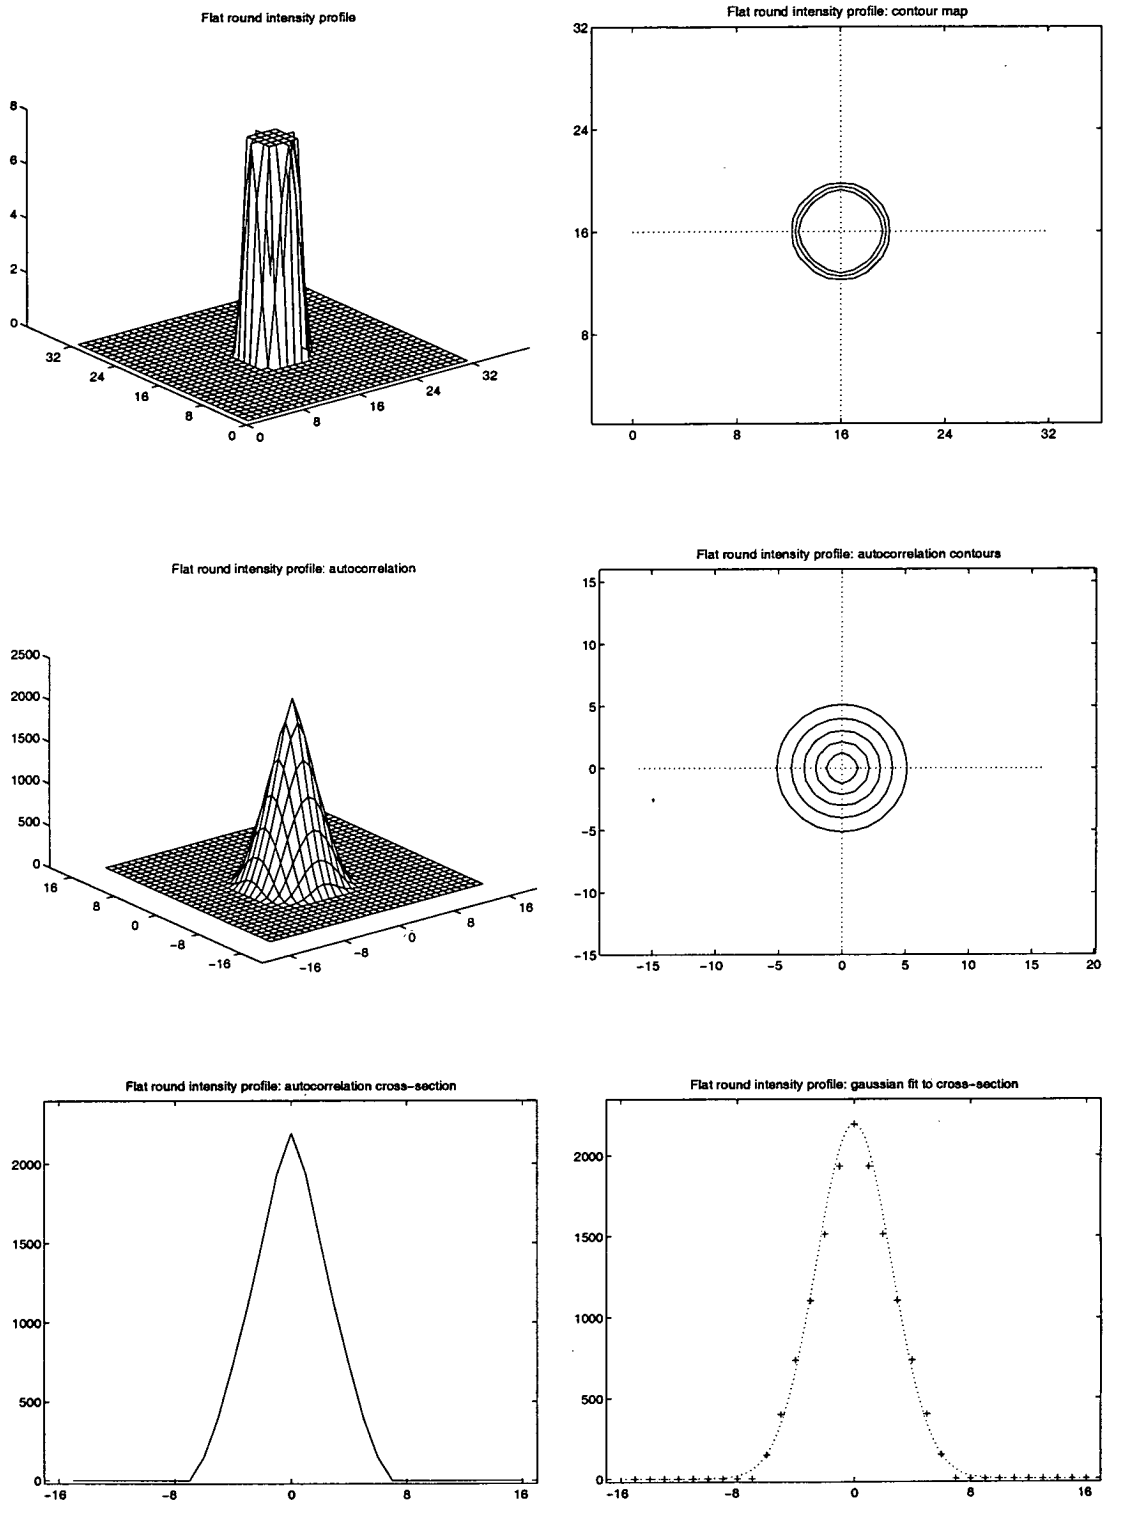


Figure 6.8: Circular images of uniform intensity: profile in the image field, contours in the image field, profile of the autocorrelation plane, contours in the autocorrelation plane, cross-section through the autocorrelation plane along $\Delta y = 0$ pixels and a Gaussian fit to data points 1 pixel apart.

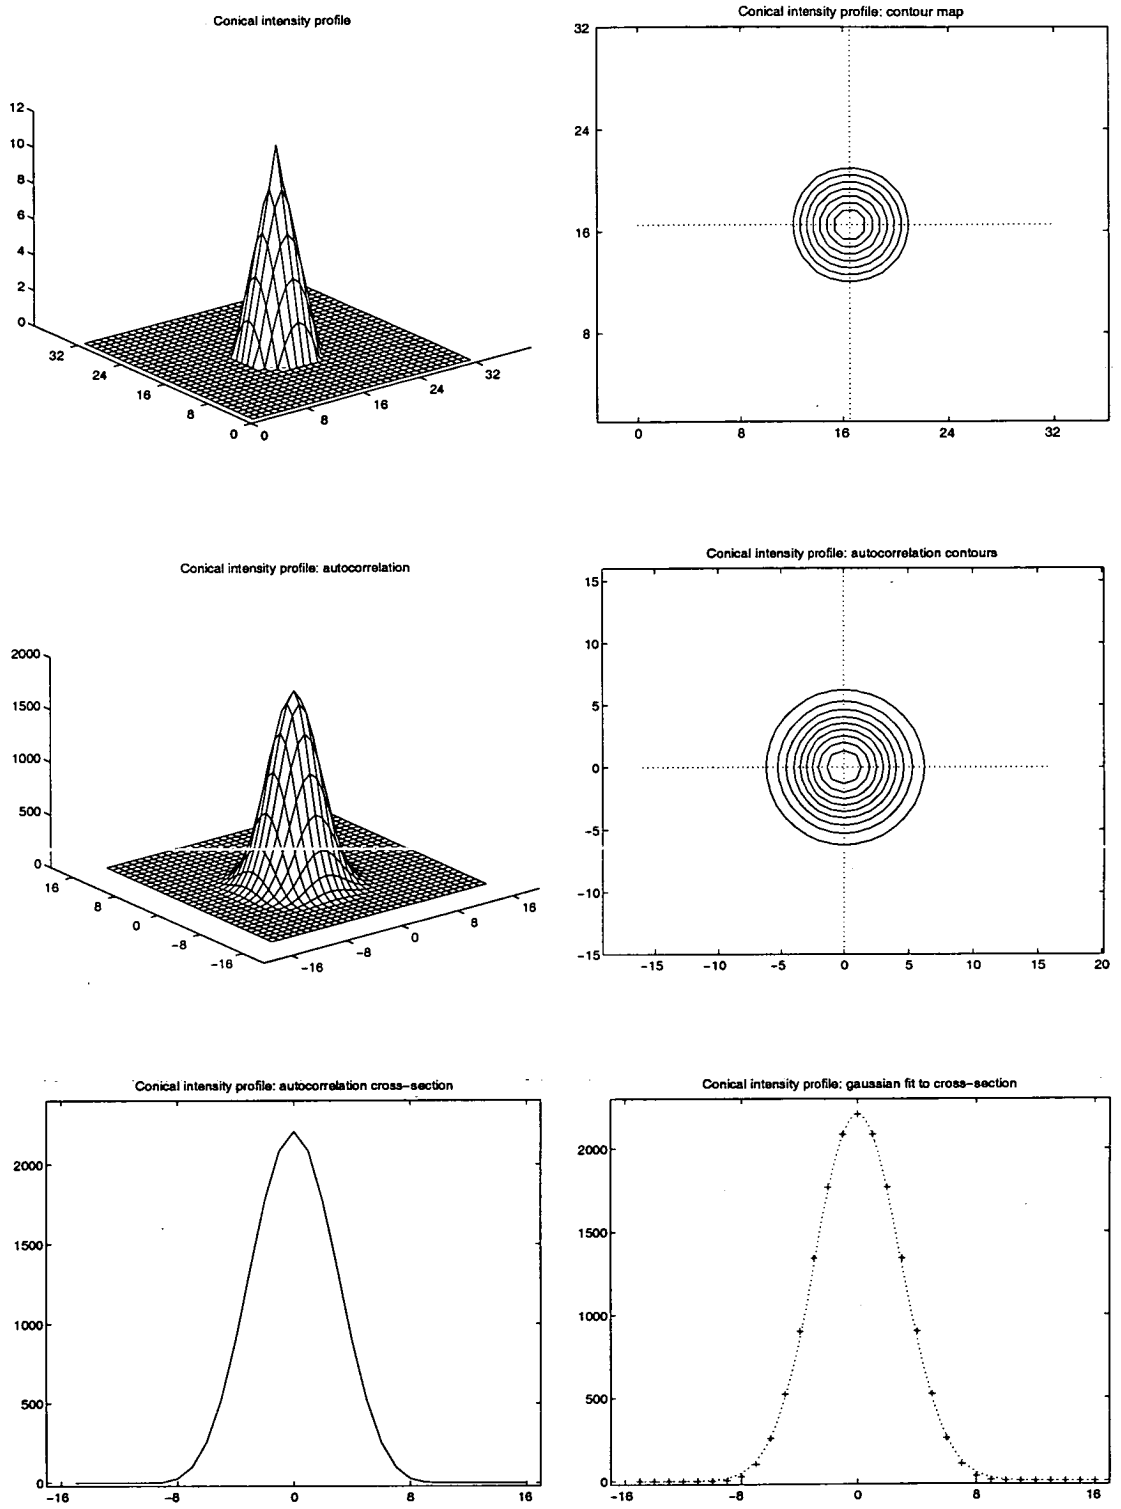


Figure 6.9: Images with a conical intensity profile: profile of the image intensity, contours of the image intensity, profile of the autocorrelation plane, contours in the autocorrelation plane, cross-section through the autocorrelation plane along $\Delta y = 0$ pixels and a Gaussian fit to discrete data points 1 pixel apart.

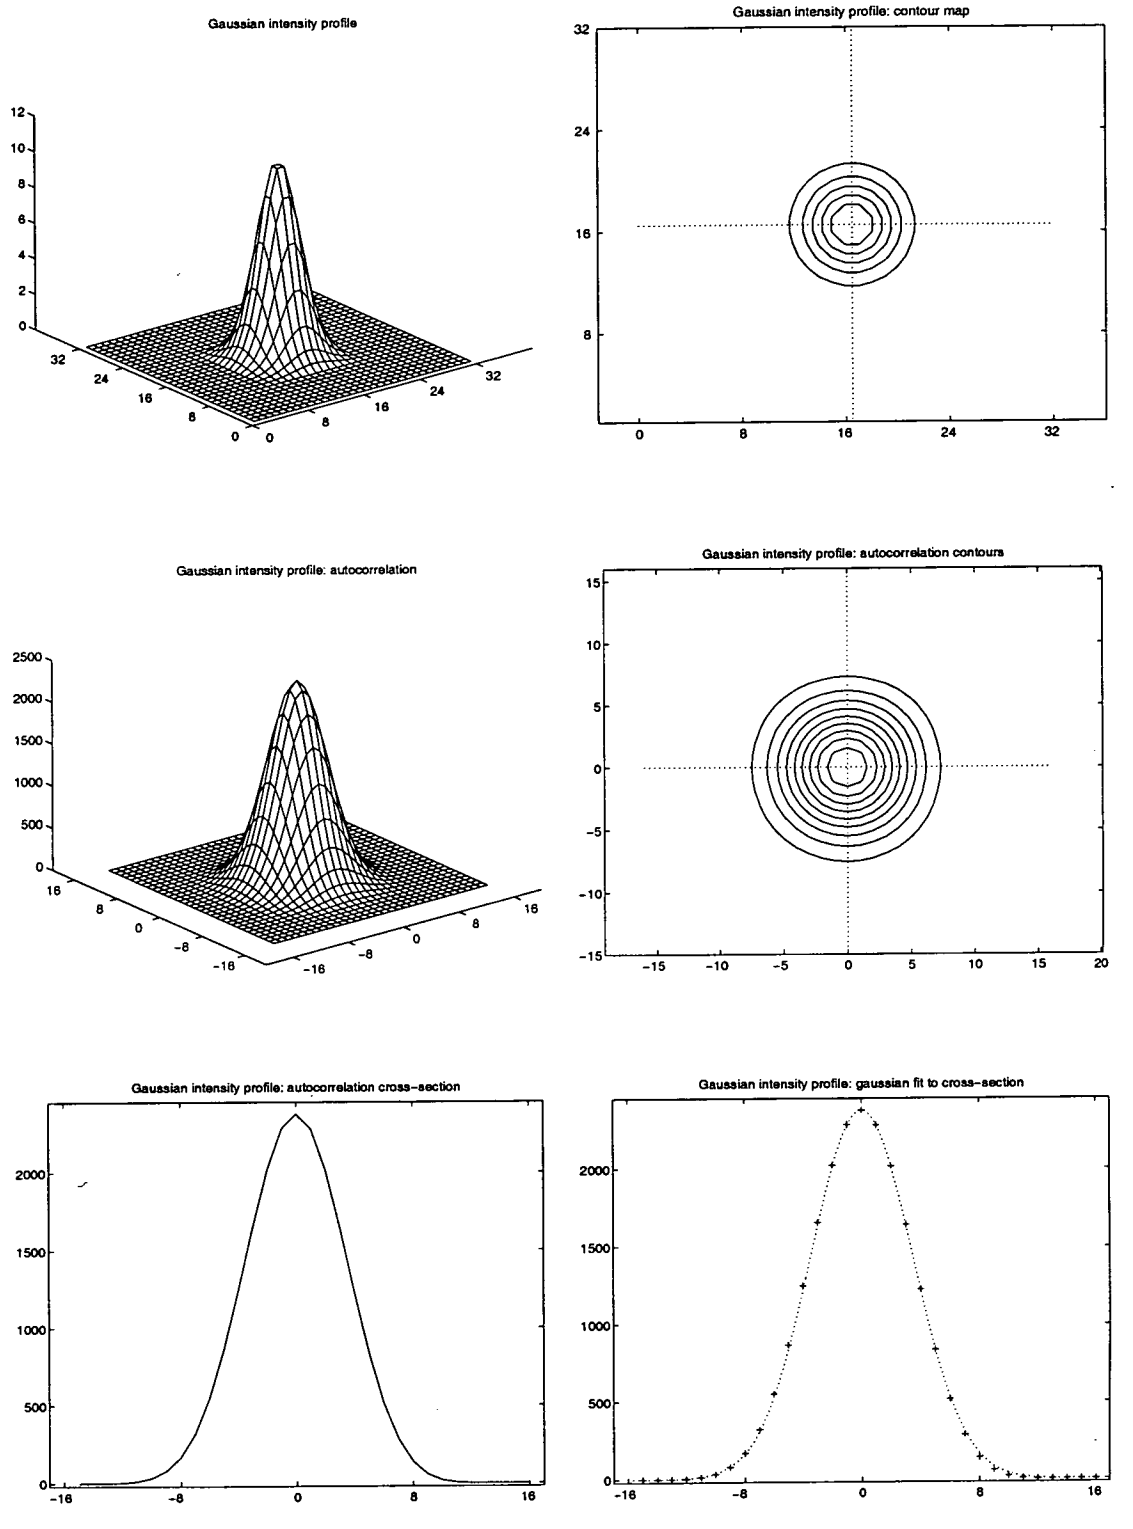


Figure 6.10: Images with a Gaussian intensity profile: profile in the image field, contours in the image field, profile of the autocorrelation plane, contours in the autocorrelation plane, cross-section through the autocorrelation plane along $\Delta y = 0$ pixels and a Gaussian fit to discrete data points 1 pixel apart.

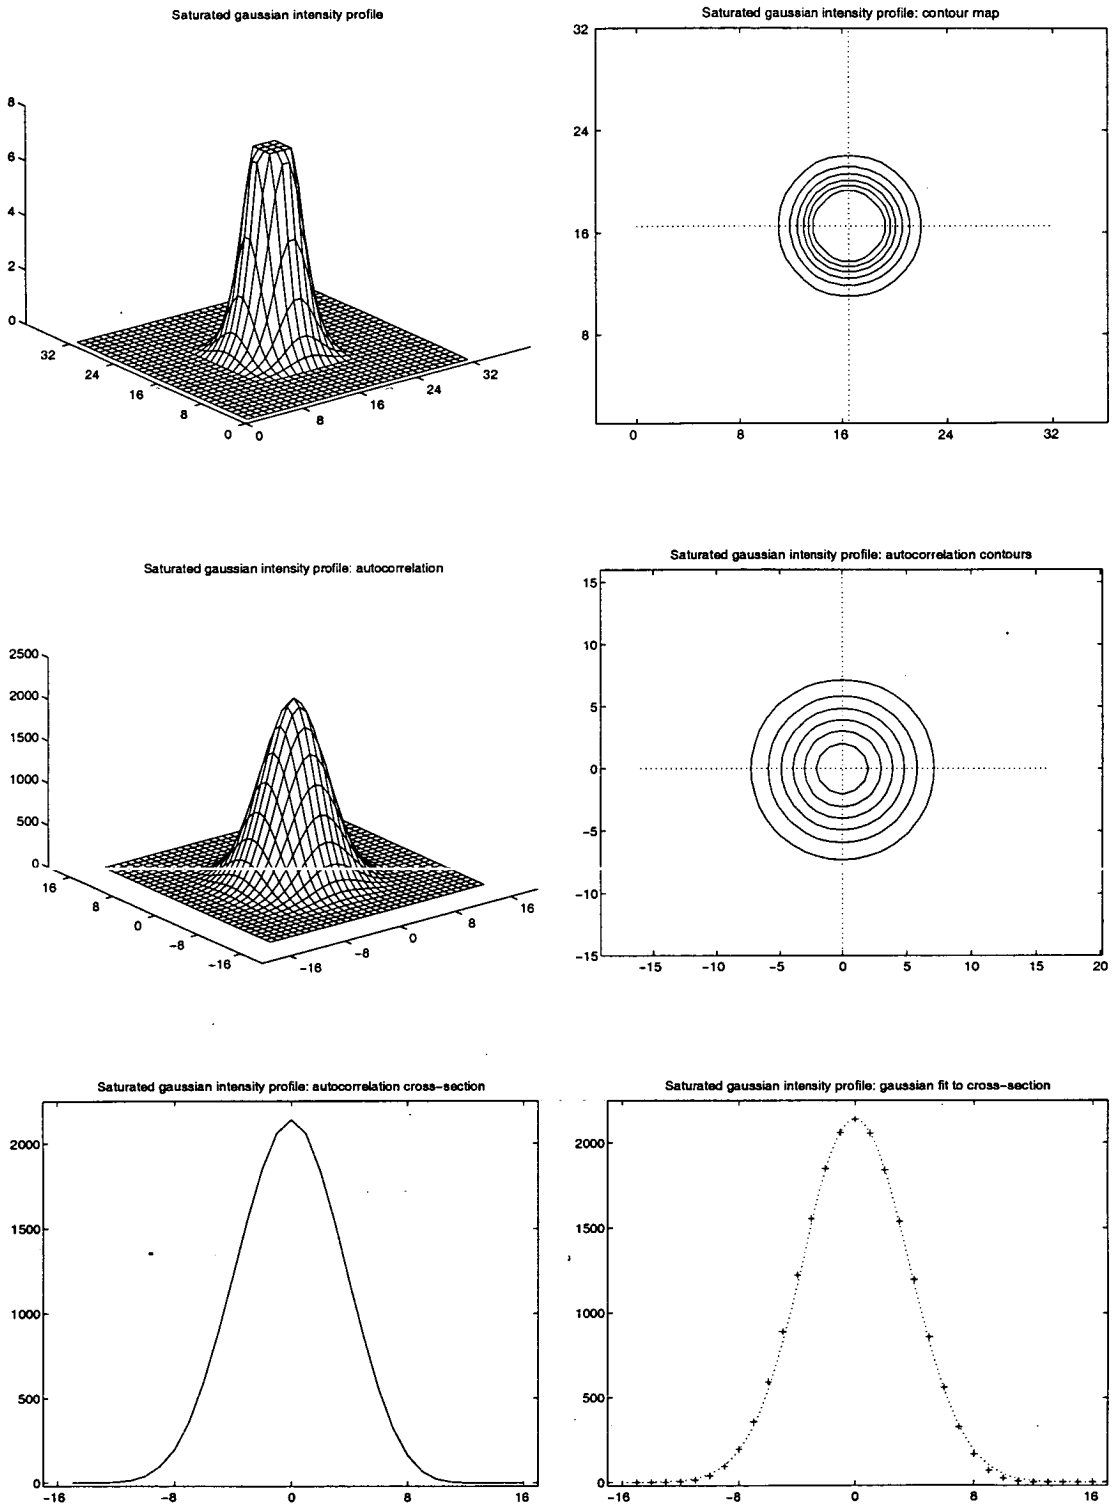


Figure 6.11: Images with a Gaussian intensity, but saturated at a value of 62.5% of the possible maximum: profile of the image intensity, contours of the image intensity, profile of the autocorrelation plane, contours in the autocorrelation plane, cross-section through the autocorrelation plane along $\Delta y = 0$ pixels and a Gaussian fit to discrete data points 1 pixel apart.

peaks are obtained, regardless of whether the image intensity profile is linear or Gaussian (Figures 6.9 and 6.10).

6.3.3 Conical peaks

The analysis of the expected displacement in the presence of wrap-around peaks is easiest if the autocorrelation peaks are conical, or at least have a triangular cross-section. This will be a reasonable assumption if the images have a uniform intensity profile, *e.g.* for fully-saturated square or circular images.

Suppose that three images of each particle are recorded at horizontal intervals of approximately 64/3 pixels. Analysis using 64×64 pixel interrogation areas gives rise to an autocorrelation plane, $A(\Delta x, \Delta y)$. A vertical cross-section through the $\Delta y = 0$ plane will reveal apparently-triangular peaks, making this case simple to analyse. Let each peak be two pixels in diameter, for the sake of argument.

The expected ratio of spurious wrap-around peaks to valid correlation peaks is 1:4 for this case, where $N = M = 3$. So, for each conical peak of magnitude 4 units centred on a position $(d_1, 0)$ in the autocorrelation plane, there will be a corresponding wrap-around peak one unit in magnitude, located at $(d_2, 0)$. The distance d_2 is given by $d_2 = 64 - 2\,d_1$. The fundamental displacement peak, being based on four times as many image correlations, will be four times steeper than the wrap-around peak and will dominate the summation. The situation is illustrated in Figure 6.12 for the value $d_1 = 21.1$ pixels.

Mathematically, the conical autocorrelation functions shown in Figure 6.12 can be represented by the sum of two individual functions A_1 and A_2:

$$
\begin{aligned}
A\,(\Delta x, 0) &= A_1\,(\Delta x) + A_2\,(\Delta x) \\
\text{where } A_1\,(\Delta x) &= \begin{cases} 4\,(1 - |(\Delta x - d_1)|) & \text{if } d_1 - 1 < \Delta x < d_1 + 1 \\ 0 & \text{otherwise} \end{cases} \\
\text{and } A_2\,(\Delta x) &= \begin{cases} 1 - |(\Delta x - d_2)| & \text{if } d_2 - 1 < \Delta x < d_2 + 1 \\ 0 & \text{otherwise.} \end{cases}
\end{aligned}
$$

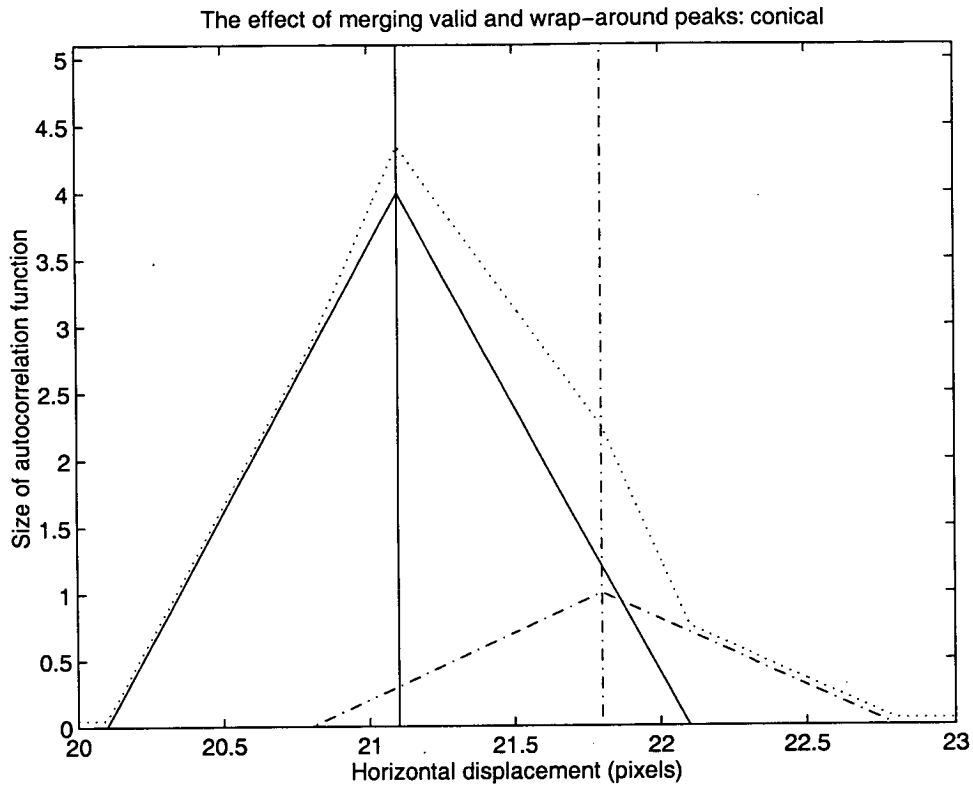


Figure 6.12: For the case of three images of each particle, separated by entirely horizontal displacements, analysed using interrogation areas measuring 64×64 pixels, consider the effect of combining a conical fundamental correlation peak of magnitude 4 units (solid lines) with a conical wrap-around peak of magnitude 1 unit (dash-dot lines). A summation function can be obtained (dotted lines). The vertical lines indicate the location of the maximum of each function.

For displacements in the range $d_1 < \Delta x < d_2$, if d_1 and d_2 are close enough together for the peaks to overlap (*i.e.* neither $A_1(\Delta x)$ nor $A_2(\Delta x)$ is zero), the above expression can be written as:

$$\begin{aligned} A(\Delta x, 0) &= 4\,(1 - (\Delta x - d_1)) + (1 - (d_2 - \Delta x)) \\ &= 4 - 4\,\Delta x + 4\,d_1 + 1 - d_2 + \Delta x \\ &= (5 + 4\,d_1 - d_2) - 3\,\Delta x \end{aligned}$$

The summation function obtained by adding the two triangular peaks in the only region where wrap-around errors can arise (*i.e.* $d_1 < \Delta x < d_2$) is therefore a strictly decreasing function from $(d_1, 0)$ to $(64 - 2d_1, 0)$, with no possible turning point. No further calculations are necessary, as the situation is straightforward.

The greatest value of the summation function can only be that found at $(d_1, 0)$, so autocorrelation analysis will probably detect the correct image displacement anyway. Therefore, if the autocorrelation peaks are conical (or some other shape with a triangular cross-section), wrap-around phenomena will actually have little effect on the value of the measured displacement.

6.3.4 Parabolic peaks

If instead the peaks in the autocorrelation plane are strictly-decreasing non-linear functions, it is possible that the sum of the intended peak at $(d_1, 0)$ and the spurious peak at $(64 - 2\,d_1, 0)$ will have a maximum in an intermediate position.

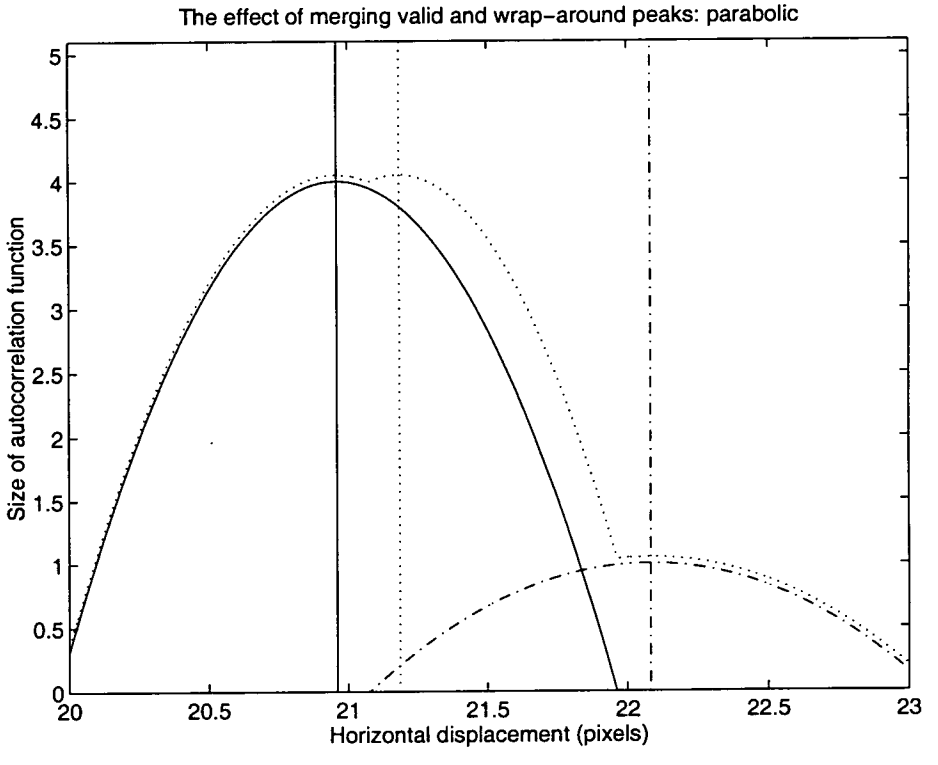


Figure 6.13: As in the previous figure, except this time a parabolic fundamental correlation peak (solid lines) at $d_1 = 20.961$ pixels is added to a parabolic wrap-around peak (dash-dot lines) to obtain a summation peak (dotted lines). The magnitude of the summation peak has been increased by 0.05 pixels for clarity.

Figure 6.13 was obtained by adding two parabolic peaks, centred on positions given by $d_1 = 20.961$ pixels. Parabolic peaks may or may not arise in practice, but do at least give rise to an analytical solution. The total width of each peak is,

once again, taken to be 2 pixels; outside a zone 1 pixel in radius, the contribution of each peak to the overall autocorrelation function is zero.

An image separation of $d_1 = 20.961$ pixels produces the largest overestimate of the true displacement for this autocorrelation peak width (and hence particle size). This error is 0.224 pixels. The summation curve (dotted lines) has three local maxima here. The peak at $(21.185, 0)$ is slightly taller than the one at $(20.961, 0)$, but both are much larger than the wrap-around peak at $(22.078, 0)$.

If d_1 is decreased slightly from this limiting value of 20.961 pixels, d_1 and d_2 move further apart. The magnitude of the central summation peak of Figure 6.13 decreases, allowing the peak at $(d_1, 0)$ to become the global maximum. This means that a search of the autocorrelation plane for the global maximum should now detect the correct displacement, so the error should undergo a step change down to zero. As d_1 decreases further, the central maximum drops lower and lower until it eventually disappears, leaving only the peaks at $(d_1, 0)$ and $(d_2, 0)$.

At slightly larger values of d_1, the central maximum grows and becomes the global maximum. The centres of the fundamental peak (solid vertical line) and the wrap-around peak (dash-dot vertical line) move closer together and the error is reduced. The error in the measurement falls to zero when $d_1 = 64/3$ pixels. The two outermost maxima can disappear if d_1 increases far enough towards a value of 64/3 pixels (Figure 6.14, left); the condition for this is that the tail of the wrap-around peak must overlap with the maximum of the fundamental displacement peak, *i.e.* $d_2 - d_1 < 1$ pixel, the radius of the peaks in the autocorrelation plane.

If d_1 is increased to slightly more than 64/3 pixels, the wrap-around peak will be closer to the origin than the fundamental peak, so the displacement will be under-read. The summation function will have only one maximum if d_1 remains close to 64/3 pixels, but two more will appear at d_1 and d_2 if $d_1 - d_2 > 1$ pixel.

The greatest under-read error is found when $d_1 = 21.706$ pixels (Figure 6.14, right), at which point the central summation peak and the intended correlation peak $A_1(d_1)$ once again have almost the same amplitude. At larger displacements,

$A_1(d_1)$ will become larger than the middle summation peak, which will gradually diminish in size then disappear. At image displacements larger than 21.706 pixels, therefore, the measured displacement should once again be d_1 and the error will undergo a step change back up to zero.

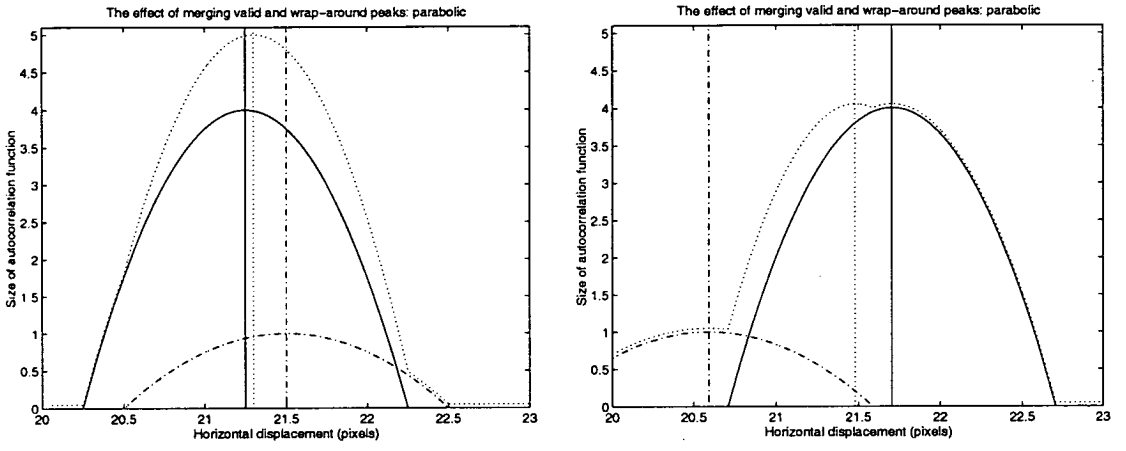


Figure 6.14: As in the previous figure, except the simulated displacement has been increased to $d_1 = 21.25$ pixels (left) or 21.706 pixels (right).

Mathematically, the equation for parabolic cross-sections of radius R_a is:

$$
\begin{aligned}
A(\Delta x, 0) &= A_1(\Delta x) + A_2(\Delta x) \\
\text{where } A_1(\Delta x) &= \begin{cases} 4\left(1 - (\Delta x - d_1)^2 / R_a^2\right) & \text{if } |\Delta x - d_1| < R_a \\ 0 & \text{otherwise} \end{cases} \\
\text{and } A_2(\Delta x) &= \begin{cases} 1 - (\Delta x - d_2)^2 / R_a^2 & \text{if } |\Delta x - d_2| < R_a \\ 0 & \text{otherwise,} \end{cases}
\end{aligned}
$$

with $d_2 = 64 - 2d_1$ for displacements close to 64/3 pixels, as before. When both A_1 and A_2 are non-zero:

$$
\begin{aligned}
A(\Delta x, 0) &= 4\left(1 - \frac{(\Delta x - d_1)^2}{R_a^2}\right) + \left(1 - \frac{(\Delta x - d_2)^2}{R_a^2}\right) \qquad (6.3) \\
&= 4 - 4\,\frac{\Delta x^2 - 2\,d_1\,\Delta x + d_1^2}{R_a^2} + 1 - \frac{\Delta x^2 - 2\,d_2\,\Delta x + d_2^2}{R_a^2} \\
&= \left(5 - \frac{4\,d_1^2 + d_2^2}{R_a^2}\right) + 2\,\frac{4\,d_1 + d_2}{R_a^2}\,\Delta x - \frac{5}{R_a^2}\,\Delta x^2. \qquad (6.4)
\end{aligned}
$$

It is not obvious where the maximum of this function lies. However, the equation can easily be differentiated in order to search for turning points:

$$\begin{aligned} \frac{dA}{d\Delta x} &= 2\,\frac{4\,d_1 + d_2}{R_a^2} - 10\,\frac{\Delta x}{R_a^2} = \frac{2}{R_a^2}\,(4\,d_1 + d_2 - 5\,\Delta x) = 0 \\ \Rightarrow \Delta x &= \frac{4\,d_1 + d_2}{5} = d^*. \end{aligned} \tag{6.5}$$

This means that $A(\Delta x, 0)$ has a turning point in the region $d_1 < \Delta x < d_2$, provided that d_1 and d_2 are close enough together for both A_1 and A_2 to be greater than zero throughout this region. This turning point occurs at $(d^*, 0)$, which is a point one fifth of the way from the fundamental displacement peak $(d_1, 0)$ to the wrap-around peak at $(d_2, 0)$. As the second derivative $d^2A/d\Delta x^2 = -10/R_a^2$ which is negative, this point must be a maximum.

In addition, as shown above, secondary local maxima can also occur at both d_1 and d_2, if the two peaks are far enough apart. This is not apparent in the above mathematics, however, as it was assumed that both A_1 and A_2 were valid at the point in question. Nevertheless, it is obvious that, if A_2 is zero near $\Delta x = d_1$ because the two peaks are more than R_a apart, then the sum of A_1 and A_2 will be at a maximum at any point where A_1 on its own is at a maximum.

The upshot of this is that wrap-around peaks can disrupt the measurement of a true image displacement of d_1, in such a way that the measured displacement (d^*) is actually $(4\,d_1 + d_2)/5$. As $d_2 = 64 - 2\,d_1$, this can also be written as:

$$d^* = \frac{2\,d_1 + 64}{5}, \tag{6.6}$$

i.e. the measured displacement will be 3/5 of the way from the true value to a fixed 'attractor point', $(64/3, 0)$. Clearly the error is zero when $d_1 = 64/3 = d_I/3$, when the fundamental displacement peak and the wrap-around peak overlap completely: $d_1 = d_2$. In general, the error in the measured displacement is given by:

$$\text{Error} = d^* - d_1 = \frac{4\,d_1 + d_2}{5} - \frac{5\,d_1}{5} = \frac{d_2 - d_1}{5} = \frac{64 - 3\,d_1}{5}. \tag{6.7}$$

This result is independent of the diameter of each autocorrelation peak, because R_a became a common factor after differentiating to obtain equation 6.5 and could later be cancelled out. A larger autocorrelation peak diameter means that the width of the region in which the two autocorrelation signals overlap will increase, so equation 6.7 will be valid for a wider range of values of d_1. This will cause the greatest errors obtained to be worse for wide autocorrelation peaks.

It was noted earlier that the greatest over-read and under-read errors predicted using parabolic peaks occurred at $d_1 = 20.961$ or 21.706 pixels, or $64/3 \pm 0.372$ pixels. These points are interesting, as they mark the lower and upper limits of d_1 for which equation 6.7 is valid. At these values, the amplitude of the central maximum at d^* is the same as the amplitude of $A_1(\Delta x)$ at $\Delta x = d_1$. Setting $A(\Delta x, 0)$ equal to 4 in equation 6.4, requiring $A(\Delta x, 0)$ to be at a maximum by substituting $\Delta x = (4d_1 + d_2)/5$ (equation 6.5) and eliminating d_2 leads to a quadratic equation in d_1. This has roots $d_1 = 64/3 \pm \sqrt{5}\, R_a/6$ pixels, which agrees with the numerical observation when $R_a = 1$ pixel. Thus, equation 6.7 is valid if:

$$\frac{64}{3} - \frac{\sqrt{5}\, R_a}{6} < d_1 < \frac{64}{3} + \frac{\sqrt{5}\, R_a}{6}.$$

The theoretical behaviour of parabolic autocorrelation peaks, as predicted by this model, has been compared to data from some of the simulations presented in the previous chapter. In Figure 6.15, data from three-particle simulations using fully-padded FFTs have been subtracted from the corresponding results for unpadded FFTs, in an attempt to cancel out the variable errors in the unpadded data and leave only the wrap-around signal ($\times$ signs, dash-dot curve). The assumption of parabolic peaks in the autocorrelation plane gave rise to equation 6.7 above, which has been plotted using a solid line (with no validity limits imposed).

It can be seen that the attempt to cancel out the systematic oscillations in the error due to bias has been reasonably successful, as the simulated results are fairly flat around $d_1 = 19$ pixels. However, the constant offset error for the fully-padded run was slightly larger than in the unpadded simulation, so the resulting difference was slightly greater than zero. On the understanding that the points plotted in Figure 6.15 could all have been lowered by about 0.01 pixels by taking

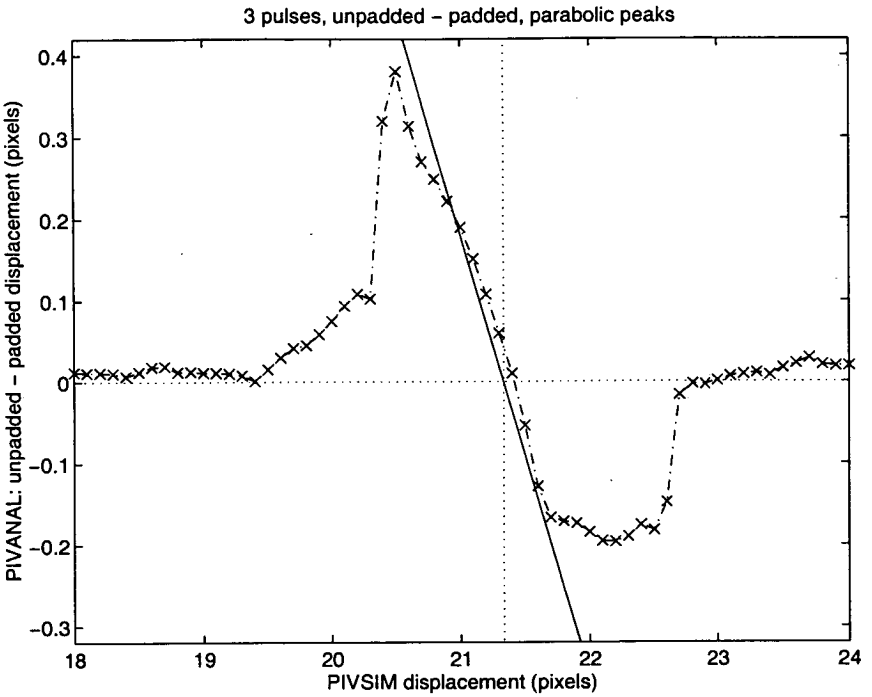


Figure 6.15: A comparison of simulated PIV results against the predictions which result from assuming parabolic autocorrelation peaks. Simulations were carried out using three images of each particle and analysis both with and without FFT padding, for particles with an e^{-2} image diameter of 4 pixels. The padded results have been subtracted from the unpadded results, to produce a plot of the measured error against the simulated displacement (dash-dot lines, $\times$ symbols). The predicted wrap-around error for parabolic peaks is plotted using a continuous line.

this into account, it appears that the estimate of the error which arises from the assumption of parabolic autocorrelation peaks is a reasonable fit to the simulated results, in the region about 0.5 pixels on either side of $d_1 = 64/3$ pixels.

One drawback to the parabolic peak model is that true displacements at either end of the effected range ($d_1 = 64/3 \pm \sqrt{5}\, R_a/6$ pixels) give rise to two maxima in the summation curve, of equal magnitude. For example, $R_a = 1$ pixel meant that the left-hand summation peak at $(d_1, 0)$ was greater than the central summation peak at d^* when d_1 was 20.960 pixels, but the reverse was true when d_1 was increased marginally to 20.961 pixels. This means that a step change in the wrap-around error should be found at points such as this. The simulated data shown in Figure 6.15 show a more gradual change in error in these regions (*e.g.* around $d_1 = 20$ pixels). This suggests that the autocorrelation peaks are less steeply curved at their edges than is assumed in the parabolic model.

An alternative type of autocorrelation function which might therefore be considered is a Gaussian model. It was noted earlier that Gaussian curves appear to give good fits to the autocorrelation functions which arise when the particle image intensity varies significantly from pixel to pixel (*e.g.* conical or Gaussian images).

6.3.5 Gaussian peaks

If the autocorrelation peaks are Gaussian with an e^{-1} radius of R_a pixels, they will obey the following equation along $\Delta y = 0$:

$$A(\Delta x, 0) = 4e^{-(\Delta x - d_1)^2/R_a^2} + e^{-(\Delta x - d_2)^2/R_a^2}. \tag{6.8}$$

These autocorrelation peaks have an e^{-2} radius of $\sqrt{2}\,R_a$ and the above expression is valid for all Δx. Further, as $e^x \approx 1 + x$ for small x, it is apparent that the parabolic and Gaussian expressions will be very similar if $\Delta x \approx d_1 \approx d_2$, which will be the case when all three distances approach 64/3 pixels. Thus, a Gaussian peak model should retain the accuracy of the parabolic model in this limit whilst hopefully predicting a more gradual change in error for other values of d_1.

Taking $R_a = 1$ pixel and then differentiating the above expression:

$$\begin{aligned} \frac{dA}{d\Delta x} &= 4e^{-(\Delta x - d_1)^2}.-2(\Delta x - d_1) + e^{-(\Delta x - d_2)^2}.-2(\Delta x - d_2) \\ &= 2e^{-(d_2 - \Delta x)^2}(d_2 - \Delta x) - 8e^{-(\Delta x - d_1)^2}(\Delta x - d_1). \end{aligned}$$

On setting the above derivative to zero, only minor simplifications can be made. It is clear that a general expression for the maximum of equation 6.8 will not be as easy to obtain, solve or manipulate as was the case for parabolic peaks.

If the behaviour of the Gaussian model is to be tested over a wide range of values of d_1 and Δx, 'brute force' numerical methods can be used. For a given d_1, a computer can evaluate $A(\Delta x, 0)$ for a range of values of Δx using equation 6.8,

then these values can be examined one by one, to search for the largest. A sample plot of the resulting autocorrelation plane is shown in Figure 6.16, in which d_1 was fixed at 21.0 pixels whilst Δx was varied in steps of 0.005 pixels. The simplest way to increase the accuracy of this method is to use a smaller step size.

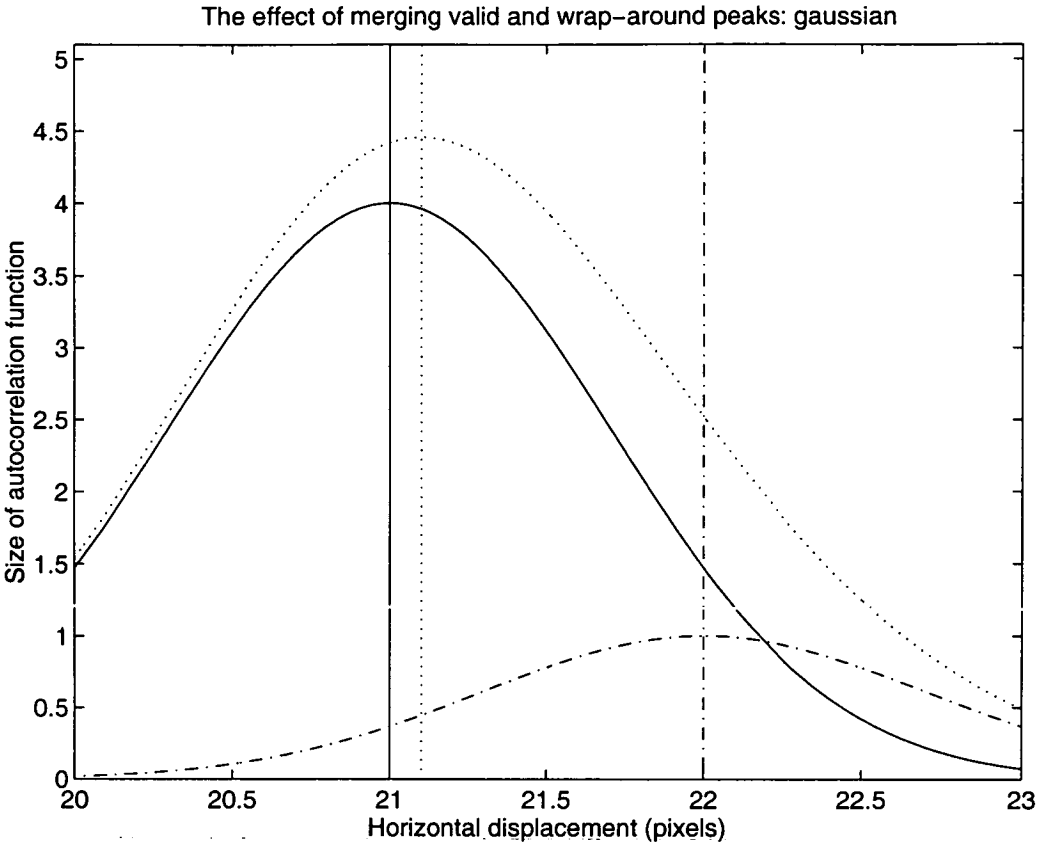


Figure 6.16: A section of the autocorrelation plane along $\Delta y = 0$, plotted in order to assess the predictions which result from assuming that autocorrelation peaks are Gaussian. Three images of each particle were simulated, 21.0 pixels apart. The fundamental displacement peak (solid curve) has four times the amplitude of the wrap-around peak (dash-dot curve). The summation peak (dotted curve) has a maximum at 21.10 pixels.

A test of the Gaussian model of autocorrelation peaks is shown opposite, in Figure 6.17 (top). The simulated data from the previous chapter is the same as that shown in Figure 6.15, except an extra 0.01 pixels was subtracted from the error points obtained using PIV simulations, to bring them closer to zero. In the Gaussian model, a ratio of four valid image pairs per wrap-around pair was used, along with an e^{-2} autocorrelation peak radius of 3.4 pixels (fitted by trial and error). Equation 6.7, the expected error for parabolic peaks, is also plotted.

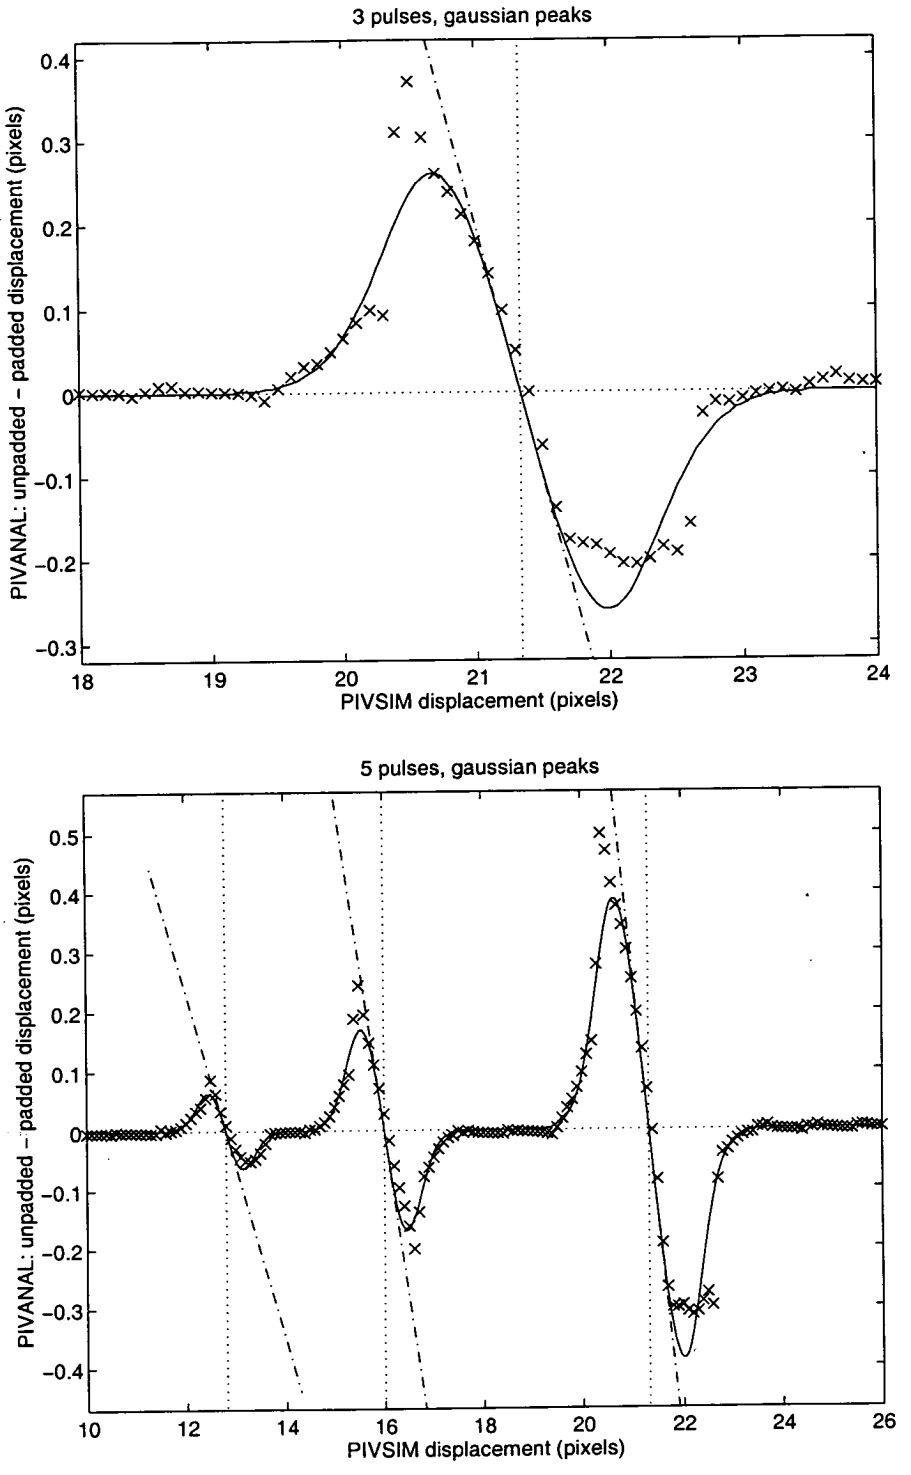


Figure 6.17: A comparison of simulated PIV results against the predicted behaviour of Gaussian autocorrelation peaks. Simulations were analysed both with and without FFT padding. The padded results plus 0.01 pixels were subtracted from the corresponding unpadded results, then the wrap-around error was plotted against the simulated displacement ($\times$ symbols). The numerically-predicted wrap-around error for Gaussian peaks is plotted using a continuous line, with the result for parabolic autocorrelation peaks also given for comparison (dash-dot line). Top: three images of each particle; Bottom: five images of each particle.

The Gaussian model with an e^{-2} autocorrelation peak diameter of 6.8 pixels gave quite a good fit to most of the simulated PIV analyses from the previous chapter, in which the e^{-2} particle image diameter was 4 pixels. The agreement at simulated displacements of about 20 pixels was pleasing, but deviations were apparent near the points of greatest over-read and greatest under-read error. As expected, the Gaussian and parabolic models converged as the simulated displacement d_1 approached 64/3 pixels. Even if autocorrelation peaks do not usually have parabolic cross-sections, the work using parabolic peaks has still proved useful.

In a further test, data from earlier PIV simulations using five images of each particle were compared to the predictions of this section. Once again, the results from fully-padded analyses were subtracted from those using unpadded analysis (Figure 6.17, bottom). For this number of images, the ratios of spurious wrap-around correlations to valid particle image correlations at displacements of 64/3, 64/4 and 64/5 pixels are 1:8/3, 1:6 and 1:16 respectively (equation 6.2). The e^{-1} autocorrelation peak radius which gave the best results was 2.35 pixels; this is equivalent to an e^{-2} radius of 3.32 pixels (very close to the value which gave the best results in the upper graph, for three images of each particle). The width of peaks in the autocorrelation plane should depend on the particle diameter in the interrogation area, but not on the number of images of each particle.

The errors predicted using parabolic peak models were:

$$\text{Error (pixels)} = \begin{cases} \frac{64-3\,d_1}{1+8/3} & \text{at } d_1 \approx 64/3 \text{ pixels} \\ \frac{64-4\,d_1}{1+6} & \text{at } d_1 \approx 64/4 \text{ pixels and} \\ \frac{64-5\,d_1}{1+16} & \text{at } d_1 \approx 64/5 \text{ pixels,} \end{cases}$$

from which a general pattern can very easily be hypothesised. It can be seen that the Gaussian and parabolic models are very similar at displacements close to the key values of 64/3, 64/4 and 64/5 pixels, as expected.

The Gaussian peak model produces a good estimate of the wrap-around error over a wide range of displacements, with the only significant deviations occurring close to the turning points predicted by the Gaussian model. This agreement

validates the Gaussian peak model and suggests that the ratio of wrap-around correlations to 'good' particle image correlations has been determined correctly.

It is possible that the deviations between the data from the previous chapter and the predictions from this model arise simply because different methods were used to locate the maxima to sub-pixel accuracy. With the Gaussian model, the values were examined one by one in a direct search for the global maximum, as equation 6.8 can be evaluated at any desired displacement. Accuracy was improved simply by reducing the step size by which Δx was varied.

In contrast, PIVANAL analyses interrogation areas containing image intensity data at intervals of 1 pixel. The resulting autocorrelation plane is also discrete, again having intervals of one pixel. To obtain sub-pixel accuracy, the program fits a Gaussian surface to a 3×3 grid of discrete data points centred on a peak in the autocorrelation plane; it then takes the vertex of this surface to be the measured displacement. The method used by PIVANAL might be affected by the addition of two adjacent peaks producing a distorted pseudo-Gaussian summation peak, whereas the 'brute force' method used here would have found the global maximum whatever the shape of the summation function.

Consider the shape of the summation function shown in Figure 6.16. The low displacement side of the summation function is very similar to the component signal $4e^{-(\Delta x - d_1)^2/R_a^2}$ and so must be almost Gaussian. The right-hand side of the summation function is somewhat flatter, however, so a good Gaussian fit to the points where $\Delta x < d^*$ will probably underestimate those points where $\Delta x > d^*$.

This is tested in Figure 6.18. In these two graphs, the summation function from Figure 6.16 (equation 6.8) is plotted. Seven points at intervals of 0.5 pixels are also shown, to emphasize the fact that PIVANAL must fit a Gaussian curve through discrete data points such as these. In the left-hand side of Figure 6.18, a Gaussian fit is made with its maximum at the same point as the global maximum: $(21.10, 0)$. The continuous curve clearly underestimates all points to the right of the maximum and so will not be favoured under a least-squares-fit calculation. A

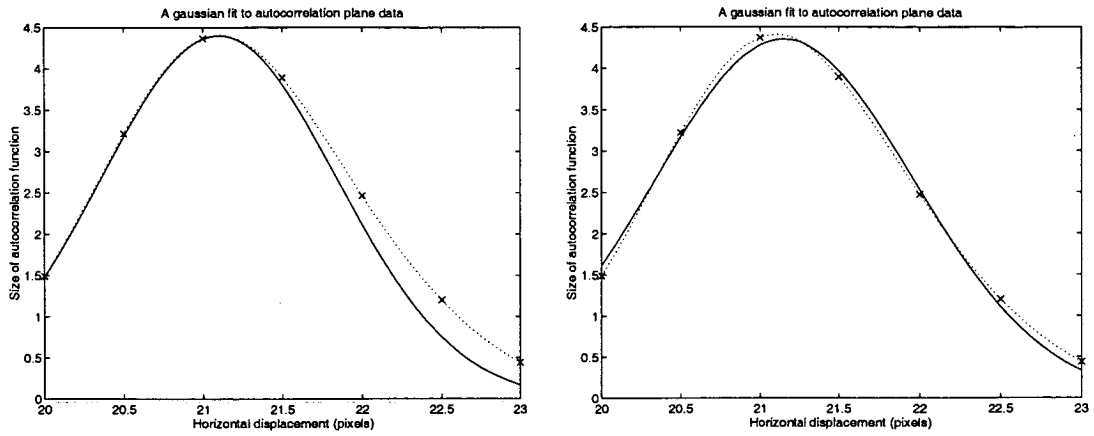


Figure 6.18: Gaussian fits to the predicted autocorrelation peak data from Figure 6.16, where the true image displacement was 21.0 pixels (dotted curves, $\times$ symbols). Left: the predicted autocorrelation plane values are fitted by a Gaussian function with a maximum value at the global maximum point (21.1 pixels) predicted by the Gaussian peak model. Right: the predicted data are fitted by a broader Gaussian function, whose maximum value occurs at 21.15 pixels.

better fit to the autocorrelation data points can be obtained if the best-fit function is moved slightly to the right of the true maximum, as shown in the right-hand side of Figure 6.18 (although that solution is not necessarily optimal).

Thus it is expected that displacements close to but less than $d_I/3$ will be over-read if the autocorrelation analysis does not use zero-padding:

1. because the addition of the intended autocorrelation peak and the wrap-around peak produces a summation signal with a global maximum at an intermediate position, closer to the point $(d_I/3, 0)$; and

2. because the Gaussian fit made by PIVANAL will tend to overestimate the displacement at which the maximum of this summation signal occurs, thus increasing the amount by which the displacement is over-read.

In contrast, displacements close to but slightly greater than $d_I/3$ should be under-read by PIVANAL, to an even greater extent than that predicted by the Gaussian peak model. This is because the Gaussian fit to the autocorrelation data involves positioning the summation peak closer to $d_I/3$, compounding the error due to the fact that the summation peak is too close to the origin.

Thus, it is plausible that, for displacements close to but slightly less than $d_I/3$, points calculated by PIVANAL by fitting a Gaussian surface to the autocorrelation plane have a greater error than that predicted by this Gaussian peak model (as shown in Figure 6.17). However, it is not clear why displacements slightly greater than $d_I/3$ appear to have *less* error than is predicted by this model.

6.4 Useful images from other flow records

The above analyses for 'horizontal' displacements were possible because the image displacements were parallel to one of the sides of the interrogation area, thus simplifying the geometry. Although a number of flows *will* produce images which satisfy this condition, it will not always be the case. Without wishing to delve in any great detail into the complications caused by displacements which are not parallel to any side of the interrogation area, it is nevertheless possible to show that this 'good image' method can be adapted for more general situations.

A relatively simple case to consider is shown in Figure 6.19, which examines the case of five images of each particle and a displacement of $(16, 16)$ pixels between consecutive illumination pulses. Once again, it is assumed that the interrogation areas are 64×64 pixels. The construction involves marking a number of 16×16 squares on the region around (and including) the interrogation area. It so happens that all area elements in this case are 16×16 squares.

It can be seen that the probability of finding four images from a particle with this displacement is rather low. Track A will give rise to only one image in the interrogation area, whereas tracks B and E will give rise to two images and track D will produce three (Figure 6.19, left). Something along the lines of track C is required if the full complement of four images is to be obtained (the greatest possible number in this case).

In the right-hand side of Figure 6.19, the possibility of the first image from each track lying in each square is considered. The number of images from the

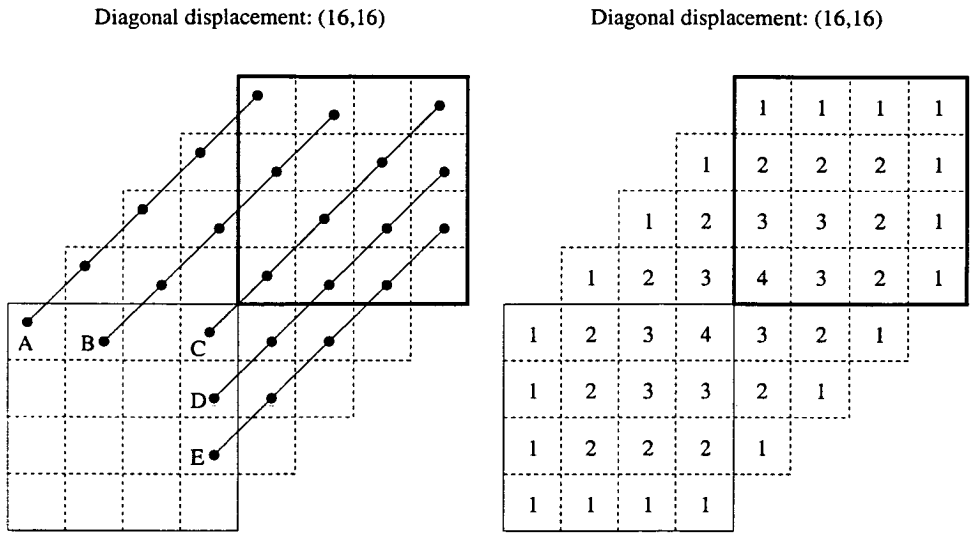


Figure 6.19: For the measurement of a diagonal displacement (16, 16) using 64×64 pixel interrogation areas, five-images tracks can cause between 1 and 4 particle images to be found in the interrogation area (thick square). Left: sample places in which the first image of the track (darker circles) can lie, in order that at least one image from the track lies within the interrogation area; Right: the number in each square indicates the number of images which will be found within the given interrogation area, if the first image of the five-image track lies in that square.

track which will lie inside the interrogation area is noted in the centre of each square. It can be seen that only two possible starting zones enable four images to be found. There is a total of 44 possible 16×16 squares in which the first image of a particle track must lie, if at least one image from the track is to be found within the interrogation area. Thus, the probability of four images being found is only 2/44. In contrast, the probability that only a single, unpaired image will be obtained is 20/44, *i.e.* almost 0.5. In the simulations of the previous chapter, it was noted that yields of valid vectors were poor for these diagonal displacements.

The left-hand half of Figure 6.20 shows a more general situation, where the displacement is $(12, 8)$. A number of significant 12×8 rectangles can be marked in and around the interrogation area. However, unlike in the $(16, 16)$ case, rows and columns of even five of these areas cannot cover the full width or the full height of the interrogation area. So, a region within the interrogation area measuring $(64 - 4 \times 12) \times (64 - 4 \times 8)$, *i.e.* 16×32, can be identified in the bottom left-hand corner of the interrogation area; the first image of a track must lie in this region, if all five images are to be found in the interrogation area (track C).

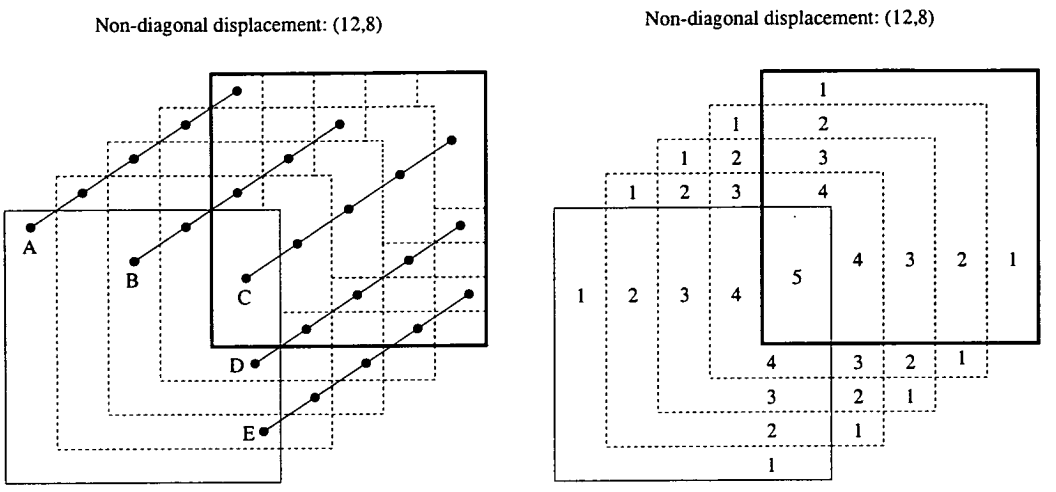


Figure 6.20: For the measurement of a displacement $(12, 8)$ using 64×64 pixel interrogation areas, tracks which are five images long can cause between 1 and 5 particle images to lie in the interrogation area. Left: some sample places in which the first image of a particle track (darker circles) can be situated, in order that at least one image from the track lies within the interrogation area (dark square); Right: the number within each square indicates the number of images from a track of five images which will lie within the given interrogation area, if the first image of the track is found in the given square.

A simplified diagram is shown in the right-hand half of Figure 6.20, in which the number of images found in the interrogation area is shown for each possible locus of the first image. It can be noted that the diagram is rotationally symmetric about the centre of the region labelled '5', the starting region from which five images will be obtained. This point can be thought of as the centre of gravity of the two continuous-sided squares shown in this figure, which represent the interrogation area (thick-edged square) and an image of the interrogation area translated by a displacement $(-48, -32)$, *i.e.* $-4\times$ the particle displacement (the thin-edged square). The advantage of this definition of the point of rotational symmetry is that it applies even when the displacement is large (*e.g.* for the $(16, 16)$ displacements in Figure 6.19, where the two squares had no area of overlap).

A justification for this symmetry can be obtained by considering what the diagram would look like if the displacement was rotated through $180°$ to become $(-12, -8)$. In the $(-12, -8)$ case, the rôles of the two 64×64 squares in Figure 6.20 would be reversed and the particle images would be taken to have arrived in the

reverse order. Because displacements of $(12, 8)$ and $(-12, -8)$ cannot be distinguished by this PIV system, the yield of valid vectors must be the same in each case, as must the probability of finding (say) exactly two images of a particle in the interrogation area. Rotational symmetry should be expected.

Following the simplified diagram in the right-hand side of Figure 6.20, a generalised diagram can be proposed which illustrates the number of images in the interrogation area which results from a particle track starting in each suitable locus. Suppose that N images of each particle are recorded and that consecutive particle images are a displacement $\boldsymbol{d}$ apart. The generalised diagram consists of the interrogation area itself and $N-1$ intersecting copies of the interrogation area, each produced by translating the interrogation area through a displacement $-\boldsymbol{d}$, $-2\boldsymbol{d}$, $-3\boldsymbol{d}$, $\ldots$, $-(N-1)\boldsymbol{d}$. The intersection of the edges of the various squares divides the relevant area into a number of smaller sub-regions. A particle track which starts in the outermost layer of sub-regions in this diagram will give rise to only one image in the interrogation area. A track which begins in the next outermost layer of sub-regions will give rise to two images in the interrogation area, *etc.* Hopefully the observations in this section will prove to be a useful starting point for further investigation of this topic.

Reconsidering the work given earlier in this chapter, it can be seen that the near-horizontal displacements of Figures 6.2, 6.3 and 6.6 could well have been analysed differently. All of the starting zones shown in these figures could have been obtained by drawing the interrogation area and various copies obtained by translations. The earlier analysis could perhaps be reconsidered in the light of this more general method.

Chapter 7

Experimental apparatus

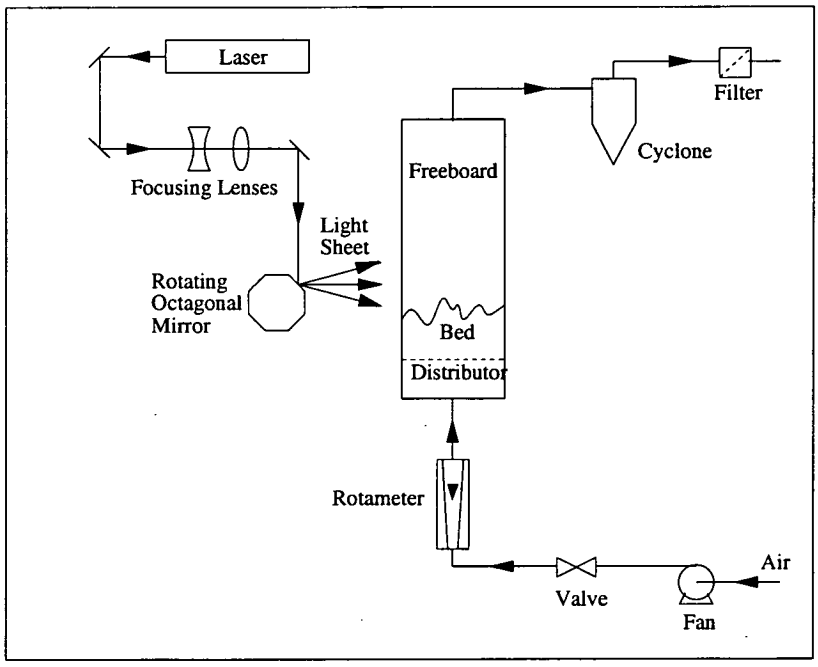


Figure 7.1: A schematic diagram of the apparatus used in the initial work.

The apparatus used is shown in Figure 7.1. A laser beam was directed at a rotating mirror, to produce a sheet of high-intensity light. This was shone at a glass section above the fluidised bed (the 'test section'). Air passed through a bed of glass particles, then discharged into the lab after passing through a cyclone with a filter on the exit stream. The practical details of this arrangement, however, varied according to the scale of the experimental rig.

7.1 Initial small-scale work

Early work on a 3 inch square bed used a bench-top system. This apparatus could easily be dismantled between runs or removed when not in use. The fluidised bed was a modification of one used in a previous fluidisation project, by Lang (1986).

The fluidised bed consisted of an air box and porous-plate distributor unit, a number of 3" × 3" × 10" bed sections and a tapered exit section (at an angle of 15°). Pressure tappings had been made into the air box (below the distributor) and one of the bed sections. All of these components were attached to flanges, so the rig could easily be bolted together.

Two of the three bed sections were made of perspex. The third section was made of glass, which is optically superior to perspex and hence more suitable for PIV work; this section was the one used for PIV measurements. On top of these sections sat a reducing duct with a square cross-section, in which the area available for flow was gradually reduced from 7.6 cm × 7.6 cm until it met a 2.8 cm × 2.8 cm block, into which a 1.25 cm-diameter exit nozzle was set. Gas leaving the fluidised bed passed through a cyclone and then a fine filter bag, before discharging into the lab. The filter bag trapped particles bigger than about 1 μm.

The source of illumination was a continuous-wave Spectraphysics 171 Argon Ion laser, which emits a blue-green beam with a wavelength of about 490 nm. This laser needs to be used with a water-conditioning unit, which filters and de-ionises 'clean' water in a closed laser-cooling circuit, whilst transferring the considerable heat given out by the laser to a second, external water circuit. This laser emits a power of about 0.6 W at its minimum operational current. The power can be reduced to a few milli-Watts by means of aperture and attenuator controls, or boosted up to about 20 W by increasing the current.

Because of the small scale, the laser was housed on a sturdy optical bench and mirrors were used to direct the beam along the relatively short path to the rotating mirror. It was found that the light sheet could be focused down to a

narrow waist (about 0.5 mm) if the beam passed through a diverging lens followed by a converging lens. The lenses were chosen so that the converging lens had a slightly shorter focal length than the diverging lens. A final mirror deflected the beam onto an octagonal rotating mirror, the speed of which could be carefully controlled. The time interval between successive illumination pulses could be kept constant at a value in the range 0.5–12.5 ms. Light from the rotating mirror was not collimated into a parallel sheet using a parabolic mirror in this early work, but was simply allowed to diverge into an expanding sheet. The fluidised bed could be held in place close to the rotating mirror, so the expansion was not a problem.

For safety, a hardboard hutch was assembled around the optical components, allowing only the expanding light sheet to emerge into the lab. This reduced the risk of accidental exposure to the high-intensity laser beam. A beam with a light intensity of more than 5 mW can cause instant blindness if it is somehow reflected into an eye, because damage occurs before the eye's blink reflex can take effect. A scanning light sheet is less dangerous than a steady laser beam of the same power, because the scanning action distributes light energy over a wide area and reduces the time in which damage can take place. Nevertheless, the Class 4 laser in this project was used at powers of up to 12 W, so extreme caution was required with both the expanding light sheet and the raw beam from the laser.

Other precautions included:

- carrying out all experiments in a designated, locked, blacked-out room, with an interlocked sign outside the door warning others when the laser was on;
- using a beam stop after the light had passed through the test section;
- aligning the beam at minimum power and only increasing the power after the alignment was complete;
- manipulating the beam only after removing watches and other jewellery;
- wearing eye-protecting goggles when working with the beam at high power; as well as

- various other recommendations made by the University of Edinburgh (1993) as part of its Health and Safety Policy.

Fluidisation was carried out using air, compressed by a small fan. This was capable of delivering a superficial velocity of 0.03–0.25 $\mathrm{m\,s^{-1}}$ (12–100 $\mathrm{l\,min^{-1}}$). The air flow was measured using a suitable rotameter and adjusted using a hand valve.

The bed fills consisted of lead glass ballotini. Two different fractions were used, with sizes in the ranges 53–106 μm ('AG grade') and 150–325 μm ('Grade 10').

Drag co-efficient calculations for spherical particles, assuming a typical glass solid density of $\rho_s \sim 2500\,\mathrm{kg\,m^{-3}}$, suggested that the terminal velocities of these particles were approximately $U_t = 0.4\,\mathrm{m\,s^{-1}}$ and $4\,\mathrm{m\,s^{-1}}$ respectively. Using the Ergun equation, minimum fluidising velocities were estimated to be $0.0067\,\mathrm{m\,s^{-1}}$ and $0.070\,\mathrm{m\,s^{-1}}$ respectively. Experimental measurements with the larger, Grade 10 ballotini found $U_{mf} = 0.0715\,\mathrm{m\,s^{-1}}$, which was a good agreement (Figure 7.2).

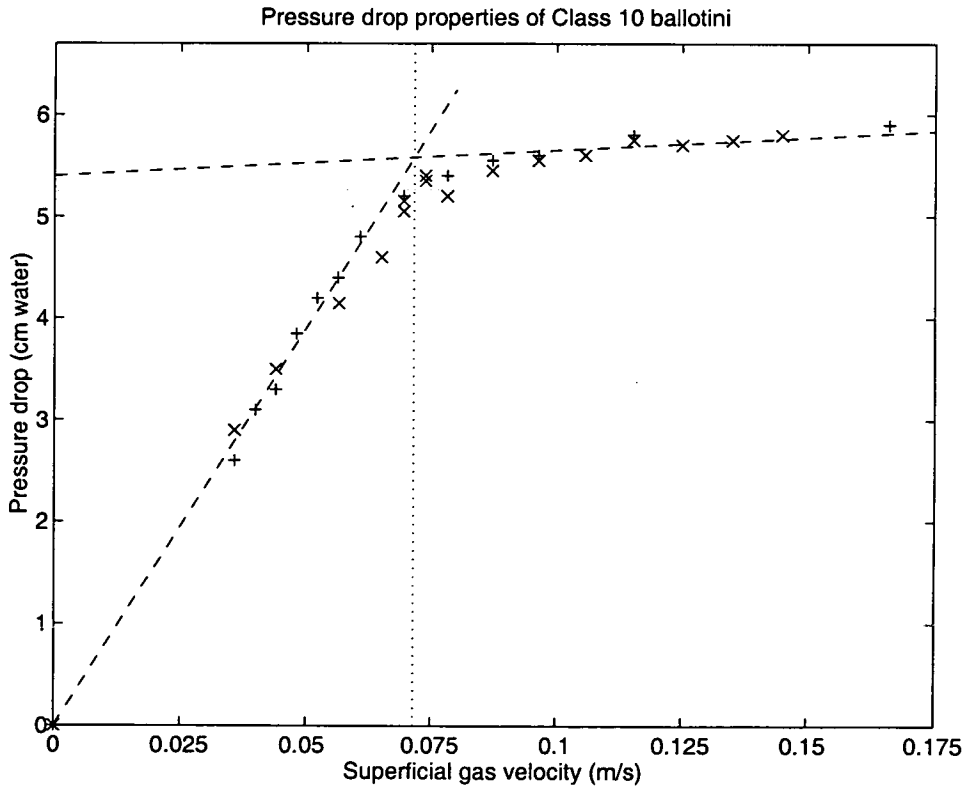


Figure 7.2: Pressure drop characteristics of Grade 10 ballotini. The two dashed lines intersect at a value $U_{mf} = 0.0715\,\mathrm{m\,s^{-1}}$ (dotted line).

In most experiments, small 'Sphericel' hollow glass spheres were added to the bed, in order to boost the fines concentration in the bed and to seed the gas flow in the freeboard. These seeding particles have a stated median diameter of $8\,\mu$m, with 30% of the particles greater than $10\,\mu$m and 5% greater than $20\,\mu$m. The hollow nature of the seeding particles reduces their effective density to about $1100\,\mathrm{kg\,m^{-3}}$ (less than half that of solid glass). Together with their smaller size, this means that the Stokes Law terminal velocities of the seeding particles were about $2.1 \times 10^{-3}\,\mathrm{m\,s^{-1}}$ (based on the median diameter of $8\,\mu$m) to about $4.8 \times 10^{-3}\,\mathrm{m\,s^{-1}}$ (assuming that the finest of these particles are more readily elutriated, so that the mean diameter of the remaining particles is about $12\,\mu$m).

The value of terminal velocity for the median-diameter hollow glass spheres ($2.1 \times 10^{-3}\,\mathrm{m\,s^{-1}}$) is about 200 times lower than that of the solid 53–106 μm AG grade particles, and about 2000 times lower than that of the Grade 10 ballotini. These seeding particles should be easily elutriated from the apparatus. When seeding a gas flow of $0.07\,\mathrm{m\,s^{-1}}$, *i.e.* when $U = U_{mf}$, this terminal velocity works out to be only 3% of the flow velocity. The velocity relaxation time of these spheres (τ) was calculated to be 0.20 ms (section 4.1.4). The time between successive illuminations was always several milli-seconds, so the particles should be able to follow variations in the local fluid velocity between illumination pulses.

These seeding particles were mixed into the bed with the air flow off, then each experiment was run as a batch process. If the initial seeding concentration within the fluidised bed was somewhat higher than the ideal, it was likely that the seeding concentration would fall to a satisfactory level for an adequate period of time, at some point during the course of each experimental run. In PIV work, the concentration of seeding particles is quite important: too much seeding means that a cloud of particles between the camera and the light sheet acts as a fog obscuring the light sheet. (This fog could be a combination of airborne particles within the bed and particles which have become attached to the walls.) On the other hand, an inadequate seeding concentration in some regions of the flow means that the displacement correlations calculated by the analysis method have a low signal-to-noise ratio, so the calculated velocities are prone to error.

In the initial small-scale work, partly dismantling the fluidised bed was an easy task. This allowed seeding particles to be added to the bed, whilst the inside surfaces of the test section could also be cleaned. With the lower sections of the bed in place but some of the upper sections removed, it was possible to wrap a cotton thread and Blu-tack 'plumb line' around a rod, which could be rested on the upper flange of the highest remaining bed section. With the laser operating at minimum power, the light sheet could be directed at the plumb line, in approximately the centre of the bed. This made it possible to check if the light sheet was vertical.

Lining up the fluidised bed test section with the light sheet was fairly straightforward. The procedure ran along the following lines:

- The fluidised bed could be made vertical by aligning it with the plumb line.
- When the incident light sheet passed through the glass in the wall of the test section which lay nearest to the rotating mirror, part of it was reflected back towards the rotating mirror (eventually being stopped by the wall of the protective hardboard hutch). If the bed was rotated about a vertical axis so that this light reflected back into its incident plane without being deflected to either side, a position could be found in which the light sheet entered the bed perpendicular to the plane of two of the bed walls.
- Without rotating the bed further, it was then possible to translate the bed perpendicular to the illumination plane, until the light sheet entered the bed at the required position (*e.g.* through the centre of a side wall). Reference was made to various marks added to the outside of the bed.
- The camera could now be placed in position. Ideally, the camera would be levelled on its stand or tripod, using spirit levels. Looking through the camera at the light sheet, two vertical lines could be observed, due to light scattering from the glass surfaces of the test section on both the entry and exit sides. The camera could be moved parallel to the light sheet, until the two lines were seen in the centre of the field of view.

- All would now be well if the camera was looking at right angles to the test section. This could be tested by moving the plumb line to the centre of the light sheet, where it would be co-planar with vertical marker lines on the front and back walls. Looking through the camera with the light sheet on and the lab lights off, the plumb line could hopefully be observed to be coincident with its reflection in the back wall of the test section.

- In some situations, it was helpful to direct a spotlight at the front of the camera and then, in an otherwise darkened room, focus the camera on the image of itself reflected in the front wall of the test section. This would also reveal whether the camera image plane was parallel to the light sheet, *e.g.* rotations of the camera about a horizontal axis would be detected, by the camera's image appearing to point up or down.

- With the vertical light sheet now passing through the centre of the bed and entering or leaving the test section at right angles to the side glass walls, a camera (viewing along a normal to the light sheet) could now be focused on the plumb line. This was found to be an easy way of achieving an accurate focus, with all types of camera. Focusing was usually easiest if a spotlight was directed at the plumb line, with the laser illumination blocked off.

The position of the camera relative to the light sheet could then be checked. A focused camera placed too close to the light sheet will not be able to photograph the whole of the intended area of the light sheet. If need be, the camera could be moved closer to or further away from the light sheet, before repeating the above steps. In fact, the whole alignment process is rather iterative, as adjustments need to continue until all of the above criteria are met simultaneously.

Provided that the optical components and the camera were not disturbed after the focusing had been carried out, the plumb line could be removed and the rest of the fluidised bed reassembled, with the camera still focused on the position of the light sheet.

Mounting a piece of dark card on the outside of the back wall of the fluidised bed was helpful, as it was found that the seeding particles were more visible if they were photographed against a dark background. Similarly, it was worth ensuring that the camera only received light from the test section, by giving it 'blinkers'.

7.2 Later, larger-scale work

In later work, a larger fluidised bed was used, in a busy fluid dynamics lab. Various minor modifications to the above arrangement were required.

The laser was now housed in a hut within the lab. The door of the lab was self-locking and the door of the hut could also be locked. In this large multi-user lab, there were long beam paths between the (fixed) laser and each experimental rig, with many right-angle bends required. Ensuring good alignment (and thus good beam propagation) in such a system would have been difficult: even a small error in the orientation of one mirror would cause large deflections 20 m and several more mirrors down the line. Further, each mirror would dissipate 5–10% of the incident light, so the overall transmission efficiency would not be good.

Instead, a fibre-optic delivery system was used, based on a fibre-optic cable enclosed within a protective 30 m length of plastic tubing. The main alignment task at the laser end was to align the fibre's entry head with the beam from the Argon Ion laser. This procedure is rather slow and delicate, but only needs to be done once if the laser and the entry head are securely clamped in place. Optimal alignment gave transmission efficiencies of about 60–70% overall. The laser could only be used at powers of less than 12 W, to avoid damaging the fibre.

At the delivery side, the output from the fibre was focused using a converging lens, then reflected off a rotating octagonal mirror and collimated using a parabolic mirror. The optical components and the mount for the fibre-optic delivery head were held in place inside a scanning-beam box (Figure 7.3). This meant that the light sheet could be moved to a different location in a matter of minutes.

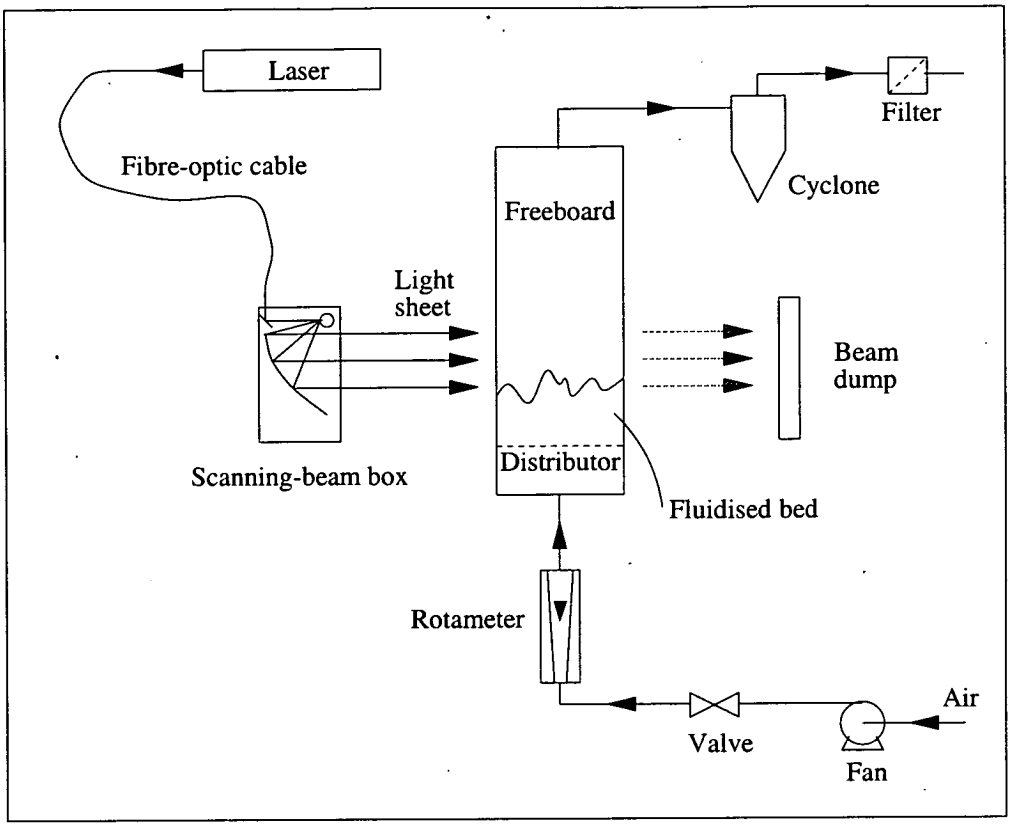


Figure 7.3: Schematic diagram of the apparatus used in the larger-scale work.

The larger fluidised bed was square, with sides 12" wide. There were a number of stacked glass sections, all of which were suitable for use in PIV. One of these sections was 0.6 m tall, whilst the others were 0.9 m tall. It was originally thought that these sections might be interchanged periodically, as was done using the small-scale rig, if it was considered necessary to take PIV photographs at a position which, in the original arrangement, was partly obstructed by the flanges joining two bed sections. This need was never felt to be sufficiently acute to justify the effort, so the sections were not actually rearranged in the course of this work.

With a total height of about 5 m, it was not practical to dismantle the larger rig between runs, in order to clean the glass surfaces. Instead, a 15 cm-diameter manhole was cut into the back wall of each test section, to provide access to the inside of the bed. Each test section could be cleaned in this way.

At the top of the bed, a 0.5 m-tall section was used to reduce the area available for air flow. In this section, the corners of the square apparatus were progressively

filled in until the area available for flow was 30 cm in diameter, then an aluminium cone gradually reduced the diameter to 10 cm. The cone angle was about 15°. A length of 4"-diameter flexible hosing carried the exit air stream down to about head height and then into a cyclone.

In order to support it, the rig was enclosed within a framework of scaffolding. A wooden platform was added, which could be moved to any required height. The camera, the scanning-beam box and the other equipment required for PIV could conveniently be housed on this platform.

Sheets of hardboard could be bolted onto the sides of the scaffolding. The main reason for this was to ensure that the light sheet was directed at a beam stop and enclosed as far as possible, reducing the risk of dangerously-powerful beams of stray light escaping into the lab. However, an added advantage was that external light was excluded, thus allowing high-contrast photographs to be taken.

Air to this new rig was supplied by a compressor in an adjacent laboratory. This compressor was capable of developing high pressures but delivered only moderate flowrates (up to about $1000\,\mathrm{l\,min^{-1}}$) for a device of its size. By fully-charging a pressurised reservoir immediately downstream of the compressor and running experiments with the compressor on, it was possible to pass air at a superficial velocity of $0.2\,\mathrm{m\,s^{-1}}$ through a 0.6 m-deep fluidised bed for about 5 minutes.

In order to align the camera with the light sheet, the plumb line method was adapted, with a 'T' type arrangement proving successful. One length of cotton thread acted as the plumb line, with a blob of Blu-tack acting as a weight at one end whilst the other end was tied into a loop. A second length of thread passed through this loop; each end of this thread was attached to a piece of Blu-tack, so that these ends could be attached to opposite inner surfaces of any test section. This formed a movable support rail, along which the plumb line could move.

The particles used in this bed were a standard grade of 'impact beads'. The nominal size range of these particles was 180–300 μm, with about 0–10% of the

particles expected to be smaller than 180 μm and about 0–10% expected to be larger than the nominal maximum (in a 300–355 μm size band). These proportions were determined by the manufacturer, using a sieve analysis. The fluidisation properties of these beads were shown earlier in Figure 2.2, where the minimum fluidising velocity was found to be about 0.070 m s^{-1}.

The particles were run through the department's Malvern particle size analyser, which uses a diffraction method. Four runs were performed and analysed, to see if the distribution followed a normal or a log-normal pattern (Figure 7.4). An RRSB model was also tried, but this gave a poorer fit.

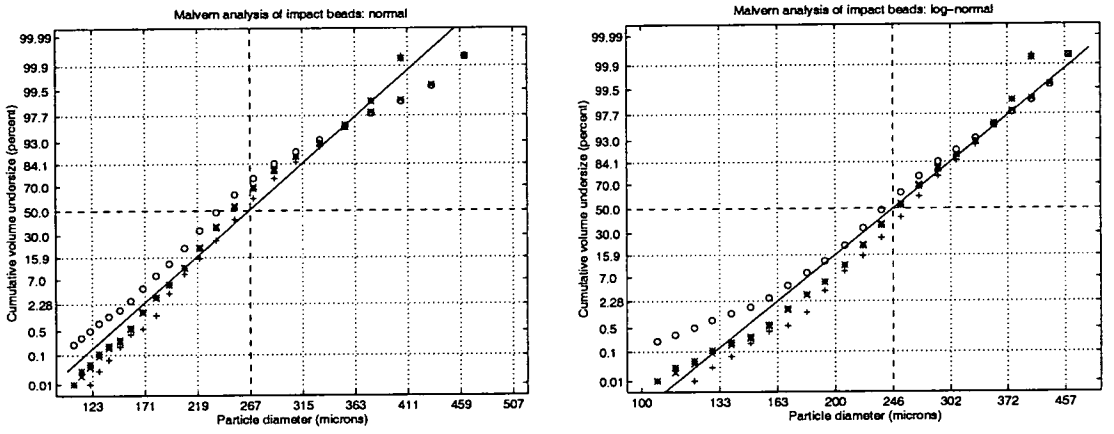


Figure 7.4: Normal and log-normal fits to size distribution data for 180–300 μm impact beads. The measurements were made using four separate runs on a Malvern particle size analyser. The median size is indicated using a dashed line.

The test for a normal distribution is that a straight line is obtained when the cumulative fractional mass undersize is plotted on a probability scale, against the particle size (on a linear scale). In the left-hand half of Figure 7.4, it can be seen that the normal distribution gives a reasonable (but not entirely convincing) straight-line fit, based on all of the data points. A median diameter of 267 μm was obtained, along with a standard deviation of 48 μm. However, in the central region of the plot, where the percentage mass undersize is between 7 and 93% ($\pm$1.5 standard deviations), the data does not appear to be particularly linear and the overall best-fit line misses the data points closest to the 50% mark.

The requirement for a log-normal distribution is that a straight line is obtained when the fractional mass of all particles smaller than a given size is plotted on

a probability scale, against the logarithm of the chosen size. This model gave a somewhat better fit to the data (Figure 7.4, right), particularly if it was possible to ignore the 2% by mass at either end of the distribution. It can be noted that this material would be more than two standard deviations away from the median. The result was a median size of about 246 μm, with the standard deviation found to be a factor of about 1.23 (meaning that diameters of $1.23 \times 246\,\mu$m and $(1/1.23) \times 246\,\mu$m are both one standard deviation away from the median).

These results agree with the assertion that fewer than 10% of the particles are smaller than 180 μm, but suggest that a larger proportion than specified are bigger than 300 μm (16%, rather than 10%). The agreement is probably as good as can be expected from two quite different methods of size analysis. Perhaps the sieve trays allowed some large particles to pass through which should have been retained, *e.g.* due to the occasional oversize pore or particle attrition whilst sieving; perhaps the Malvern analyser calculated excessively large particle sizes if two particles happened to be unusually close together as they passed through the laser beam.

In some experiments, image-shifting was used. As well as providing a scanning-beam box, the University of Edinburgh spin-off company Optical Flow Systems supplied a ISS–200 image-shifter unit. This consists of a rotating mirror on a turntable, a mount for a camera, a control box and various control leads.

When not in use, the turntable of this image-shifter unit returned to its rest position ('45°'), to allow alignment to take place. When primed for action, the turntable rotated back a short distance and waited for a trigger signal. When this came, the turntable accelerated, rotated for a short distance at constant speed, then decelerated and returned to its cocked-back position, ready for the next photograph. If the camera exposure time was chosen to be 1/30 s, it seemed reasonable to use a mirror sweep time of about 300 ms, to help achieve a constant rotational speed. The mirror still needed to decelerate and return to its primed position, so it was not really possible to capture much more than one or two frames per second with this device. Nevertheless, this was an improvement on an

earlier model, which rotated continuously at a constant speed and allowed only one exposure to be taken per revolution (*e.g.* every 36 s when rotating at $10^\circ\,s^{-1}$).

An RS–232 serial link from the image-shifting control box could be plugged into a PC, which acted as the user interface. The user gave the control box values for the exposure time, the shutter delay, the rotating-mirror sweep time, the angular rotational speed (minimum: $10^\circ\,s^{-1}$), *etc.* The control box then arranged for the rotating mirror to pass through its rest position, at constant rotational speed, at the exposure centre point (*i.e.* at a time exactly halfway through the exposure). Measured experimentally, the shutter delay is the time interval between a trigger signal reaching the camera and the shutter starting to open.

One input lead to the image-shifter allowed a pulse to be sent to the control box, which could cause a photograph to be taken a given time after the arrival of the pulse. This signal could be linked to some action in the experimental apparatus, *e.g.* a signal sent to a wavemarker. Alternatively, an image-shifted photograph could be triggered by typing characters on the keyboard (as was done in this work).

Chapter 8

Experimental results with the smaller-scale apparatus

Table 8.1 lists the main rolls of film taken with the smaller fluidised bed:

Bed material size (μm)	Bed depth (cm)	Viewing height (cm)	U/U_{mf}	Film number	Frames analysed
53–106	5	18	30	3	none
53–106	5	18	19	3	none
180–300	3	18	1	39	8–24
180–300	3	33	1	37	13–24
180–300	5	17	1	41	8–23
180–300	5	33	1	33	0–5, 11–16

Table 8.1: Summary of films taken using PIV to record the flow through the centre of the freeboard space above the smaller fluidised bed.

Early results from these runs have appeared in print, in refereed papers by McCluskey *et al.* (1994) and Rix *et al.* (1996) which are appended to this thesis (with permission). Each film has been re-analysed here, using larger interrogation areas than before in order to achieve significantly better yields of vectors.

8.1 Fluidised beds of AG grade ballotini

The first attempts at PIV were made with the AG grade ballotini. These particles had a size range of 53–106 μm according to the manufacturer, a Stokes Law terminal velocity of about $0.4\,\mathrm{m\,s^{-1}}$ and a minimum fluidising velocity of about $0.0067\,\mathrm{m\,s^{-1}}$ (according to the Ergun equation and a force balance).

No seeding particles were added to this grade of ballotini. In the first part of film three, an air flow of $0.19\,\mathrm{m\,s^{-1}}$ was used. The ratio U/U_{mf} was approximately 30, so the fluidisation was very vigorous. The exact depth of bed was not recorded as this was only intended to be a trial run, but was probably about 5 cm or so. Some sample photographs under these conditions are given in Figures 8.1 and 8.2, as well as in the top half of Figure 8.3.

In these photographs, for example in film 3 frame 4 (Figure 8.1, top), a number of features can be noted. Light was scattered by particles in the flow and also by the inner walls of the test section (causing two bright vertical lines to be seen). As this test section was 7.6 cm×7.6 cm in horizontal cross-section, it was easy to determine the image:object magnification in the photographs and the area of the test section recorded in each photograph (about 7.6 cm horizontally by about 8.8 cm vertically). In fact, *two* distinct vertical lines can be seen at each wall, consisting of a bright line (due to the scattering at an inner air-glass interface, possibly helped by reflections from particles stuck to the wall) and a less bright line (perhaps due to light being scattered from the outer glass-air interface). Moreover, as the inside of the back wall was not painted black and so could reflect incident light, the camera also recorded the reflections of these vertical lines in the back wall. These lines can be seen approximately 5 and 11 mm across the photograph, measuring from the brightest vertical line at one wall towards the other.

A change in background material can be seen near the top of the photographs. A piece of black card was placed behind the glass back wall of most of the test section, as a dark background. However, an aluminium flange connected the test section to the other sections of the bed; this surrounded the glass at either end

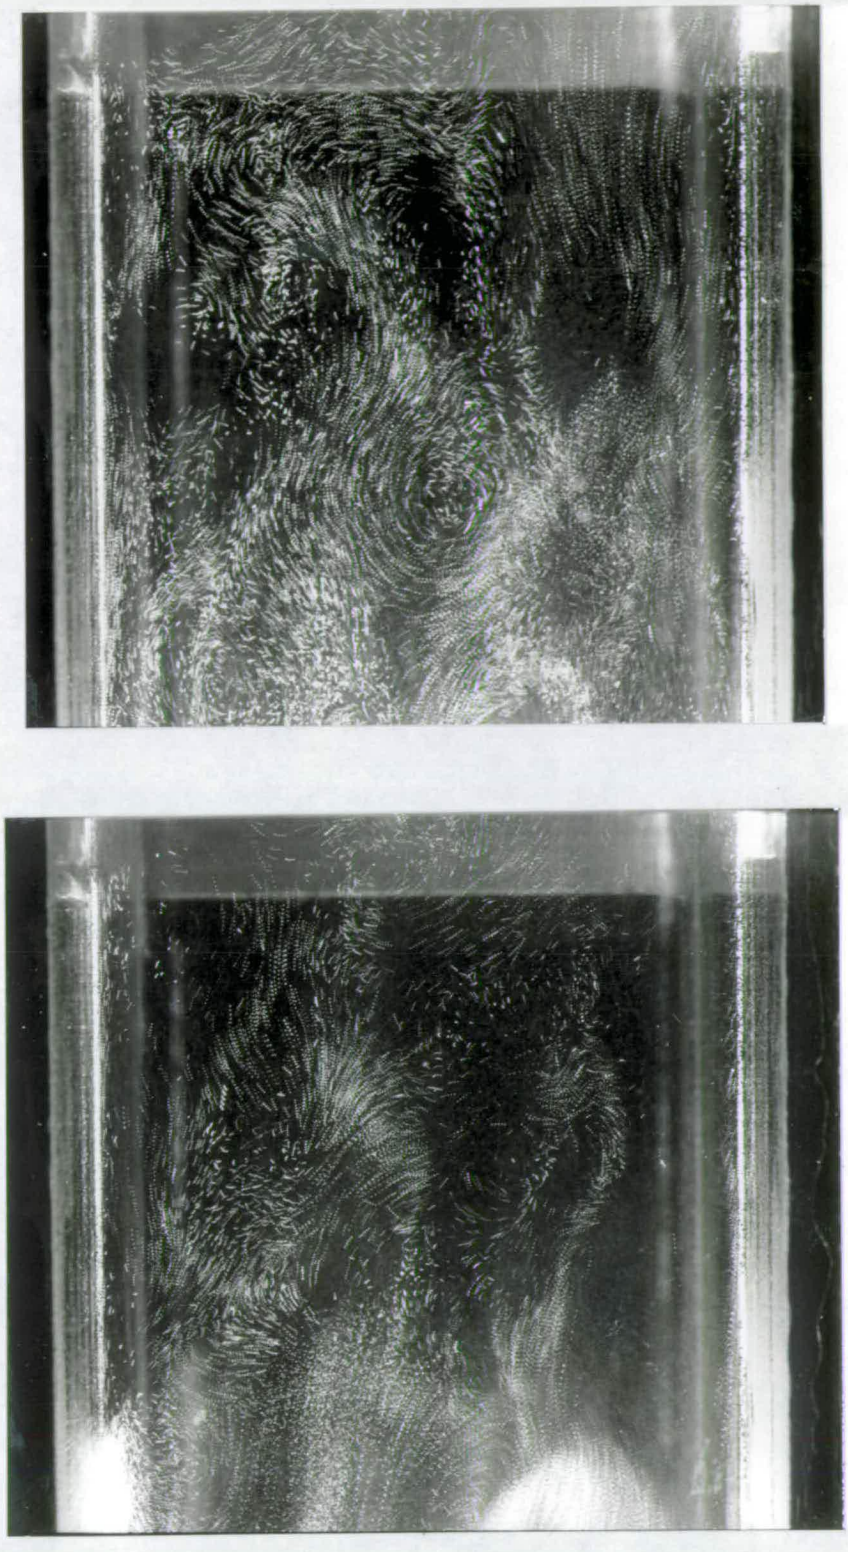


Figure 8.1: PIV photographs for the flow about 10 cm above a bed of AG grade glass particles (film 3, frames 4 (top) and 5 (bottom). $U/U_{mf} \approx 30$.

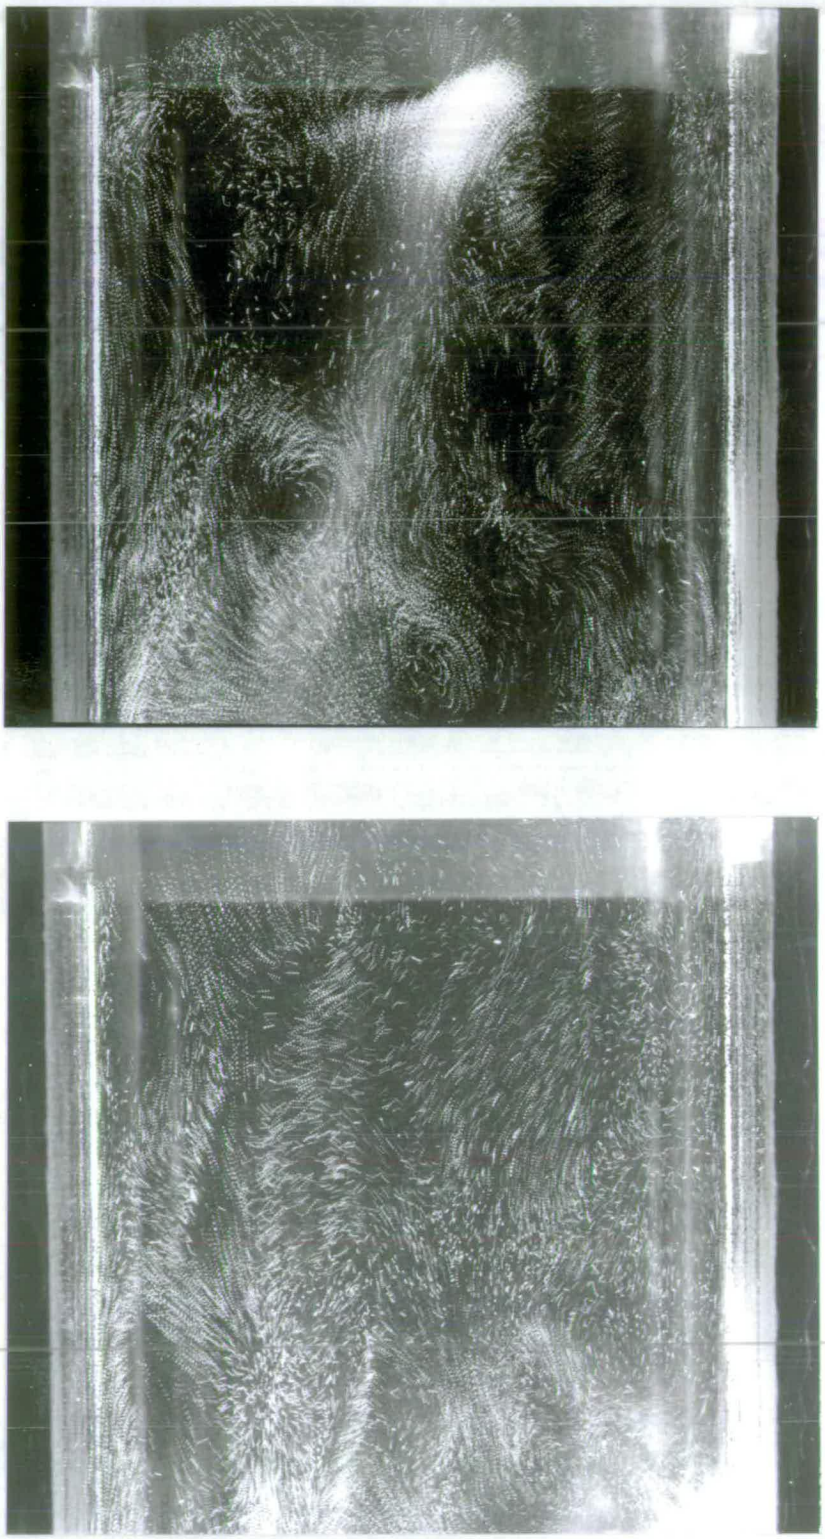


Figure 8.2: PIV photographs for the flow about 10 cm above a bed of AG grade glass particles (f3f6 and f3f7). $U/U_{mf} \approx 30$.

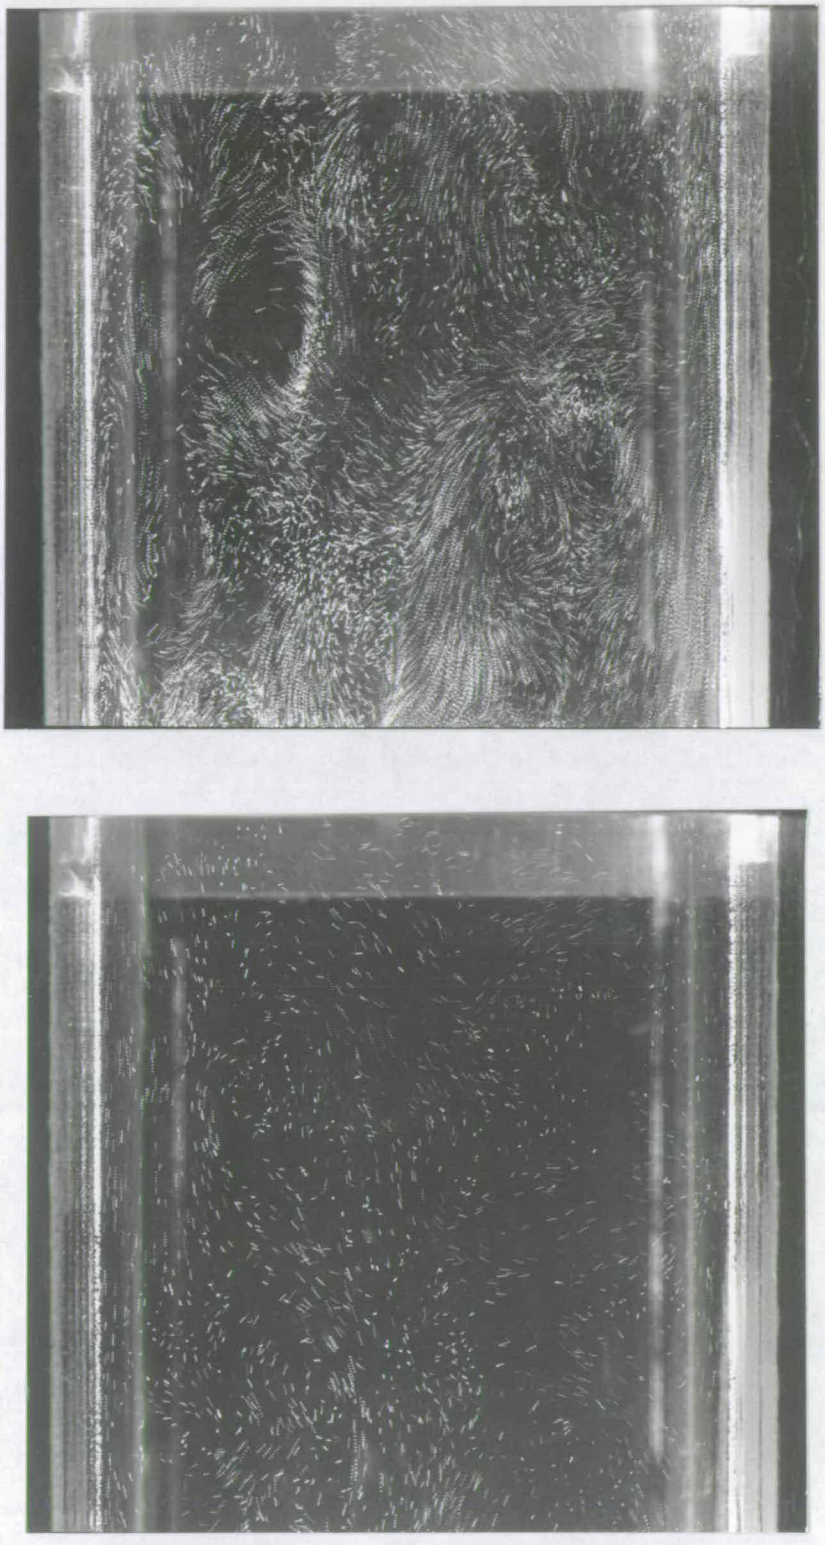


Figure 8.3: PIV photographs for the flow about 10 cm above a bed of AG grade glass particles. Top: f3f8, for which $U/U_{mf} \approx 30$; bottom: f3f11, where U/U_{mf} was about 19.

and provided a more reflective background. Where the glass and aluminium came into contact, it would have been worth blackening the inner surface of the test section if many photographs were to be taken at this camera position, in order to improve the visibility of airborne particles. With hindsight, the aluminium flanges could have been anodised before the glass test section was assembled.

At some viewing heights above the distributor, it was not possible to avoid seeing part of this aluminium flange, as all three bed sections had the same 25 cm length. Instead, as there were no compelling reasons to take PIV photographs at heights affected by these flanges, it was decided that subsequent photographs would be taken in positions where the flanges did not cause difficulties. It was noted that problems caused by flanges being in awkward positions could have been avoided if the three sections of the fluidised bed had been designed to have different lengths, so that one of the possible arrangements of the bed sections would have given the camera an unimpeded view of the flow.

The actual photographs taken are interesting in that they provide detailed snapshots of the highly chaotic flow in the freeboard at a number of different times. A number of observations can be made about the nature of the particle flow in the freeboard. However, it would not be wise to speculate on the nature of air flow in the freeboard, due to the large size of the bed particles. The particles had a Stokes Law terminal velocity which was approximately double the superficial gas velocity. As a result, some particles would have had a long residence time in the freeboard and, whilst they were suspended in the air above the fluidised bed, a mean velocity which was significantly lower than the mean air velocity.

In Figure 8.1, it can be seen that the particle concentration in the freeboard is by no means uniform. In frame 4 (the upper photograph), there are two particle-free regions halfway up each photograph, close to either wall. There are also regions where so many particles are present that individual images cannot be discerned. The particle motion in frame 4 appears to be predominantly vertical only in the top right-hand corner of the photograph and near the walls; the motion is largely horizontal in the top left-hand corner. It is possible that the particles which

move vertically near the walls are falling back down towards the bed; however, the direction of this motion cannot be determined for certain from these photographs, due to the directional ambiguity of the PIV system. In frame 5, two eruptions from the bubbling bed can be seen (Figure 8.1, bottom) and the particle concentration is much lower than in the other frames.

Figure 8.2 is also interesting. In frame 6 (upper photograph), a cluster of particles can be seen, rising up the centre of the bed like a comet. The cluster appears to be about 1 cm wide and about 2 cm long, although its length might be exaggerated by movement during the course of the exposure (about 1/60 s). Again, particle downflow may occur close to the walls, although this is not certain. In frame 7 (lower photograph), there are numerous intersecting particle trajectories near the centre of the photograph, providing further evidence of the confusing nature of the flow. A similar story is told by the upper photograph of Figure 8.3, in which no obvious pattern of particle movement can be discerned.

In the lower half of Figure 8.3, the superficial velocity was reduced to $0.125\,\mathrm{m\,s^{-1}}$ ($19\,U_{mf}$), significantly reducing the particle concentration.

The various exposures in film 3 were interesting and showed that PIV could, in principle, be useful when researching the flow in the freeboard. The photographs were well focused, suggesting that the experimental technique had been implemented successfully. However, a number of problems were identified, leading to the conclusion that PIV was of only limited use under conditions of such vigorous fluidisation. Whilst it was relatively easy to obtain focused pictures of the freeboard, thus 'visualising' the flow, it was not possible to extract significant amounts of velocity information from the photographs, for a number of reasons:

- the particle concentration was variable, with some regions of the film not recording any particle images whilst other regions were saturated.
- the concentration of bed particles in the freeboard was a strong function of U, so the amount of velocity information which could be collected by PIV photographs of these particles was limited.

move vertically near the walls are falling back down towards the bed; however, the direction of this motion cannot be determined for certain from these photographs, due to the directional ambiguity of the PIV system. In frame 5, two eruptions from the bubbling bed can be seen (Figure 8.1, bottom) and the particle concentration is much lower than in the other frames.

Figure 8.2 is also interesting. In frame 6 (upper photograph), a cluster of particles can be seen, rising up the centre of the bed like a comet. The cluster appears to be about 1 cm wide and about 2 cm long, although its length might be exaggerated by movement during the course of the exposure (about 1/60 s). Again, particle downflow may occur close to the walls, although this is not certain. In frame 7 (lower photograph), there are numerous intersecting particle trajectories near the centre of the photograph, providing further evidence of the confusing nature of the flow. A similar story is told by the upper photograph of Figure 8.3, in which no obvious pattern of particle movement can be discerned.

In the lower half of Figure 8.3, the superficial velocity was reduced to $0.125\,\mathrm{m\,s^{-1}}$ ($19\,U_{mf}$), significantly reducing the particle concentration.

The various exposures in film 3 were interesting and showed that PIV could, in principle, be useful when researching the flow in the freeboard. The photographs were well focused, suggesting that the experimental technique had been implemented successfully. However, a number of problems were identified, leading to the conclusion that PIV was of only limited use under conditions of such vigorous fluidisation. Whilst it was relatively easy to obtain focused pictures of the freeboard, thus 'visualising' the flow, it was not possible to extract significant amounts of velocity information from the photographs, for a number of reasons:

- the particle concentration was variable, with some regions of the film not recording any particle images whilst other regions were saturated.
- the concentration of bed particles in the freeboard was a strong function of U, so the amount of velocity information which could be collected by PIV photographs of these particles was limited.

- the trajectories of the particles which were photographed were not always linear, suggesting that the velocities were not always constant during the exposure. In fact, the displacements of the particles between consecutive images often changed in both magnitude *and* direction.

- some particles, which were quite close together in space and moving in the same direction, had significantly different speeds.

- some particle trajectories even cross over one another, occasionally intersecting at angles as large as 90°. It is not at all obvious what velocity *should* be measured under these circumstances.

- no overall pattern to the flow could be discerned, *e.g.* it was not clear that all particles in the centre of the freeboard were rising whilst all particles near the walls were falling back down towards the bed. PIV seems to work well when it is used to make accurate velocity measurements in a flow which is already reasonably well understood.

- the PIV analysis system did not record the 'quality' of each vector (*i.e.* any measure of how confident the system was that the correct displacement had been detected), so it was very difficult to select the correct vectors in the output vector maps because they were buried in a forest of incorrect vectors (see also the previous point). Ideally, one would like to see an image of each interrogation area alongside each vector, so that it would be possible to judge the success of each interrogation. This would, however, become a very slow operation if there were several thousand vectors to validate.

- there was considerable variation in the magnitudes of the particle velocities throughout the flow, *e.g.* some particles moved very rapidly whilst others were almost stationary. A PIV analysis system has a finite dynamic range, as explained in earlier chapters, so it would not be possible to measure all of the velocities on display without using an image-shifting technique.

- even when the angle of motion and velocity magnitude was determined correctly for a given interrogation area, it was not clear whether the particles were moving predominantly upwards or predominantly downwards, due to

the directional ambiguity of a direct PIV system. This problem could be solved, to some extent, by the use of image-shifting or by comparing the measured vectors to an overall circulation pattern or a video recording.

It was therefore decided that better results would be obtained using slightly lower air flowrates and the larger grade of ballotini, aiming for milder fluidisation conditions (and hence fewer bed particles in the freeboard). Seeding particles were added to the bed, in order to ensure that fine particles would be elutriated. PIV was used to measure air velocities in the freeboard (or, at any rate, the velocities of elutriable particles), rather than the velocities of the relatively-few bed particles in the freeboard. It was thought that air velocities would be easier to measure, because there was a clear, upward direction to the air flow.

8.2 Fluidised beds of Grade 10 ballotini

8.2.1 Bed depths of 3 cm

An investigation was carried out into the freeboard flow above a bubbling bed of Grade 10 ballotini, with the air velocity $U \approx U_{mf}$. Two different bed depths were used: 3 cm and 5 cm. A small dose of seeding particles (Sphericel hollow glass spheres) was mixed into the bed between runs. Two different freeboard heights were used, which were about 18 cm and about 33 cm above the position of the unfluidised bed surface. The light sheet was directed at the centre of the bed.

A number of PIV photographs were taken of the flow in each situation. A camera with a motor drive was used, allowing about five high-quality photographs to be taken each second. In each exposure, an area slightly larger than the test section was recorded. The flow records were digitised, either by scanning prints made from photographic negatives or by scanning the negatives directly. The resulting bitmap images covered approximately 3400 × 2350 pixels. These images were cropped, so that only pixels representing points within the flow were retained.

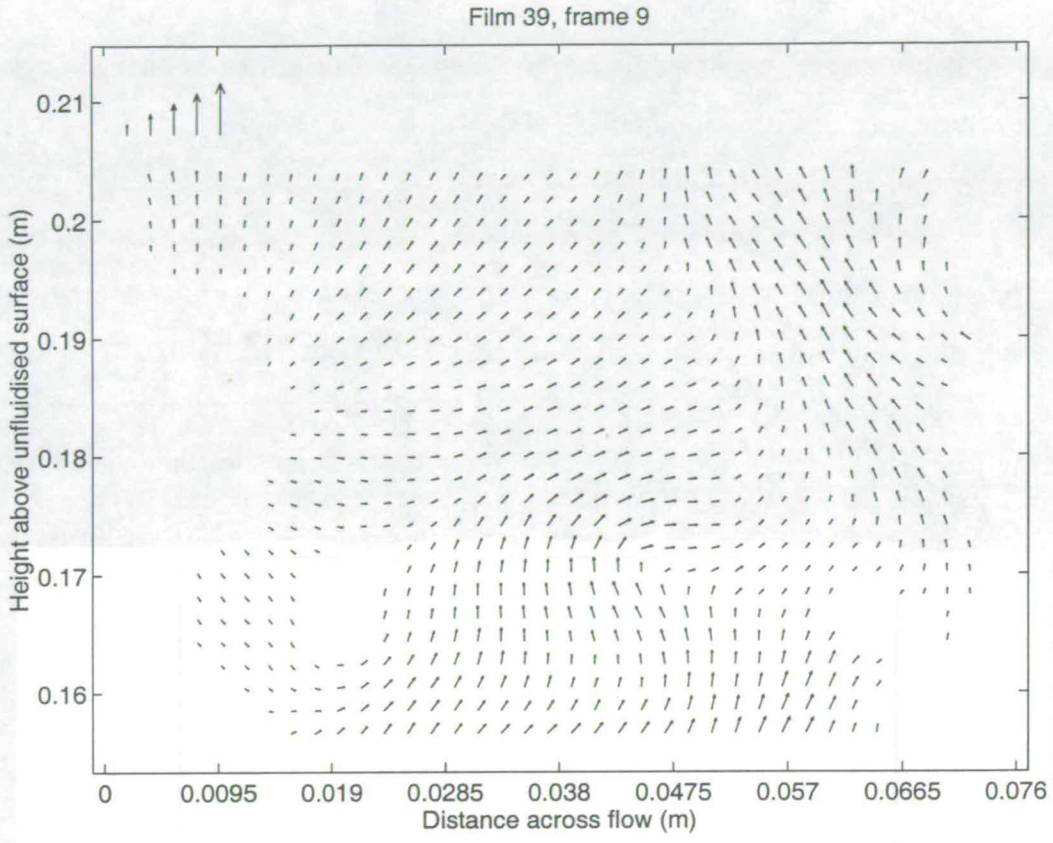
The horizontal width of each bitmap then corresponded to the 7.6 cm width of the test section, so the image:object magnification could easily be determined. These images were analysed using a digital autocorrelation program (PIVANAL).

Each analysis used 128 × 128 pixel interrogation areas, centred on a grid of points about 85 pixels apart. This gave rise to an array of about 27 × 40 velocity vectors, each recording the average velocity in a 3.0 mm×3.0 mm region of the flow. Compared to the early results using 64 × 64 interrogation areas, the larger interrogation area size improved the yields of valid vectors (particularly in the regions of fastest flow) and produced a more practical array of vectors (which could be plotted in full on a half-page). In these vector maps, points are identified by means of abscissa and ordinate co-ordinates, given in the laboratory frame.

Five additional vectors have been added to the top left-hand corner of each vector map, on a row of their own. These vectors are vertical and have magnitudes of 0.1, 0.2, 0.3, 0.4 and 0.5 $m\,s^{-1}$. The length of these reference vectors may vary from vector map to vector map, as some sets of vectors have been scaled up in order to increase the visibility of the smaller vectors.

Within each vector map, the length of each arrow is proportional to the measured displacement. The absence of a vector usually means that no valid measurement was made at that point (*e.g.* because the image separation was too small or because there were insufficient image pairs in the interrogation area). For example, in Figure 4.18 of Chapter 4, the actual displacements added by image-shifting were calculated and then plotted, throughout the field of view. The central image-shifting displacement was then subtracted from all values, resulting in error vectors which were often very small. In Figure 4.19, it is difficult to see a vector at each point, simply because the errors are very small in the middle of the field of view.

When the bed depth was 3 cm, the flow was rather variable at a height of about 18 cm above the bed surface (this height being measured when the air flow was off). In Figure 8.4, the flow appears to be somewhat wavy. In the upper vector



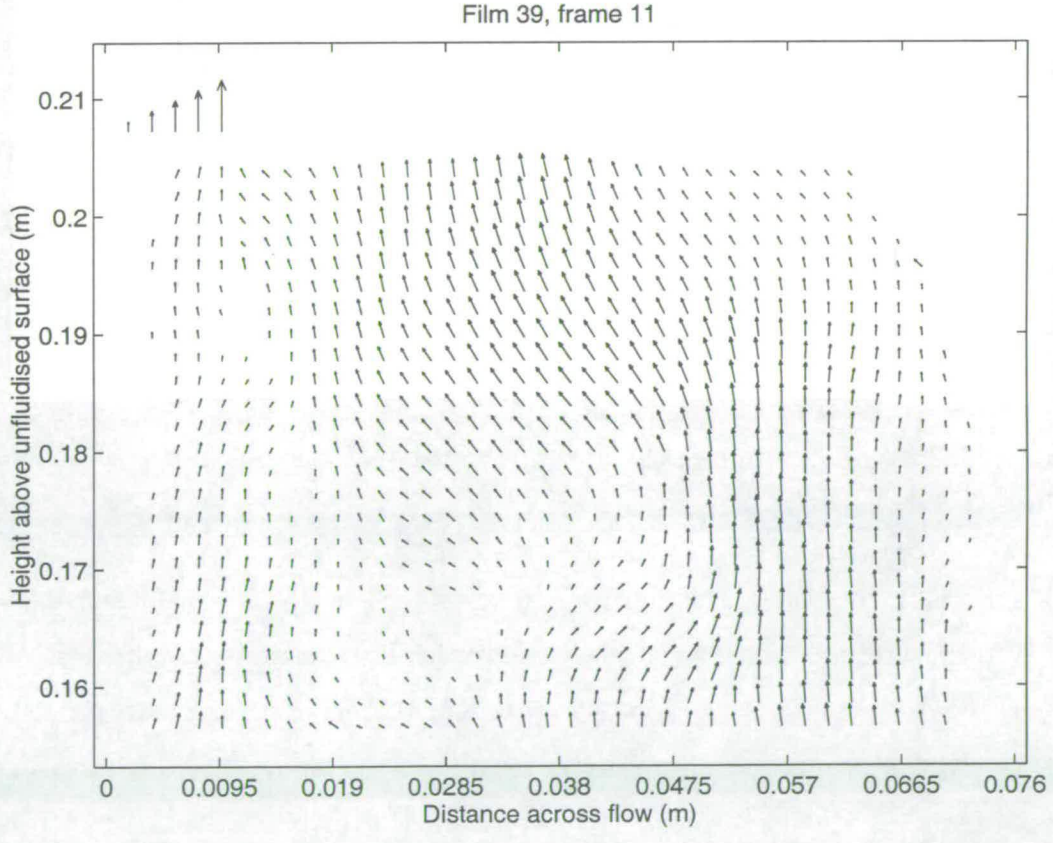


Figure 8.4: PIV vector maps for the flow about 18 cm above a 3 cm bed of particles. $U \approx 1.0\, U_{mf}$

map (film 39, frame 9), there are regions where the air velocity appears to be mainly horizontal (*e.g.* on the left of the image, at a height of about 0.18 m) or even downward (*e.g.* at the point (0.01, 0.17) in the co-ordinates shown). There are few vectors in the left-hand-most 0.01 m, due to a lack of distinct, paired images in this region; this could be because the velocity was so low that images overlapped, or because motion perpendicular to the plane of the light sheet meant that only one image of each particle could be recorded before the particle moved out of the light sheet. In film 39 frame 11, a jet of rapid air upflow can be seen, running from the bottom right corner of the photograph to the top centre.

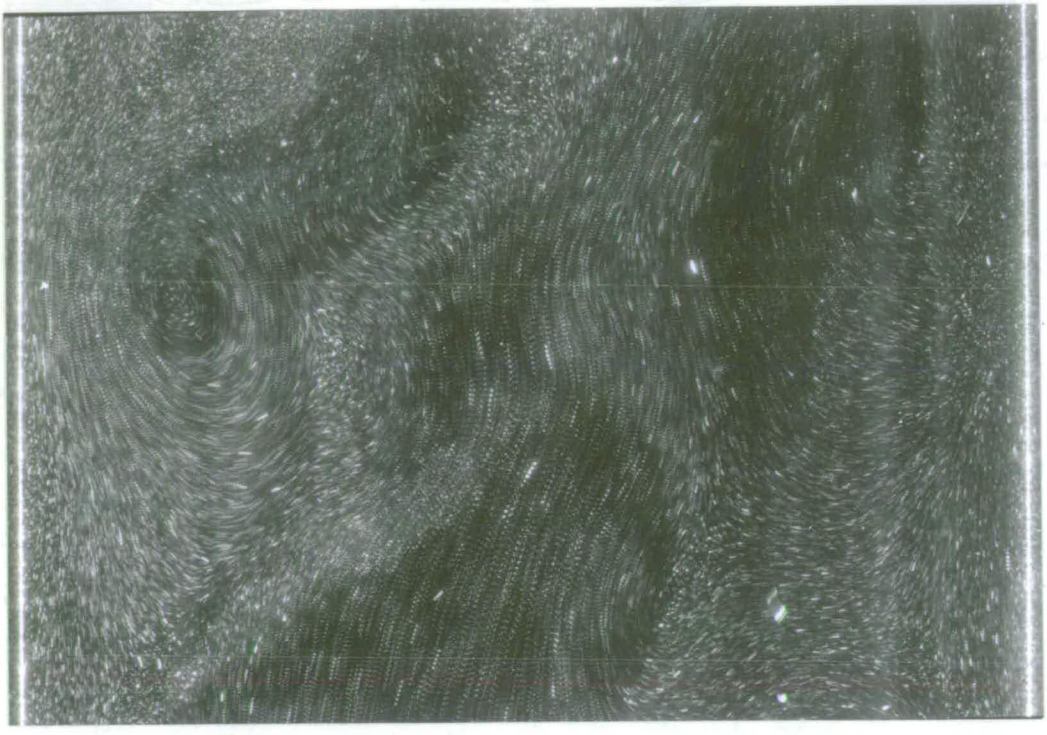


Figure 8.5: PIV photograph of the flow about 18 cm above a 3 cm bed of ballotini (film 39, frame 23).

The flows in the rest of this film were also interesting. Much variation was observed from frame to frame. At the very end of the film, a large recirculation region was observed (sample photograph in Figure 8.5; vector maps in Figure 8.6). Unfortunately, this sequence of frames was terminated by the end of the film, which also caused the right-hand-most 0.01 m of frame 24 to be spoiled. From frame 23 to frame 24, the centre of the recirculation region rose from about 0.185 m to about 0.19 m and also moved closer to the left-hand wall. The total air downflow

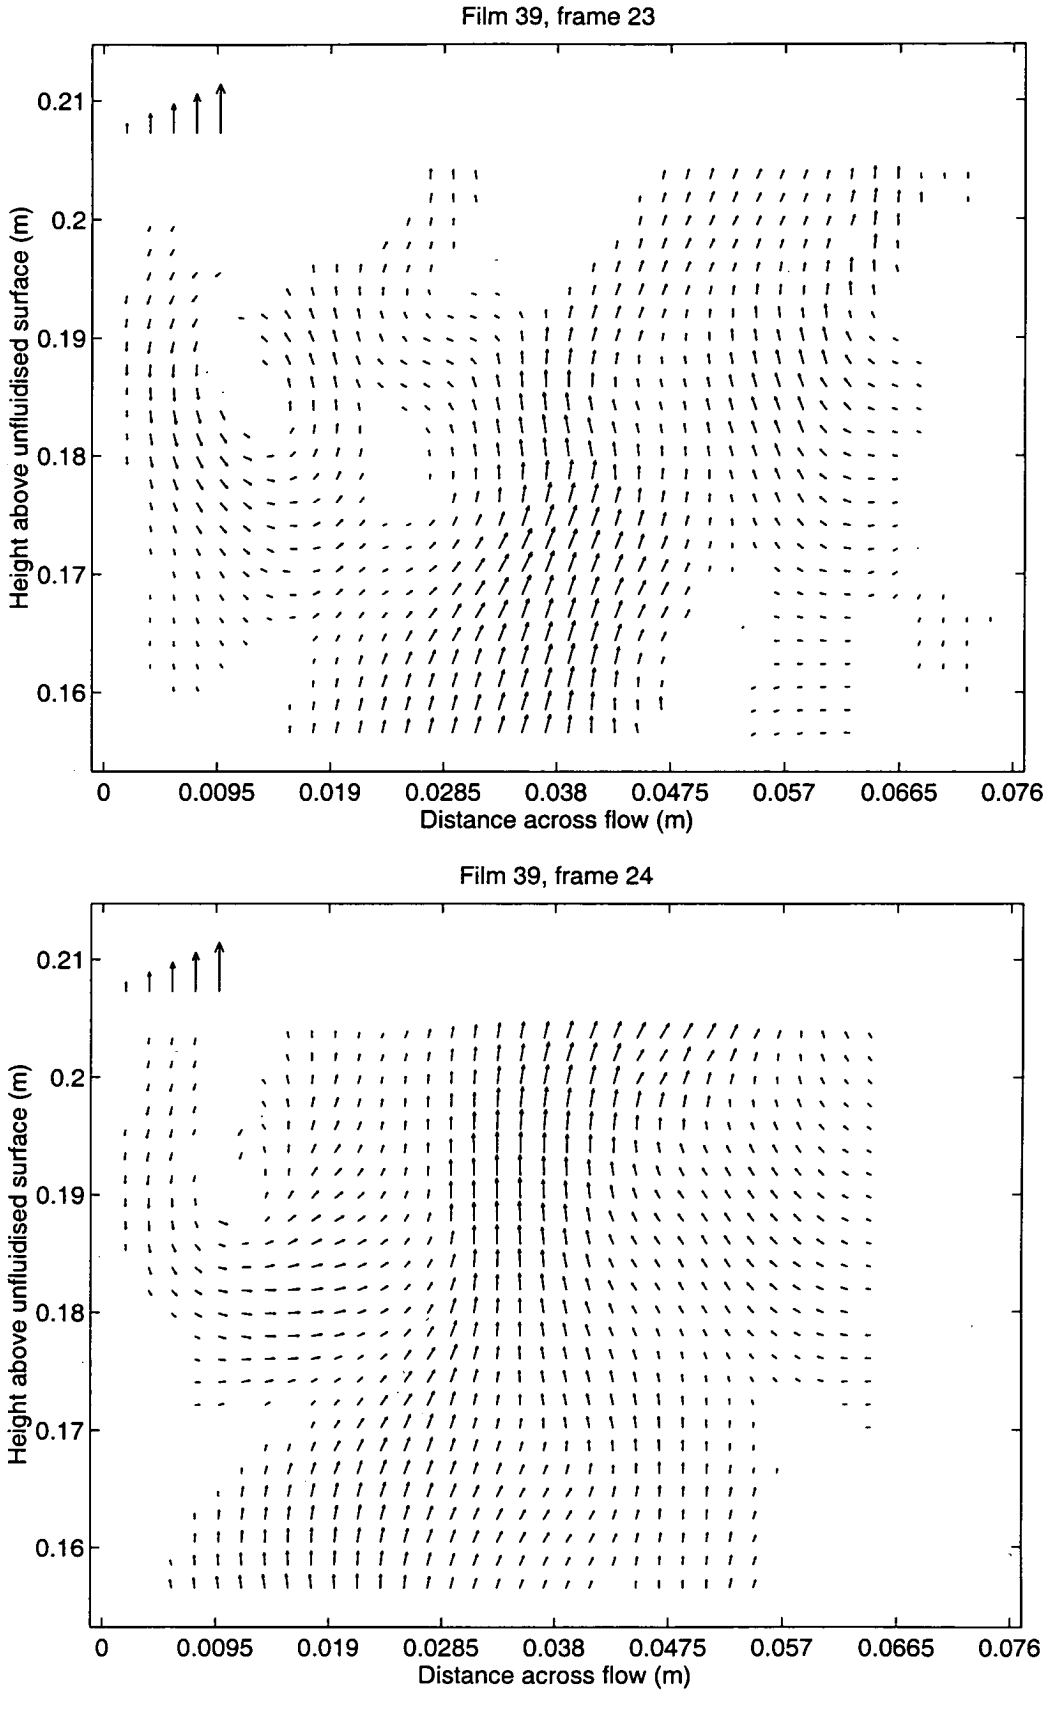


Figure 8.6: PIV vector maps for the flow about 18 cm above a 3 cm bed of particles. $U \approx 1.0\, U_{mf}$

between the centre of each recirculation zone and the left-hand wall should approximately cancel out the total air upflow in the region about 0.01 m to the right of each vortex. If the net air upflow in these regions is minimal, continuity states that greater than average velocities should be found elsewhere. In both cases, rapid air upflow was noted near the centre of the test section.

It could be argued that a second recirculation zone is visible in Figure 8.6, near the right-hand walls. This would appear to vindicate the 'ghost bubble' theory of Pemberton & Davidson (1984), in that apparently-paired vortices are seen. Although interesting in their own right, frames 23 and 24 were very much 'one-off' results, however: similar features were *not* observed in the rest of the film. In fact, during the course of this research, little other evidence has been found to support the ghost bubble theory.

It is not always easy to tell if the lack of paired images in a given flow record (*e.g.* the left-hand edge of film 39 frame 9) is the result of out-of-plane motion or is instead due to overlapping images. In the former case, no pattern will be discernible in each interrogation area, despite the presence of a number of distinct particle images. In the latter case, a low-velocity region might be identified if non-circular image streaks are found, consisting of a number of superimposed particle images a short distance (*e.g.* about 0.2 particle diameters) apart. Further, there can be clues in the neighbouring velocity vectors, which may get smaller as the low-velocity region is approached. The two effects may well occur together, if the direction of the velocity becomes entirely perpendicular to the light sheet. In fact, given that some of the vectors in frame 9 of film 39 appear to be horizontal and predominantly in-plane, it is reasonable to expect other vectors in this three-dimensional flow to be horizontal and predominantly out-of-plane.

Some statistics for these (and other) frames from film 39 are given in Table 8.2, in which v_x denotes a horizontal velocity and v_y an upward vertical velocity.

These values were obtained by taking the validated vector maps (those used to create Figures 8.4 and 8.6), interpolating to fill in as many of the missing vectors as possible, plotting suitable size distribution functions and calculating mean values.

Frame	Yield (%)	Mean v_y ($m\,s^{-1}$)	Max v_y ($m\,s^{-1}$)	Min v_y ($m\,s^{-1}$)	Mean v_x ($m\,s^{-1}$)	Max v_x ($m\,s^{-1}$)
f39f8	79	0.080	0.19	−0.10	+0.017	−0.14
f39f9	78	0.060	0.17	−0.06	+0.027	+0.14
f39f10	74	0.093	0.23	−0.06	−0.009	−0.15
f39f11	87	0.118	0.26	+0.02	−0.022	−0.13
f39f12	82	0.115	0.32	−0.06	−0.013	−0.13
f39f13	77	0.099	0.28	−0.06	−0.005	−0.10
f39f14	80	0.090	0.25	−0.06	+0.020	−0.11
f39f15	73	0.117	0.24	−0.00	−0.034	−0.12
f39f16	85	0.105	0.21	+0.01	+0.016	+0.13
f39f17	87	0.107	0.20	−0.08	+0.022	+0.13
f39f18	82	0.116	0.20	+0.02	+0.003	+0.09
f39f19	78	0.107	0.24	−0.05	−0.006	−0.09
f39f20	80	0.102	0.25	−0.11	+0.008	+0.09
f39f21	78	0.095	0.22	−0.08	+0.015	+0.12
f39f22	66	0.050	0.20	−0.09	+0.005	−0.10
f39f23	74	0.057	0.22	−0.15	+0.003	+0.11
f39f24	75	0.076	0.22	−0.11	+0.010	+0.13
mean:	79	0.093	0.23	−0.06	+0.003	~ 0.12

Table 8.2: Summary of velocity values for film 39, where the bed depth was 3 cm and the mean viewing height was 18 cm above the unfluidised bed surface.

When evaluating the average value of the greatest horizontal velocity, the plus and minus signs in the right-hand-most column were ignored and only the magnitude was considered.

Not all of the missing vectors could be filled in, however, so the averaging process counted up the number of non-zero velocity vectors and used this number when calculating the mean, rather than the total number of vectors (typically $27 \times 40 = 1080$). In frame 24, the right-hand-most 0.01 m of the flow gave rise to no valid vectors. The interpolation method could fill in missing vectors within regions bounded by non-zero vectors, *i.e.* where there were vectors between which an interpolation could be performed, but would not extrapolate into the region from 0.065 m to 0.075 m across the flow (where there were no vectors at all). The air upflow in this region *was* relatively low, but it was ignored in the averaging analysis. This perhaps explains why the mean upward velocity in frame 24 is somewhat larger than that in frame 23 and, more generally, may account for some of the variations in the mean vertical velocity from frame to frame.

From Table 8.2, it is apparent that the mean vertical velocity was often about $1.5 \times U$, the superficial air velocity. For comparison, this method of averaging should give a mean of $1.44\,U$ for a PIV section through fully-developed laminar flow in a square-sided duct (see Section 2.4.6). This is quite a good agreement, even though it cannot be argued that the 'laminar' flow was well developed here.

The magnitude of the greatest vertical velocity component was about 2.5–4 U. For comparison, the bubble rise velocity in the fluidised bed was calculated using correlations from the literature. Assuming velocity values of $U = 0.077\,\mathrm{m\,s^{-1}}$ and $U_{mf} = 0.070\,\mathrm{m\,s^{-1}}$ in order to ensure that $U > U_{mf}$, the correlation of bubble diameter by Mori & Wen (1975) gave $D_b = 0.0032\,\mathrm{m}$, whereas work by Rowe (1976) suggested that $D_b = 0.0034\,\mathrm{m}$ (in good agreement). Rowe & Partridge (1965) found that the single-bubble rise velocity was given by $U_{b,\infty} = K\sqrt{g\,D_b/2}$, for some constant K which is generally close to 1. These estimates of the bubble diameter allow $U_{b,\infty}$ to be estimated to be $0.13\,\mathrm{m\,s^{-1}}$. The absolute bubble rise velocity will exceed this value by approximately $(U - U_{mf})$, so the absolute bubble rise velocity (in the laboratory frame) will be expected to be about $0.14\,\mathrm{m\,s^{-1}}$ (approximately $2\,U$) near the top surface of this 3 cm fluidised bed.

These estimates of the bubble diameter, however, were somewhat smaller than those observed in experiments using the smaller fluidised bed. Possibly the porous distributor in this fluidised bed was damaged, so that the gas flow through the bed was not evenly distributed (allowing larger bubbles to form). It was thought that a bubble diameter of about 1 cm was more reasonable. This would predict a bubble rise velocity of about $0.22\,\mathrm{m\,s^{-1}}$ (or about $3\,U$), which can be taken to be a rough estimate of the air velocity in the freeboard region, close to the bed surface. These velocities are somewhat greater than the largest vertical velocity expected in fully-developed laminar flow in a square duct, which is a velocity of about $2.10\,U$ a long way above the bed surface.

Significant horizontal velocity components were also present, approximately equal in peak magnitude to the mean vertical velocity. These horizontal velocities could be either to the left or to the right, as indicated by the signs in Table 8.2.

No non-zero horizontal velocities will be found in a truly laminar flow, so these horizontal velocity components should diminish at greater freeboard heights.

A vector map for frame 23 of film 39 (Figure 8.5) was given in Figure 8.6, using 128 × 128 interrogation areas. The corresponding result for analysis with 64 × 64 interrogation areas is shown in Figure 8.7. The velocities were very similar to those given in Table 8.2, but the yield of valid vectors fell from 73% to 69%. Overall, the average yields when using 128 × 128 and 64 × 64 interrogation areas were 79% and 70% respectively, because the larger interrogation areas gave better yields in regions of high velocity. In frame 23, however, this advantage was partly offset by the smaller interrogation areas being better at analysing regions of high velocity gradient, *e.g.* the recirculation region near the left-hand wall.

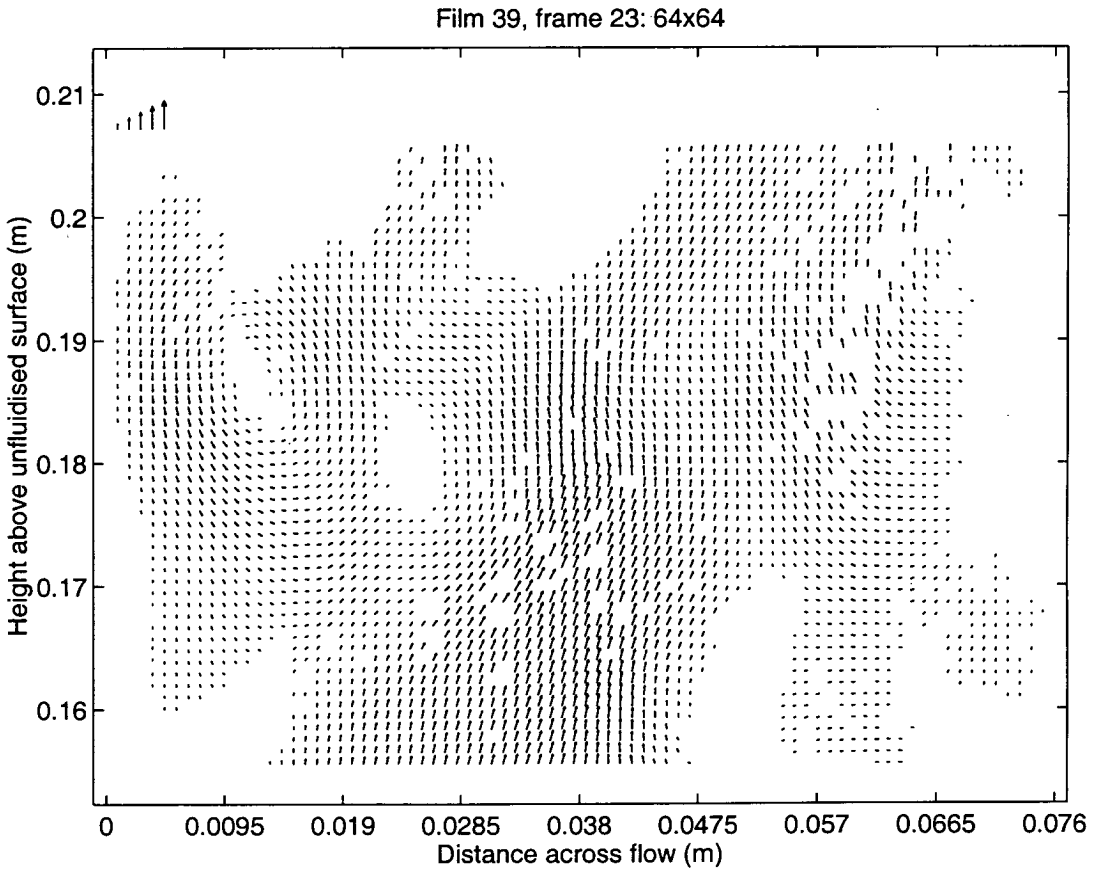


Figure 8.7: Film 39 frame 23, analysed using 64 × 64 interrogation areas.

Keeping the bed depth at 3 cm, results further up the freeboard showed a more universally-upward character (film 37; Figures 8.8 and 8.9). In some cases, *e.g.* frame 15, the bulk of the upflow took place in two jets of rapidly-rising air,

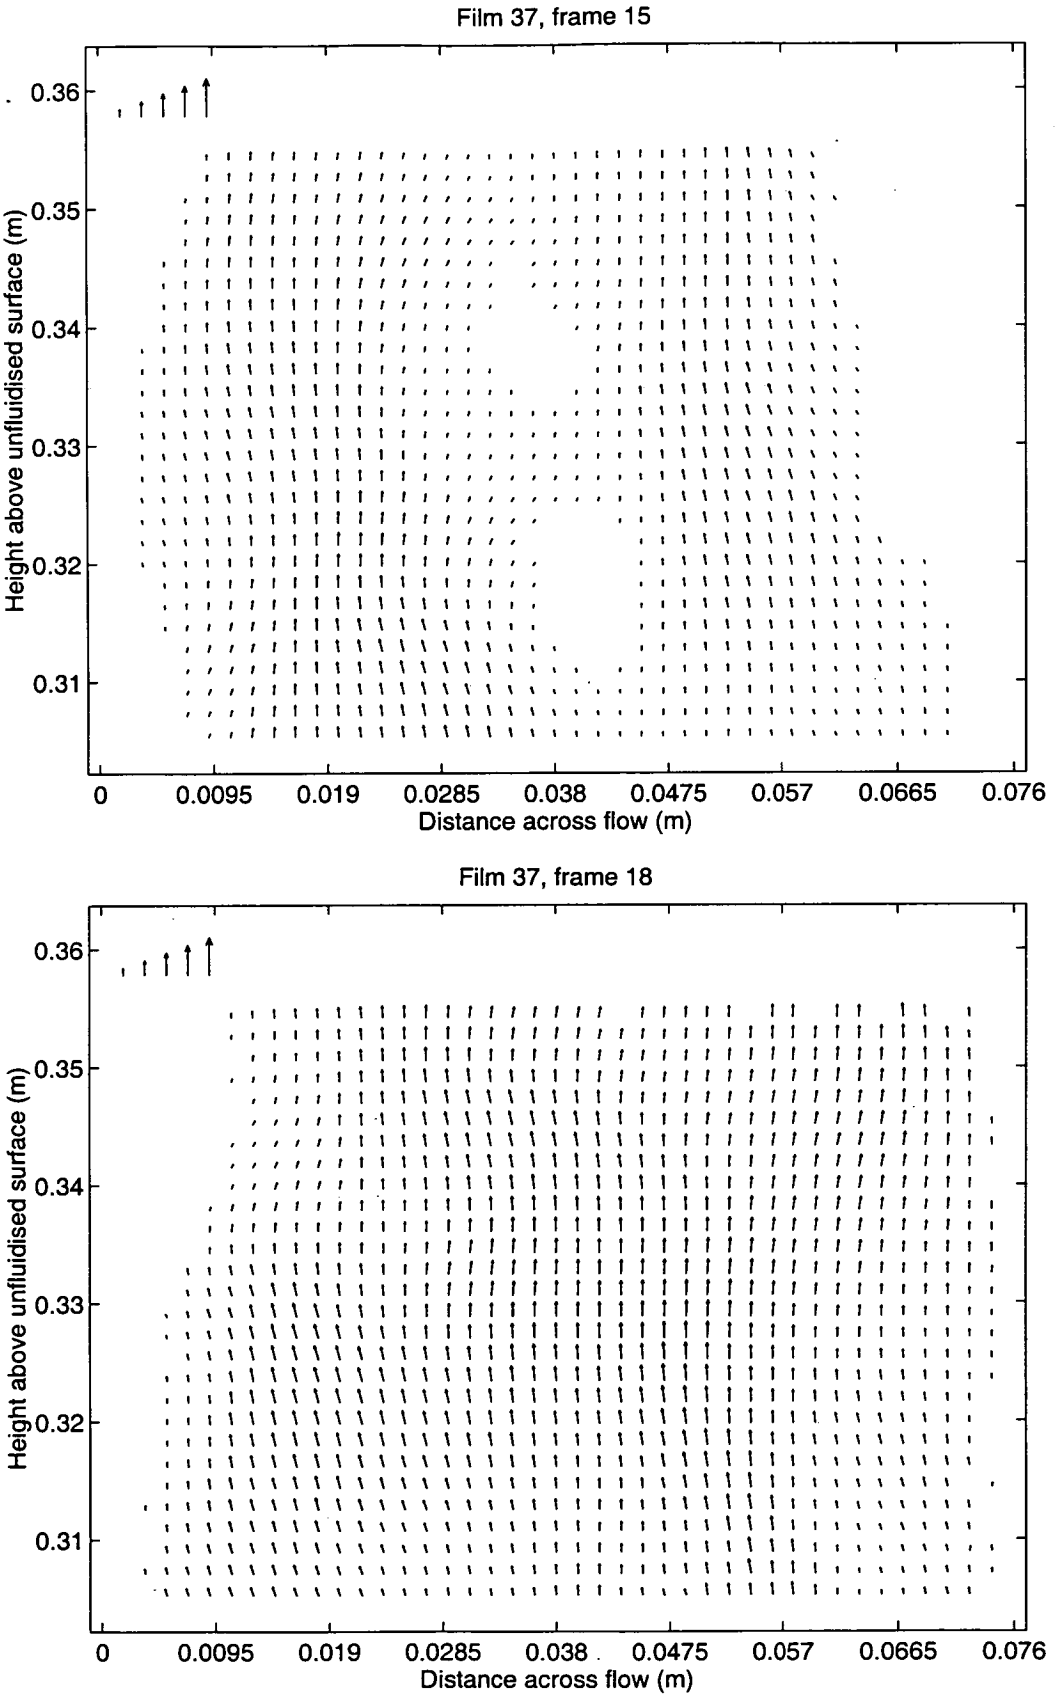


Figure 8.8: PIV vector maps for the flow about 33 cm above a 3 cm bed of particles. $U \approx 1.0\, U_{mf}$

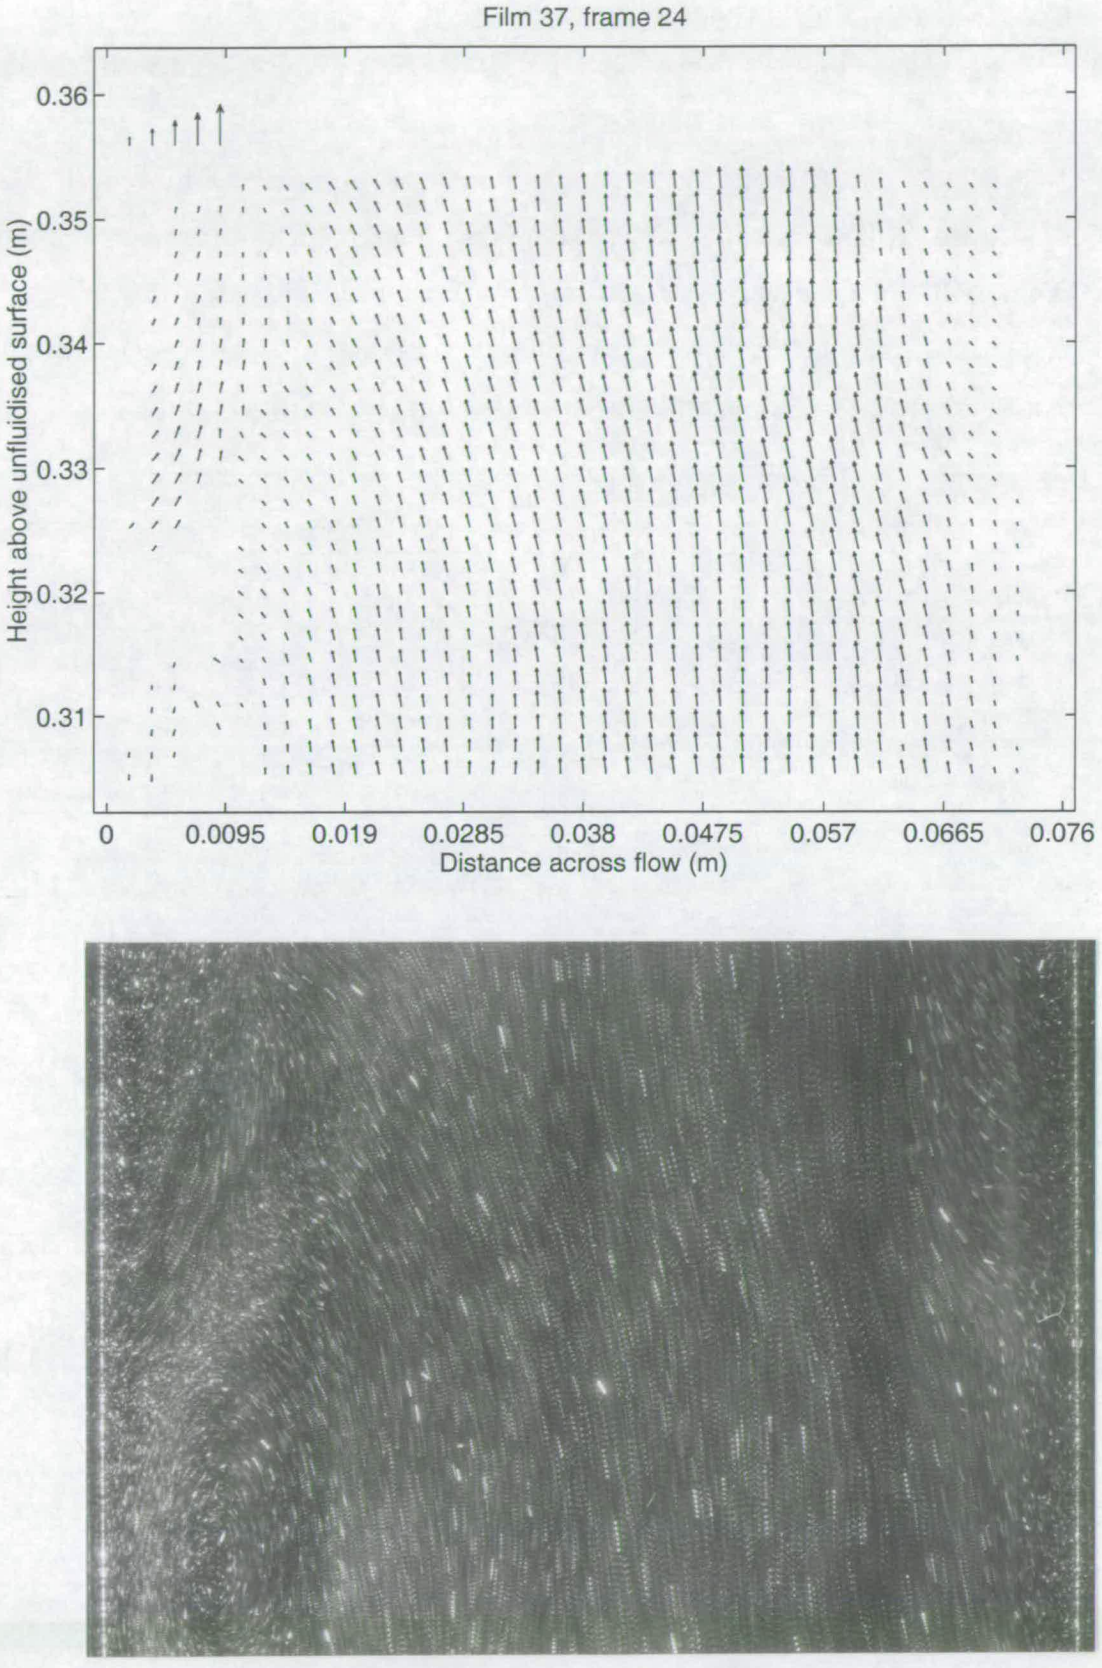


Figure 8.9: PIV vector map (top) and photograph (bottom) for the flow about 33 cm above a 3 cm bed of particles. $U \approx 1.0\,U_{mf}$.

with slower-moving regions found (a) between the two jets and (b) between each jet and the nearest wall. A few stagnant regions and the occasional recirculation zone were observed even at this height above a rather shallow bed, as shown in the bottom left corner of frame 24 (vector map and photograph in Figure 8.9). The more uniform upflow of frame 18 was, however, more typical of the vector maps from this roll of film (Figure 8.8, bottom). With no two PIV vector maps looking particularly similar in this chaotic situation, it can be tempting to concentrate on the more exotic flow records and thus lose track of the relatively mundane.

Frame	Yield (%)	Mean v_y ($\mathrm{m\,s^{-1}}$)	Max v_y ($\mathrm{m\,s^{-1}}$)	Min v_y ($\mathrm{m\,s^{-1}}$)	Mean v_x ($\mathrm{m\,s^{-1}}$)	Max v_x ($\mathrm{m\,s^{-1}}$)
f37f13	74	0.095	0.20	+0.02	−0.000	−0.08
f37f14	74	0.082	0.19	+0.02	−0.006	−0.08
f37f15	78	0.095	0.17	+0.01	−0.004	±0.05
f37f16	83	0.110	0.20	+0.01	−0.017	−0.08
f37f17	88	0.129	0.21	+0.03	−0.021	−0.09
f37f18	90	0.144	0.24	+0.02	−0.009	−0.06
f37f19	89	0.131	0.23	−0.07	−0.016	−0.08
f37f20	81	0.149	0.31	+0.01	−0.021	−0.06
f37f21	85	0.169	0.30	+0.02	−0.027	−0.08
f37f22	88	0.166	0.32	−0.07	−0.017	−0.07
f37f23	89	0.172	0.33	+0.01	−0.018	−0.11
f37f24	91	0.153	0.29	−0.09	−0.027	±0.08
mean:	84	0.133	0.25	−0.001	−0.015	∼ 0.08

Table 8.3: Summary of velocity values for film 37, where the bed depth was 3 cm and the mean viewing height was 33 cm above the unfluidised bed surface.

Some velocity statistics are given in Table 8.3. Compared to the values at a mean height of 18 cm in the earlier run, the mean upward velocity seems to be somewhat larger at the greater height in the freeboard, although the greatest detected upward velocities are comparable. As time goes on, all velocities seem to increase, *e.g.* in the final six exposures, the mean vertical velocity was about 2–2.3 U and the greatest velocity about 4–5 U.

This may be due to oscillations in the air flow, with these photographs recording only a part of the cycle in which the velocities in the illumination plane happened to be relatively large. A test for this could have been made if further frames had been taken: would the mean vertical velocity have decreased over the next few

exposures? Looking back to the table of results for film 39 and taking frames 9–18 in isolation, it might appear that there is an increase in the mean vertical velocity as time goes by (Table 8.2). However, frame 8 and frames 19–23 suggest that frames 9–18 were merely the higher flowrate part of a time-dependent oscillation in the mean centre-section velocity. Such oscillations may also also occur at the greater viewing height used in film 37, if the time period of these oscillations velocity was longer than the time taken to use up the 24-exposure roll of film (about 5–10 s).

Alternatively, and perhaps more plausibly, the increase in mean velocity seen in film 37 could be the result of a genuine increase in the air flow through the fluidised bed. A rotameter reading of 5 cm corresponded to an air superficial velocity of $0.070\,\mathrm{m\,s^{-1}}$ and an expected section-average gas velocity of $0.101\,\mathrm{m\,s^{-1}}$. The rotameter reading would have had to drift as high as 9 cm to account for an increased superficial velocity of $0.106\,\mathrm{m\,s^{-1}}$ and hence an increased average of $0.152\,\mathrm{m\,s^{-1}}$ across the illuminated section. This is too great a difference to be accounted for by reading errors, but a sustained drift towards higher velocities might have occurred during the course of each run, because it was not easy to keep an eye on the rotameter whilst taking photographs and because it was important to take photographs as soon as possible after fluidisation began (before the walls of the test section became coated by seeding particles). As the fan warmed up, its characteristic curve might have changed and more air might have passed through the fluidised bed.

In this run, the greatest horizontal velocity components at a height of about 33 cm were closer to zero than they were at the lower freeboard height. This was due to a decrease in the standard deviation of the horizontal velocity size distribution, at the greater freeboard height. Another difference was found in the mean values, which were either to the left or to the right in film 39 (hence giving an average value close to zero) but almost always pointed to the left in film 37 (as can be seen in frame 24, shown in Figure 8.9).

8.2.2 Bed depths of 5 cm

Moving back down to a mean viewing height of 17 cm above the unfluidised bed surface and increasing the bed depth to 5 cm, more turbulence was apparent (Figures 8.10 and 8.11). In Film 41 frame 11, a strongly-recirculating zone can be seen in the bottom left corner of the picture (but this was not part of a pair). In frames 12 and 13, the strength of this vortex appeared to decrease significantly, whilst a jet of rapidly-rising air developed to the right of the bed axis.

Some limitations of direct PIV were apparent in the older results from this film. There were a number of regions where the velocities were either too small to measure at all (*e.g.* the various vector-free regions) or too large to measure consistently (the yield of valid vectors in the most rapidly-rising regions was poor). This is a result of the finite dynamic range of a PIV system: particle images must be far enough apart to allow distinct images to be recorded, yet not so far apart that loss-of-pairs effects reduce the yield. Using 64×64 interrogation areas, the average yield of valid vectors was only 56%.

Frame	Yield (%)	Mean v_y ($m\,s^{-1}$)	Max v_y ($m\,s^{-1}$)	Min v_y ($m\,s^{-1}$)	Mean v_x ($m\,s^{-1}$)	Max v_x ($m\,s^{-1}$)
f41f8	58	0.081	0.21	−0.02	−0.019	−0.12
f41f9	61	0.081	0.25	−0.08	−0.014	−0.11
f41f10	67	0.097	0.23	−0.08	−0.014	±0.13
f41f11	89	0.121	0.30	−0.13	−0.034	−0.14
f41f12	85	0.124	0.34	−0.09	−0.038	−0.14
f41f13	71	0.117	0.34	−0.11	−0.017	−0.10
f41f14	76	0.079	0.32	−0.14	−0.020	−0.13
f41f16	66	0.079	0.32	−0.11	+0.015	+0.12
f41f17	66	0.075	0.30	−0.09	+0.022	+0.17
f41f18	59	0.071	0.23	−0.05	−0.024	−0.12
f41f19	62	0.081	0.21	−0.08	+0.011	+0.13
f41f20	86	0.116	0.25	+0.01	−0.021	−0.13
f41f21	87	0.112	0.26	−0.02	−0.027	−0.10
f41f22	72	0.087	0.25	−0.10	−0.025	−0.10
f41f23	66	0.082	0.26	−0.11	−0.013	−0.07
mean:	71	0.094	0.27	−0.08	−0.015	∼ 0.12

Table 8.4: Summary of velocity values for film 41, where the bed depth was 5 cm and the mean viewing height was 17 cm above the unfluidised bed surface.

Things can be improved by using PIV with image-shifting, but this technology had not previously been used in this department at the time this early work was carried out and only became well-established in the University of Edinburgh Fluid Dynamics group about a year or two later. Instead, the frames were re-analysed using larger, 128×128 interrogation areas. This improved the yields of the larger vectors but did not help in regions of low velocity. As the velocity properties at this bed depth show, the average yield of valid vectors rose to 71% (Table 8.4).

For the greater bed depth, the correlation of Mori & Wen (1975) suggested that $D_b = 0.0051\,\mathrm{m}$, in good agreement with work by Rowe & Partridge (1965) which proposed that $D_b = 0.0050\,\mathrm{m}$. Once again, however, these values seriously underestimated the bubble diameter observed in this small fluidised bed. The diameters obtained by correlation suggest that the absolute bubble rise velocity near the top of this 5 cm bed will be about $0.165\,\mathrm{m\,s^{-1}}$ (or about $2.4\,U$). On the other hand, a more realistic bubble diameter of about 2 cm would predict that bubble rise velocities as large as $0.32\,\mathrm{m\,s^{-1}}$ ($4.6\,U$) could be found.

The mean vertical velocity appeared to go through two full cycles of oscillations. Assuming, rather questionably, that the time interval between consecutive photographs was the same in all runs, a tentative conclusion would be that the time period of these oscillations is shorter than that found above the shallower bed. In film 39, oscillations lasted for 14 exposures, but the cycle appeared to repeat every 8 exposures in film 41. The mean velocities were very similar to those above the 3 cm bed at the comparable viewing height, but the greatest vertical velocities were larger (3–5 U, consistent with the increased bubble rise velocity).

For film 33, the bed depth was 5 cm and the viewing height was 33 cm above the initial position of the bed surface (Figures 8.12 and 8.13). Two separate runs were made on this roll of film, using the first 10 and then the last 14 exposures. The latter half of each mini-run has been discarded, because the concentration of seeding particles in the test section became too low.

The flow again appeared to be channeled in zones, *e.g.* there was a slow-moving region about 0.06 m across the flow in frame 11, whilst there were strong upward

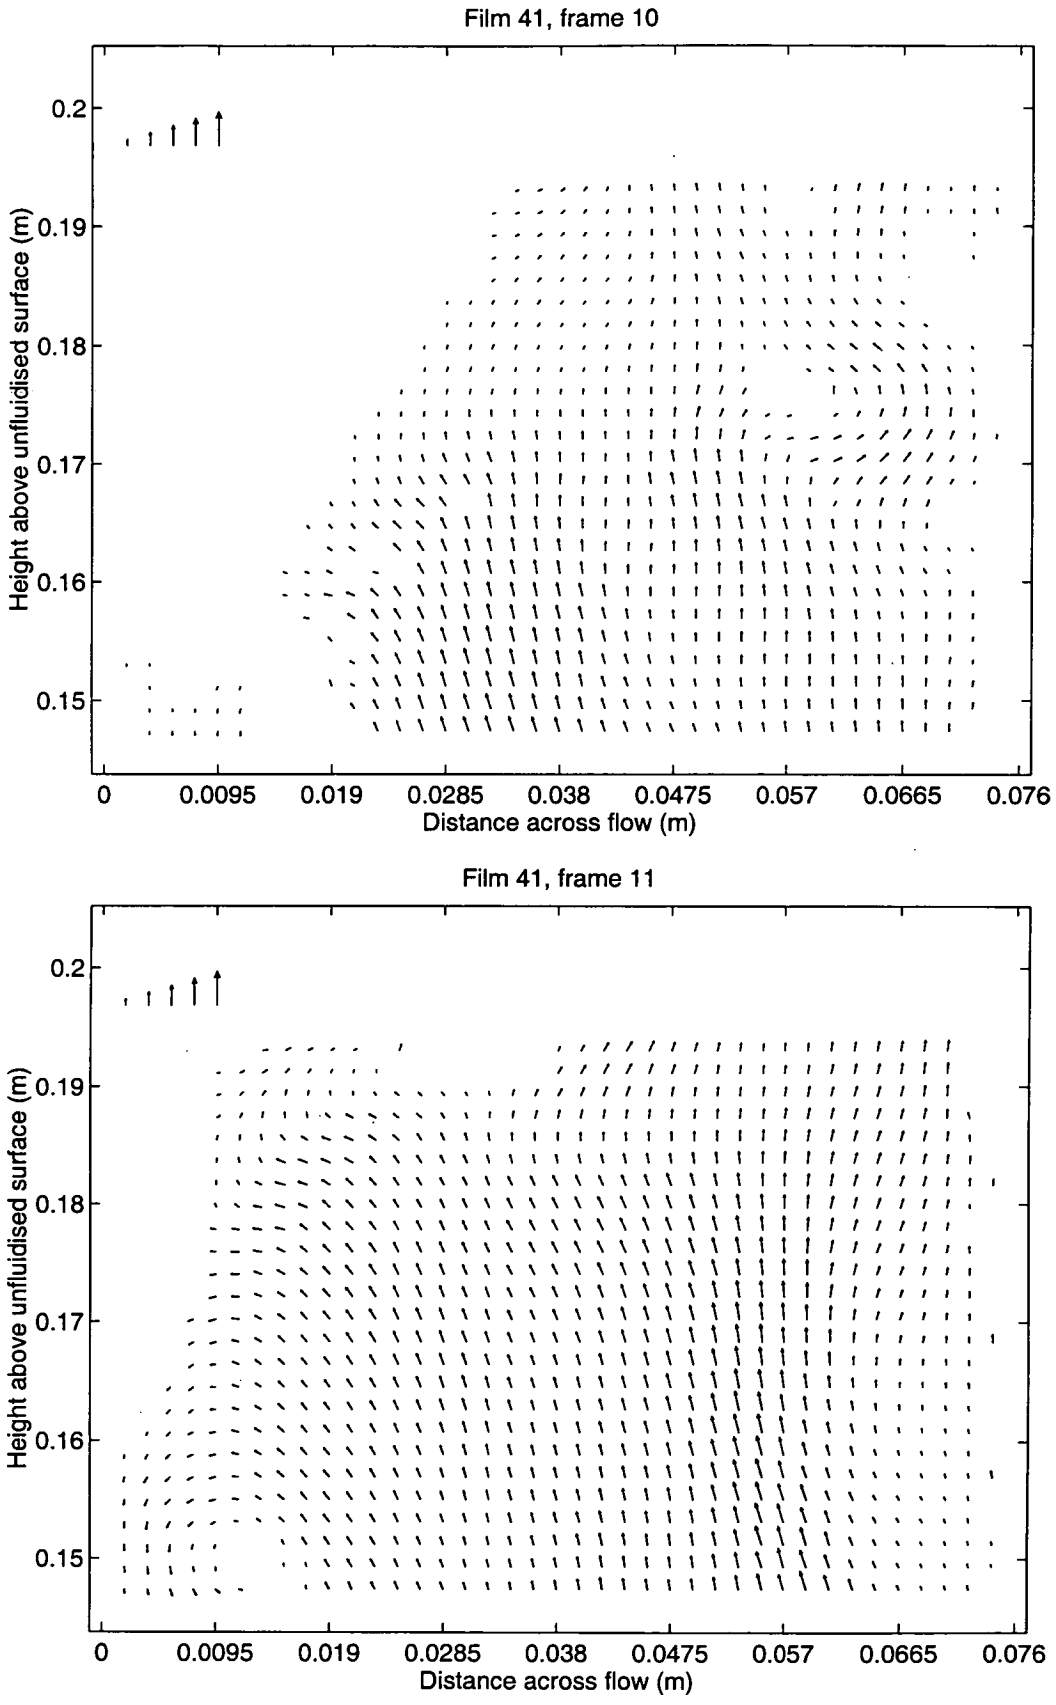
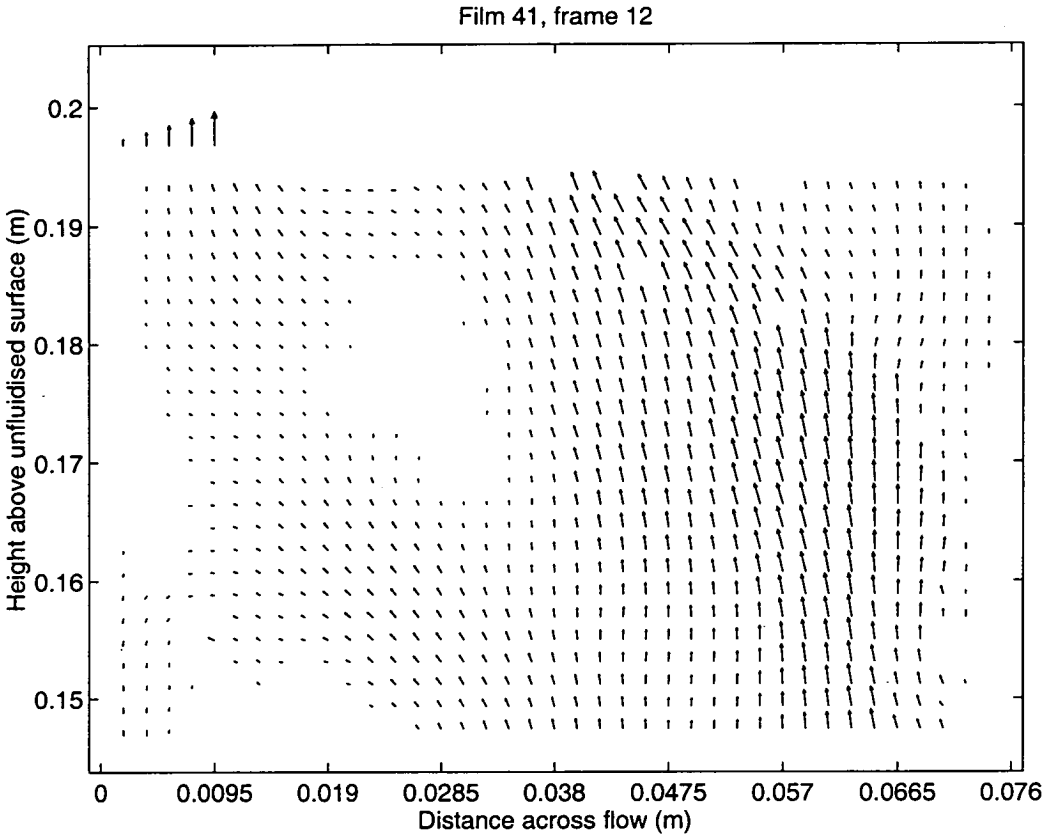


Figure 8.10: PIV vector maps for the flow about 17 cm above a 5 cm bed of particles. $U \approx 1.0\, U_{mf}$



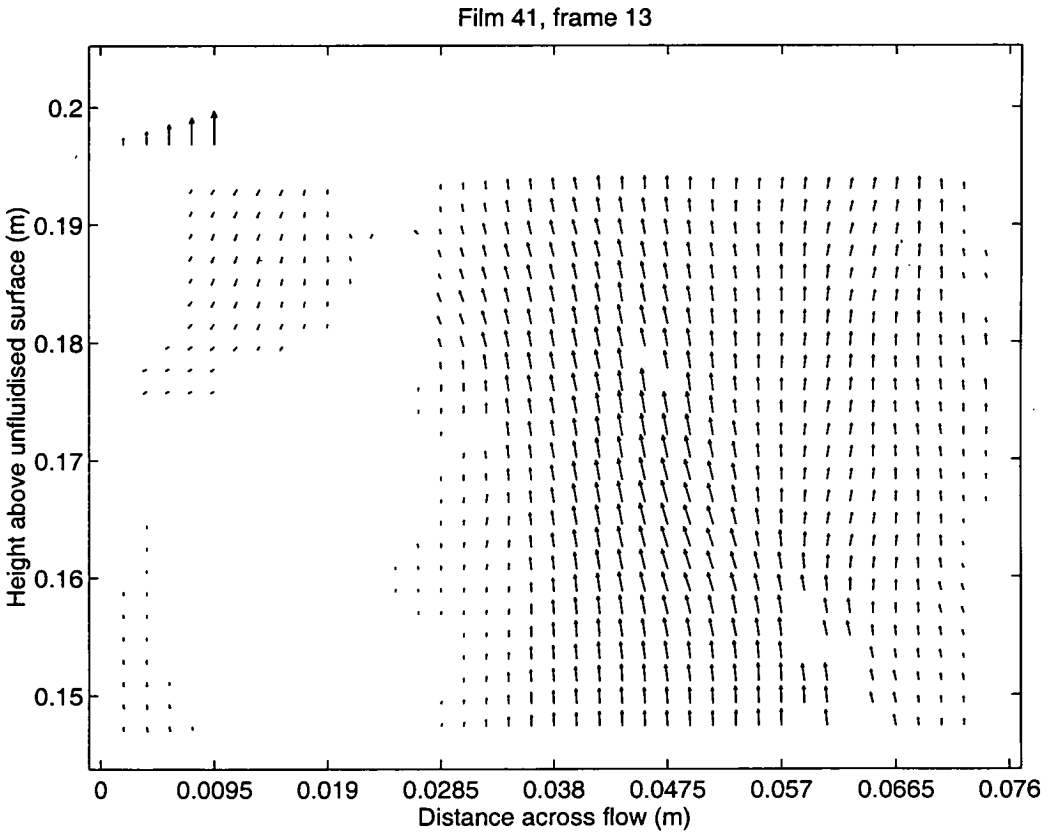


Figure 8.11: PIV vector maps for the flow about 17 cm above a 5 cm bed of particles. $U \approx 1.0\,U_{mf}$

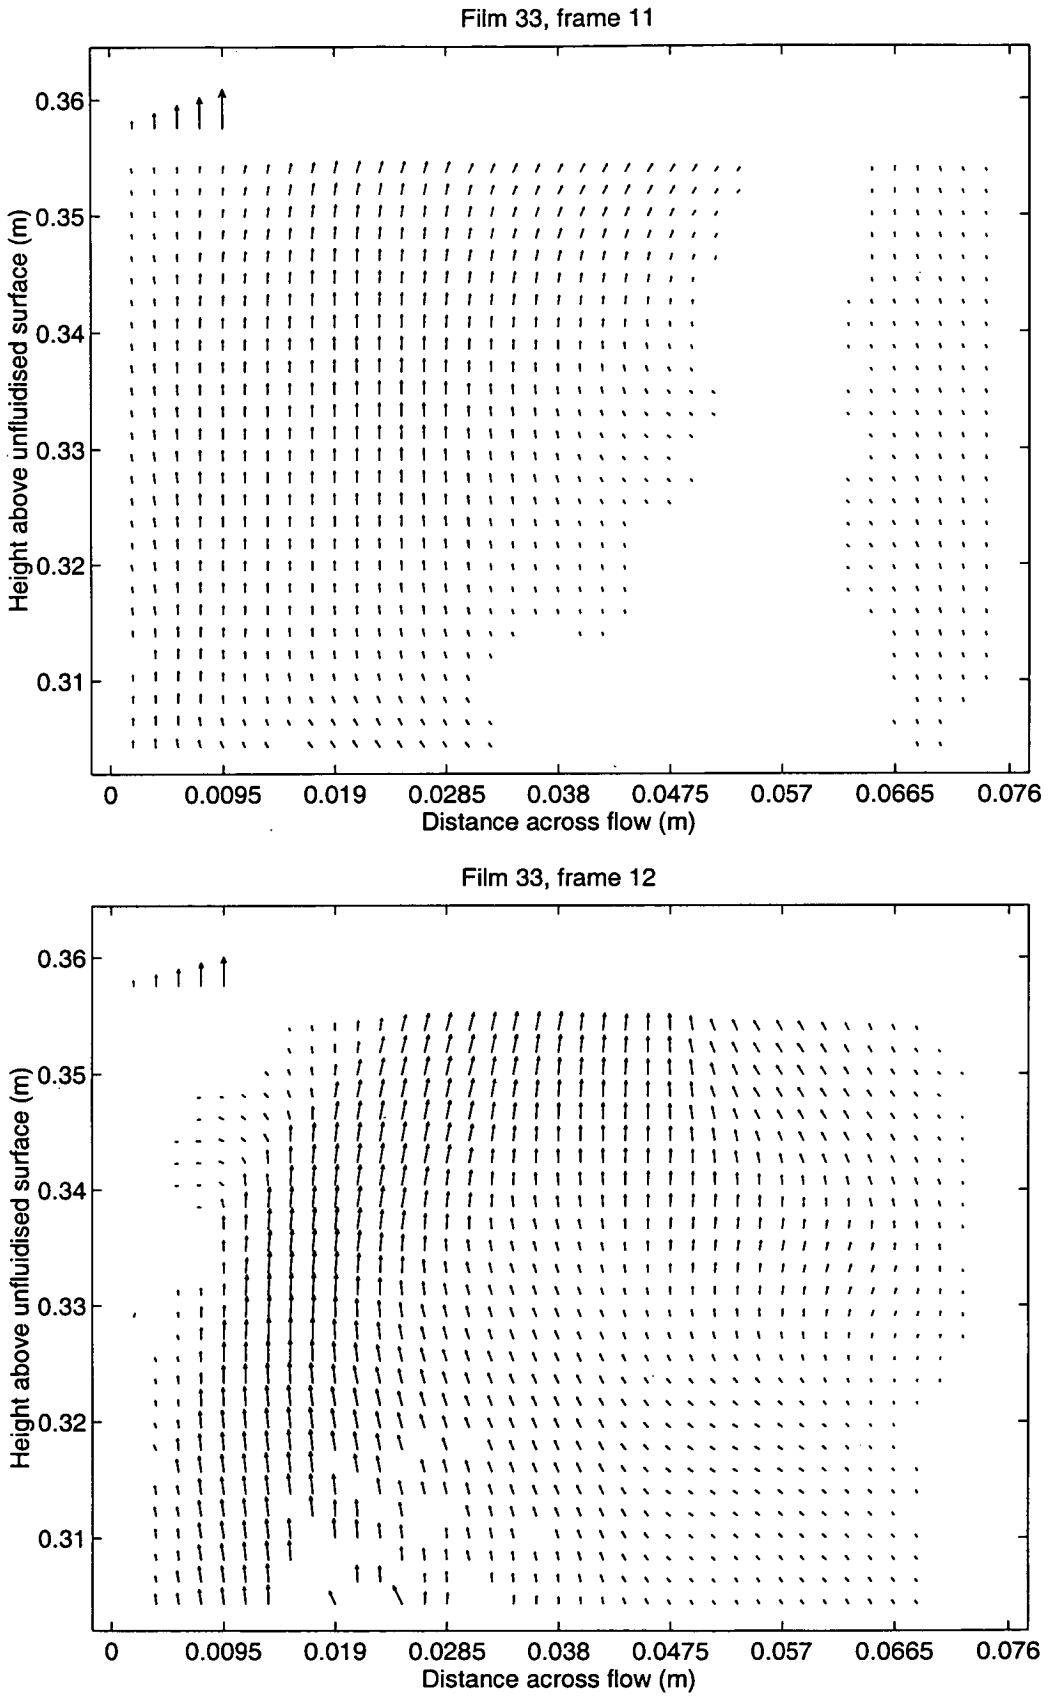


Figure 8.12: PIV vector maps for the flow about 33 cm above a 5 cm bed of particles. $U \approx 1.0\, U_{mf}$

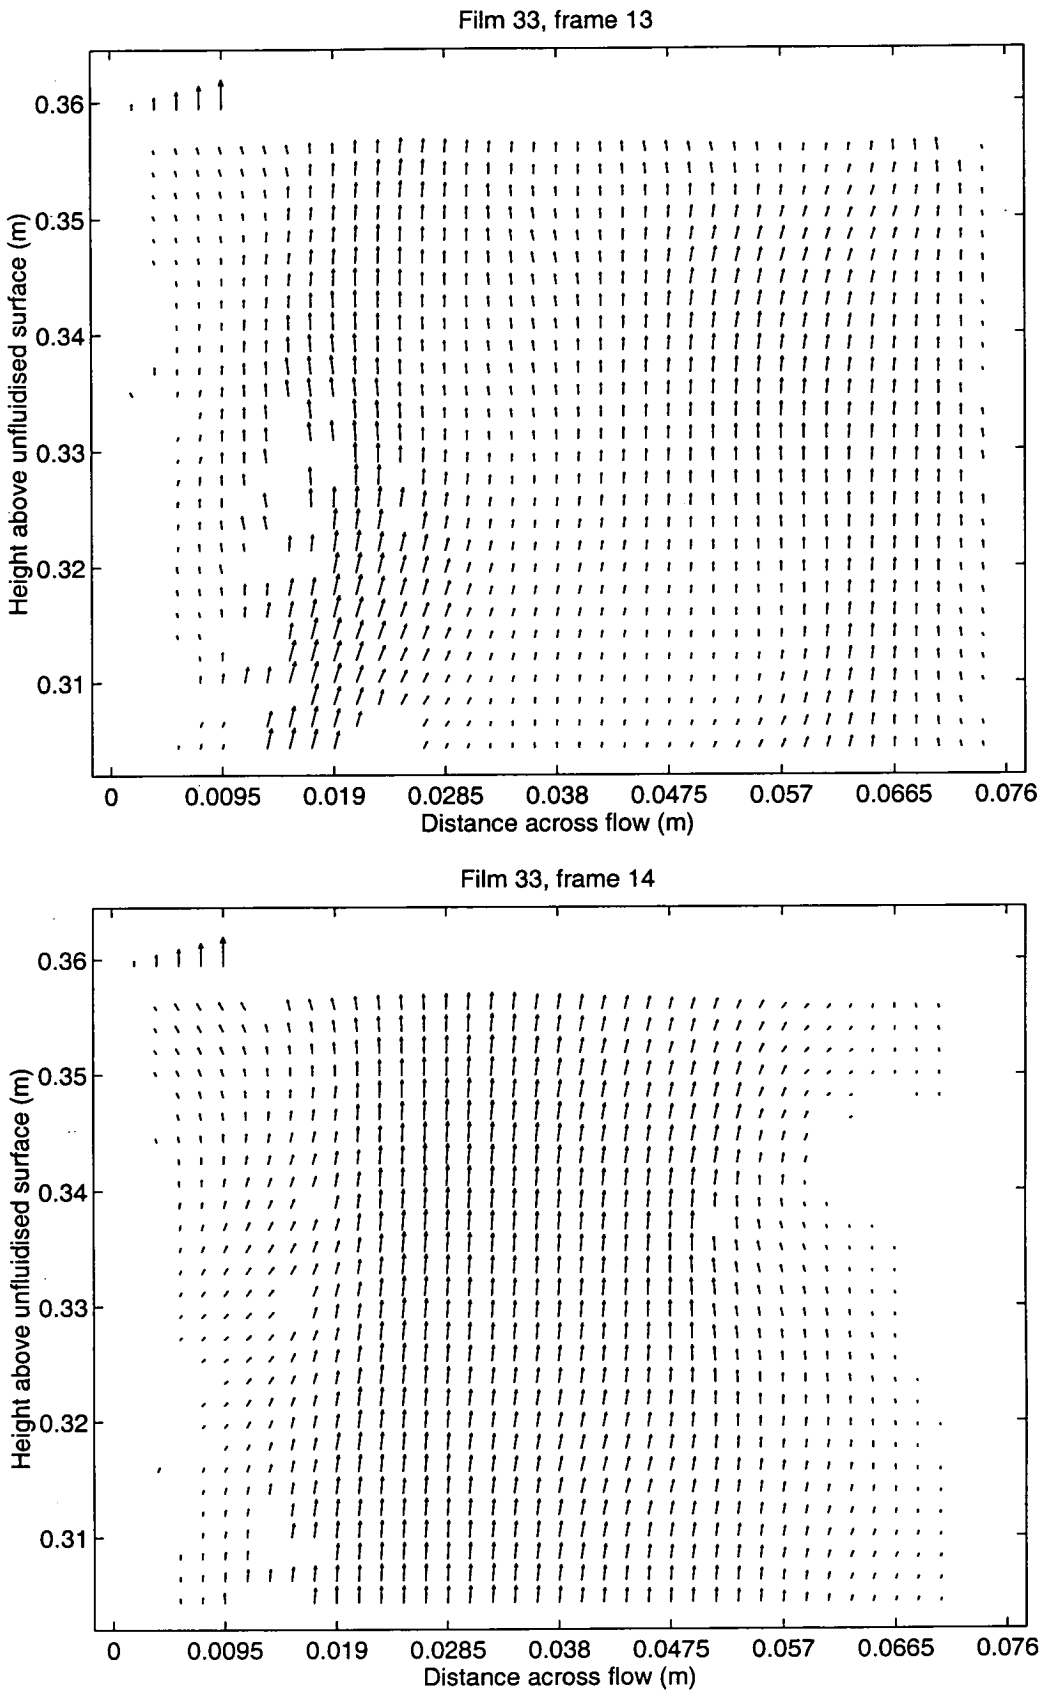


Figure 8.13: PIV vector maps for the flow about 33 cm above a 5 cm bed of particles. $U \approx 1.0\, U_{mf}$

velocities to the left of centre in frame 12 and in the centre of frame 14. Some velocity statistics are given in Table 8.5.

Frame	Yield (%)	Mean v_y ($m\,s^{-1}$)	Max v_y ($m\,s^{-1}$)	Min v_y ($m\,s^{-1}$)	Mean v_x ($m\,s^{-1}$)	Max v_x ($m\,s^{-1}$)
f33f0	95	0.081	0.14	+0.005	−0.026	−0.08
f33f1	75	0.086	0.21	+0.020	−0.005	−0.08
f33f2	94	0.132	0.22	+0.030	+0.010	+0.06
f33f3	89	0.134	0.23	+0.020	+0.026	+0.08
f33f4	88	0.123	0.24	−0.100	+0.009	+0.05
f33f5	90	0.137	0.21	+0.030	−0.014	−0.06
f33f11	76	0.081	0.19	+0.005	−0.007	+0.07
f33f12	86	0.176	0.38	−0.025	−0.024	−0.11
f33f13	91	0.186	0.38	+0.030	+0.013	+0.12
f33f14	84	0.192	0.36	+0.010	+0.023	+0.10
f33f15	94	0.173	0.32	+0.010	+0.005	+0.10
f33f16	82	0.218	0.38	−0.080	+0.021	−0.08
mean:	87	0.143	0.27	−0.004	+0.001	~ 0.08

Table 8.5: Summary of velocity values for film 33, where the bed depth was 5 cm and the mean viewing height was 33 cm above the unfluidised bed surface.

Comparing the two runs with a 5 cm bed depth, the horizontal velocity standard deviation was reduced higher up in the freeboard (*i.e.* the greatest horizontal velocities were smaller). There appeared to be a net horizontal flow to the left of $0.015\,m\,s^{-1}$ at the lower viewing height, but no net horizontal flow was found at the greater height; the opposite was observed above the shallower bed (Table 8.6).

Film number	Bed depth (cm)	Viewing height (m)	Mean v_y ($m\,s^{-1}$)	Typical max v_y ($m\,s^{-1}$)	Typical min v_y ($m\,s^{-1}$)	Mean v_x ($m\,s^{-1}$)	Typical max v_x ($m\,s^{-1}$)
f37	3	0.33	0.133	0.25	−0.001	−0.015	0.08
f39	3	0.18	0.093	0.23	−0.06	+0.003	0.12
f33	5	0.33	0.143	0.27	−0.004	−0.001	0.08
f41	5	0.17	0.094	0.27	−0.08	−0.015	0.13

Table 8.6: Summary of average velocities for the films in this chapter. 'Mean v_y' is the mean of the various section-average vertical velocities from the analysed frames of each film, whilst 'Typical max v_y' refers to the average of each frame's value of greatest upward vertical velocity, *etc.*

The largest vertical velocity could be slightly larger at the greater freeboard height, reaching about 2–6 U at a mean height of 33 cm. A more significant increase was found in the minimum vertical velocity, which was about 0.06–

$0.08\,\mathrm{m\,s^{-1}}$ downward at the lower viewing height but almost zero at the greater freeboard height (above both beds). This resulted in the mean vertical velocities becoming somewhat larger as the viewing height increased; the values of about 1.2–3 U found 33 cm above the 5 cm bed were the largest of all four sets of results.

For a superficial air velocity of $0.070\,\mathrm{m\,s^{-1}}$ in each case, the expected values for fully-developed laminar flow are a peak vertical velocity of $0.1468\,\mathrm{m\,s^{-1}}$ and a section-average of $0.1008\,\mathrm{m\,s^{-1}}$. The measured mean velocities were either about 7% less (at the lower viewing height) or about 35% greater (at the higher position) than these theoretical values. The peak vertical velocities were about 70% greater.

Comparing the two runs at a freeboard height of about 17 cm, the horizontal velocity components were similar in terms of both peak magnitude and the absolute value of mean velocity in a typical vector map. The vertical velocity components were amplified at the greater bed depth: the maximum velocity became more upward and the minimum velocity became more downward. The section-average velocities were comparable, however. Similar trends were observed when the two runs at a freeboard height of about 33 cm were compared.

In summary, increasing the bed depth appears to produce greater vertically-upward and vertically-downward velocities, without significantly altering the mean values. On the other hand, increasing the viewing height increases the section-average flow through the centre of the freeboard (by reducing the magnitude of the downward velocities, without changing the upward velocities).

8.2.3 Ensemble averages

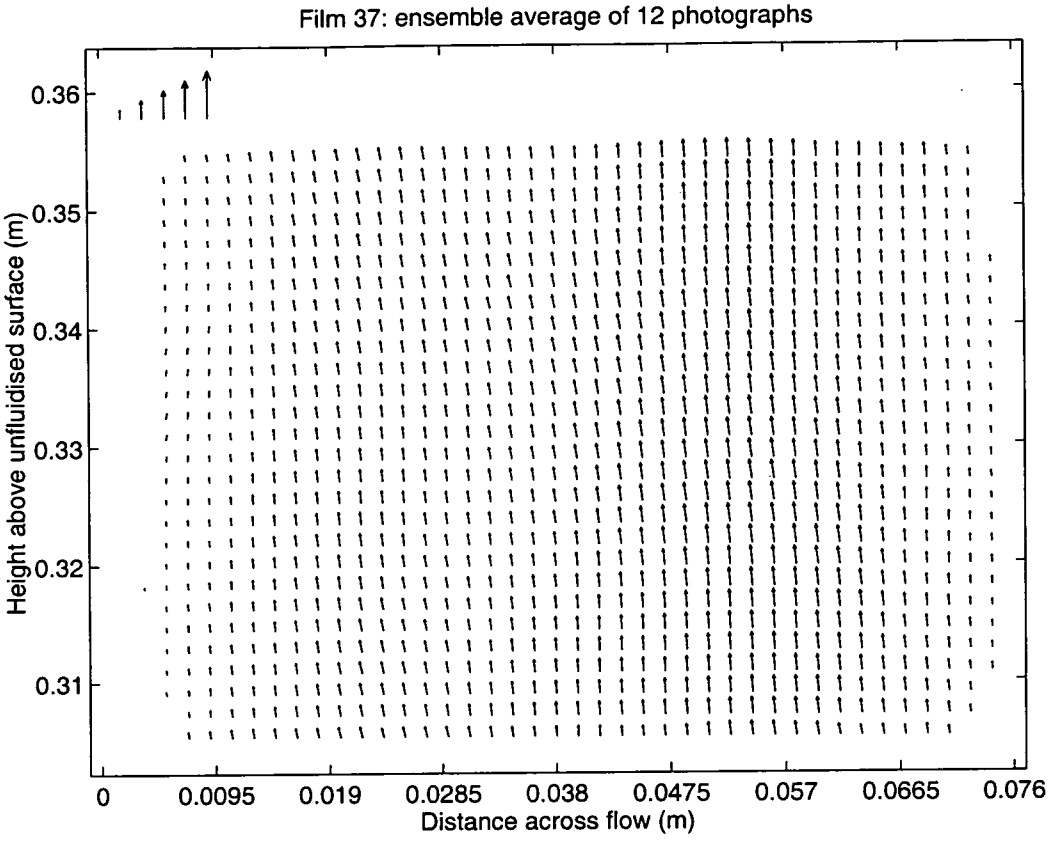
After obtaining the above vector maps, it was possible to combine many interpolated vector maps into a single, ensemble-average vector map. At each point, the vectors from the various frames were added together and then divided by the number of non-zero vectors. The averages near the walls were often derived from only some of the frames, however, so results in those regions were rather

erratic (even after discarding average velocities based on only 1–4 non-zero vectors). Ensemble-averaged vector maps are shown in Figure 8.14 (for a bed depth of 3 cm) and Figure 8.15 (when the bed depth was 5 cm).

Table 8.6 suggested that the average vertical velocity was somewhat larger at the greater viewing height. Perhaps the flow immediately above an evenly-bubbling shallow bed is relatively uniform throughout the cross-section, whereas the velocities further up the freeboard become somewhat larger along the bed axis and somewhat lower near the walls (especially in the corners), as laminar flow becomes established in the square freeboard (Figure 2.8). The total volumetric gas upflow must be the same at each height, but the *distribution* of gas flow across a horizontal cross-section should change as a function of height: air must be drawn towards the centre of the bed, as the laminar velocity profile develops.

However, Figure 8.16 suggests that this was not always happening at the heights investigated using the smaller apparatus. In this figure, vertical velocity profiles across the central cross-section are plotted, using data from six different heights within each ensemble-averaged vector map (Figures 8.14 and 8.15). The use of three different symbols when plotting vertical velocities in Figure 8.16 suggests that there is no dramatic difference between the six heights plotted. Points marked with '+' are vertical velocities from horizontal sections near the top of each ensemble-averaged vector map, those marked by 'o' are more central and those indicated by '×' signs are from nearer the bottom of each vector map.

The ideas outlined two paragraphs ago are supported by the experimental results above the deeper bed (Figure 8.16, right). At the lower viewing height, the velocity profile was skewed and also included some negative velocities. The profiles at the greater viewing height were reasonably laminar, but suggested that the actual air flow was greater than that measured (*cf* the large velocities found in Table 8.5). Thus, the velocity profile did appear to become more laminar as the viewing height was increased, by means of a transfer of gas flow towards the centre of the bed. However, it is not possible to be sure of this, as these results are from only one thin vertical slice through the flow.



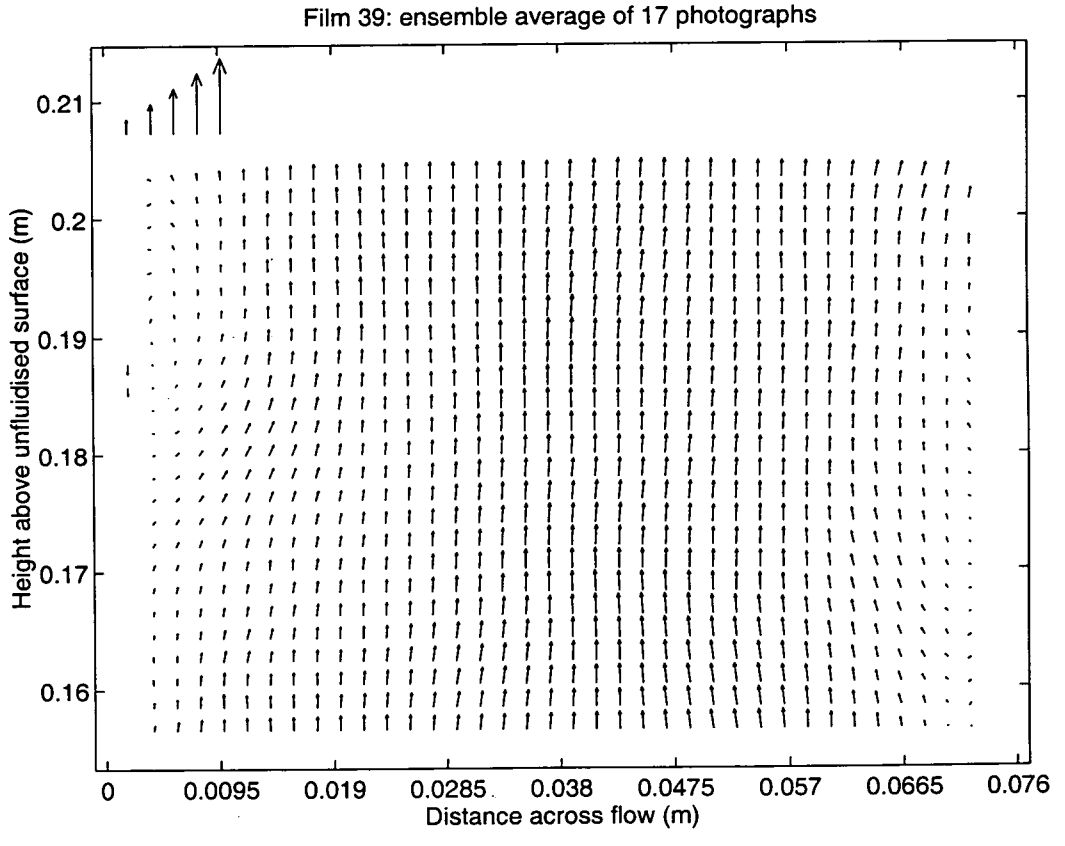


Figure 8.14: Ensemble-averaged PIV vector maps for the flow about 33 cm (top) and about 18 cm (bottom) above a 3 cm bed of particles. $U \approx 1.0\,U_{mf}$; Re = 350.

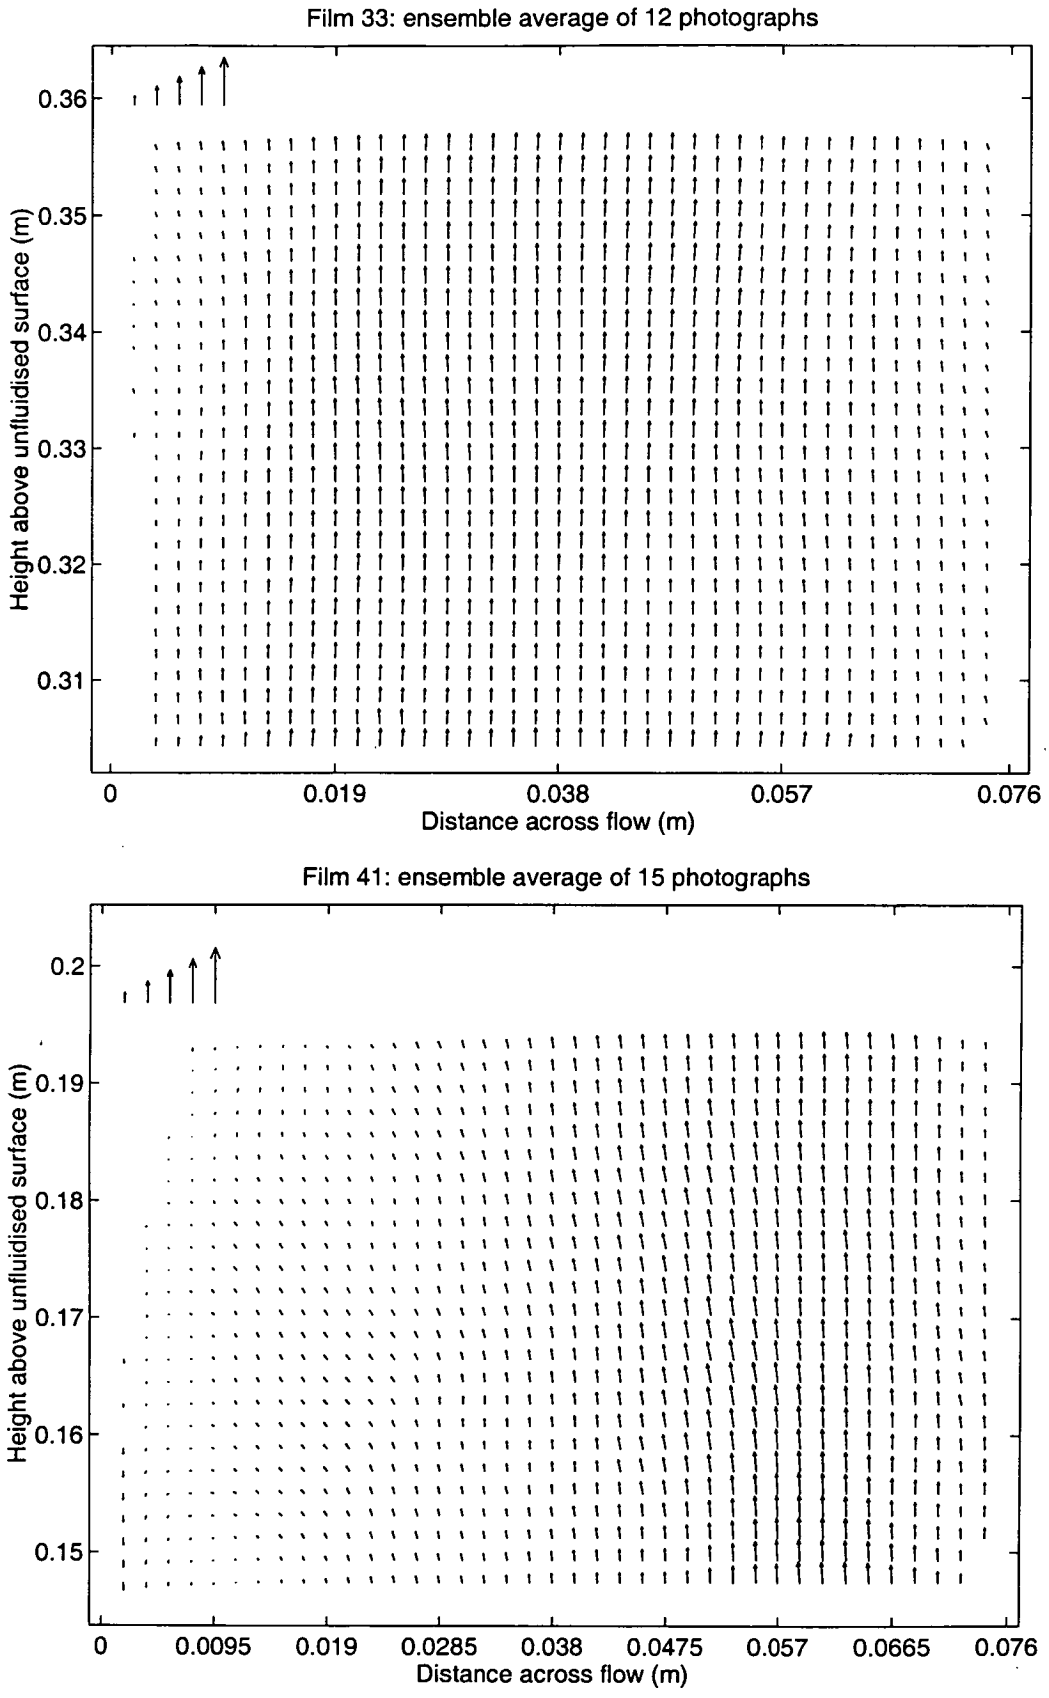


Figure 8.15: Ensemble-averaged PIV vector maps for the flow about 33 cm (top) and about 17 cm (bottom) above a 5 cm bed of particles. $U \approx 1.0\,U_{mf}$; Re = 350.

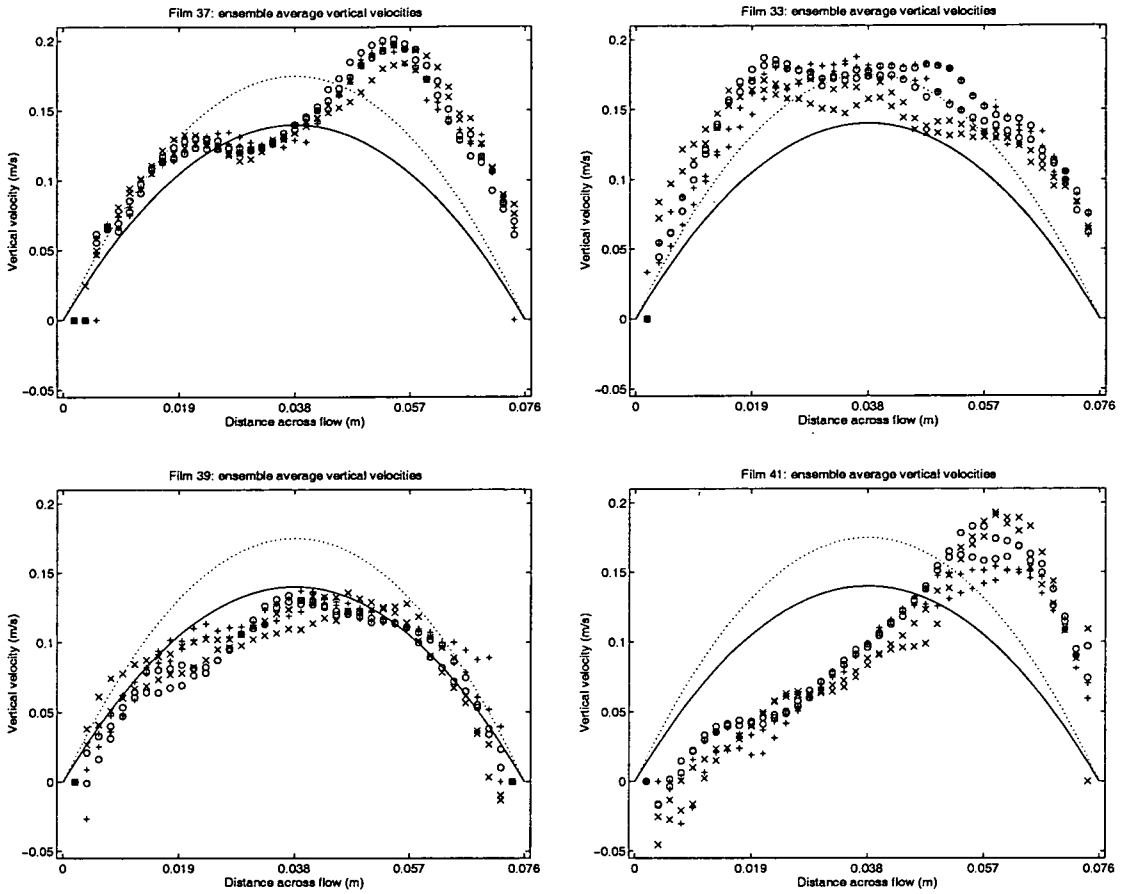


Figure 8.16: Vertical velocity profiles across the freeboard of the smaller fluidised bed, at viewing heights of 33 cm (top) and about 17 cm (bottom) above beds of particles which were 3 cm (left) or 5 cm (right) deep. The curves represent the laminar velocity profile based on the given flowrate (continuous curves) and 25% more than the given flowrate (dotted curves). $U \approx 1.0\,U_{mf}$; Re = 350.

However, the trend appeared to be reversed for the shallower bed (Figure 8.16, left). At the lower viewing height, the ensemble-averaged velocities were actually quite well described by a laminar profile, even though the individual vector maps were not particularly laminar (Figures 8.4 and 8.6). The equivalent profiles at the greater viewing height were less parabolic; greater velocities than expected were found and the profiles were not symmetric. However, better agreement might well have been obtained if several hundred frames had been analysed when calculating each ensemble-average vector map, rather than just 12–17.

It seems that the laminar velocity profile given in Section 2.4.6 is a surprisingly reasonable estimate of the time-average vertical velocity profile in this apparatus. (If something like 20 bed widths are required for laminar flow to develop, then

this means that a freeboard height of 1.5 m would be required above this bed!) Nevertheless, for slow-moving particles which spend a long time in the freeboard, the average upward local gas velocity could well be something like 2–2.5 U.

However, elutriation is more likely to involve particles being swept upwards by a jet of rapid gas upflow, with a low residence time in the freeboard. Ensemble-average velocity profiles may not be relevant, because the instantaneous vector maps were found to be quite different. Upward gas velocities of about 4 U were found in most frames; this seems to be a more appropriate figure when performing elutriation calculations.

Chapter 9

Experimental results with the larger-scale apparatus

Numerous PIV photographs were taken using the larger fluidised bed, the cross-section of which measured 30 cm × 30 cm. A number of parameters were varied:

- two bed depths were used: 31 cm and 60 cm.
- three main viewing heights were chosen, each using a different section of the freeboard as the test section. These heights were about 40 cm, about 140 cm and about 200 cm above the surface of the unfluidised bed (*i.e.* measured when the air flow was off).
- three different air flowrates were used, which converted to superficial gas velocities of about $1.25\,U_{mf}$, $1.75\,U_{mf}$ and $2.25\,U_{mf}$.
- in most cases, photographs were taken of the flow in two different vertical planes in the test section. One of these planes ran through the centre of the bed, whilst the other lay halfway between the front wall and the centre of the bed (*i.e.* one quarter of the way from the front wall of the test section to the back wall). Each plane was parallel to the front and back walls, and perpendicular to the two side walls.

This produced a total of 36 possible sets of experimental conditions, most of which were recorded on a separate roll of film. Each film had up to 36 exposures, so more than 1200 individual vector maps could have been obtained if every exposure on every roll of film had been analysed. Even presenting two vector maps and flow statistics from 5–10 analysed frames for each permutation would take up a large amount of space in this thesis. Instead, the PIV results through the centre of the freeboard above a 31 cm-deep fluidised bed will be presented in some detail, with brief comments accompanying the flows found at the greater bed depth and at the nearer viewing position.

These results included some situations which were very well suited to the direct PIV technique (*i.e.* PIV without the aid of image-shifting) and some situations where image-shifting appeared to be absolutely essential. Unless otherwise stated, the analyses took place using 128 × 128 interrogation areas and no zero-padding.

Various fluidisation properties can be estimated for this larger fluidised bed, using correlations from the literature. For example, the bubble diameter at the bed surface depends on both the bed depth and the superficial velocity. It was found for the smaller bed that estimates of the bubble diameter using equations 2.7 (Mori & Wen, 1975) and 2.8 (Rowe, 1976) produced very similar results; this seemed to be the case here too. Because of its greater ease of use, Rowe's result was preferred. The resulting bubble diameters seemed reasonable. They then allowed bubble rise velocities to be estimated, using $K = 1$ in Rowe & Partridge's result that $U_{b,\infty} = K\sqrt{g\,D_b/2}$ (equation 2.9) and then equation 2.11. Hence Table 9.1:

	U/U_{mf} =	1.25	1.75	2.25
31 cm bed	$\Rightarrow D_b$ =	0.031 m	0.054 m	0.070 m
	$\Rightarrow U_b$ =	0.40 m s^{-1}	0.57 m s^{-1}	0.67 m s^{-1}
60 cm bed	$\Rightarrow D_b$ =	0.051 m	0.088 m	0.114 m
	$\Rightarrow U_b$ =	0.51 m s^{-1}	0.71 m s^{-1}	0.84 m s^{-1}

Table 9.1: Bubble diameters (Rowe, 1976) and rise velocities (Rowe & Partridge, 1965) in the larger fluidised bed.

These velocities should be checked against the rise velocity of a slug in this bed. Equations 2.10 and 2.11 give $U_b = 0.60\,\mathrm{m\,s^{-1}} + (U/U_{mf} - 1) \times 0.07\,\mathrm{m\,s^{-1}}$ in the laboratory frame, which suggests that the bubble rise velocities listed in Table 9.1 are similar to the corresponding slug velocities. As a result, the values for the deeper bed at the higher flowrates may be too large.

Horio *et al.* (1980) gave a correlation for the TDH: $z_{TDH} = 14\sqrt{D_b/g}$. Baron *et al.* (1988) suggested that $z_{TDH} = (7.5\,U_b)^2/2g$ for a slugging fluidised bed, which seems more appropriate than their bubbling-bed result given the velocities found here. Fournol *et al.* (1973) found that $g\,z_{TDH}/U^2 = 1000$, but caution is needed because their result was based on experiments using fine powders. The values of TDH predicted by these relationships are listed in Table 9.2.

$U/U_{mf} =$		1.25	1.75	2.25
31 cm bed				
Horio *et al.*:	$z_{TDH} =$	0.77 m	1.0 m	1.2 m
Baron *et al.*:	$z_{TDH} =$	0.46 m	0.9 m	1.3 m
Fournol *et al.*:	$z_{TDH} =$	0.78 m	1.5 m	2.5 m
60 cm bed				
Horio *et al.*:	$z_{TDH} =$	1.0 m	1.3 m	1.5 m
Baron *et al.*:	$z_{TDH} =$	0.7 m	1.4 m	2.0 m
Fournol *et al.*:	$z_{TDH} =$	0.8 m	1.5 m	2.5 m

Table 9.2: Transport Disengaging Heights predicted for the larger fluidised bed by Horio *et al.* (1980), Baron *et al.* (1988) and Fournol *et al.* (1973).

The agreement is very good in some situations, *e.g.* when using a superficial velocity $U = 1.75\,U_{mf}$ in the deeper bed. It seems that the lowest freeboard height of 0.40 m should definitely be below the TDH, at all flowrates. The intermediate freeboard height of 1.3 m is probably close to the TDH when $U/U_{mf} = 1.75$ (or possibly even 2.25, in the shallower bed). The highest viewing position (2.1 m) should be above the TDH in most runs, although maybe not by much when $U = 2.25\,U_{mf}$.

9.1 31 cm bed depth, centre of freeboard

For results obtained in the centre of the freeboard using the shallower bed depth, a list of some of the main film numbers and flow conditions is given in Table 9.3.

Bed depth (cm)	Viewing height (cm)	U/U_{mf}	Film number	Image-shifting used?
31	220	1.16	90	No
31	220	1.27	105	Yes
31	220	1.69	88	No
31	220	2.25	86	No
31	140	1.27	82	No
31	130	1.27	133	Yes
31	150	1.70	80	No
31	150	2.67	77	No
31	150	2.67	78	No
31	42	1.55	71	No

Table 9.3: Summary of films taken using PIV to record the flow through the centre of the freeboard space above a 31 cm fluidised bed.

9.1.1 Viewing heights of 220 cm

Starting at the top of the freeboard, at a viewing height of about 210–230 cm above the bed surface, the intermediate air flow ($U = 1.69\,U_{mf}$) gave exceptionally good results with the direct PIV method. An illumination time interval of 10 ms was used in film 88, giving three images of each particle throughout most of the flow record when the camera was set to 1/30 of a second.[1] After analysis, the vectors were found to vary from about 7–20 pixels in magnitude, indicating that the dynamic range was low. In addition, all vectors were approximately vertical, which meant that the displacements were roughly parallel to two of the sides of

[1] An illumination time interval of about 10.3 ms would have been marginally better, producing exactly three images of each particle in the light sheet.

each interrogation area. In fact, conditions were so good for PIV that the average yield of valid vectors was over 99%. Vector alignment was straightforward, so there was nothing to be gained by the use of image-shifting. Sample vector maps are given in Figure 9.1. A photograph of the flow in frame 11 is given in Figure 9.2, along with some velocity statistics.

The intermediate air flowrate was about $640\,\mathrm{l\,min^{-1}}$, which gave a superficial velocity of $0.118\,\mathrm{m\,s^{-1}}$. The greatest vertical velocity was almost always somewhat larger than the superficial velocity. However, the mean vertical velocity found by PIV was about $0.086\,\mathrm{m\,s^{-1}}$, obtained by averaging across each vertical section through the centre of the freeboard. Thus PIV found section-averaged velocities which were lower than the superficial velocity, whereas the model of laminar flow in the square freeboard predicts that the PIV-average velocity should be about $1.44\,U$ (Section 2.4.6). Even if turbulent flow is found instead, the section-average velocity obtained from PIV measurements should still exceed U if the light sheet passes through the centre of the bed. Given $U = 0.118\,\mathrm{m\,s^{-1}}$ in this apparatus, a check on the Reynolds number reveals that $\mathrm{Re} = 2350$ or so, which is somewhat larger than the upper limit of $\mathrm{Re} = 2000$ for laminar flow (*NB*: in a circular pipe). As the flow immediately above the bed surface will certainly not be laminar, it seems unlikely that the flow will switch into the laminar regime in the freeboard of the larger fluidised bed if U is greater than about $0.11\,\mathrm{m\,s^{-1}}$.

When the gas velocity was increased to $U = 0.157\,\mathrm{m\,s^{-1}} = 2.25\,U_{mf}$, reasonable results were obtained in film 86 using an illumination time interval of 8 ms. It was again helpful that the gas velocities were all nearly vertical. There were four images of each particle in most of the flow record and, thanks to the reduced Δt, the image separations remained satisfactory (increasing slightly to about 5 to 21 pixels). A couple of vector maps for this flow are given in Figure 9.3.

The mean air flowrate under these conditions was about $850\,\mathrm{l\,min^{-1}}$, hence the superficial velocity was about $0.157\,\mathrm{m\,s^{-1}}$. Once again, the mean vertical velocity found by PIV was somewhat lower than this, at about $0.130\,\mathrm{m\,s^{-1}}$ (Table 9.4). As with the intermediate flowrate (film 88, above), there was an apparent 'slip

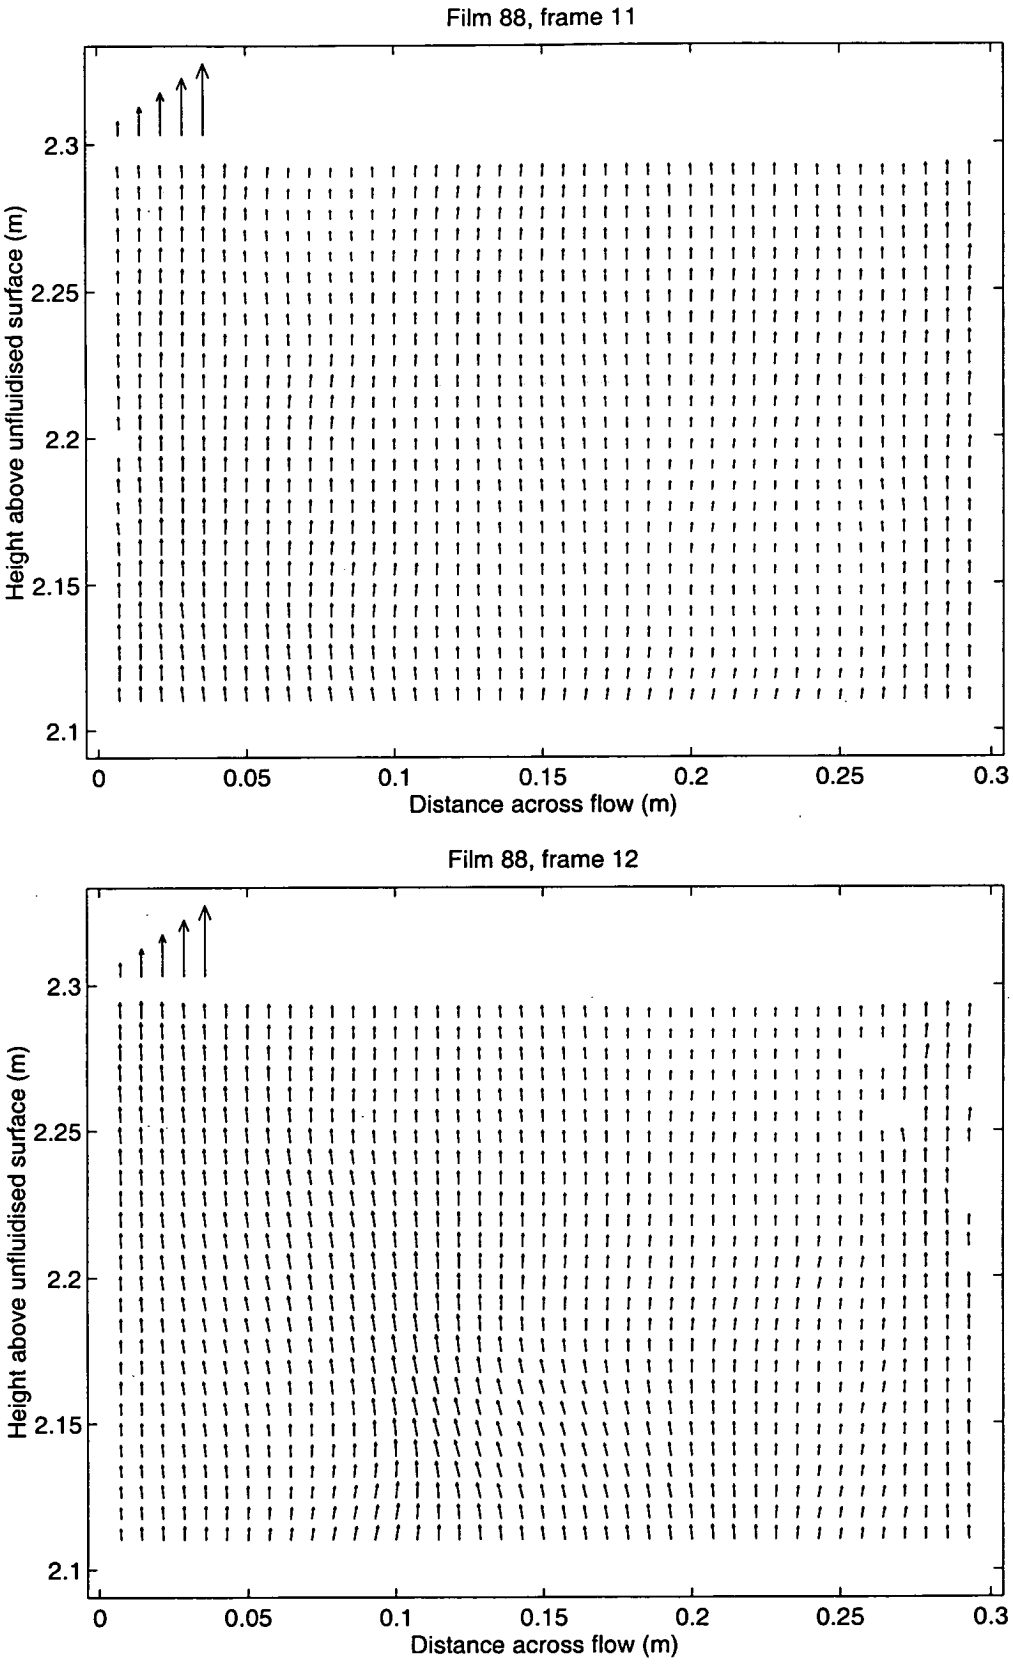
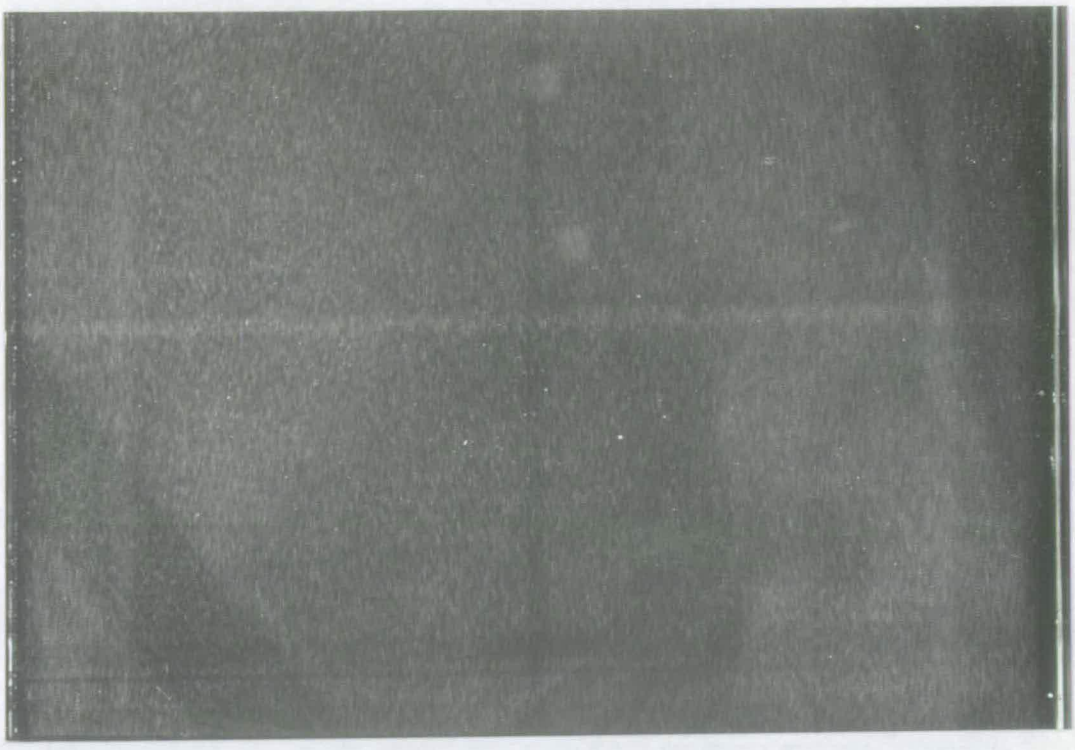


Figure 9.1: PIV vector maps for the flow about 220 cm above a 31 cm bed of particles. $U \approx 1.7\,U_{mf}$. The marker vectors in the top left corner are 0.1–0.5 m s^{-1} in magnitude.



Frame	Yield (%)	Mean v_y ($\mathrm{m\,s^{-1}}$)	Max v_y ($\mathrm{m\,s^{-1}}$)	Min v_y ($\mathrm{m\,s^{-1}}$)	Mean v_x ($\mathrm{m\,s^{-1}}$)	Max v_x ($\mathrm{m\,s^{-1}}$)
f88f10	100	0.0821	0.115	0.050	−0.0012	−0.030
f88f11	100	0.0823	0.125	0.045	+0.0003	±0.025
f88f12	99	0.0909	0.140	0.040	−0.0042	−0.040
f88f13	100	0.0994	0.140	0.060	−0.0047	−0.045
f88f14	100	0.0900	0.125	0.050	−0.0031	−0.030
f88f15	99	0.0855	0.135	0.035	+0.0001	+0.040
f88f16	99	0.0793	0.125	0.035	+0.0010	+0.040
f88f17	100	0.0824	0.125	0.040	+0.0030	±0.030
f88f18	100	0.0839	0.130	0.040	+0.0007	−0.030
f88f19	98	0.0742	0.120	0.040	−0.0022	−0.040
f88f20	100	0.0915	0.165	0.050	−0.0009	+0.045
f88f21	100	0.0779	0.140	0.040	−0.0007	±0.035
f88f22	99	0.0886	0.140	0.040	−0.0010	±0.040
f88f23	99	0.0911	0.145	0.035	+0.0037	+0.040
f88f25	100	0.1038	0.145	0.050	−0.0016	±0.030
f88f26	100	0.0924	0.130	0.050	−0.0058	−0.040
f88f27	100	0.0783	0.115	0.050	−0.0038	−0.040
f88f28	100	0.0766	0.110	0.045	−0.0058	−0.040
mean:	99	0.0861	0.132	0.044	−0.0015	~ 0.037

Figure 9.2: Top: PIV photograph from frame 11 of film 88; Bottom: velocity values for this film, using $U = 0.118\,\mathrm{m\,s^{-1}} = 1.69\,U_{mf}$. The bed depth was 31 cm and the mean viewing height was about 220 cm above the unfluidised bed surface.

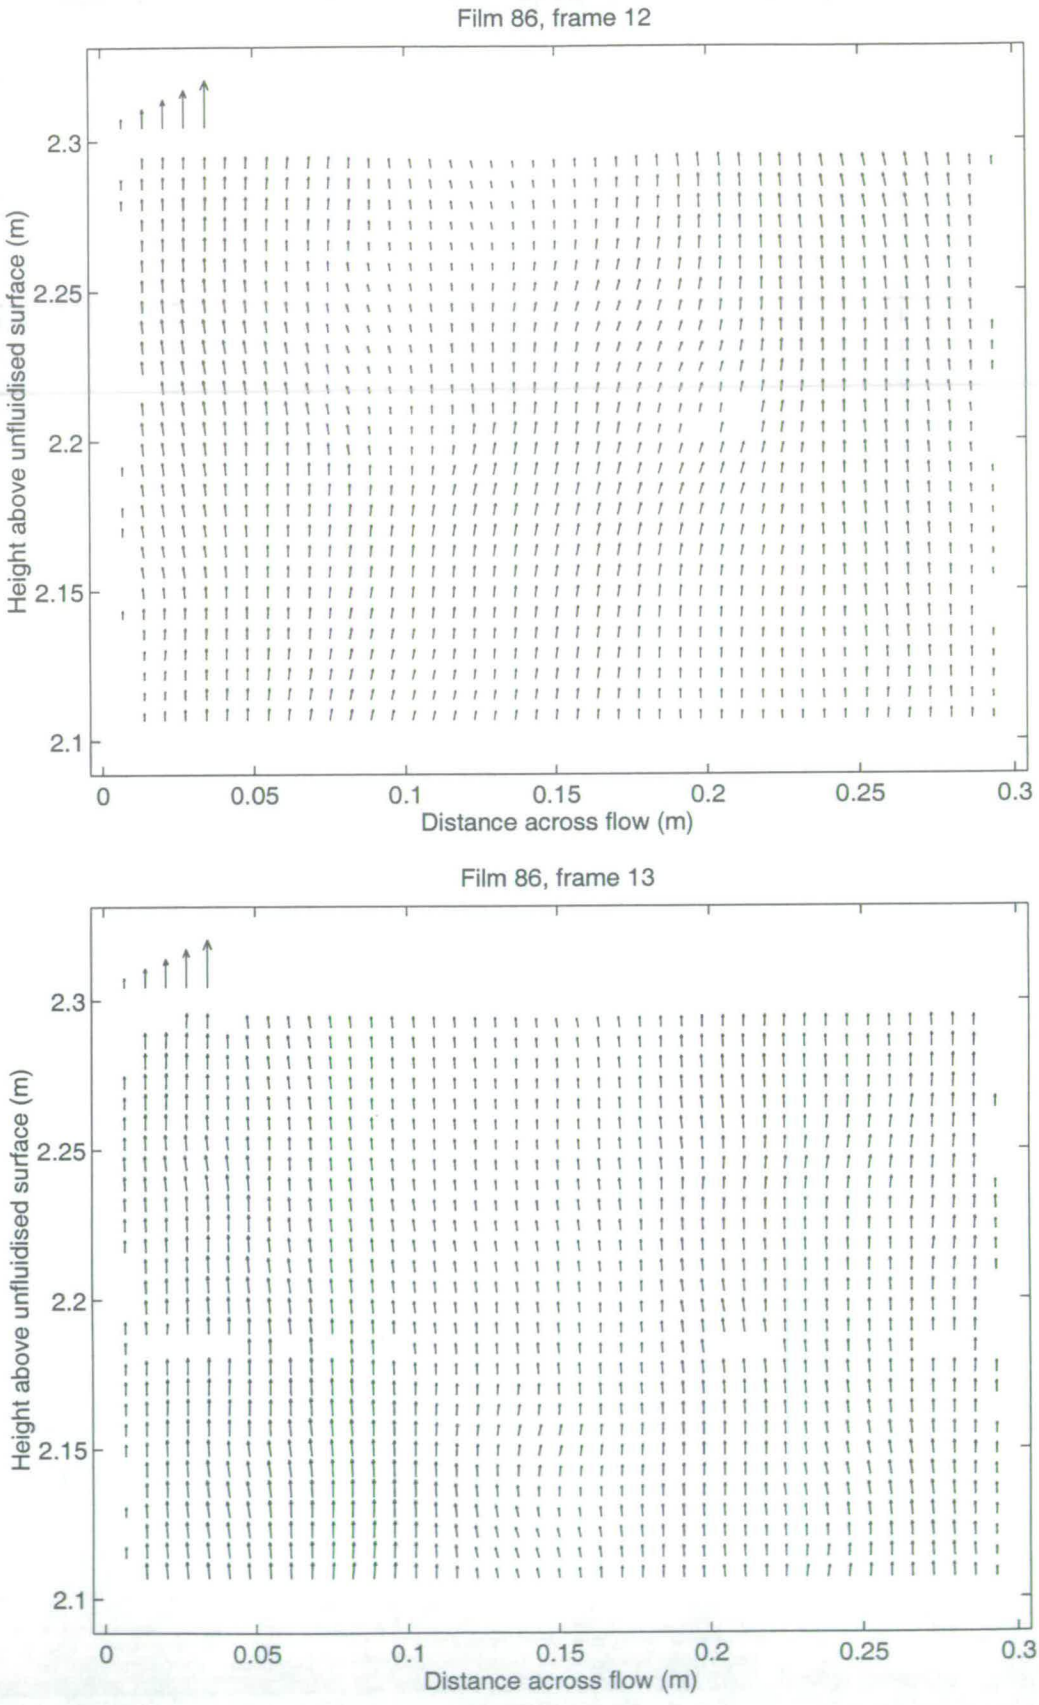


Figure 9.3: PIV vector maps for the flow about 220 cm above a 31 cm bed of particles. $U \approx 2.25\, U_{mf}$.

Frame	Yield (%)	Mean v_y ($m\,s^{-1}$)	Max v_y ($m\,s^{-1}$)	Min v_y ($m\,s^{-1}$)	Mean v_x ($m\,s^{-1}$)	Max v_x ($m\,s^{-1}$)
f86f10	97	0.1384	0.195	0.075	−0.0027	−0.070
f86f11	96	0.1320	0.185	0.055	−0.0046	−0.060
f86f12	97	0.1128	0.155	0.055	+0.0027	+0.050
f86f13	97	0.1345	0.200	0.075	−0.0051	−0.040
f86f14	95	0.1277	0.195	0.065	−0.0025	−0.060
f86f15	92	0.1376	0.220	0.075	+0.0025	+0.050
f86f16	95	0.1287	0.195	0.065	+0.0019	+0.050
f86f17	95	0.1266	0.200	0.065	+0.0030	+0.050
mean:	96	0.1298	0.193	0.066	−0.0006	~ 0.054

Table 9.4: Summary of velocity values for film 86, using $U = 0.157\,m\,s^{-1} = 2.25\,U_{mf}$. The bed depth was 31 cm and the mean viewing height was about 220 cm above the position of the bed surface (measured with the air off).

velocity' of about $0.03\,m\,s^{-1}$ between the air flow (measured by the rotameter) and the cloud of fine seeding particles.

The yields of valid vectors were about 96% in film 86. These yields are very good, even though they are not quite as impressive as those at the intermediate flowrate. Because the illumination time intervals were different, the recorded image displacements in films 86 and 88 were similar even though the flow velocities were different. As the vectors were mostly vertical, similar yields were expected. The reduced yield at the higher flowrate can be attributed to the following effects:

- a scanner problem meant that a spurious horizontal band was added near the centre of the bitmap images. This spoiled the results in a few rows of vectors (*e.g.* at viewing heights of about 2.19 m in frame 13 (Figure 9.3)).
- the reduced illumination time interval meant that the amount of light energy given to each particle during each illumination was reduced by about 20%, so the particle images were less bright.
- the displacements in the digitised images were slightly greater in film 86. The yield of valid vectors worsens as the image separation increases.
- there were a number of 'large particle' images in the freeboard at the larger velocity, but very few were observed when $U = 1.69\,U_{mf}$. These im-

ages could be from bed particles ejected into the freeboard or from seeding particles clustering together (*e.g.* due to electrostatic effects or after picking up moisture from the air). It was noted in Table 9.1 that the highest viewing position was close to the TDH at the highest flowrate (hence many bed particles were found here) but probably above it at the intermediate flowrate (few bed particles were recorded in film 88). Whilst these bed-particle images could be filtered out to some extent by passing the bitmap images through a large-image detection program (PIVSPLIT) before the PIV analysis, the deletion of these images left behind 'holes' in the bitmap recordings. These holes cannot have helped the detection of the correct peak. The application of PIV to a flow is made more difficult if two or more phases are present, although McCluskey *et al.* (1994) noted some success.

Nevertheless, these errors are not too bad, because the missing vectors occur in a thin band near the centre of the vector maps, with correct vectors found above and below. Under these circumstances, it is possible to fill in the missing information by interpolation and to have confidence in the interpolated results.

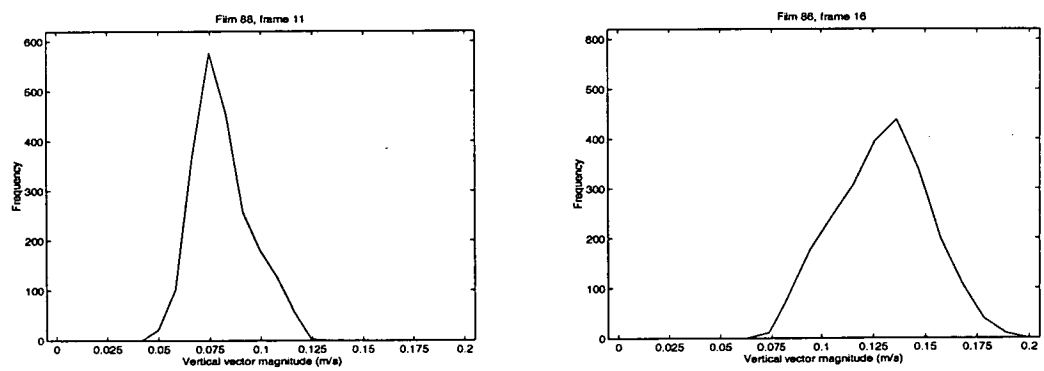


Figure 9.4: Histograms of the distribution of vertical velocities in two sample photographs from film 88 ($U \approx 1.69\,U_{mf}$, left) and film 86 ($U \approx 2.25\,U_{mf}$, right).

The vertical velocities at the intermediate and greatest flowrate are contrasted in Figure 9.4. There was a relatively-narrow distribution of vertical velocity at the intermediate flowrate (film 88; left-hand histogram). A drift towards more rapid velocities was observed at the greater flowrate (film 86; right-hand histogram). The frequency peak at the greater flowrate is lower, because the velocities (in $\mathrm{m\,s^{-1}}$) were more spread out and because the yield of valid vectors was lower.

Different problems were encountered when the air flow was turned down to $440\,\mathrm{l\,min^{-1}}$ ($U = 0.0815\,\mathrm{m\,s^{-1}} = 1.16\,U_{mf}$) for film 90. The measured velocities were generally somewhat lower than this (Figure 9.5, left). The most commonly-occurring image displacement was about 7 pixels and few vectors were larger than 12 pixels (Figure 9.5, right). It is highly likely that a number of low-velocity vectors were not detected, due to problems of overlapping particle images.

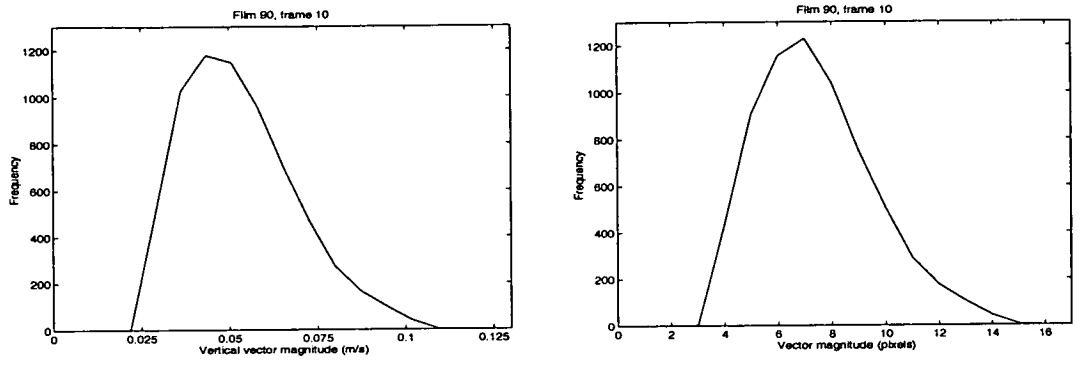


Figure 9.5: Histograms of the velocity distributions obtained from frame 10 of film 90 ($U \approx 1.16\,U_{mf}$). Left: a histogram of the vertical velocity, in $\mathrm{m\,s^{-1}}$; Right: a histogram of the magnitude of each displacement vector, in pixels.

The obvious solution to this bunching of displacements at the lower end of the scale was to extend the illumination time interval beyond the value of 12 ms used in film 90. This was not easy to achieve, however, as the longest illumination time interval provided by the given scanning-beam box was 12.5 ms. A time of nearer 20 ms would have been more appropriate. In order to achieve this, a slower-speed motor could have been purchased, four alternate faces of the octagonal rotating mirror (inside the scanning-beam box) could have been blacked out or an optical chopper could have been used.

The effects of these low displacements can be seen in Figure 9.6. In this re-analysed version of frame 10, there were large regions in which no velocity information could be determined (although it would be safe to guess that the missing vectors were rather small). Consequently, the average yield of valid vectors was about 70%. Film 90 was originally analysed using 64 × 64 interrogation areas. It was not considered worthwhile to re-analyse all the bitmap files from this film using larger interrogation areas, because these would not have helped resolve the low image displacements.

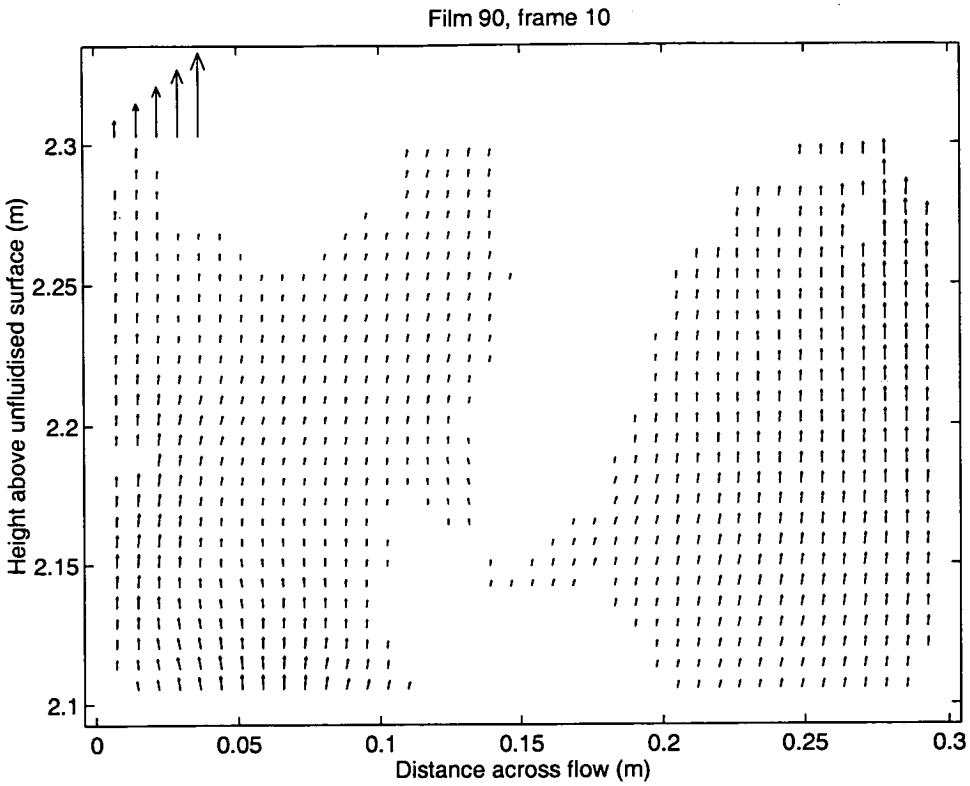


Figure 9.6: PIV vector maps for the flow about 220 cm above a 31 cm bed of particles. $U \approx 1.2\,U_{mf}$.

Problems with low velocities and directional ambiguities can be avoided by the use of image-shifting, so the technique was applied in film 105. A slightly greater air flowrate of $480\,\mathrm{l\,min^{-1}}$ was used, which gave $U = 0.0889\,\mathrm{m\,s^{-1}} = 1.27\,U_{mf}$. To achieve image-shifting, the flow was photographed via a mirror which rotated at $10^\circ\,\mathrm{s^{-1}}$ during each exposure. The mean image-shifting displacement was about 16.2 pixels to the left. An interrogation area size of 128×128 pixels was necessary, in order to avoid wrap-around errors: with four images of each particle and 64×64 interrogation areas, wrap-around errors could arise at near-horizontal displacements approximately 64/4 and 64/3 pixels in magnitude if the vertical displacement was small. Some vector maps are given in Figures 9.7–9.9.

In the vector maps obtained from film 105, the yields of valid vectors were typically about 91–94%. The main cause of missing vectors was the vertical lines from light scattered at either wall. Without image shifting, a single line was seen at each wall and it was easy to crop the bitmap images so as to retain only the area between these lines. With image-shifting, however, these vertical lines were

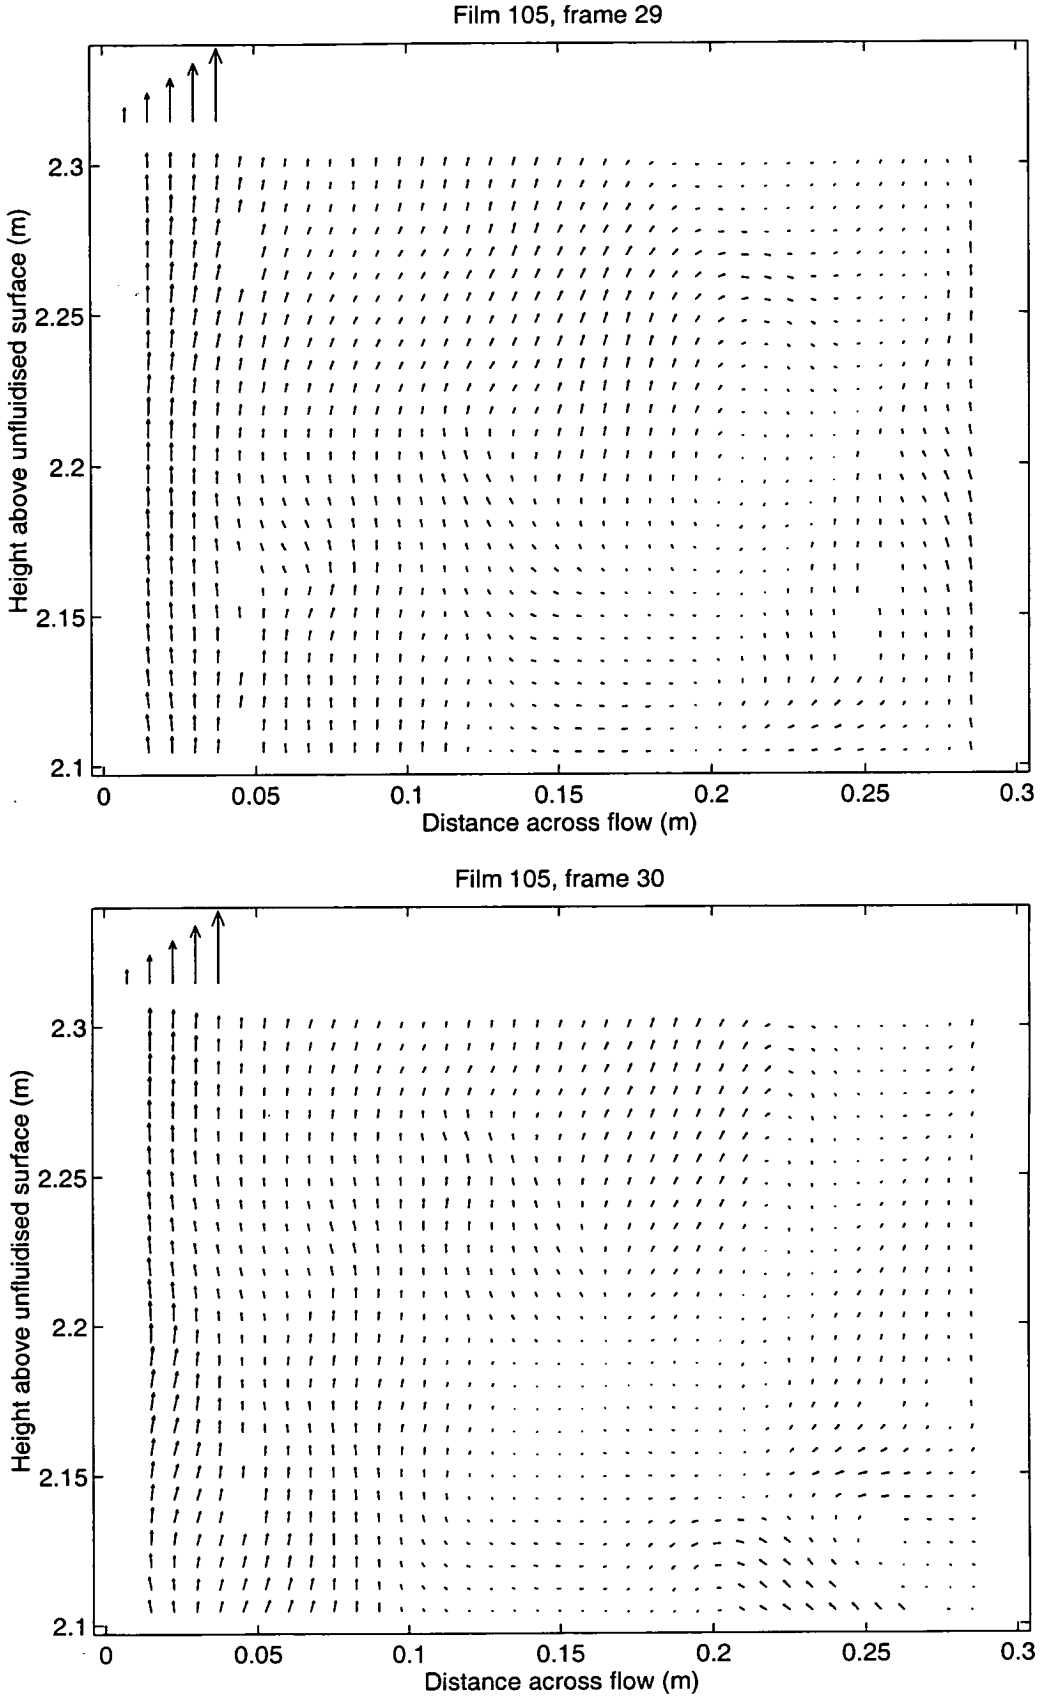


Figure 9.7: PIV vector maps for the flow about 220 cm above a 31 cm bed of particles, with image-shifting used. $U \approx 1.27\,U_{mf}$.

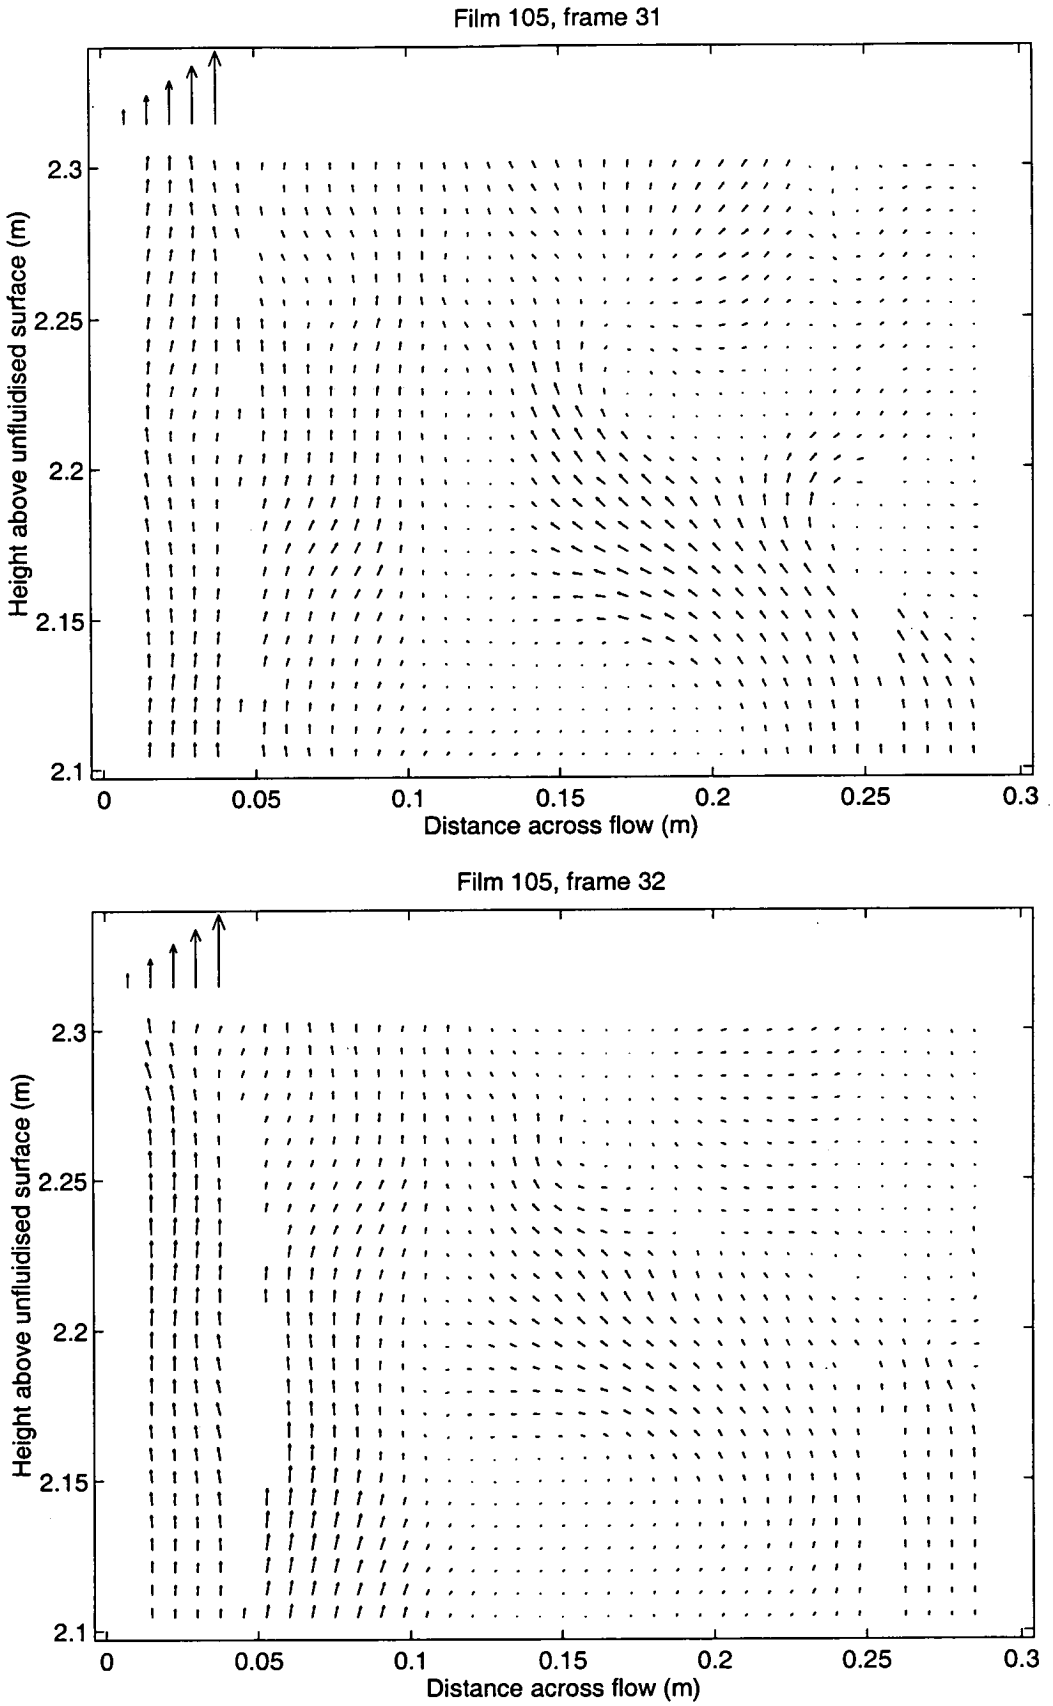


Figure 9.8: PIV vector maps for the flow about 220 cm above a 31 cm bed of particles. Image-shifting was used. $U \approx 1.27\,U_{mf}$.

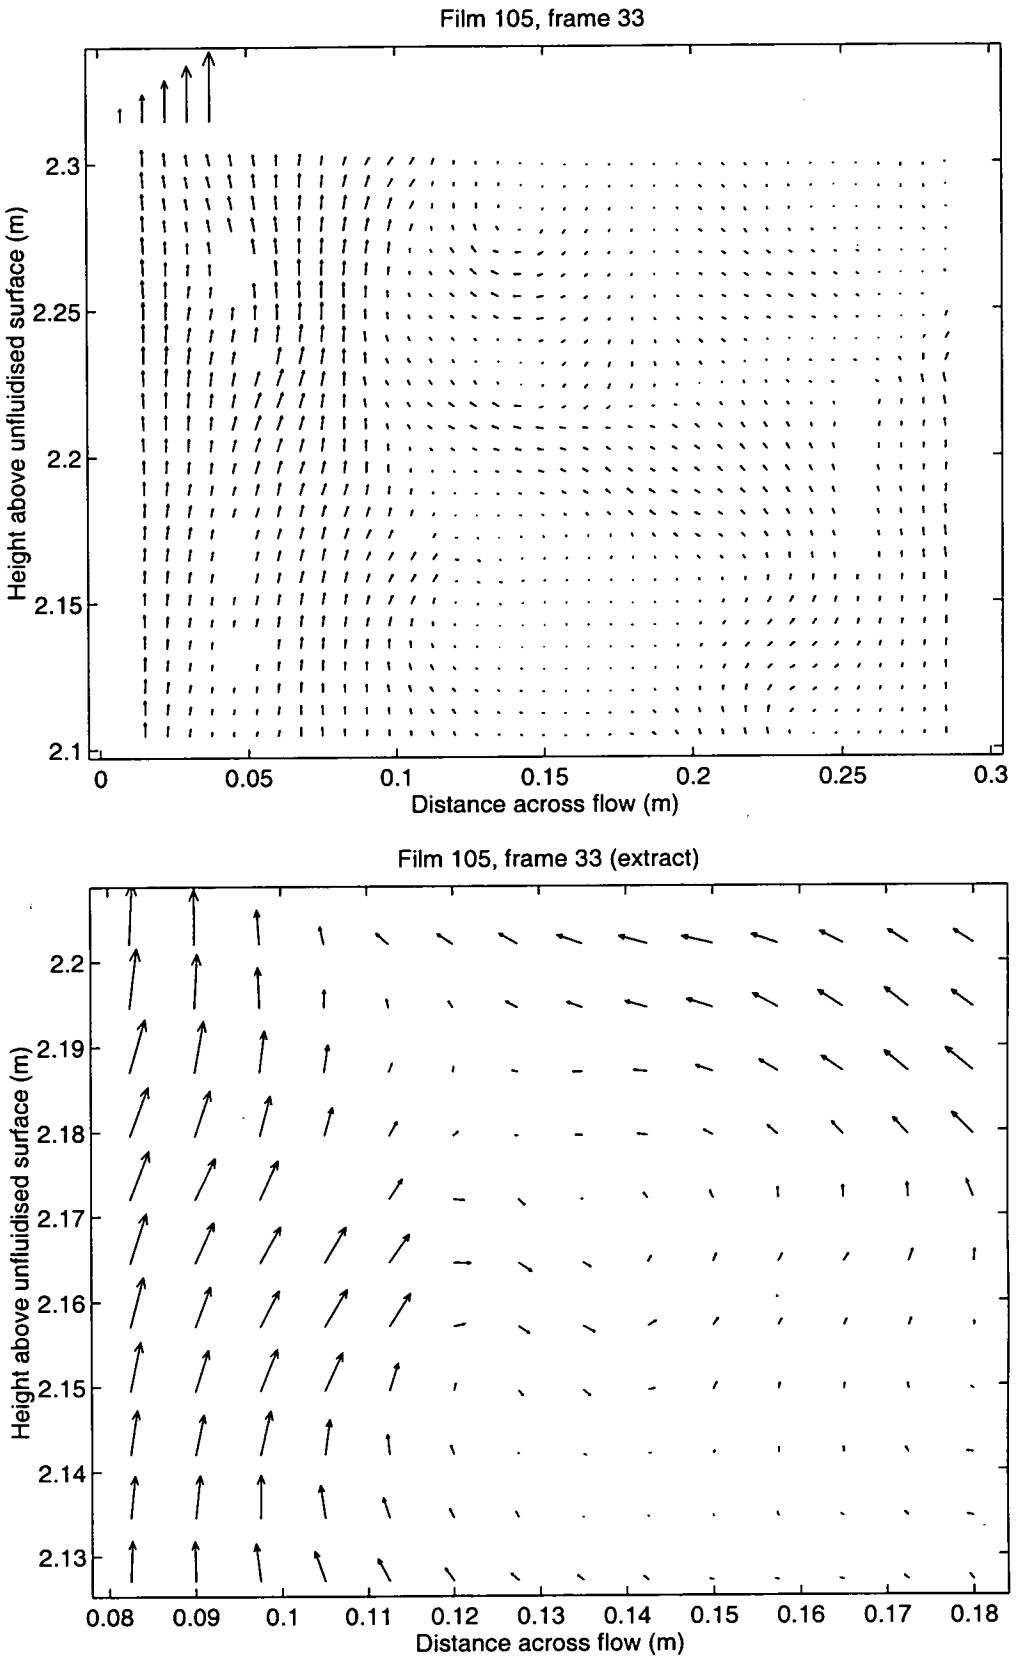


Figure 9.9: PIV vector maps for the flow about 220 cm above a 31 cm bed of particles, using image-shifting. $U \approx 1.27 U_{mf}$. Top: the full width of the test section; Bottom: a close-up view near the centre of the full vector map.

photographed in three or four different positions, due to the movement of the rotating mirror between exposures. As a result, wide regions of the negative were affected near each wall. Half of these vertical bands were deliberately retained when cropping the images before analysis, so as to continue to use bitmap files which were the digital images of 0.3 m-wide regions of the test section (so the image:object magnification could easily be calculated). This, however, meant that the left-hand-most and right-hand-most columns of vectors were wrong.

A second problem, as illustrated in Figure 9.10, was that the bright lines from the glass-air interfaces could also be reflected into the camera via the back wall of the test section. These images increased the background light level in a couple of narrow vertical bands within the flow record and often caused an erroneous vector to be detected in the affected interrogation areas. The reflected lines were out of focus, so wide areas of the film were affected. These lines were visible in some earlier films but did not cause serious problems; they were more prominent in film 105, however, causing some loss of yield. These reflections could have been avoided by sticking matt black tape to the inside of the back wall in the appropriate places, as was done in later work (film 106 and onwards).

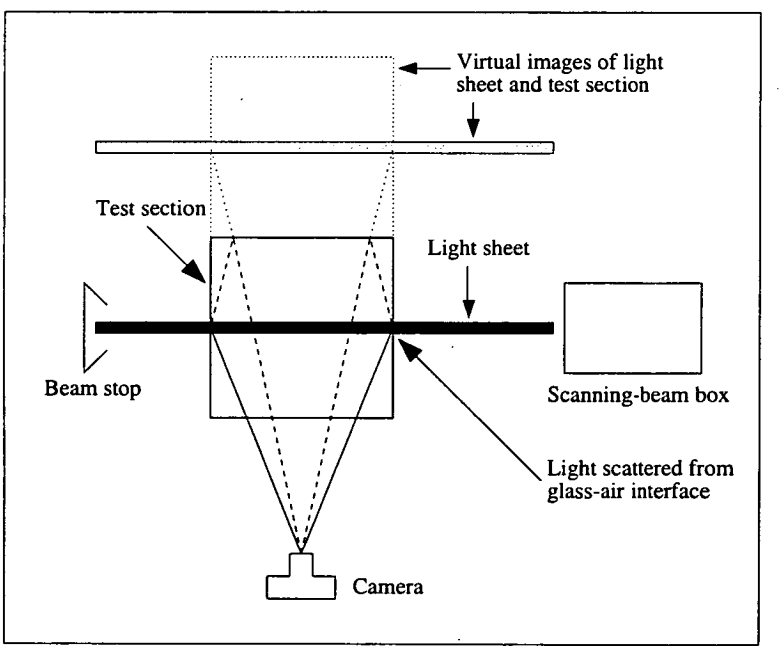


Figure 9.10: The scattering of the light sheet as it enters or leaves the test section. Images can be viewed either directly (continuous-line rays) or after reflection in the back wall of the test section (dashed lines).

A sample PIV photograph from this run clearly shows the multiple images of the glass-air interface, as well as the vague reflections of these lines in the back wall (Figure 9.11). This print is an enlargement of frame 31. The corresponding vector map was given earlier, in the top half of Figure 9.8.

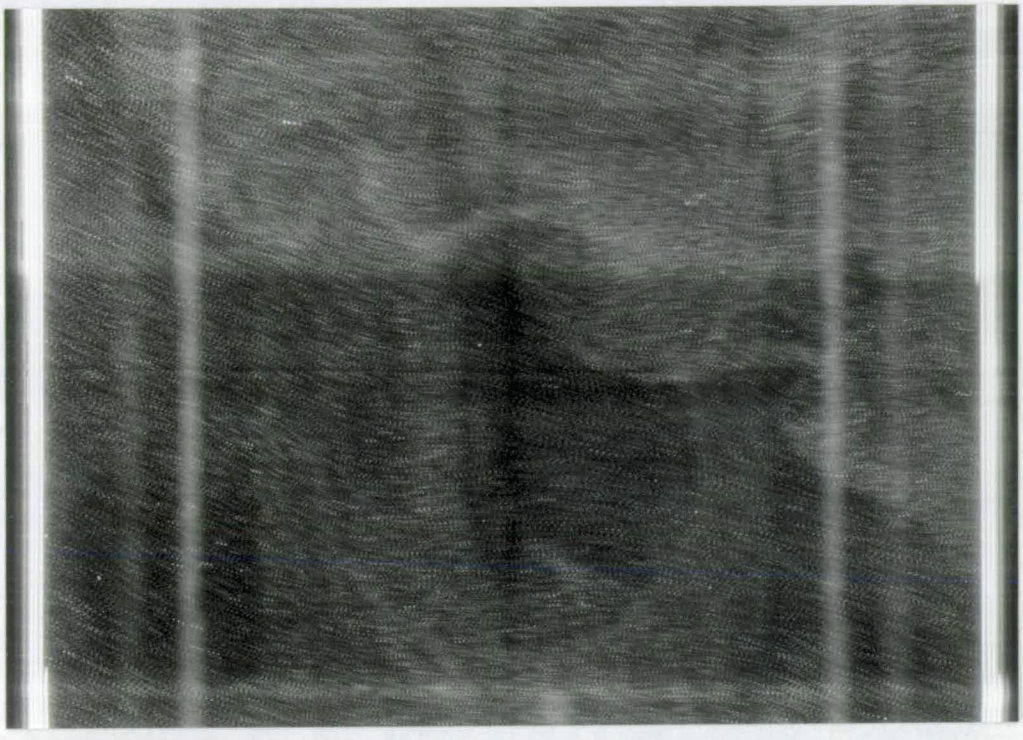


Figure 9.11: PIV photograph of the flow recorded in film 105, frame 31.

The illumination time interval in film 105 was 8.8 ms. This was an unfortunate choice, because some regions of each frame recorded four images of each particle, whilst others managed only three (*e.g.* the top 40% and the bottom 10% of frame 31 recorded four images). A more even illumination would have been obtained if Δt had been reduced to about 8 ms (allowing four complete scans to be made in the 1/30 s exposure time) or increased to about 11 ms (recording exactly three images of each particle, at a higher contrast). This practical consideration might impose an extra constraint when choosing PIV parameters for an experiment.

The vector maps shown in Figures 9.7–9.9 consisted of 27 rows and 39 columns, but the extra, direct images of the bright glass-air interface meant that the columns of vectors nearest each side wall contained no valid measurements. Thus, about $2/39 = 5\%$ of the vectors were not valid, simply because of these directly-

viewed images of the walls. Along similar lines, the reflections of the glass-air interfaces in the back wall were also recorded in about four different positions, spoiling some of the results in a further two columns. For example, it can be seen in frames 31 and 32 (Figure 9.8) that columns of missing vectors were found near either wall and also at distances of 0.045 m and 0.255 m across the flow section. Due to these limitations, the yield of valid vectors was unlikely to exceed 95%, so the yields obtained in practice were quite good under the circumstances.

In frame 29 (Figure 9.7, top) and frame 31 (Figure 9.8, top), it can be seen that almost all of the lower-velocity vectors were resolved correctly, with the missing vectors occurring mainly at random locations in regions of large upward velocity. Concerning the features of each vector map, the largest vertical velocities were always found close to the left-hand wall. The flow near the right-hand wall was also mainly vertically upward, but a number of very low velocity or even downward regions were found near the centre of the bed. This can be seen most clearly in the close-up partial vector map shown in the lower part of Figure 9.9.

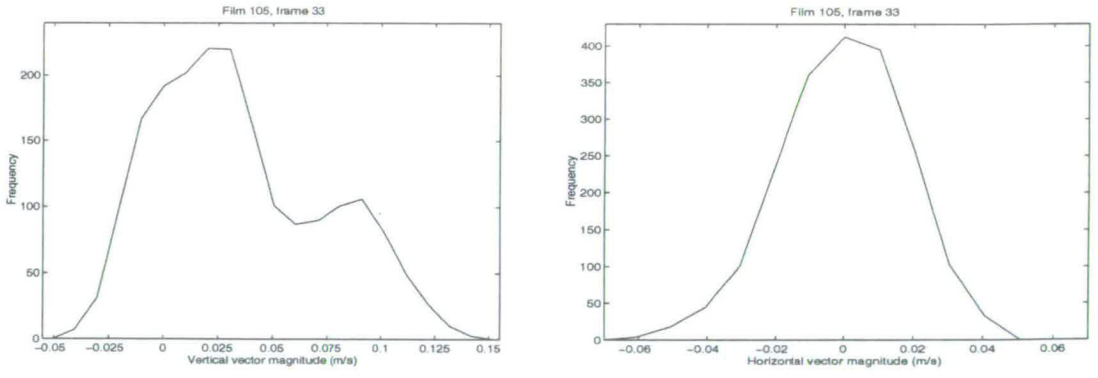


Figure 9.12: Histograms of the distribution of the vertical (left) and horizontal (right) velocity information obtained from frame 33 of film 105 ($U \approx 1.27\,U_{mf}$).

Histograms of the vertical and horizontal velocity components in frame 33 are shown in Figure 9.12. The vertical velocity distribution was bi-modal, with peaks at $0.025\,\mathrm{m\,s^{-1}}$ and $0.09\,\mathrm{m\,s^{-1}}$. This ties in with the vector map from this frame (Figure 9.9): there were a number of stagnant regions (*e.g.* just below the centre of the frame) and a number of rapidly-rising regions (*e.g.* near the left-hand wall). In contrast, the horizontal velocities were approximately normally-distributed, with a mean of $-0.0008\,\mathrm{m\,s^{-1}}$ (*i.e.* almost zero).

An attempt is made in Figure 9.13 to show the extra PIV measurements which can be obtained only if image-shifting is used. The vector map in this figure is based on frame 33, which was shown in full in Figure 9.9. Vectors from this frame were deleted, for the purposes of Figure 9.13, if their magnitude was less than 2 pixels (after subtracting the image-shifting displacements). These vectors would certainly not have been resolved without image-shifting. Thus, this new vector map suggests the results which would have been obtained from the flow recorded in this frame, if the direct PIV recording conditions of film 90 had been used instead (*cf* Figure 9.6). The yield of valid vectors fell from about 93% when image-shifting was used to about 70% when it was 'not', so about 23% of the displacements in this frame were smaller than 2 pixels in magnitude.

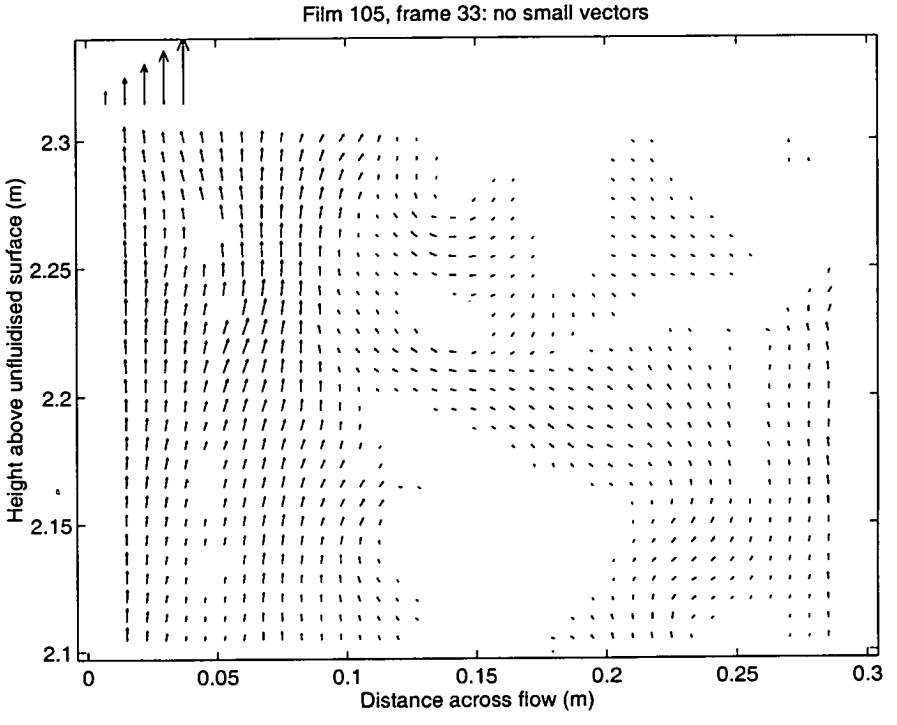


Figure 9.13: A PIV vector map for the flow about 220 cm above a 31 cm bed of particles. Image-shifting was used, then the smallest vectors were deleted, to illustrate the gains which can be made by using image-shifting. $U \approx 1.27\,U_{mf}$.

Some velocity statistics from the low-flowrate run with image-shifting are given in Table 9.5. The mean vertical velocity was approximately $0.0363\,\mathrm{m\,s^{-1}}$, which is considerably less than the superficial velocity. These velocity statistics are from interpolated vector maps, which were calculated using 128 × 128 interrogation areas and unpadded FFTs. Analyses with 64 × 64 interrogation areas and fully-

padded FFTs were also tried; although the velocities were similar, the smaller interrogation areas gave disappointingly low yields (about 80–85%).

Frame	Yield (%)	Mean v_y ($\mathrm{m\,s^{-1}}$)	Max v_y ($\mathrm{m\,s^{-1}}$)	Min v_y ($\mathrm{m\,s^{-1}}$)	Mean v_x ($\mathrm{m\,s^{-1}}$)	Max v_x ($\mathrm{m\,s^{-1}}$)
f105f29	93	0.0452	0.15	−0.030	+0.0052	±0.055
f105f30	94	0.0378	0.155	−0.040	+0.0027	−0.070
f105f31	92	0.0382	0.115	−0.035	−0.0019	−0.090
f105f32	91	0.0383	0.13	−0.035	−0.0013	−0.070
f105f33	93	0.0355	0.15	−0.050	−0.0008	−0.070
f105f34	92	0.0317	0.15	−0.040	−0.0039	−0.070
f105f35	92	0.0290	0.15	−0.035	−0.0013	−0.070
f105f36	91	0.0346	0.18	−0.035	+0.0014	−0.060
mean:	92	0.0363	0.148	−0.038	+0.0001	~ 0.069

Table 9.5: Summary of velocity values for film 105, where the superficial velocity U was $0.0889\,\mathrm{m\,s^{-1}} = 1.27\,U_{mf}$ and image-shifting was used. The bed depth was 31 cm and the mean viewing height was about 220 cm above the unfluidised bed surface. Yields were measured before interpolation.

It was hoped that the seeding-particle slip velocity could be assessed in some way, using data recorded in the first part of film 104. In this run, the bed was fluidised gently in order to drive some seeding particles into the freeboard, then the air was shut off. Using image-shifting, the movement of the seeding particles was photographed as they fell back down towards the bed. Histograms of the velocity components (Figure 9.14) suggest that the most common vertical velocity was about $-0.05\,\mathrm{m\,s^{-1}}$ (*i.e.* downwards), with vertical velocities as low as $-0.09\,\mathrm{m\,s^{-1}}$ and as high as $+0.11\,\mathrm{m\,s^{-1}}$ also being found. The flow seemed to be based on a cloud of particles falling downwards near the centre of the bed, with some quite large upward velocities found near the walls (Figure 9.15).

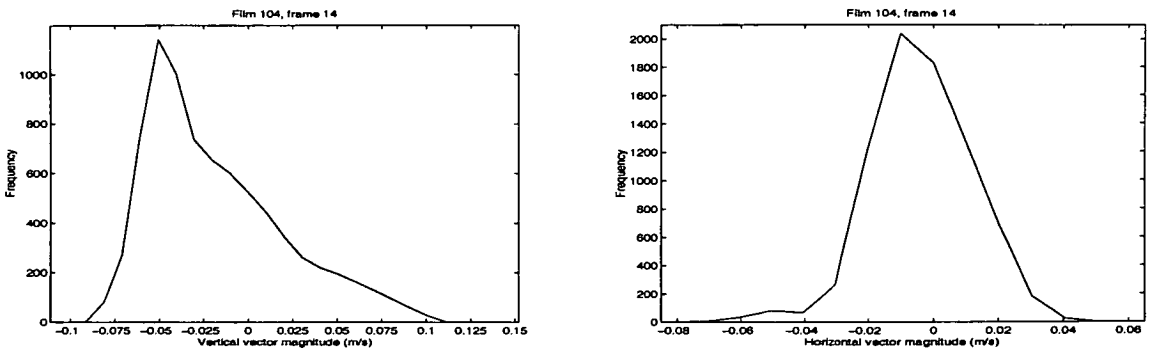


Figure 9.14: Histograms of the vertical (left) and horizontal (right) velocity distributions from film 104, frame 14.

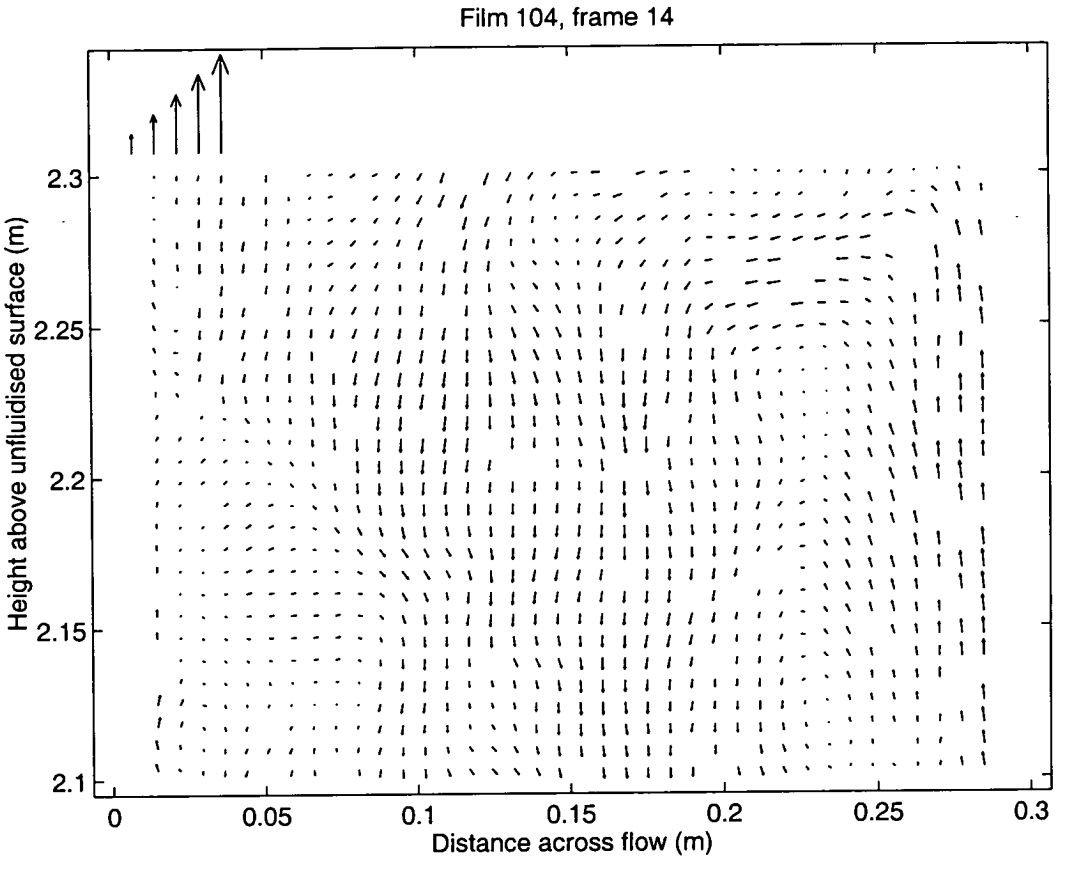


Figure 9.15: PIV vector map for the flow about 220 cm above a 31 cm bed of particles, with image-shifting used.

Applying Stokes' Law, the terminal velocity of the seeding particles is somewhere between $0.0021\,\mathrm{m\,s^{-1}}$ (based on the mean particle diameter of $8\,\mu$m) and $0.0065\,\mathrm{m\,s^{-1}}$ (based on $d_p = 14\,\mu$m, the particle diameter representing a 90% undersize value). The smaller value is more typical of the supplied seeding particles, but the larger particles might remain in the bed for longer if the finer particles are more easily elutriated. Thus, the concentration of 'large' seeding particles in the bed after a period of operation could be somewhat greater than that in a fresh sample. The velocities of airborne particles falling in the freeboard seem to be several times greater than their terminal velocities, which however are based on isolated particles falling in a stagnant medium of infinite extent.

One explanation for this behaviour is that fine particles first start to fall back down in the centre of the freeboard and drag some air down with them. There is no net air flow, however, so this downward air flow in the centre must be balanced

by an upward air flow near the walls. The downward air flow in the centre of the freeboard allows the seeding particles to fall faster than their terminal velocity, whilst the upward air flow near the walls drags some seeding particles back up to greater freeboard heights. Thus, the flow recorded in film 104 resembled a natural convection pattern, rather than uniform particle free-fall.

However, these particles were driven into the freeboard using a relatively-low air flowrate. The air velocities in the freeboard would have been similar to those recorded in film 105 (Figures 9.7–9.9) just before the air was shut off. Afterwards, the seeding particles may simply have continued to follow the down-at-centre, up-at-wall circulation pattern which had *already* been established. A test for this would be to elutriate the seeding particles using $U \sim 2\,U_{mf}$ before switching the air off, to achieve a more uniform initial flow pattern (*cf* Figures 9.1 and 9.3).

Another explanation is that the seeding particles were not uniformly distributed throughout the freeboard. Bubble eruptions near the centre of the bed might have produced a high seeding particle concentration near the axis. Static electricity might have attracted seeding particles to the walls, reducing the local airborne concentration. The exit arrangement gradually reduced the freeboard area from the walls towards the axis, forcing particles inwards at the top of the freeboard.

There would typically have been about 20 g of particles, in a freeboard of volume $0.3\,\mathrm{m} \times 0.3\,\mathrm{m} \times 4\,\mathrm{m} = 0.36\,\mathrm{m}^3$, *i.e.* the particle bulk density was about $0.06\,\mathrm{kg\,m^{-3}}$. If $\rho_f = 1.2\,\mathrm{kg\,m^{-3}}$, then variations in particle concentration can increase the air-particle density by about 5%. The convection could be equivalent to that found in air at 300 K, within vertical pipe walls heated to 315 K. Tests for this could involve varying the initial mass of seeding particles in the freeboard and using smaller seeding particles (*e.g.* 1 μm styrene particles).

So, down-at-centre, up-at-wall patterns definitely arose in the freeboard above this fluidised bed, both with air flow and without air flow. What is not clear is whether the fluidisation conditions caused these to arise (*i.e.* with or without the seeding particles) or whether the particle loading in the freeboard affected the gas flow (which would otherwise have been greatest near the axis).

Engineers should note that the feed rate of pulverised coal particles to a fluidised bed combustor might affect the flow pattern in the freeboard and hence the completeness of combustion. Too many coal particles could encourage a down-at-centre flow pattern, with elutriation taking place via a rapidly-rising gas stream near the wall. On the other hand, a density difference might not actually arise above an industrial coal combustor, as the slow-moving particles near the centre of the freeboard would have plenty of time in which to burn out.

It is believed that the circulation pattern shown in Figure 9.15 is caused by particle descent near the centre of the section and particle-air drag forces, rather than convective effects due to the 12 W laser beam heating up the air near each wall. One check which can be made is to set up the flow pattern with the light sheet blocked off, allowing the apparatus to 'cool down' to room temperature. The result is that the circulation pattern is seen as soon as illumination is restored, suggesting that it exists even when the supposed heat source is removed.

9.1.2 Viewing heights of about 140 cm

The highest flowrate used at the intermediate viewing height was $1010\,\mathrm{l\,min^{-1}}$, which gave a superficial air velocity of $0.187\,\mathrm{m\,s^{-1}}$ ($2.67\,U_{mf}$). Two similar runs were made at this flowrate, with no image-shifting used in either case.

In film 77, the illumination time interval was about 4 ms and the camera was set to 1/30 s, so about 8 images were recorded of each particle. When analysed, the most common image displacement was about 7 pixels, with magnitudes lying in the range 3–14 pixels. These values were rather low and there were probably too many images of each particle. Nevertheless, a number of frames were analysed successfully and the yield of valid vectors was very good (usually over 95%). Some sample vector maps are given in Figure 9.16. Histograms of the vertical and horizontal velocity components from frame 10 are plotted in Figure 9.17, above a table of velocity statistics for this run.

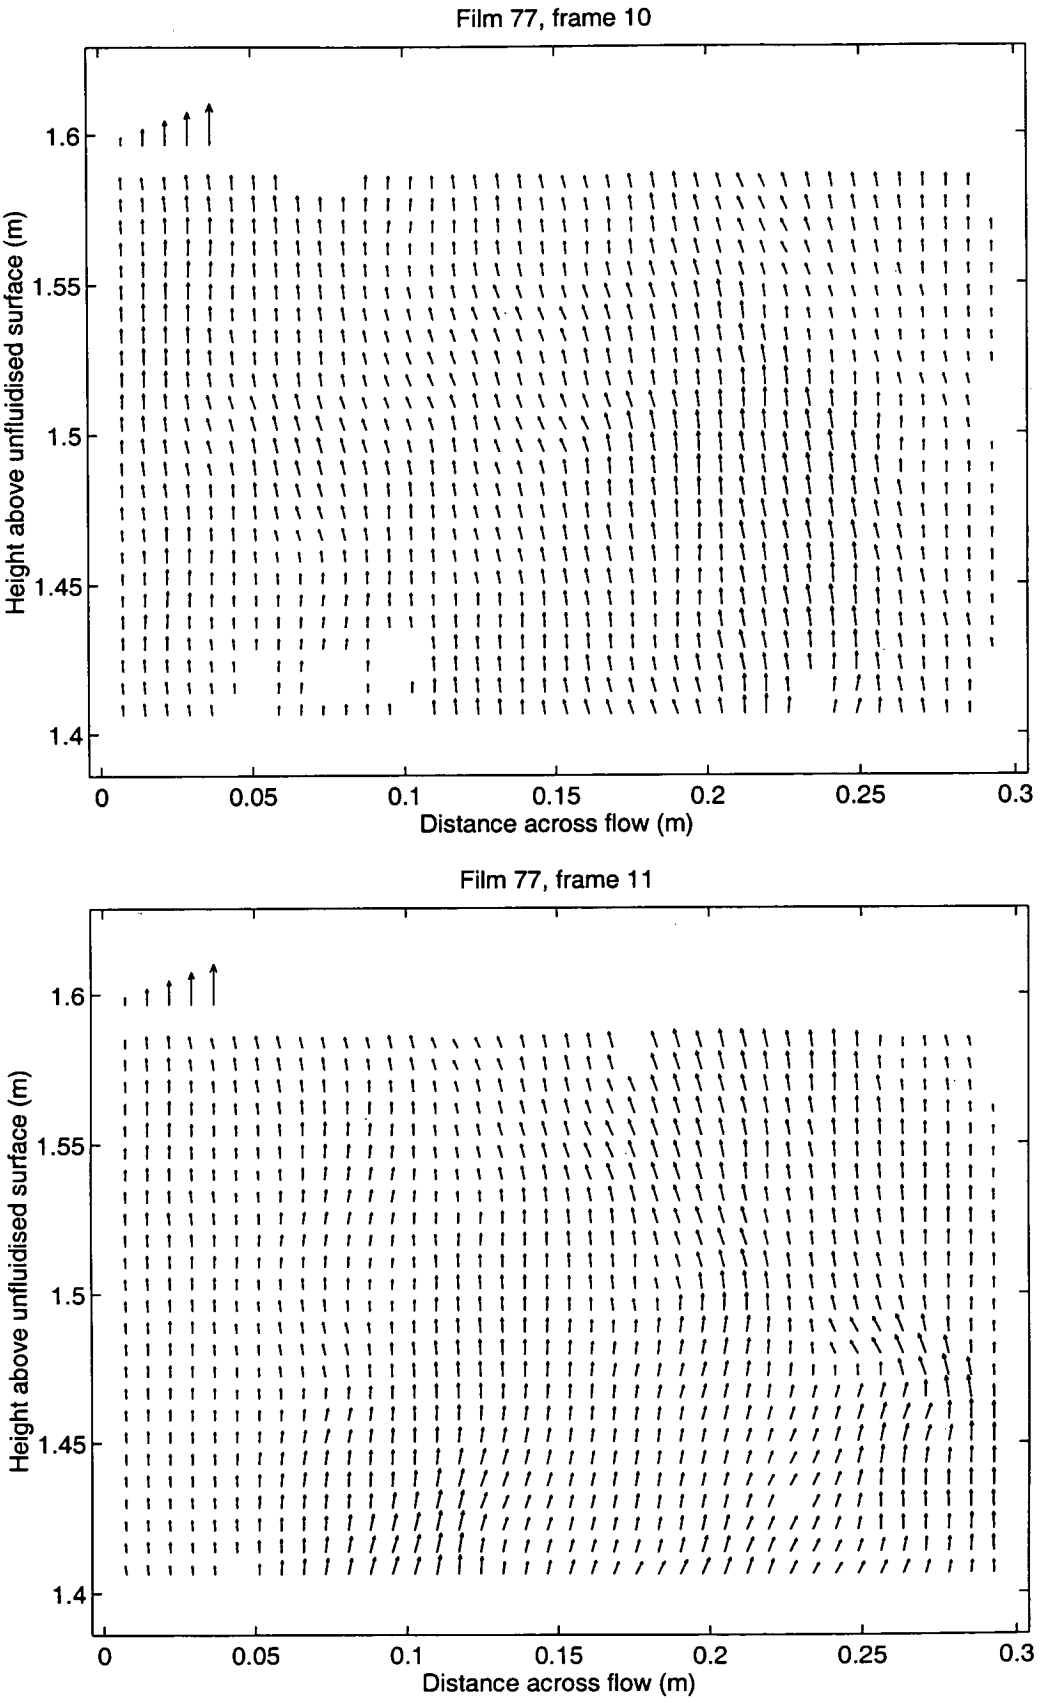
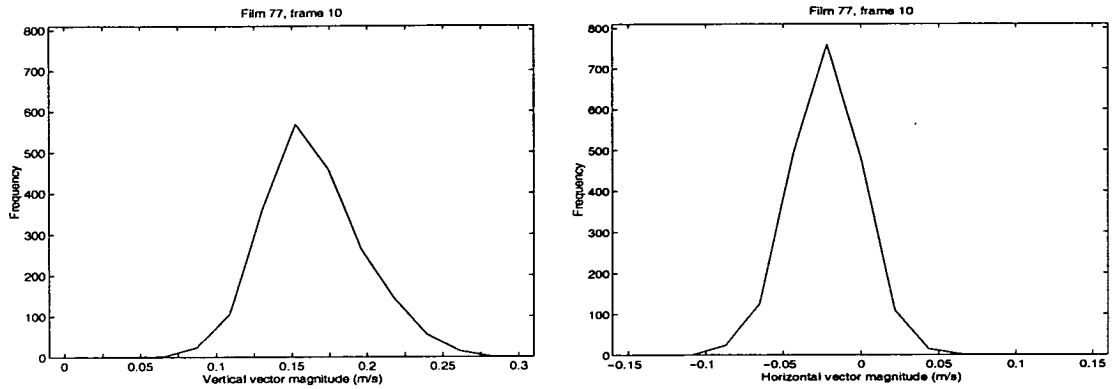


Figure 9.16: PIV vector maps for the flow about 150 cm above a 31 cm bed of particles. $U \approx 2.67\, U_{mf}$.



Frame	Yield (%)	Mean v_y ($\mathrm{m\,s^{-1}}$)	Max v_y ($\mathrm{m\,s^{-1}}$)	Min v_y ($\mathrm{m\,s^{-1}}$)	Mean v_x ($\mathrm{m\,s^{-1}}$)	Max v_x ($\mathrm{m\,s^{-1}}$)
f77f7	99	0.1653	0.27	0.070	+0.0038	+0.105
f77f8	97	0.1648	0.26	0.065	+0.0057	+0.090
f77f9	96	0.1723	0.26	0.065	−0.0108	−0.105
f77f10	97	0.1646	0.28	0.060	−0.0213	−0.105
f77f11	99	0.1608	0.26	0.065	−0.0012	−0.120
f77f12	94	0.1676	0.28	0.065	+0.0214	+0.120
f77f13	89	0.1607	0.28	0.060	+0.0131	+0.105
f77f14	91	0.1428	0.28	0.050	−0.0164	−0.120
f77 mean:	95	0.1624	0.27	0.063	−0.0007	~ 0.11
f78f4	97	0.1322	0.19	0.070	−0.0036	±0.055
f78f5	98	0.1320	0.20	0.070	+0.0088	+0.10
f78f6	99	0.1320	0.24	0.060	+0.0135	±0.08
f78f7	99	0.1553	0.24	0.070	+0.0086	±0.07
f78f8	98	0.1328	0.24	0.070	−0.0015	−0.065
f78f9	99	0.1341	0.24	0.060	−0.0055	−0.065
f78f10	97	0.1648	0.31	0.030	−0.0080	±0.10
f78f11	97	0.1563	0.27	0.070	−0.0103	−0.095
f78f12	93	0.1510	0.25	0.060	−0.0152	+0.12
f78 mean:	97	0.1434	0.24	0.060	−0.0015	~ 0.08
Average:	96	0.1523	0.255	0.061	−0.0011	~ 0.094

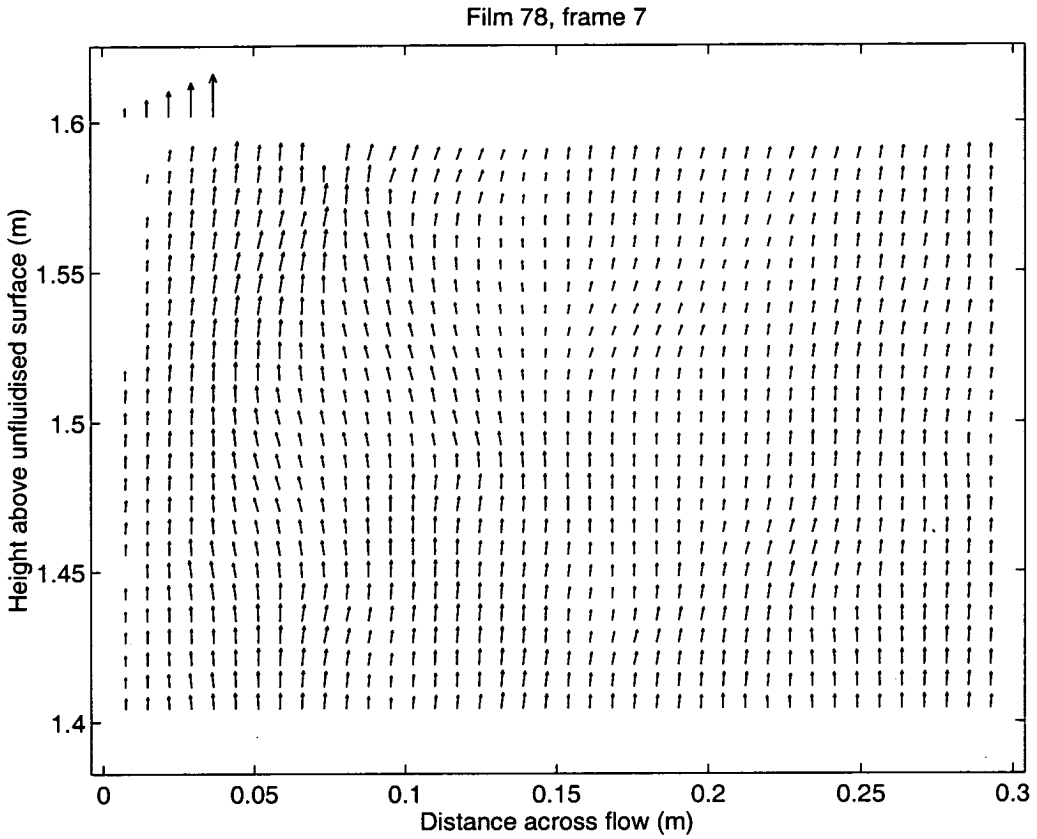
Figure 9.17: Top: Histograms of the distribution of the vertical (left) and horizontal (right) velocity information obtained from frame 10 of film 77. Bottom: Velocity statistics for films 77 and 78. In both runs, $U = 0.187\,\mathrm{m\,s^{-1}} \approx 2.67\,U_{mf}$, the bed depth was 31 cm and the mean viewing height was 150 cm above the bed surface.

The run was repeated in film 78, with an increase in the illumination time interval (to 6 ms) being the only change. There were now about 5 images of each particle. Image displacements of 4–20 pixels were obtained, with the mode value increased to 11 pixels. The yields of valid vectors were again very good (averaging 97%). Velocity statistics for film 78 were tabulated in Figure 9.17. A couple of typical vector maps are plotted in Figure 9.18.

Compared to the superficial air velocity of $0.187\,\mathrm{m\,s^{-1}}$, the section-average values from PIV were $0.162\,\mathrm{m\,s^{-1}}$ for the eight analysed frames from film 77 and $0.143\,\mathrm{m\,s^{-1}}$ for the nine results from film 78. The difference may be due to chance, or it may reflect a systematic over-reading of displacement in film 77. Errors might have arisen if the image separation was below the minimum displacement for accurate resolution using this PIV system.

At an intermediate air flowrate of $640\,\mathrm{l\,min^{-1}}$, the superficial air velocity was $0.119\,\mathrm{m\,s^{-1}}$ ($1.69\,U_{mf}$). The flow was recorded in film 80, using an illumination time interval of 6.5 ms. As the vector maps in Figure 9.19 show, the flow was again found to be upward everywhere and the yields of valid vectors were reasonable (usually over 97%). One slight criticism is that a longer illumination time interval (*e.g.* about 8–10 ms) could have been used, as the magnitude of the most common image displacement in the bitmap file was only about 7 pixels. However, it did not seem worth repeating the run just to make this minor change.

Some velocity statistics are given in Table 9.6. The mean upward velocity found by PIV, averaging across each section, was $0.0924\,\mathrm{m\,s^{-1}}$ (about $0.026\,\mathrm{m\,s^{-1}}$ less than the superficial velocity). The peak upward velocity was rather less than that found at the higher flowrate, being typically about $0.17\,\mathrm{m\,s^{-1}}$ compared to about $0.255\,\mathrm{m\,s^{-1}}$ at the higher flowrate. This ratio of greatest vertical velocities is similar to the appropriate ratio of superficial velocities: $2.67\,U_{mf}/1.69\,U_{mf} = 1.58$; $0.255\,\mathrm{m\,s^{-1}}/0.17\,\mathrm{m\,s^{-1}} = 1.50$. Alternatively, one could say that the ratio Max v_y/U was reasonably constant at this height, at about 1.54.



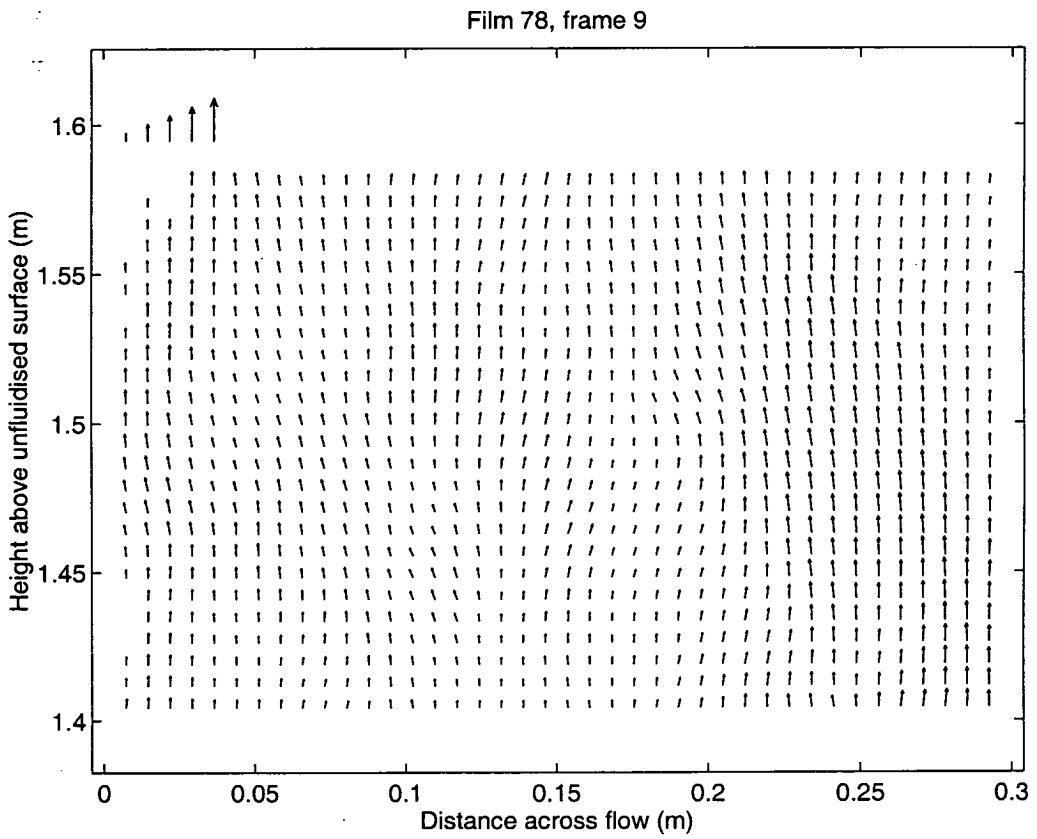


Figure 9.18: PIV vector maps for the flow about 150 cm above a 31 cm bed of particles. $U \approx 2.67\,U_{mf}$.

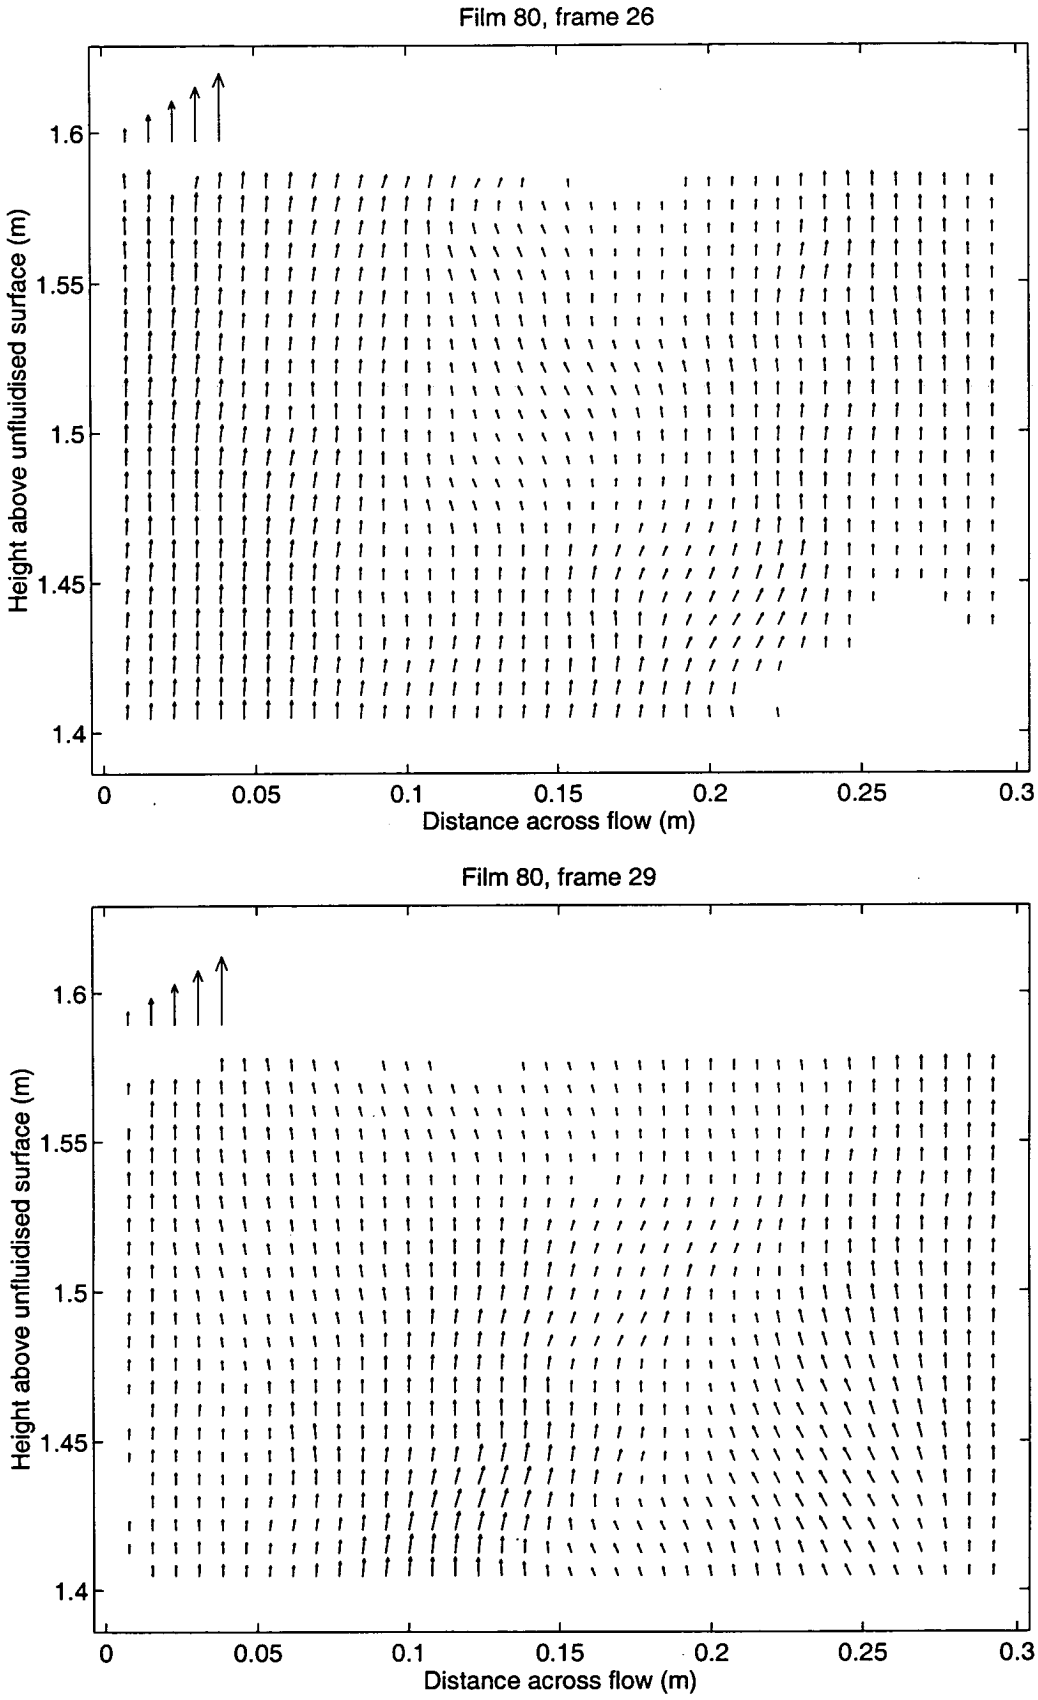


Figure 9.19: PIV vector maps for the flow about 150 cm above a 31 cm bed of particles. $U \approx 1.69\, U_{mf}$.

Frame	Yield (%)	Mean v_y ($m\,s^{-1}$)	Max v_y ($m\,s^{-1}$)	Min v_y ($m\,s^{-1}$)	Mean v_x ($m\,s^{-1}$)	Max v_x ($m\,s^{-1}$)
f80f26	95	0.0950	0.18	0.030	+0.0031	±0.050
f80f27	98	0.0939	0.17	0.030	+0.0021	+0.065
f80f28	98	0.0896	0.15	0.035	+0.0007	±0.050
f80f29	98	0.0875	0.16	0.035	−0.0037	−0.065
f80f30	99	0.0888	0.15	0.030	−0.0013	−0.065
f80f31	99	0.0929	0.16	0.030	−0.0023	±0.050
f80f32	96	0.0844	0.15	0.030	−0.0051	−0.050
f80f33	94	0.0854	0.16	0.025	−0.0050	±0.050
f80f34	98	0.0914	0.18	0.020	−0.0094	−0.075
f80f35	96	0.1045	0.18	0.025	−0.0174	−0.090
f80f36	97	0.0909	0.18	0.025	+0.0003	−0.065
f80f37	98	0.1045	0.18	0.040	−0.0114	−0.065
mean:	97	0.0924	0.17	0.030	−0.0041	∼ 0.062

Table 9.6: Summary of velocity values for film 80, where the superficial velocity was $0.119\,m\,s^{-1}$, or about $1.69\,U_{mf}$. The bed depth was 31 cm and the mean viewing height was about 150 cm above the unfluidised bed surface.

Results were also obtained at a lower air flowrate, with the viewing height reduced by 10 cm to become 130–150 cm above the unfluidised surface. An air flowrate of $480\,l\,min^{-1}$ gave a superficial air velocity of $0.0889\,m\,s^{-1}$ ($1.27\,U_{mf}$). Image-shifting was not used in the first run at this flowrate and viewing height (film 82). The illumination time interval of 12 ms and the longer camera exposure time of '1/15' of a second meant that about 5 images of each particle were recorded.

The flow at this air velocity of $0.0889\,m\,s^{-1}$ was more suitable for direct PIV in film 82 than it was at the greater freeboard height (2.2 m in film 90). Regions of unmeasurably low velocity were still present, but had a much reduced area. Almost all the valid vectors were upward. The image displacements were rather low, with the mode lying between 7 and 9 pixels. A longer illumination time interval was suggested, but this could not have been achieved using this apparatus. Nevertheless, results were reasonable even without image-shifting, *e.g.* using 64 × 64 interrogation areas, the yields of valid vectors were between 68 and 87% (compared to 50–75% at the highest viewing position). This film has not been re-analysed using 128 × 128 interrogation areas, because image-shifting gave much better results.

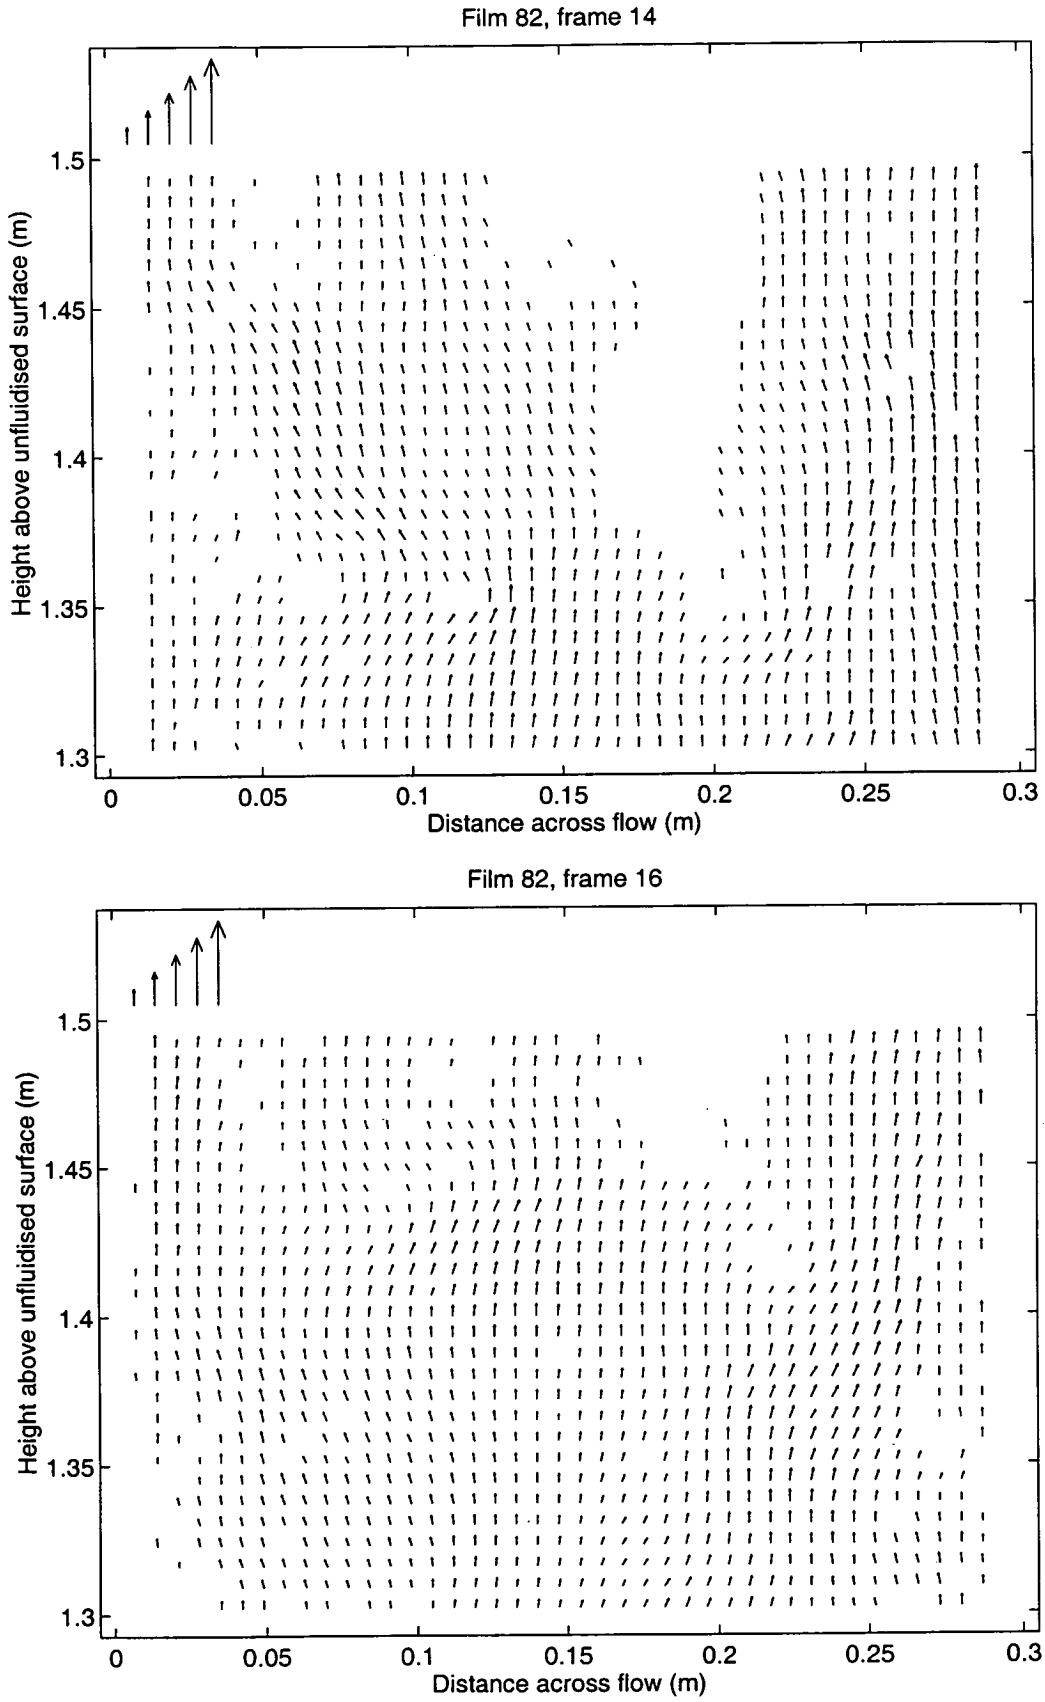


Figure 9.20: Direct PIV vector maps for the flow about 140 cm above a 31 cm bed of particles. $U \approx 1.27\,U_{mf}$.

As the two vector maps in Figure 9.20 show, much of the variation in yields in film 82 can be attributed to the size of the slowest-moving regions of the flow. In frame 14, the yield was 79% (a typical value for this film) with most of the missing vectors lying in the slow-moving region found about 0.15–0.20 m across the flow (Figure 9.20, top). Frame 13 was slightly worse, as this slow-moving region covered the full height of the area photographed, reducing the yield to 74%. By the time frame 16 was taken, however, this region had moved up about 0.1 m and so covered considerably less of the area photographed; the yield consequently increased to 87%. An inference which can be drawn here is that the missing low-magnitude vectors are upward, because the regions of unresolved vectors could be seen to drift upwards from frame to frame. It appeared that direct PIV could be adequate for this flow, given a long enough illumination time interval.

Keeping the same air velocity of 0.0889 m s^{-1} but moving slightly further down to a viewing height of 120–140 cm above the bed surface, the run was repeated using image-shifting (film 133). The illumination time interval was reduced to 10.4 ms, which gave almost exactly three images of each particle. Four images were found in a small rectangular region on some photographs, however. An examination of the photographic negatives allowed the height of this four-image region in the object plane to be determined and then compared to the 0.5 m height of the light sheet. The actual exposure time was thus estimated to be $3.01 \times 10.4\,\text{ms} = 31.3\,\text{ms}$, compared to the nominal value of 1/30 s.

The analysis calculations used 128 × 128 pixel interrogation areas, with no zero-padding. The image-shifting displacement was 20.0 pixels, but it was shown in Section 5.7.7 that wrap-around errors would not arise here, because only three images were recorded of each particle.

The resulting arrays of vectors had 28 rows and 41 columns. As was often found with image-shifting in this work, the outermost two columns of each vector map were spoilt by the inclusion of images of the bright glass-air interface at either wall and yielded no valid vectors. Typical yields were 94–95%, which meant that almost all of the other vectors were resolved correctly. Sample vector maps are given in Figure 9.21. All vectors in this run of frames were upward.

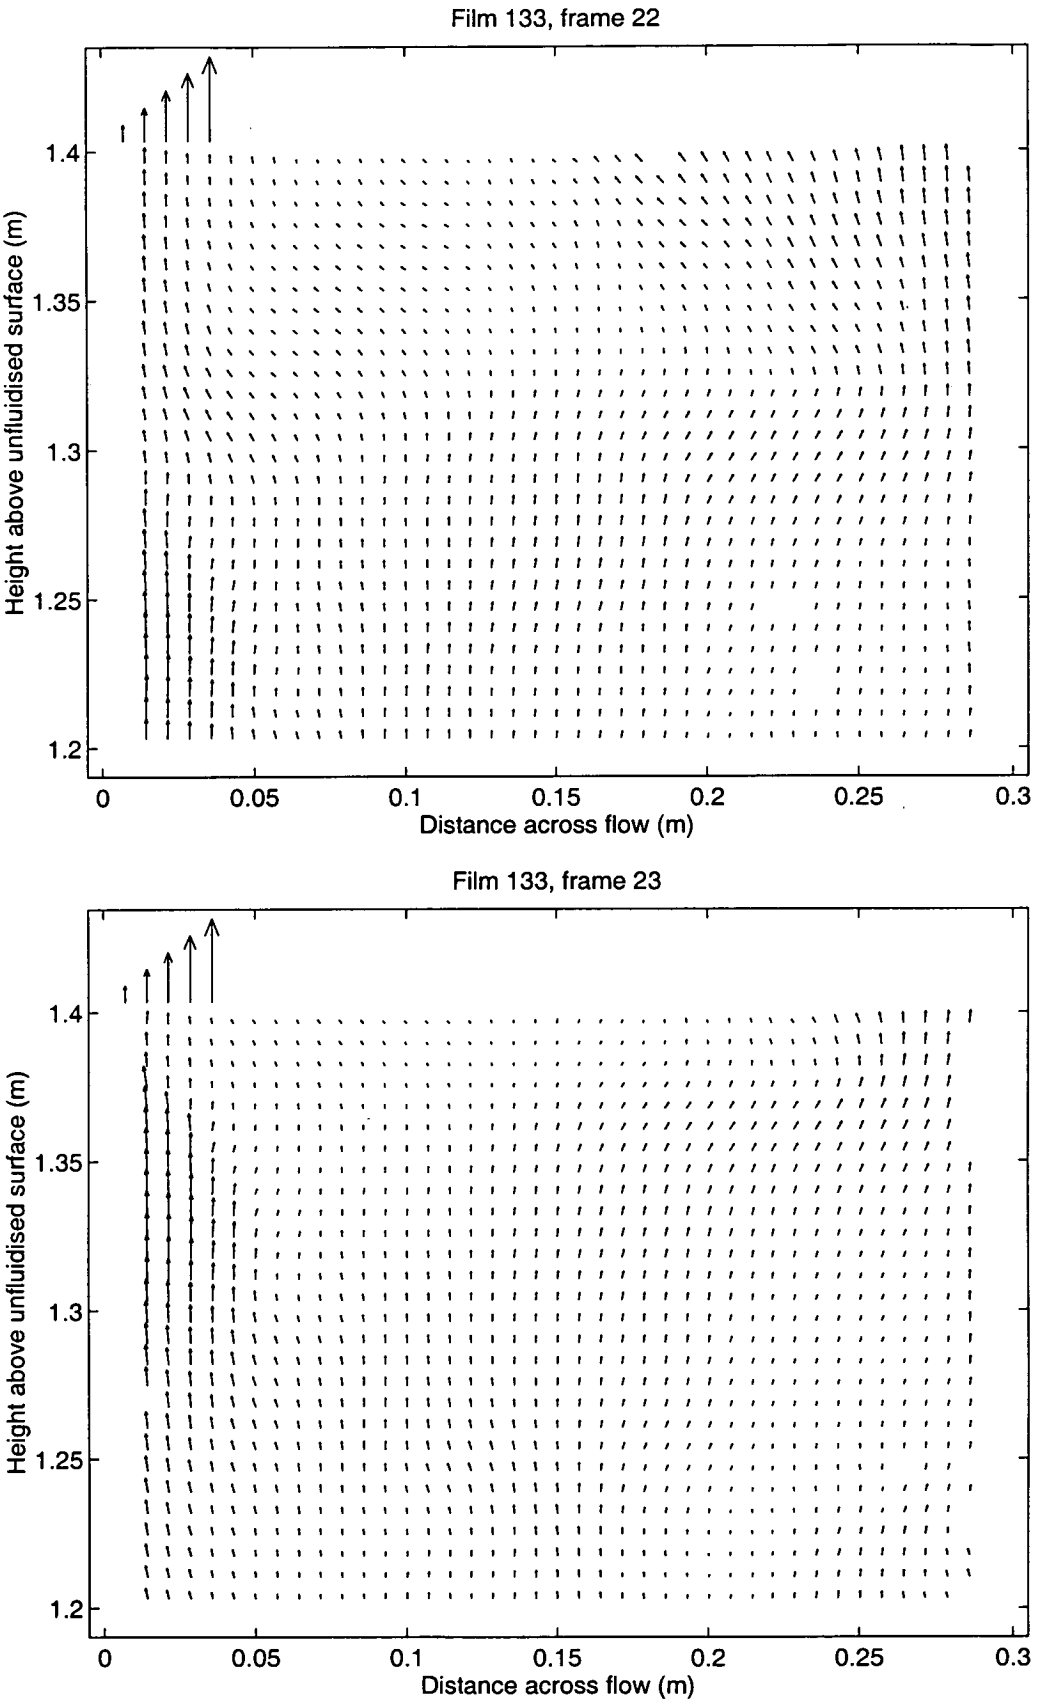


Figure 9.21: PIV vector maps for the flow about 130 cm above a 31 cm bed of particles, after subtracting the image-shifting velocity. $U \approx 1.27\,U_{mf}$.

Some velocity statistics from the data obtained using film 133 are presented in Table 9.7. The mean vertical velocity was found to be $0.0445\,\mathrm{m\,s^{-1}}$ and the greatest vertical velocity was about three times larger than this ($0.139\,\mathrm{m\,s^{-1}}$). The section-average vertical velocities quoted in this table underestimate the flow in these experiments, however, because the upward velocities were greatest near the wall. There were no vectors in the outermost two columns of each vector map, so velocities in these columns could not be included in the means calculated in Table 9.7. If 39 columns of vectors with a mean vertical velocity of $0.0445\,\mathrm{m\,s^{-1}}$ are combined with two additional columns with a somewhat greater upward velocity of about $0.10\,\mathrm{m\,s^{-1}}$, say, the result will be an increased section-average vertical velocity of about $0.047\,\mathrm{m\,s^{-1}}$ (still well below the superficial velocity, however).

Frame	Yield (%)	Mean v_y ($\mathrm{m\,s^{-1}}$)	Max v_y ($\mathrm{m\,s^{-1}}$)	Min v_y ($\mathrm{m\,s^{-1}}$)	Mean v_x ($\mathrm{m\,s^{-1}}$)	Max v_x ($\mathrm{m\,s^{-1}}$)
f133f17	95	0.0450	0.135	0.010	−0.0053	−0.033
f133f18	94	0.0420	0.110	0.005	−0.0013	−0.030
f133f19	92	0.0425	0.135	0.005	+0.0008	+0.033
f133f20	94	0.0458	0.140	0.000	−0.0033	−0.040
f133f21	94	0.0439	0.135	0.010	−0.0068	−0.045
f133f22	95	0.0451	0.140	0.000	−0.0037	−0.045
f133f23	94	0.0435	0.160	0.000	+0.0019	+0.033
f133f24	94	0.0399	0.135	0.000	+0.0029	±0.033
f133f25	95	0.0446	0.135	0.000	+0.0051	+0.040
f133f26	94	0.0501	0.165	0.000	+0.0057	+0.040
f133f27	94	0.0472	0.135	0.005	+0.0010	+0.040
mean:	94	0.0445	0.139	0.003	−0.0003	∼ 0.037

Table 9.7: Summary of velocities for film 133, where the superficial velocity U was $0.0889\,\mathrm{m\,s^{-1}}$ (about $1.27\,U_{mf}$). The bed depth was 31 cm and the mean viewing height was about 130 cm above the bed surface. Image-shifting was used.

It was also possible to take a vector map from this sequence and delete all vectors smaller than 3 pixels. The resulting array of vectors then included some apparently-unresolved low-velocity regions, which could not have been measured successfully using direct PIV. In Figure 9.22, it can be seen that the regions where no vectors were resolved resemble those found under similar conditions when image-shifting was not used (*i.e.* film 82; Figure 9.20). Image-shifting thus allowed velocity vectors to be determined throughout the flow and also confirmed that the missing vectors in film 82 were upward and of low velocity.

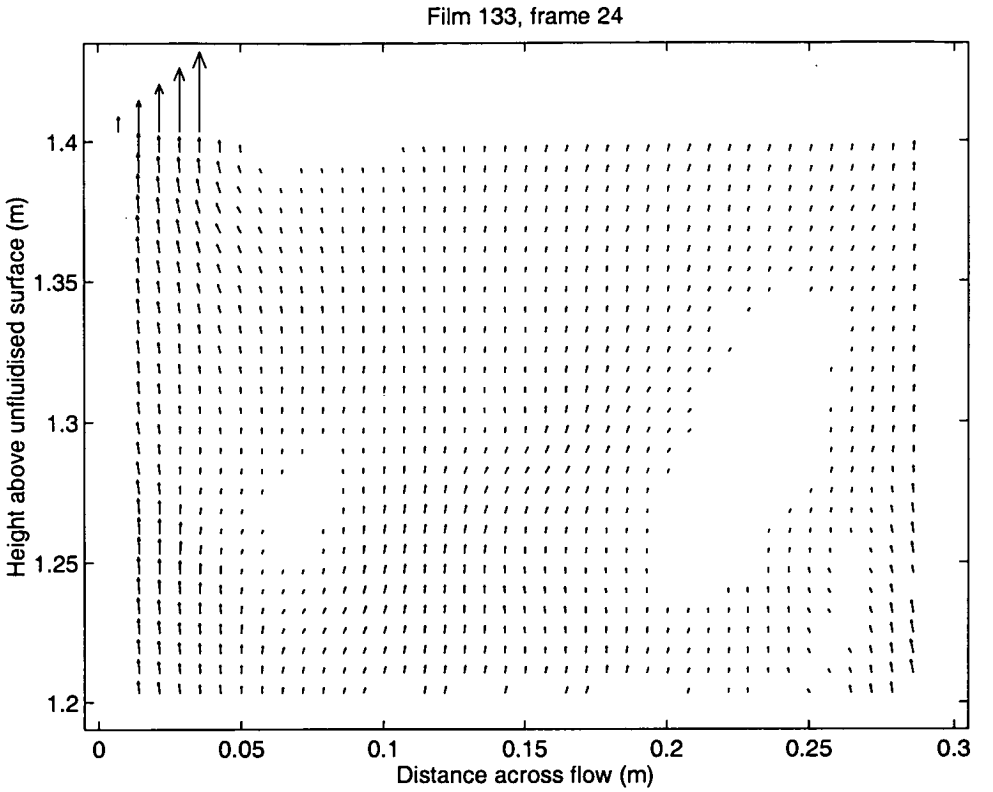


Figure 9.22: A PIV vector map for the flow about 130 cm above a 31 cm bed: film 133 frame 24, but with the smallest vectors deleted. $U \approx 1.27\,U_{mf}$.

9.1.3 Viewing heights of about 40 cm

The flow immediately above a bubbling fluidised bed is extremely turbulent. As a result, the lowest viewing height selected was about 30–50 cm above the level of the unfluidised bed surface. Heights closer to the bed surface were not examined, because particles were ejected into the freeboard immediately above the bed surface in great quantities, swamping the photographic records. In addition, the bed surface rises up in front of each bubble immediately before it bursts, causing further problems if the bed surface encroaches into the field of view. The viewing height chosen was well below the TDH, but was high enough above the bed surface to avoid the worst effects of the bubbling action.

Results without image-shifting were mixed. Reasonable yields of valid vectors were obtained in a number of photographs from film 71, which was taken at a flow velocity $U = 1.55\,U_{mf}$. A couple of vector maps are given in Figure 9.23.

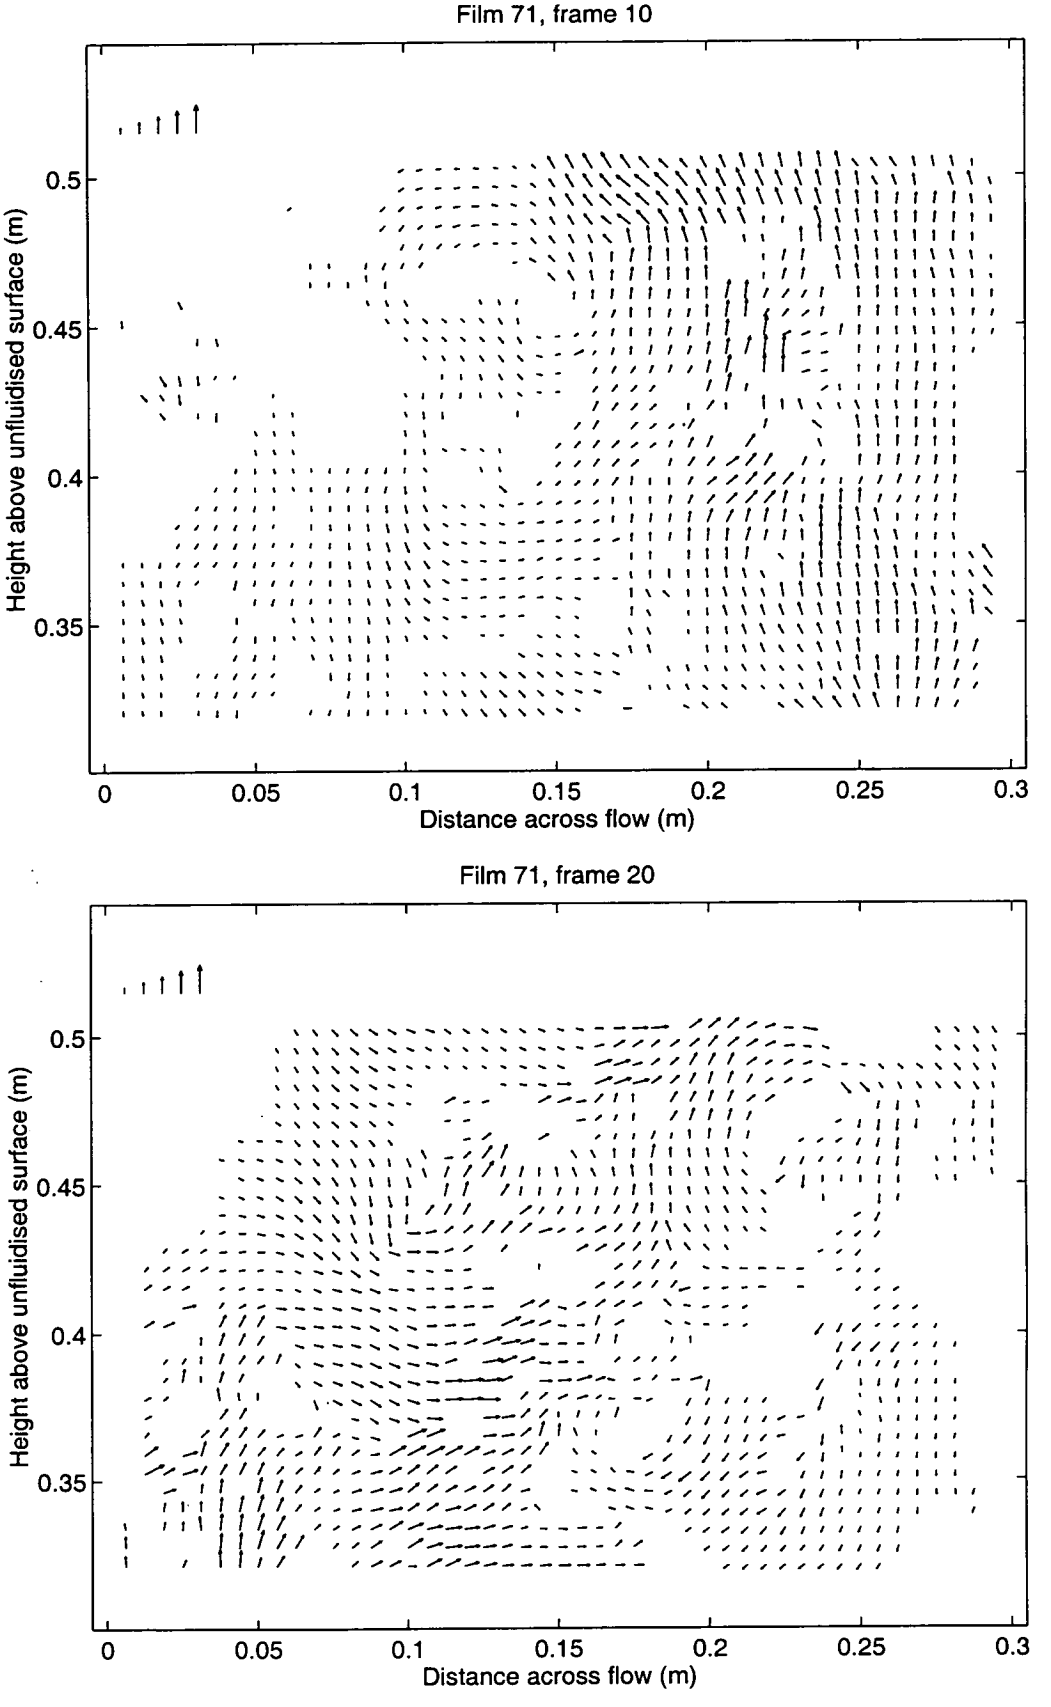


Figure 9.23: PIV vector maps for the flow about 40 cm above a 31 cm bed of particles. $U \approx 1.55\, U_{mf}$.

A problem was that it was not always easy to align the vectors, particularly when there were vector-free regions in between a number of islands of vectors, in each of which the yields were good.

Significant regions of downflow were identified, often covering half of the area photographed (*e.g.* in frames 8–11 of film 71). In this case, it is not even clear which half of the flow field contains upward velocities and which half contains downward velocities: after aligning each vector by orienting it relative to (and together with) its nearest neighbours, there is still the possibility that the direction of *every* vector is out by exactly 180°. The nature of the flow under these conditions was similar to that observed later, when the flow a similar distance above a 60 cm bed was photographed using image-shifting (film 140). Thus, although the features observed were unexpected (*e.g.* half of the vectors were downward), this is not simply a result of alignment difficulties when using direct PIV.

Frame	Yield (%)	Mean v_y ($m\,s^{-1}$)	Max v_y ($m\,s^{-1}$)	Min v_y ($m\,s^{-1}$)	Mean v_x ($m\,s^{-1}$)	Max v_x ($m\,s^{-1}$)
f71f9	54	0.1135	0.36	−0.15	+0.0128	+0.28
f71f10	71	0.0472	0.35	−0.24	+0.0052	−0.27
f71f11	69	0.0172	0.38	−0.25	−0.0409	−0.27
f71f12	70	0.0237	0.37	−0.28	−0.0123	−0.29
f71f13	77	0.1222	0.58	−0.35	−0.0914	−0.43
f71f14	69	0.0993	0.47	−0.25	−0.0104	−0.22
f71f15	69	0.0669	0.44	−0.22	+0.0094	−0.28
f71f16	63	0.0761	0.40	−0.22	+0.0058	−0.28
f71f17	63	0.0783	0.41	−0.21	−0.0028	−0.28
f71f18	60	0.0421	0.32	−0.19	+0.0525	+0.25
f71f19	70	0.0652	0.48	−0.25	−0.0078	−0.32
f71f20	71	0.0103	0.34	−0.26	+0.0610	+0.35
f71f21	65	0.0530	0.35	−0.21	−0.0308	+0.31
f71f22	61	0.0347	0.41	−0.24	+0.0582	+0.45
mean:	67	0.0607	0.40	−0.24	+0.0006	~ 0.31

Table 9.8: Summary of velocity values for film 71, where the superficial velocity was $0.108\,m\,s^{-1}$, or about $1.55\,U_{mf}$. The bed depth was 31 cm and the mean viewing height was about 40 cm above the unfluidised bed surface.

The velocity statistics in Table 9.8 show just how turbulent this flow is. The greatest upward vertical velocity was typically about $0.40\,m\,s^{-1}$ (about $4\,U$), but the lowest vertical velocity was usually about $-0.24\,m\,s^{-1}$ (*i.e.* very strongly

downward). Some large horizontal velocities were obtained in the various frames, with magnitudes occasionally reaching as high as $0.45\,\mathrm{m\,s^{-1}}$. The mean horizontal velocities, averaging all the vectors in each frame (after interpolation), ranged from $0.091\,\mathrm{m\,s^{-1}}$ to the left (in frame 13) to $0.061\,\mathrm{m\,s^{-1}}$ to the right (in frame 20).

9.2 31 cm bed depth, quarter of way in

Comparable results were also taken at a viewing position 7.5 cm beyond the front wall, *i.e.* one quarter of the way from the front wall of a test section to the back wall, remaining parallel to both of these walls. This position was halfway between the front wall and the central viewing position. Table 9.9 lists the main runs.

Bed depth (cm)	Viewing height (cm)	U/U_{mf}	Film number	Image-shifting used?
31	220	1.03	92	No
31	220	1.27	102	Yes
31	220	1.53	93	No
31	220	1.93	91	No
31	130	1.16	132	Yes
31	140	1.79	84	No
31	140	2.25	85	No
31	42	2.14	70	No

Table 9.9: A list of films taken using PIV to record the flow in a section 7.5 cm into the freeboard space above a 31 cm fluidised bed.

9.2.1 Viewing heights of about 220 cm

Starting at the top of the freeboard, the low flowrate results from film 92 (without image-shifting) were inferior to those from film 102 (with image-shifting). Large regions of unresolved vectors were found, wherever the flow velocity fell below the

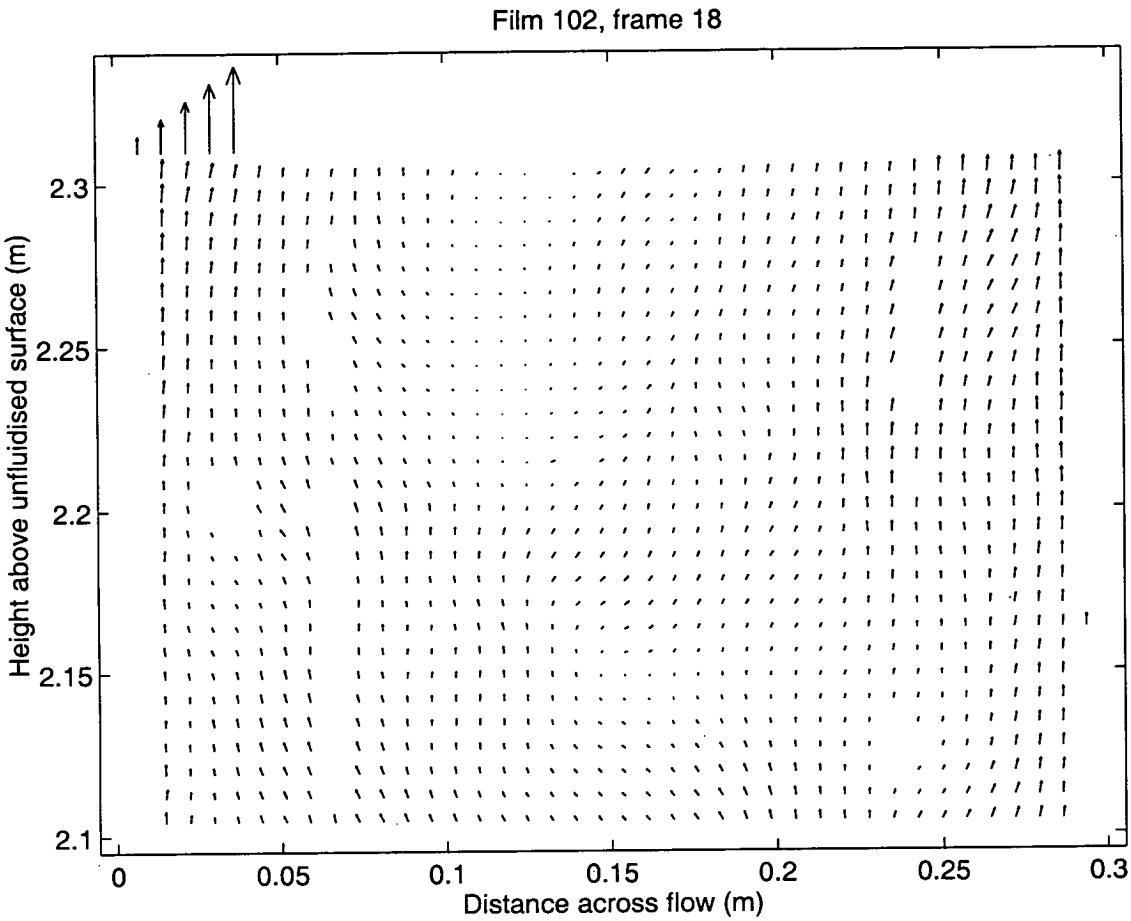
minimum resolvable value. It was also difficult to determine the correct orientation of each vector. Nevertheless, the few frames which were analysed suggested that the velocity was lowest in the centre of the bed and largest near the walls.

These conclusions were confirmed by the image-shifted results of film 102. In this run, the air flowrate was about $480\,\mathrm{l\,min^{-1}}$, so the superficial velocity U was about $0.0889\,\mathrm{m\,s^{-1}}$ ($1.27\,U_{mf}$). A sample vector map and some velocity statistics are given in Figure 9.24. Some zero-velocity vectors and some downward vectors were found in each frame. The results were similar to those obtained in the centre of the bed using a comparable air flowrate (film 105).

Film 93 recorded the results obtained using a slightly higher air flowrate of $585\,\mathrm{l\,min^{-1}}$ ($U = 0.1083\,\mathrm{m\,s^{-1}} = 1.55\,U_{mf}$). As the sample vector map from frame 15 and the velocity statistics show, the velocities measured by PIV were upward everywhere (Figure 9.25). Not all vector maps were as wavy as frame 15, however. Much of the variation in the yields of valid vectors can be attributed to the low-velocity regions which were occasionally encountered, *e.g.* about 0.24 m across the upper half of frame 15. The corresponding run in the centre of the bed (film 88) used a slightly larger superficial air velocity of $U = 0.118\,\mathrm{m\,s^{-1}}$, which appeared to be enough of an increase to avoid these low-velocity regions.

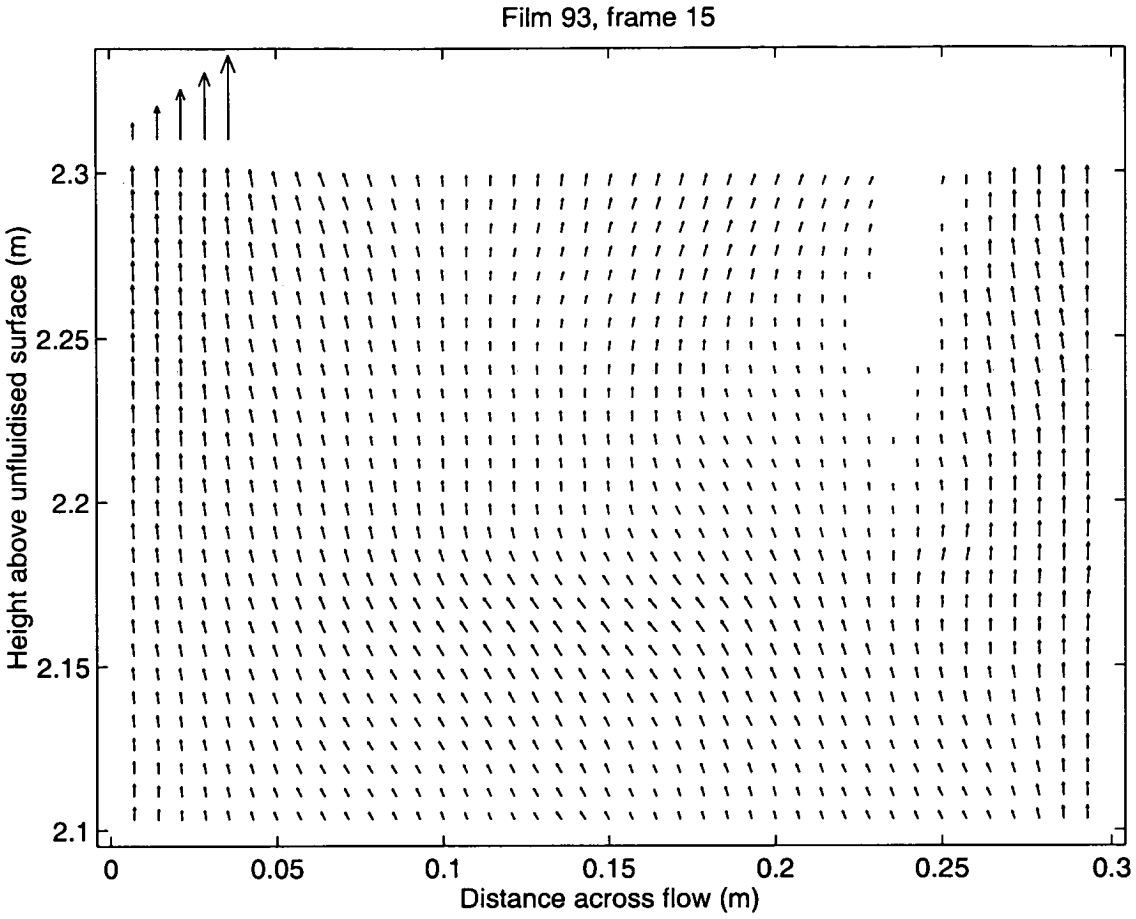
The air flowrate was increased to $735\,\mathrm{l\,min^{-1}}$ ($U = 0.136\,\mathrm{m\,s^{-1}} \approx 1.94\,U_{mf}$) for film 91. A vector map from the velocity data obtained from frame 13 is shown in Figure 9.26, along with some velocity statistics. The seeding particle velocities became much more uniform in this region of the freeboard at high air flowrates.

The vector maps from flows one quarter of the way into the freeboard at superficial velocities $U = 1.55\,U_{mf}$ (film 93, Figure 9.25) and $U = 1.94\,U_{mf}$ (film 91, Figure 9.26) were similar to those found in the centre of the freeboard, at the same viewing height and bed depth (*cf* Figures 9.1 and 9.3). In all cases, only upward velocities were found.



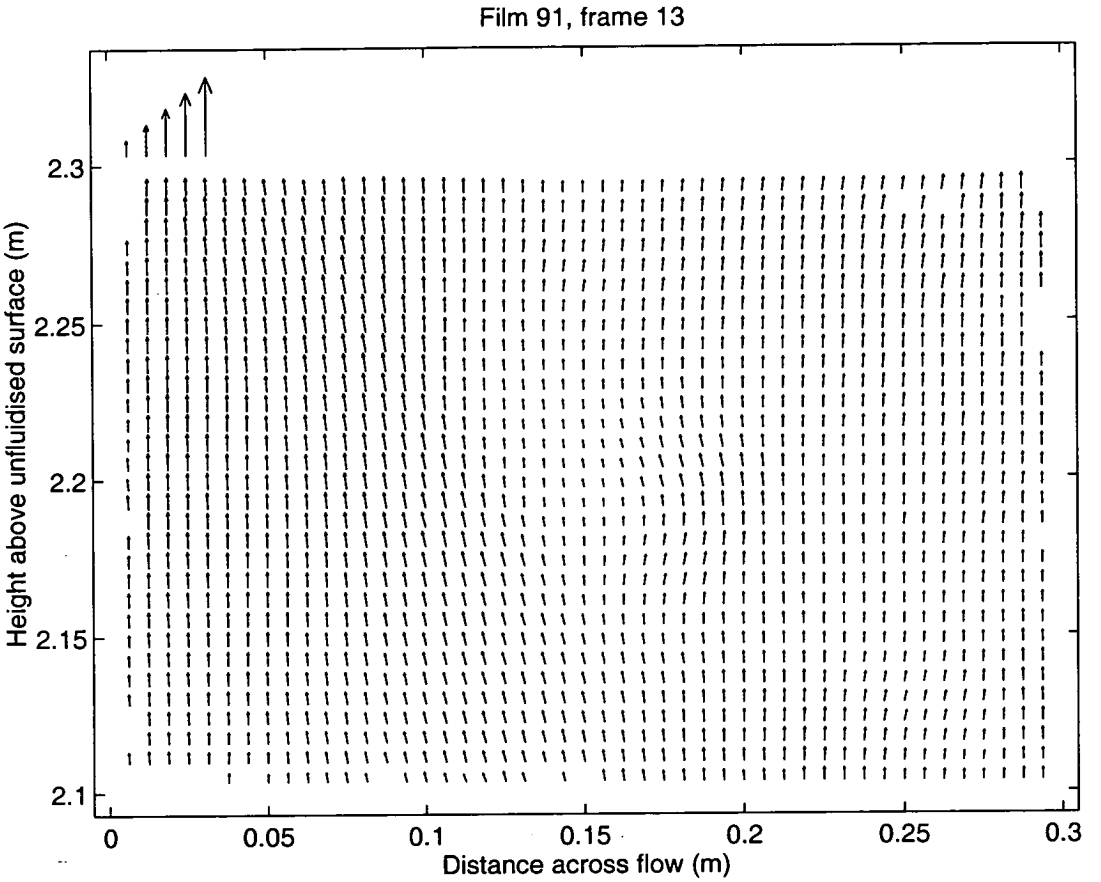
Frame	Yield (%)	Mean v_y ($\mathrm{m\,s^{-1}}$)	Max v_y ($\mathrm{m\,s^{-1}}$)	Min v_y ($\mathrm{m\,s^{-1}}$)	Mean v_x ($\mathrm{m\,s^{-1}}$)	Max v_x ($\mathrm{m\,s^{-1}}$)
f102f12	89	0.0353	0.12	−0.015	+0.0021	+0.050
f102f13	92	0.0372	0.12	−0.015	−0.0018	+0.045
f102f14	93	0.0335	0.11	−0.020	−0.0042	−0.050
f102f15	86	0.0441	0.15	−0.020	−0.0040	−0.050
f102f16	93	0.0496	0.16	−0.030	−0.0060	−0.050
f102f17	92	0.0515	0.15	−0.030	−0.0015	−0.040
f102f18	92	0.0401	0.13	−0.015	+0.0011	±0.040
f102f19	90	0.0402	0.15	−0.015	+0.0050	+0.055
f102f20	90	0.0413	0.16	−0.015	+0.0057	+0.050
mean:	91	0.0607	0.14	−0.019	+0.0004	∼ 0.048

Figure 9.24: Results for film 102, where the superficial velocity was $0.0889\,\mathrm{m\,s^{-1}}$, or about $1.27\,U_{mf}$. The bed depth was 31 cm and the mean viewing height was about 220 cm above the unfluidised bed surface, one quarter of the way into the test section. Image-shifting was used. Top: the vector map from frame 18 of this film; Bottom: the velocity statistics.



Frame	Yield (%)	Mean v_y ($\mathrm{m\,s^{-1}}$)	Max v_y ($\mathrm{m\,s^{-1}}$)	Min v_y ($\mathrm{m\,s^{-1}}$)	Mean v_x ($\mathrm{m\,s^{-1}}$)	Max v_x ($\mathrm{m\,s^{-1}}$)
f93f13	97	0.0732	0.150	0.025	−0.0040	−0.055
f93f14	98	0.0680	0.145	0.020	−0.0062	−0.055
f93f15	98	0.0650	0.130	0.020	−0.0118	−0.055
f93f16	99	0.0663	0.125	0.020	−0.0134	−0.055
f93f17	100	0.0657	0.150	0.020	−0.0087	−0.055
f93f18	100	0.0817	0.165	0.025	−0.0018	+0.055
f93f19	97	0.0772	0.165	0.010	+0.0053	+0.060
f93f20	100	0.0933	0.140	0.035	+0.0051	+0.055
f93f21	98	0.0946	0.140	0.040	−0.0016	−0.045
f93f22	96	0.0927	0.145	0.040	−0.0040	−0.045
f93f23	94	0.0889	0.150	0.040	−0.0066	−0.045
f93f24	100	0.0821	0.120	0.040	−0.0089	−0.050
mean:	98	0.0791	0.144	0.028	−0.0047	~ 0.053

Figure 9.25: Results from film 93, where $U = 0.108\,\mathrm{m\,s^{-1}} \approx 1.55\,U_{mf}$. The bed depth was 31 cm and the mean viewing height was about 220 cm above the unfluidised bed surface, one quarter of the way into the freeboard. Top: Vector map of the velocity information obtained from frame 15; Bottom: velocity statistics.



Frame	Yield (%)	Mean v_y ($\mathrm{m\,s^{-1}}$)	Max v_y ($\mathrm{m\,s^{-1}}$)	Min v_y ($\mathrm{m\,s^{-1}}$)	Mean v_x ($\mathrm{m\,s^{-1}}$)	Max v_x ($\mathrm{m\,s^{-1}}$)
f91f3	97	0.0895	0.130	0.035	+0.0049	+0.040
f91f4	97	0.0846	0.160	0.040	+0.0052	+0.040
f91f5	99	0.0959	0.160	0.050	−0.0028	−0.055
f91f6	95	0.0806	0.145	0.025	−0.0090	−0.050
f91f7	93	0.0680	0.125	0.025	−0.0078	−0.050
f91f8	96	0.0872	0.145	0.045	+0.0031	+0.065
f91f9	95	0.0852	0.130	0.025	+0.0097	+0.050
f91f10	94	0.0901	0.145	0.035	+0.0098	+0.045
f91f11	97	0.0951	0.145	0.040	+0.0050	+0.040
f91f12	97	0.0882	0.130	0.045	+0.0010	−0.035
f91f13	98	0.0855	0.130	0.045	−0.0034	−0.040
f91f14	98	0.0852	0.130	0.050	−0.0048	−0.040
f91f15	94	0.0908	0.160	0.050	−0.0039	−0.035
f91f16	97	0.0974	0.160	0.060	−0.0021	−0.035
mean:	96	0.0874	0.143	0.041	+0.0004	0.044

Figure 9.26: Results from film 91, where the superficial velocity was $0.136\,\mathrm{m\,s^{-1}}$ ($1.94\,U_{mf}$). The bed depth was 31 cm. The mean viewing height was about 220 cm above the bed surface, one quarter of the way into the freeboard. Image-shifting was not used. Top: vector map for frame 13; Bottom: some velocity statistics.

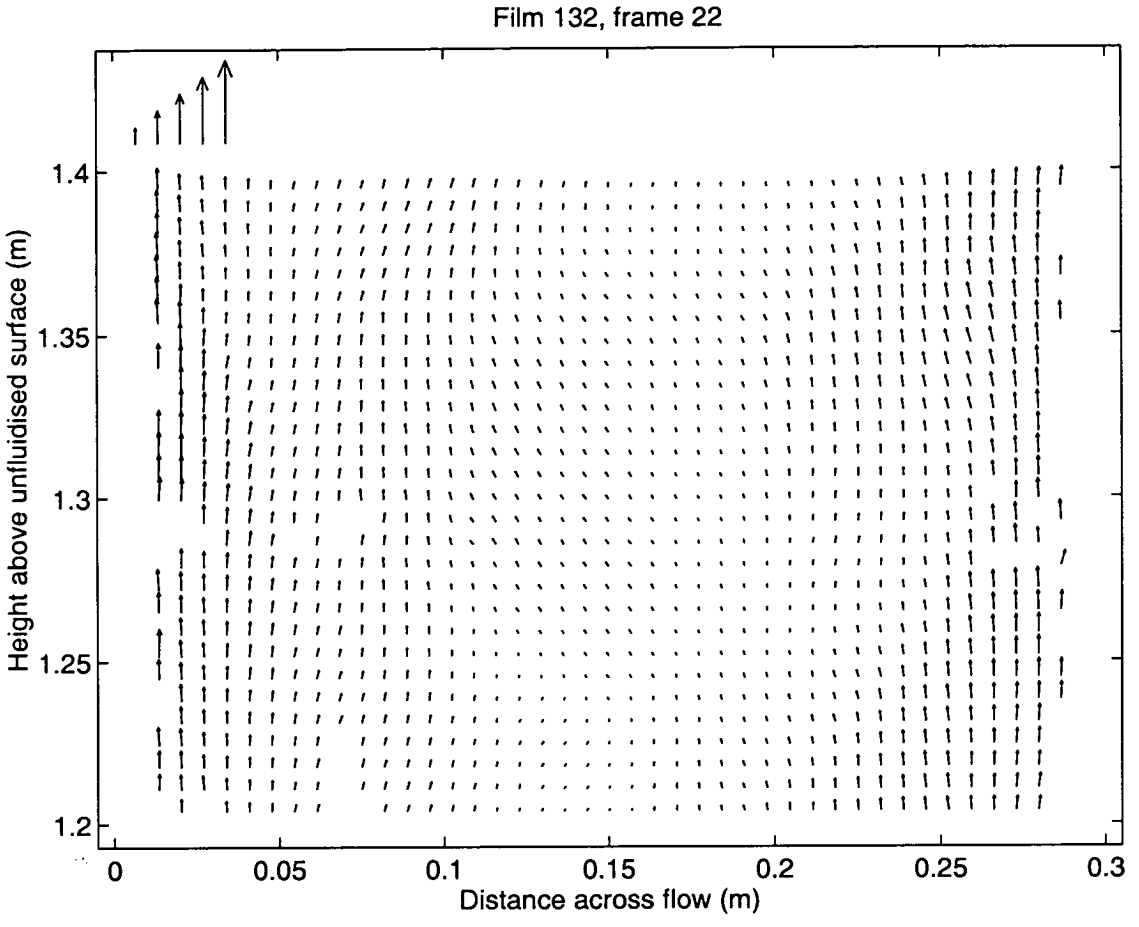
Comparing the runs when image-shifting was not used, the yield of valid vectors was significantly better in film 93 ($U = 1.94\,U_{mf}$) than in film 91 ($U = 1.55\,U_{mf}$), with both giving much better yields than film 92 ($U = 1.06\,U_{mf}$). It appears that image-shifting is not really needed at velocities greater than about $1.6\,U_{mf}$ at this height and viewing position, but becomes more and more necessary if the superficial velocity drops below this value and approaches U_{mf}.

9.2.2 Viewing heights of about 130 cm

Again, three different superficial velocities were tried at the intermediate freeboard height, at the 'quarter-in' position. Image-shifting yielded good results at the lowest air flowrate ($U = 1.12\,U_{mf}$; film 132; Figure 9.27), whilst direct PIV worked very well at the two higher flowrates ($U = 1.79\,U_{mf}$ for film 84 (Figure 9.28) and $U = 2.25\,U_{mf}$ for film 85 (Figure 9.29)).

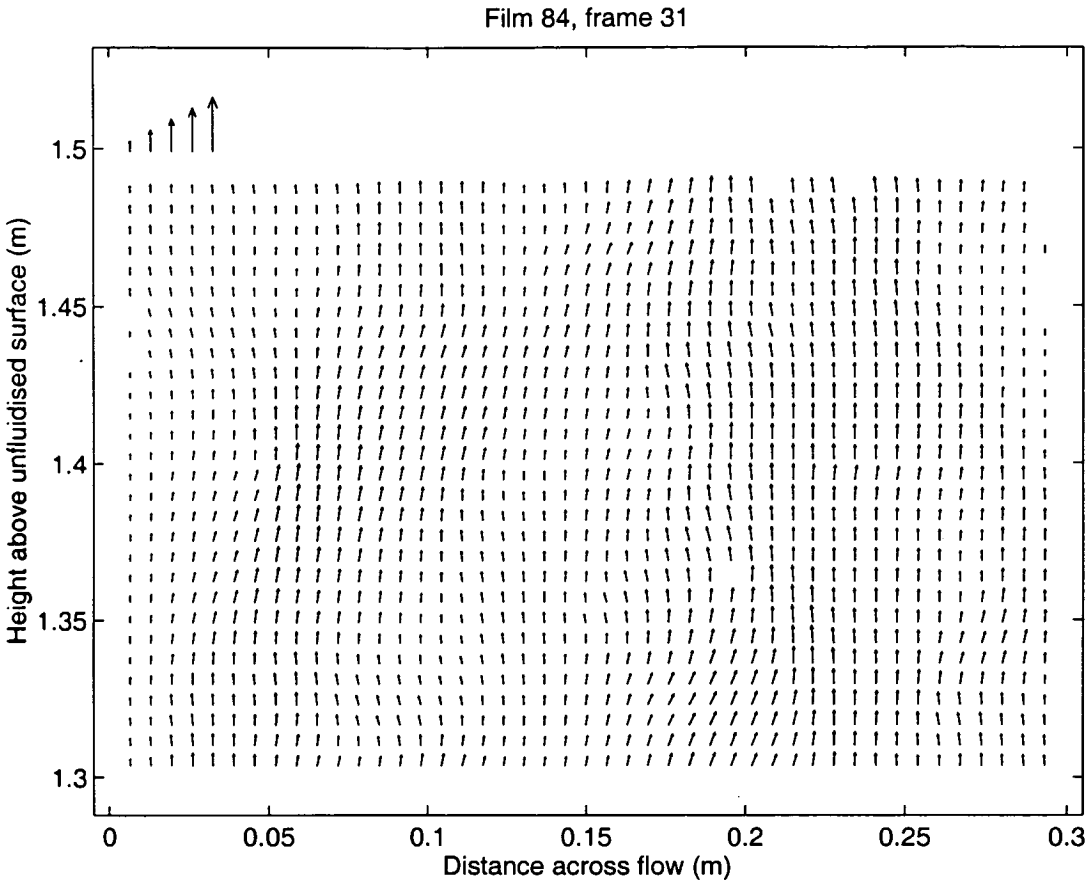
The yield of valid vectors was 91% at the lowest flowrate, when image-shifting was used. Almost all the vectors in the centre of the field of view were resolved satisfactorily, but no vectors in the outermost columns of each vector map were valid, due to the multiple bright images of the glass-air interface at either wall. The vector maps showed that the air upflow was greatest near the walls, with very low (or even downward) velocities found in the centre of the flow. A sample vector map and some velocity statistics are presented in Figure 9.27.

At the intermediate flowrate, the yields were extremely good, working out to be about 97% on average. The yields were at least 98% in 10 of the 13 analysed frames, but two frames yielded only about 90% valid vectors and brought down the average. It is possible that the two worst frames could have been analysed more successfully if the appropriate negatives had been scanned in again; however, the velocity results obtained from these two frames appeared to be consistent with those from the other frames of film 84, so repeat analyses were not deemed necessary. As the reduced yields could have been the result of regions of unusually high seeding concentration or excessive out-of-plane motion, there was no guarantee that repeat runs would have been any more successful anyway.



Frame	Yield (%)	Mean v_y ($\mathrm{m\,s^{-1}}$)	Max v_y ($\mathrm{m\,s^{-1}}$)	Min v_y ($\mathrm{m\,s^{-1}}$)	Mean v_x ($\mathrm{m\,s^{-1}}$)	Max v_x ($\mathrm{m\,s^{-1}}$)
f132f18	91	0.0469	0.155	−0.025	−0.0016	±0.030
f132f19	92	0.0515	0.170	−0.025	+0.0021	+0.025
f132f20	92	0.0512	0.160	+0.010	+0.0002	−0.030
f132f21	90	0.0498	0.145	+0.000	+0.0004	+0.040
f132f22	92	0.0556	0.165	−0.000	−0.0029	−0.045
f132f23	91	0.0532	0.140	−0.005	−0.0035	−0.040
f132f24	92	0.0447	0.125	−0.005	−0.0022	+0.050
f132f25	91	0.0446	0.125	−0.005	+0.0004	−0.035
f132f26	91	0.0478	0.155	−0.005	+0.0023	+0.035
mean:	91	0.0495	0.146	−0.007	−0.0005	~ 0.037

Figure 9.27: Results for film 132, where the superficial velocity was $0.0787\,\mathrm{m\,s^{-1}}$, or about $1.12\,U_{mf}$. The bed depth was 31 cm and the mean viewing height was about 130 cm above the unfluidised bed surface (one quarter of the way into the freeboard). Image-shifting was used. Top: vector map of the velocity information obtained from frame 22; Bottom: some velocity statistics.



Frame	Yield (%)	Mean v_y ($\mathrm{m\,s^{-1}}$)	Max v_y ($\mathrm{m\,s^{-1}}$)	Min v_y ($\mathrm{m\,s^{-1}}$)	Mean v_x ($\mathrm{m\,s^{-1}}$)	Max v_x ($\mathrm{m\,s^{-1}}$)
f84f19	89	0.1063	0.175	0.045	+0.0016	+0.055
f84f20	99	0.1054	0.175	0.050	+0.0064	+0.055
f84f21	98	0.1045	0.160	0.055	+0.0116	+0.055
f84f22	98	0.0944	0.155	0.050	+0.0122	+0.055
f84f23	99	0.0956	0.170	0.050	+0.0075	±0.060
f84f24	89	0.1115	0.200	0.045	−0.0042	±0.060
f84f25	96	0.1150	0.200	0.030	+0.0020	+0.075
f84f26	98	0.1137	0.200	0.025	+0.0094	+0.075
f84f27	99	0.1114	0.190	0.025	+0.0101	+0.065
f84f28	98	0.1147	0.190	0.035	+0.0101	+0.065
f84f29	98	0.1031	0.185	0.025	+0.0072	+0.065
f84f30	99	0.1238	0.190	0.050	+0.0038	±0.055
f84f31	99	0.1059	0.165	0.040	+0.0051	±0.060
mean:	97	0.1081	0.181	0.040	+0.0064	~ 0.062

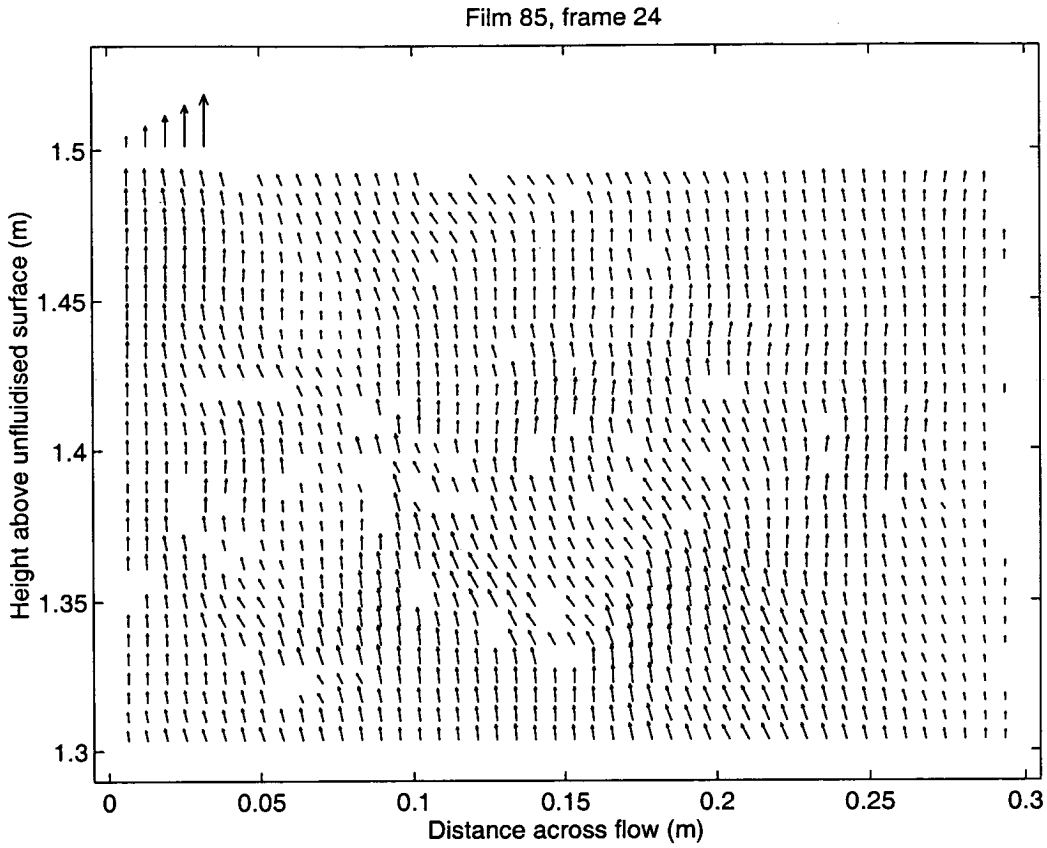
Figure 9.28: Results obtained from film 84, with $U = 0.125\,\mathrm{m\,s^{-1}}$, or about $1.79\,U_{mf}$. The bed depth was 31 cm and the mean viewing height was about 140 cm above the unfluidised bed surface (one quarter of the way into the freeboard). Image-shifting was not used here. Top: vector map of the velocity information obtained from frame 31; Bottom: the velocity statistics for this run.

A sample vector map from the intermediate flowrate run is given in Figure 9.28, along with some velocity statistics. The flow under these conditions was entirely upward, with the vertical components of velocity lying in the range $0.04\,\mathrm{m\,s^{-1}}$ to $0.18\,\mathrm{m\,s^{-1}}$. This limited range of velocities meant that direct PIV had few problems in resolving almost all of the image displacements, including near the walls. In addition, the measured vectors were all very easy to align. Under these conditions, the use of image-shifting would have been pointless.

The increased concentration of large bed particles at the higher flowrate had an impact on the yield of valid vectors ($U = 0.157\,\mathrm{m\,s^{-1}}$ for film 85). Another factor was the choice of illumination time interval, which could perhaps have been reduced by about 25% so as to record four images of each particle during each exposure, rather than three. This would have reduced the magnitudes of the image displacements on the bitmap images, which ranged between about 6 and about 27 pixels (with a median value of about 16 pixels). Nevertheless, the mean yield was about 91%, which seems quite acceptable (especially as it was dragged down by the values of 77% obtained in two of the 16 frames). A sample vector map is given in Figure 9.29, along with the velocity statistics for this film.

As in the intermediate flowrate run, the velocities at the highest flowrate were upward everywhere and the vertical velocity dynamic range was not too large (typically about 4 or 5). There was considerable variation from frame to frame in the measured velocities of film 85, however. Between frames 21 and 22, for example, the mean vertical velocity remained almost unchanged, even though there was a much greater range of velocities in frame 22 ($0.03\,\mathrm{m\,s^{-1}}$ to $0.27\,\mathrm{m\,s^{-1}}$) than in frame 21 ($0.06\,\mathrm{m\,s^{-1}}$ to $0.225\,\mathrm{m\,s^{-1}}$). Along similar lines, frames 24 and 25 had similar ranges of vertical velocities but significantly different mean vertical velocities. There were no vertical velocities slower than $0.05\,\mathrm{m\,s^{-1}}$ or faster than $0.215\,\mathrm{m\,s^{-1}}$ in either frame; the differences arose because the median velocities were at different positions within the same range.

Another observation which can be made is that the mean vertical velocities detected by PIV, averaging across each cross-section of the flow, tend to be closer



Frame	Yield (%)	Mean v_y ($\mathrm{m\,s^{-1}}$)	Max v_y ($\mathrm{m\,s^{-1}}$)	Min v_y ($\mathrm{m\,s^{-1}}$)	Mean v_x ($\mathrm{m\,s^{-1}}$)	Max v_x ($\mathrm{m\,s^{-1}}$)
f85f10	77	0.1384	0.260	−0.005	+0.0248	+0.125
f85f11	86	0.1559	0.280	+0.025	+0.0025	+0.100
f85f12	92	0.1659	0.250	+0.050	−0.0088	−0.060
f85f13	92	0.1405	0.260	+0.075	−0.0014	−0.120
f85f14	77	0.1406	0.240	+0.025	−0.0130	−0.110
f85f15	96	0.1457	0.225	+0.070	+0.0128	+0.085
f85f20	93	0.1020	0.180	+0.040	−0.0193	−0.080
f85f21	96	0.1297	0.225	+0.060	−0.0203	−0.090
f85f22	93	0.1319	0.270	+0.030	−0.0185	−0.110
f85f23	94	0.1182	0.225	+0.050	−0.0205	−0.090
f85f24	94	0.1301	0.215	+0.055	−0.0245	−0.100
f85f25	94	0.1493	0.210	+0.050	−0.0219	−0.090
f85f26	95	0.1428	0.225	+0.050	−0.0084	−0.070
f85f27	96	0.1508	0.235	+0.075	+0.0048	+0.070
f85f28	90	0.1354	0.225	+0.050	+0.0180	+0.095
f85f29	93	0.1686	0.260	+0.080	+0.0021	+0.090
mean:	91	0.1404	0.237	+0.049	−0.0057	~ 0.093

Figure 9.29: Summary of results for film 85. Conditions were the same as in film 84, but with an increased superficial velocity $U = 0.157\,\mathrm{m\,s^{-1}}$ $(2.25\,U_{mf})$. Top: the vector map from frame 24; Bottom: some velocity statistics for the run.

to the superficial air velocity in the runs taken at this intermediate viewing height of about 1.30 m than they were in the comparable runs at a viewing height of about 2.20 m. At the intermediate height, the PIV velocities averaged out to be about 0.017 m s^{-1} less than the superficial velocities in the runs at the higher flowrates (Figures 9.28 and 9.29). The differences at the viewing height of about 2.20 m were 0.030 m s^{-1} at the intermediate air flowrate (Figure 9.25) and about 0.049 m s^{-1} at the highest air flowrate (Figure 9.26). The differences at the lowest flowrates tested were greater than these values, but followed the trend of a greater apparent slip velocity being found at the highest viewing position.

9.2.3 Viewing heights of about 40 cm

No results will be presented from film 70, in which the viewing height was about 42 cm above the bed surface and the air flowrate was about 0.150 m s^{-1} (2.14 U_{mf}). In this film, which was the first film taken using a bed depth of 31 cm in the larger fluidised bed, the air flowrate was too large for PIV to work well.

The photographic negatives from film 70 appeared extremely dark, because the seeding particle concentration in the freeboard region was very high. This meant that the negatives did not scan in very well and could not be analysed successfully. The PIV parameters used here meant that about 6 images of each particle were recorded, so the camera exposure time could have been halved. Reducing the seeding particle concentration was another option. However, these changes were not investigated. Instead, more efforts were spent on the comparable runs using the centre of the test section and/or greater viewing heights.

The other runs at this viewing height used lower air flowrates (film 71, above) and possibly image-shifting as well (film 140, for a 60 cm bed; see later). It was found that the simple PIV technique used here could produce reasonable vector maps even at low viewing heights, if the PIV parameters were right.

9.3 60 cm bed, centre of freeboard

A number of runs were made using an increased bed depth of 60 cm, both in the centre of the freeboard and one quarter of the way into the freeboard cross-section. The centre-bed results are summarised in Table 9.10.

Bed depth (cm)	Viewing height (cm)	U/U_{mf}	Film number	Image-shifting used?
60	190	1.22	108	Yes
60	190	1.75	107	No
60	190	2.35	118	No
60	130	1.32	121	Yes
60	130	1.85	124	Yes
60	130	2.46	119	No
60	40	1.31	140	Yes
60	40	1.75	140	Yes

Table 9.10: Details of films taken using PIV to record the flow through the centre of the freeboard space, when the bed depth was 60 cm.

9.3.1 Viewing heights of about 190 cm

Three different air flowrates were tested at the greatest viewing height. These produced superficial velocities of $U = 0.0852\,\mathrm{m\,s^{-1}}$ (film 108), $0.122\,\mathrm{m\,s^{-1}}$ (film 107) or $0.165\,\mathrm{m\,s^{-1}}$ (film 118). Image-shifting was used only in the run at the lowest flowrate. These runs were made with the PIV apparatus at the same position above the fluidised bed distributor as it had been at the end of the last runs with the 31 cm bed. The viewing heights were therefore about 30 cm lower when measured relative to the bed surface.

A sample vector map from the lowest flowrate tested under these conditions is given in Figure 9.30. Every frame from film 108 showed upflow near the walls, usually with a peak magnitude which exceeded the superficial air velocity. However,

the section-average vertical velocity was rather low in all frames, as there were always some regions of extremely low (or even downward) velocity.

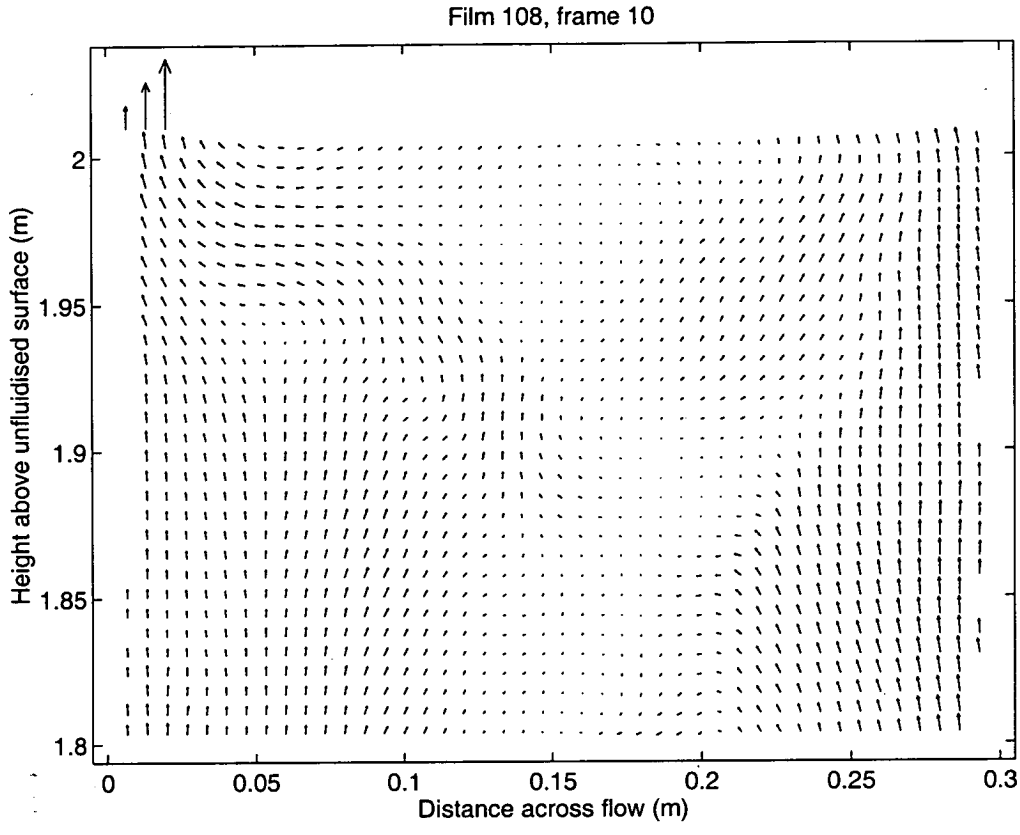


Figure 9.30: Vector map of the velocity information obtained from frame 10 of film 108, where $U = 0.0852\,\mathrm{m\,s^{-1}} \approx 1.22\,U_{mf}$. Image-shifting was used.

The velocity statistics for this lowest flowrate are given in Table 9.11. The application of an image-shifting technique allowed a very good yield of valid vectors to be obtained in the bulk of the flow (even in the regions of lowest velocity), but caused difficulties right next to the walls. This systematic under-representation of the most strongly-upward vectors in the flow meant that the mean upward velocities, obtained by averaging across each frame, were lower than the true values. A few test corrections were made, in which the corner vectors in each vector map were assumed to be identical to the vectors immediately adjacent to them when moving horizontally towards the centre of the test section, before interpolation was used to estimate the missing vectors. The result of this, however, was that the mean vertical velocities increased by only about 5–15%. Thus, the missing vectors do not appear to make a dramatic difference to the section-average velocities, which are all significantly lower than the superficial velocity.

Frame	Yield (%)	Mean v_y ($m\,s^{-1}$)	Max v_y ($m\,s^{-1}$)	Min v_y ($m\,s^{-1}$)	Mean v_x ($m\,s^{-1}$)	Max v_x ($m\,s^{-1}$)
f108f6	97	0.0189	0.105	−0.030	−0.0066	−0.040
f108f7	98	0.0207	0.090	−0.025	−0.0025	−0.040
f108f8	96	0.0200	0.090	−0.025	−0.0021	−0.040
f108f9	98	0.0229	0.090	−0.030	−0.0025	±0.040
f108f10	98	0.0260	0.090	−0.025	−0.0022	−0.050
f108f11	97	0.0287	0.105	−0.025	−0.0009	−0.045
f108f12	95	0.0293	0.115	−0.025	−0.0029	−0.050
f108f13	95	0.0238	0.090	−0.040	−0.0012	−0.050
f108f14	97	0.0222	0.085	−0.040	+0.0054	+0.040
f108f15	95	0.0193	0.105	−0.040	+0.0039	+0.040
f108f16	96	0.0162	0.105	−0.045	+0.0004	+0.040
f108f17	98	0.0135	0.090	−0.040	−0.0002	±0.033
f108f18	96	0.0084	0.080	−0.040	−0.0004	±0.033
f108f19	96	0.0081	0.070	−0.040	−0.0000	±0.033
f108f20	97	0.0186	0.095	−0.035	+0.0047	+0.040
f108f21	96	0.0232	0.090	−0.025	+0.0025	+0.035
f108f22	96	0.0220	0.090	−0.015	+0.0001	+0.033
mean:	97	0.0201	0.093	−0.032	−0.0003	∼0.040

Table 9.11: Summary of velocity values for film 108, where the superficial velocity U was $0.0851\,m\,s^{-1}$, or about $1.22\,U_{mf}$. The bed depth was 60 cm. The mean viewing height was about 190 cm above the unfluidised bed surface and through the centre of the bed. Image-shifting was used.

In the run at an intermediate flowrate, a superficial velocity of $0.122\,m\,s^{-1}$ was used ($U = 1.75\,U_{mf}$). The results from film 107 are presented by means of a sample vector map (Figure 9.31) and a table of velocity statistics (Table 9.12).

Compared to the run at the lowest flowrate, the superficial gas velocity increased by about $0.037\,m\,s^{-1}$. This increase is the same as the increase in the maximum vertical velocity (typically $0.093\,m\,s^{-1}$ at the lowest flowrate, but increased to $0.130\,m\,s^{-1}$ here). However, somewhat larger increases were found in the minimum vertical velocity (up from $-0.032\,m\,s^{-1}$ at the lowest flowrate to $+0.039\,m\,s^{-1}$ here) and in the section-average vertical velocity (up by $0.0557\,m\,s^{-1}$). The fluctuations in the mean vertical velocity in the two runs were similar (about $\pm 0.010\,m\,s^{-1}$ from frame to frame).

Features of a typical vector map at the intermediate flowrate, such as that shown in Figure 9.31, are that the velocity vectors are all upward and almost

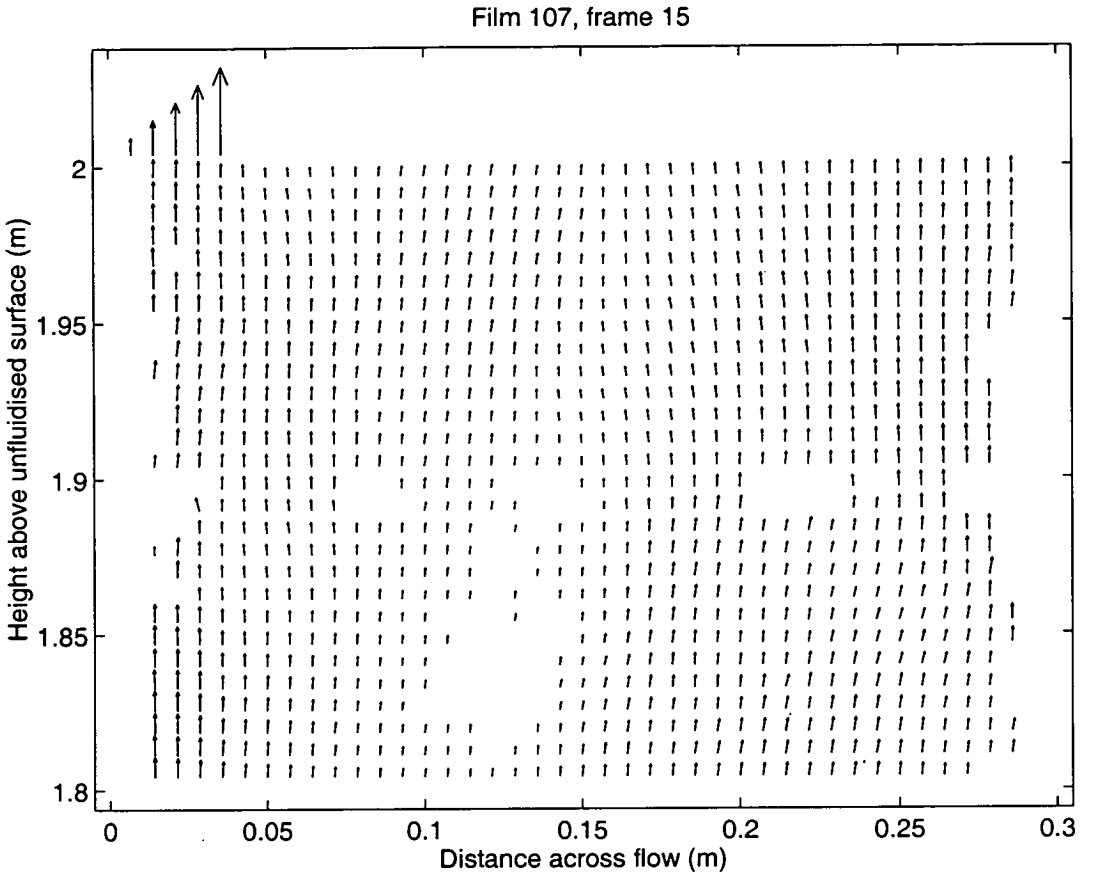


Figure 9.31: The vector map obtained from frame 15 of film 107, with $U = 0.122\,\mathrm{m\,s^{-1}}$ (about $1.75\,U_{mf}$). The bed depth was 60 cm. The mean viewing height was about 190 cm above the bed surface and through the centre of the freeboard. Image-shifting was not used.

vertical. The largest vectors were often found near the walls, but the vertical velocity profile across the flow was much flatter than it had been at the lowest flowrate tested under these conditions (Figure 9.30). Various low-velocity regions could not be resolved, which accounted for some of the variation in the yield of valid vectors from frame to frame (*e.g.* from frame 13 to frame 16).

Much of the rest of the quite substantial variation in the yield is believed to have been due to out-of-plane motion, rather than difficulties in scanning the negatives (*e.g.* frames 17–20). This belief is based on a direct visual examination of the negatives (on which no image pairs could be seen in large regions, suggesting that even optimal scanning and image-editing would have made no difference) and on the observation that the horizontal velocities were somewhat larger in these frames than in the earlier part of the film. After all, if large in-plane horizontal

Frame	Yield (%)	Mean v_y ($m\,s^{-1}$)	Max v_y ($m\,s^{-1}$)	Min v_y ($m\,s^{-1}$)	Mean v_x ($m\,s^{-1}$)	Max v_x ($m\,s^{-1}$)
f107f7	89	0.0777	0.125	0.040	−0.0003	±0.025
f107f8	88	0.0783	0.110	0.050	+0.0003	−0.025
f107f9	83	0.0743	0.110	0.045	+0.0011	±0.025
f107f10	71	0.0854	0.130	0.050	−0.0005	±0.030
f107f11	94	0.0819	0.135	0.055	−0.0018	±0.025
f107f12	94	0.0791	0.130	0.055	−0.0007	±0.025
f107f13	95	0.0821	0.130	0.050	+0.0006	+0.025
f107f14	88	0.0714	0.125	0.035	+0.0004	+0.025
f107f15	87	0.0695	0.130	0.025	+0.0030	+0.025
f107f16	83	0.0679	0.135	0.030	+0.0054	+0.025
f107f17	43	0.0707	0.135	0.025	+0.0135	+0.045
f107f18	38	0.0730	0.135	0.030	+0.0050	+0.060
f107f19	44	0.0739	0.150	0.030	+0.0022	±0.045
f107f20	49	0.0761	0.145	0.030	+0.0020	−0.040
mean:	75	0.0758	0.130	0.039	+0.0022	~ 0.032

Table 9.12: Results for film 107, with $U = 0.122\,m\,s^{-1}$ (about $1.75\,U_{mf}$). The bed depth was 60 cm. The mean viewing height was about 190 cm above the bed surface and through the centre of the bed. Image-shifting was not used. Top: the vector map obtained from frame 15; Bottom: some velocity statistics.

velocities are found, large out-of-plane horizontal velocities are also likely to be present. Some of these problems of low velocities and out-of-plane motion could have been alleviated by using image-shifting, allowing a shorter illumination time interval to be chosen (thus giving the seeding particles less chance to move into or out of the light sheet during each exposure). On the whole, however, the results from film 107 were considered to be adequate, so a repeat run with the help of image-shifting was not considered necessary.

In line with previous results, image-shifting was not needed at the highest flowrate tested at this bed depth and viewing position. For film 118, the superficial velocity was increased to $0.165\,m\,s^{-1}$ (about $2.35\,U_{mf}$). The vector maps were similar in appearance to those at the intermediate flowrate, but without the regions of low velocity. The vector map from frame 11 is given in Figure 9.32.

The main problem encountered was that the high air flowrate greatly increased the number of large particle images found in the test section. These bed particles usually moved at a somewhat lower velocity than the seeding particles and each

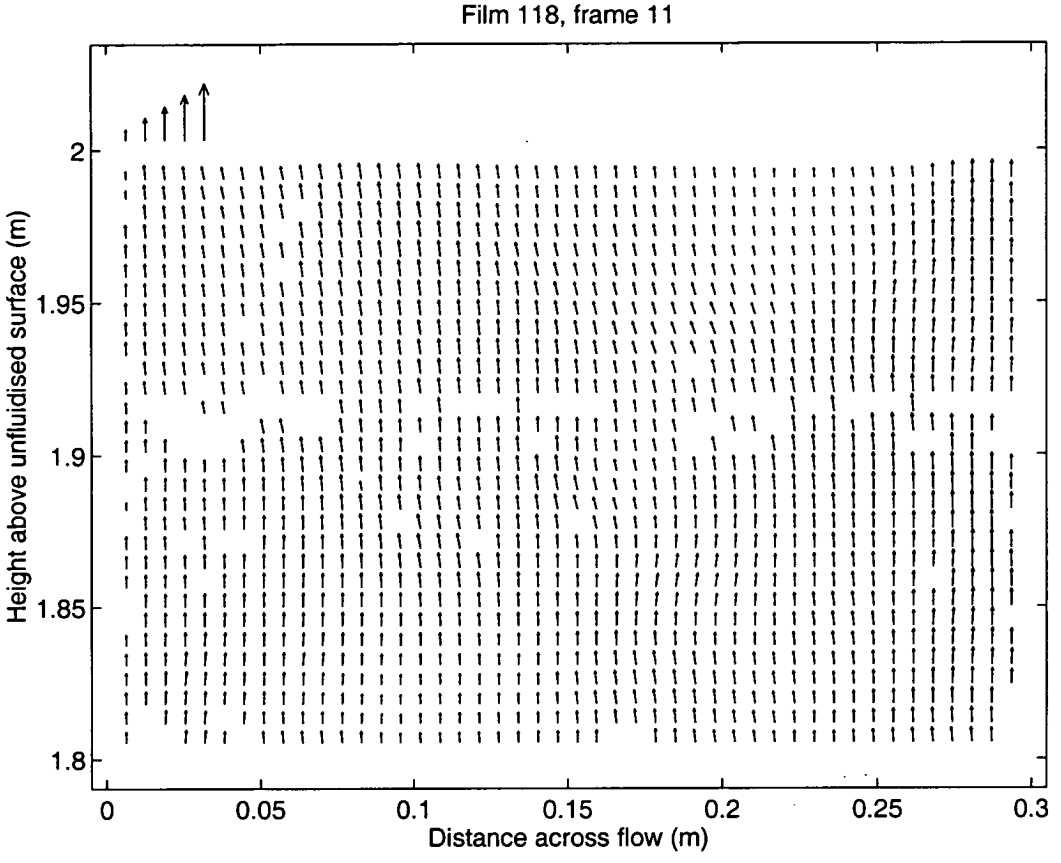


Figure 9.32: Vector map of the velocity information obtained from frame 11 of film 118, where $U = 0.165\,\mathrm{m\,s^{-1}} \approx 2.35\,U_{mf}$. Image-shifting was not used.

image was much larger and brighter, so, if the image-separation software did not achieve a full phase separation, the occasional low-velocity vector was found in a 'sea' of larger-velocity measurements based on seeding particle images. These bed-particle measurements accounted for a sizable proportion of the missing vectors.

The problems of incomplete phase separation were difficult to avoid. During operation, the computer program PIVSPLIT analyses a bitmap image file and 'deletes' any images it encounters which meet certain criteria for size, brightness, shape, *etc.* One option when deleting a bright larger-phase particle image is to set to zero the brightness of all pixels associated with that image. A drawback is that the measurement of smaller-phase image separations will be affected if the two phases have similar image displacements, as the autocorrelation background 'noise' will be lower than in the rest of the autocorrelation plane.

The alternative PIVSPLIT option, which had been recommended for use here, is to set the brightness of all pixels in a large-phase image to an average value, so that the smaller-phase data can be analysed more accurately. This, however, relies on a small-phase autocorrelation peak being obtained which is much taller than the spurious correlation peak due to the darkened large-phase images. If the small-phase seeding concentration was too low, or if the large-phase particle concentration was too high, the peak detected was likely to be that due to the larger phase, rather than the random measurement which would have been returned if the large-phase data had been completely obliterated (hence the observed large-phase displacements within the small-phase vector maps). This issue of incomplete phase separation is probably not very important, however, as these conditions of low seeding concentration and large bed-particle concentration are not likely to give rise to a valid small-phase measurement anyway.

Frame	Yield (%)	Mean v_y ($\mathrm{m\,s^{-1}}$)	Max v_y ($\mathrm{m\,s^{-1}}$)	Min v_y ($\mathrm{m\,s^{-1}}$)	Mean v_x ($\mathrm{m\,s^{-1}}$)	Max v_x ($\mathrm{m\,s^{-1}}$)
f118f2	98	0.1128	0.150	0.070	+0.0013	+0.040
f118f3	97	0.1186	0.175	0.070	+0.0018	+0.033
f118f4	93	0.1351	0.190	0.080	+0.0011	+0.045
f118f5	97	0.1260	0.190	0.065	+0.0021	+0.040
f118f6	96	0.1188	0.175	0.070	+0.0045	+0.040
f118f7	94	0.1318	0.185	0.080	+0.0022	±0.040
f118f8	93	0.1219	0.180	0.070	−0.0047	−0.040
f118f9	97	0.0931	0.145	0.040	−0.0103	−0.045
f118f10	95	0.0959	0.155	0.040	−0.0114	−0.040
f118f11	94	0.1214	0.170	0.060	−0.0085	−0.045
f118f12	87	0.1215	0.170	0.060	−0.0102	−0.050
f118f13	95	0.1149	0.170	0.070	−0.0061	−0.040
f118f14	95	0.0995	0.155	0.050	−0.0038	+0.050
f118f15	84	0.1087	0.170	0.040	+0.0066	+0.055
mean:	94	0.1157	0.170	0.062	−0.0025	~ 0.043

Table 9.13: Summary of velocities for film 118: $U = 0.165\,\mathrm{m\,s^{-1}} = 2.35\,U_{mf}$. The bed depth was 60 cm. The mean viewing height was about 190 cm above the unfluidised bed surface and through the centre of the bed.

The table of velocity statistics for film 118 (Table 9.13) suggests that the greatest vertical velocity in each frame was slightly greater than the superficial velocity (as was also found in films 108 and 107). Most vertical velocity statistics appeared to be about $0.043\,\mathrm{m\,s^{-1}}$ greater (*i.e.* more upward) than at the interme-

diate air flowrate (film 107), this being the difference between the two superficial velocities. This simple model predicts the changes in the section-average vertical velocity and the typical greatest vertical velocity rather well, but overestimates the value of the slowest upward velocity in each frame (which was only $0.023\,\mathrm{m\,s^{-1}}$ greater in film 118). There was a larger-than-expected increase in this minimum vertical velocity on moving from the slowest to the intermediate flowrate, so perhaps a smaller increase here should not be a surprise.

9.3.2 Viewing heights of about 130 cm

Again, three flowrates were tested at this viewing height. The three superficial velocities used were $0.0926\,\mathrm{m\,s^{-1}}$ (film 121, with image-shifting), $0.130\,\mathrm{m\,s^{-1}}$ (film 124, with image-shifting) and $0.172\,\mathrm{m\,s^{-1}}$ (film 119, no image-shifting). By coincidence, each of these velocities was almost exactly $0.007\,\mathrm{m\,s^{-1}}$ larger than the corresponding velocity used at this bed depth at the highest viewing position.

Frame	Yield (%)	Mean v_y ($\mathrm{m\,s^{-1}}$)	Max v_y ($\mathrm{m\,s^{-1}}$)	Min v_y ($\mathrm{m\,s^{-1}}$)	Mean v_x ($\mathrm{m\,s^{-1}}$)	Max v_x ($\mathrm{m\,s^{-1}}$)
f121f3	95	−0.0014	0.125	−0.060	−0.0095	−0.055
f121f4	95	−0.0011	0.100	−0.060	−0.0080	−0.055
f121f5	95	+0.0021	0.100	−0.060	−0.0060	−0.055
f121f6	95	+0.0018	0.105	−0.050	−0.0006	+0.055
f121f7	95	+0.0005	0.110	−0.050	−0.0012	±0.050
f121f8	95	−0.0007	0.115	−0.050	+0.0003	−0.045
f121f9	96	−0.0006	0.095	−0.045	+0.0019	−0.045
f121f10	96	−0.0015	0.130	−0.045	+0.0033	+0.050
f121f11	95	+0.0055	0.160	−0.075	+0.0072	+0.070
f121f12	95	+0.0078	0.160	−0.070	+0.0052	+0.060
f121f13	95	+0.0066	0.160	−0.060	+0.0042	−0.035
f121f14	95	+0.0014	0.130	−0.060	+0.0026	±0.035
f121f15	95	+0.0058	0.145	−0.050	+0.0006	+0.035
mean:	95	+0.0020	0.126	−0.057	−0.0000	∼ 0.050

Table 9.14: Summary of velocity values for film 121. The superficial velocity was $0.0926\,\mathrm{m\,s^{-1}}$ ($1.32\,U_{mf}$). The bed depth was 60 cm. The test section was 130 cm above the bed surface, and through the bed's centre. Image-shifting was used.

The run at the lowest flowrate gave very interesting results (Table 9.14). Although the superficial air velocity was $0.0926\,\mathrm{m\,s^{-1}}$ ($U = 1.32\,U_{mf}$) for film 121,

the net upflow measured by PIV was almost negligible. The greatest net upflow was only $0.0078\,\mathrm{m\,s^{-1}}$ (in frame 12). Some net downward flows were found, *e.g.* in frames 8–10. The yields of valid vectors were good, consisting of almost every vector in the bulk of the flow but missing most of the vectors in the two columns nearest the walls. As was often the case, these missing vectors were among the most strongly upward vectors in each vector map, so the section-average velocities should be slightly larger than those calculated by averaging all measured vectors.

Some interesting turbulence structures could be seen in the vector maps obtained from film 121, *e.g.* that for frame 12 (Figure 9.33). Two vortices can be seen near the left-hand wall, both of which rotate in a clockwise direction. A vortex rotating in an anticlockwise direction can be seen in the top right-hand corner, along with various others near the bottom centre of the vector map.

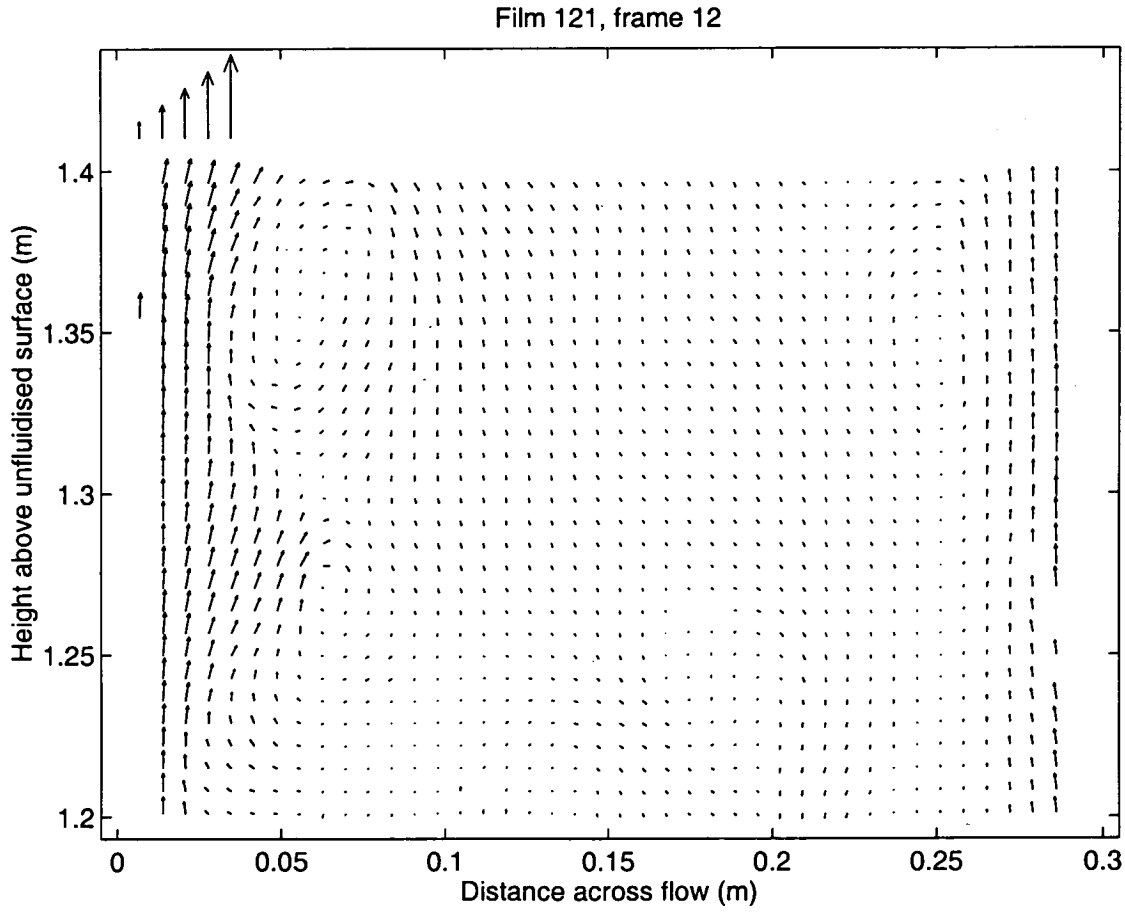


Figure 9.33: Vector map of the velocity information obtained from frame 12 of film 121, where $U = 0.0926\,\mathrm{m\,s^{-1}} \approx 1.32\,U_{mf}$. Image-shifting was used.

An increase in superficial velocity to $0.1296\,\mathrm{m\,s^{-1}}$ ($U = 1.85\,U_{mf}$) was sufficient to make all vectors in film 124 upward. A considerable amount of lateral movement was also observed, resulting in the flow appearing wavy. The vector map for frame 17 of film 124 (Figure 9.35, top) illustrates this well. A further feature of this frame is that the vertical velocity distribution function was bi-modal, with peaks at velocities of about $0.075\,\mathrm{m\,s^{-1}}$ (due mainly to the vectors in the right-hand half of the frame) and about $0.155\,\mathrm{m\,s^{-1}}$ (due to the rapidly-rising jet near the left-hand wall). In contrast, the horizontal velocity components were approximately normally-distributed, with mean and median values close to zero (Figure 9.34).

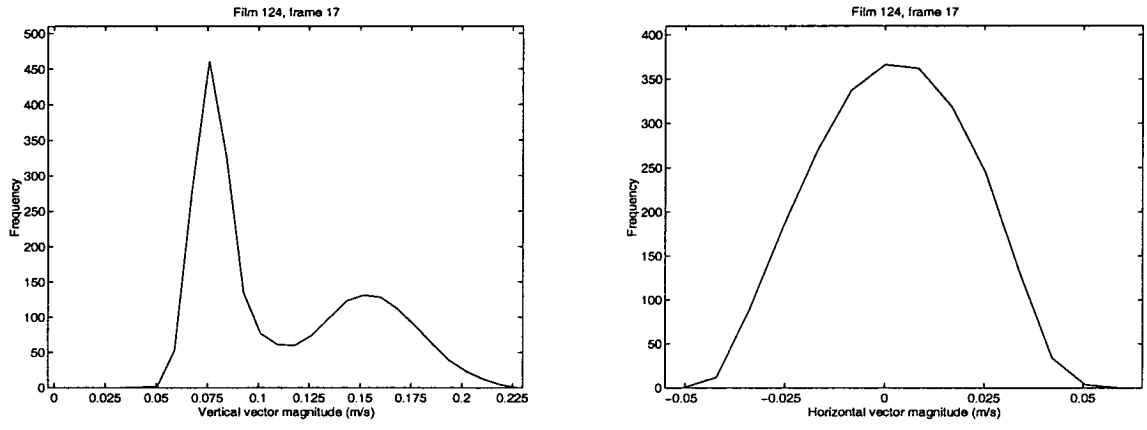
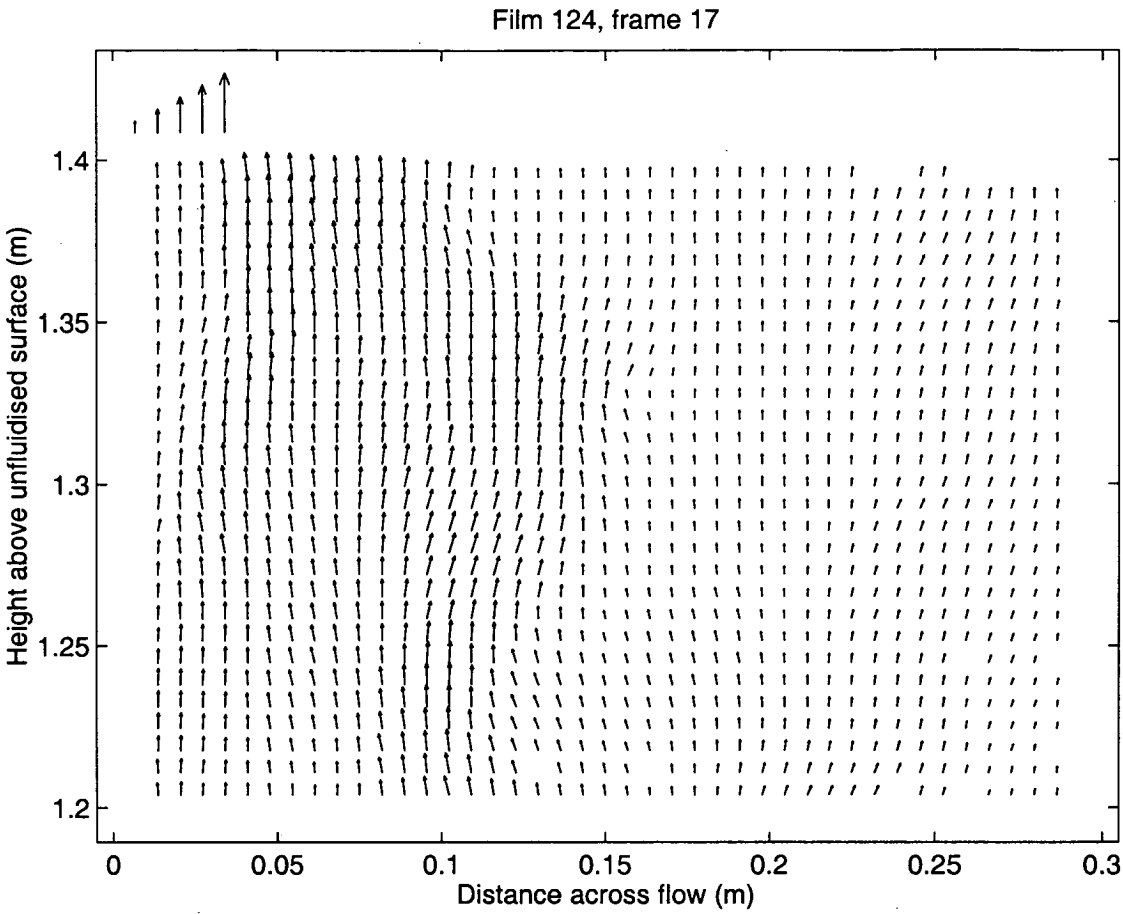


Figure 9.34: Distributions of the vertical (left) and horizontal (right) velocities in frame 17 of film 124. $U = 0.1296\,\mathrm{m\,s^{-1}} \approx 1.85\,U_{mf}$. Image-shifting was used.

As can be seen from the velocity statistics for film 124 (Figure 9.35, bottom), the increase in superficial velocity ($0.037\,\mathrm{m\,s^{-1}}$) was pretty well in line with the increase in the maximum vertical velocity, when both values are compared to those in the run at the lowest flowrate used at this viewing height and bed depth (film 121). In both cases, the maximum vertical velocity was about $0.03\,\mathrm{m\,s^{-1}}$ greater than the superficial velocity. In contrast, both the minimum measured vertical velocity and the section-average vertical velocity increased by a somewhat larger amount (about $0.095\,\mathrm{m\,s^{-1}}$) for the same increase in the superficial velocity (the same trend as observed at the greater viewing height). The yield of valid vectors was about 94% in both cases, mainly due to missing vectors in the outermost columns of each vector map.



Frame	Yield (%)	Mean v_y ($\mathrm{m\,s^{-1}}$)	Max v_y ($\mathrm{m\,s^{-1}}$)	Min v_y ($\mathrm{m\,s^{-1}}$)	Mean v_x ($\mathrm{m\,s^{-1}}$)	Max v_x ($\mathrm{m\,s^{-1}}$)
f124f5	95	0.0869	0.135	+0.040	−0.0026	−0.055
f124f6	93	0.0915	0.180	+0.040	−0.0011	−0.050
f124f7	95	0.0816	0.165	+0.025	+0.0010	+0.050
f124f8	92	0.1014	0.165	+0.050	−0.0053	−0.055
f124f9	93	0.0955	0.160	+0.025	−0.0076	−0.075
f124f10	94	0.0830	0.150	+0.025	+0.0017	−0.055
f124f11	95	0.0960	0.145	+0.040	+0.0108	+0.055
f124f12	95	0.0947	0.150	+0.045	+0.0035	±0.050
f124f13	95	0.0945	0.150	+0.045	−0.0097	−0.050
f124f14	94	0.1017	0.155	+0.050	−0.0086	−0.060
f124f15	94	0.0936	0.130	+0.040	+0.0079	+0.040
f124f16	93	0.0988	0.200	+0.040	+0.0069	+0.055
f124f17	94	0.1110	0.225	+0.050	+0.0027	+0.050
mean:	94	0.0946	0.162	+0.040	−0.0003	~ 0.054

Figure 9.35: Top: Vector map of the velocities for frame 17 of film 124, where $U = 0.1296\,\mathrm{m\,s^{-1}} \approx 1.85\,U_{mf}$; Bottom: Summary of velocity values for film 124. The mean viewing height was about 130 cm above the unfluidised bed surface and the light sheet was directed at the centre of the bed. Image-shifting was used.

At the greater superficial air velocity of $0.172\,\mathrm{m\,s^{-1}}$ ($U \approx 2.46\,U_{mf}$) used in film 119, the mean vertical velocity remained about $0.036\,\mathrm{m\,s^{-1}}$ below the superficial velocity (compared to the intermediate flowrate run). At the same time, the greatest vertical velocities increased significantly to about $0.272\,\mathrm{m\,s^{-1}}$ (about $0.10\,\mathrm{m\,s^{-1}}$ more than U and about twice the section-average vertical velocity). In contrast, the minimum vertical velocity was almost unaffected by this increase in flowrate; if anything, it was slightly lower than at the intermediate flowrate. Some large horizontal velocities were also found, occasionally having the same magnitude as the superficial velocity. Nevertheless, the average yield of valid vectors was reasonably good (86%). Some velocity statistics are given in Table 9.15. It was noted that there was a lot of variation from frame to frame.

Frame	Yield (%)	Mean v_y ($\mathrm{m\,s^{-1}}$)	Max v_y ($\mathrm{m\,s^{-1}}$)	Min v_y ($\mathrm{m\,s^{-1}}$)	Mean v_x ($\mathrm{m\,s^{-1}}$)	Max v_x ($\mathrm{m\,s^{-1}}$)
f119f2	93	0.1465	0.28	0.05	+0.0023	−0.09
f119f3	89	0.1412	0.27	0.04	+0.0174	+0.15
f119f4	84	0.1789	0.285	0.06	+0.0368	+0.19
f119f5	83	0.1609	0.32	0.04	+0.0396	+0.19
f119f6	87	0.1469	0.27	0.03	+0.0422	+0.12
f119f7	89	0.1785	0.32	0.04	+0.0158	+0.14
f119f8	92	0.1347	0.27	0.04	−0.0099	−0.09
f119f9	70	0.0793	0.23	0.02	+0.0050	+0.12
f119f10	92	0.0951	0.24	0.03	−0.0022	−0.13
f119f11	88	0.1177	0.27	0.04	−0.0196	−0.15
f119f12	79	0.0964	0.24	0.02	−0.0234	−0.14
f119f13	85	0.1248	0.27	0.03	−0.0056	+0.17
mean:	86	0.1334	0.272	0.036	+0.0082	~ 0.093

Table 9.15: Velocity statistics for film 119, where $U = 0.172\,\mathrm{m\,s^{-1}}$ ($2.46\,U_{mf}$). The bed depth was 60 cm. The mean viewing height was 130 cm above the bed surface and through the centre of the bed. Image-shifting was not used.

Frame 10 gave one of the better yields of valid vectors (Figure 9.36), although there were a number of missing vectors at a height of approximately 1.28 m above the bed surface. Similar horizontal bands of missing vectors could be observed in frames 4–6 and 11–13 of this run. These were due to a fault with the image-scanner, which produced a bright horizontal band near the centre of each bitmap image. This occasionally resulted in hardly any vectors being found within a band of 1–3 rows of vectors near the centre of the field of view. Similar problems were noted earlier, for films 86 (Figure 9.3) and 118 (Figure 9.32).

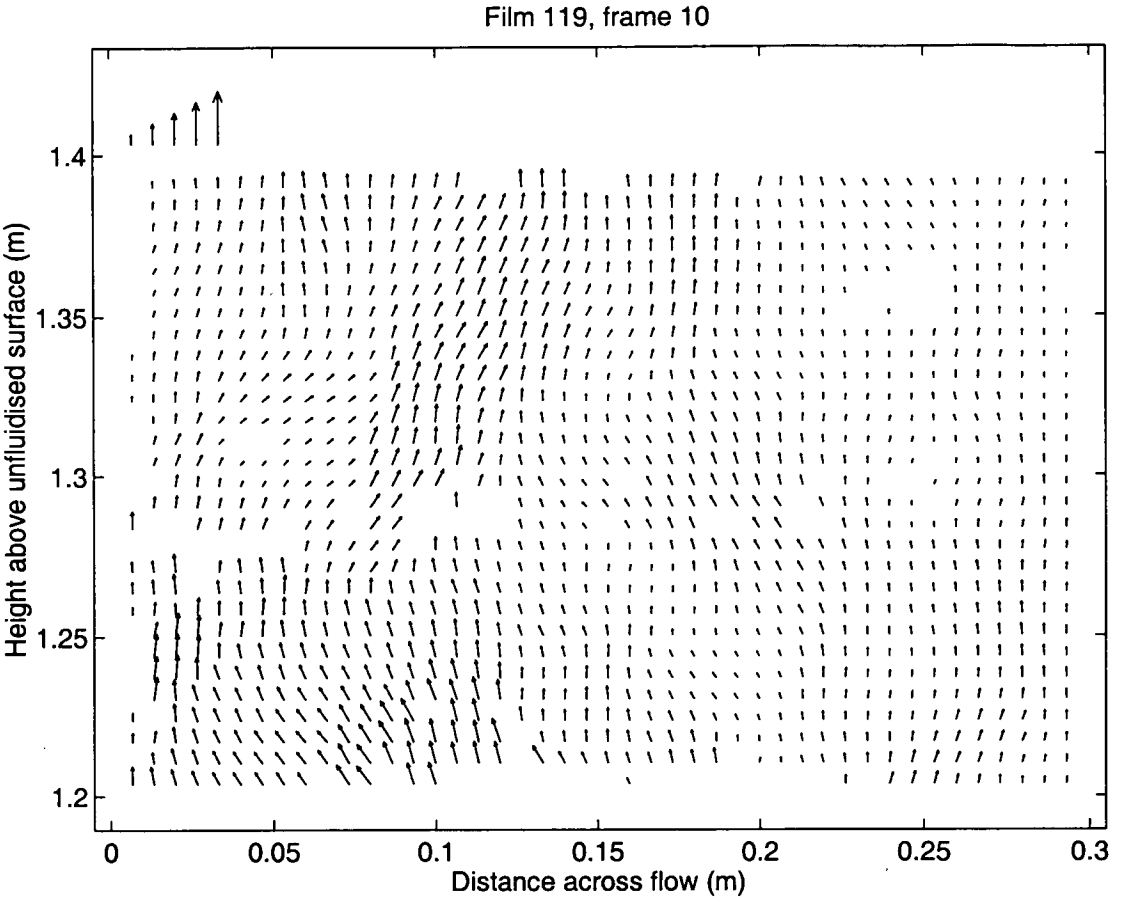


Figure 9.36: Vector map of the velocities measured in frame 10 of film 119, where $U = 0.172\,\mathrm{m\,s^{-1}}$ $(2.46\,U_{mf})$. Image-shifting was not used.

As far as the vector map in Figure 9.36 is concerned, the flow was once again rather wavy. Some low-velocity vectors were found in the right-hand half of the frame. It could therefore be argued that image-shifting would have been of some assistance even at this large a flowrate, to help reduce the dynamic range of the measured vectors. For example, the dynamic ranges of the vertical velocities at the intermediate flowrate (film 124) were about 3–6, whereas the corresponding ratios of maximum v_y to minimum v_y in film 119 were about 6–10. Another advantage of image-shifting is that it enables the direction of every vector to be determined automatically, but this would not be important here. All in all, the results from film 119 are probably as good as they are going to get. In any case, interpolation can be used to estimate reasonable values for most of the missing vectors.

9.3.3 Viewing heights of about 40 cm

The flow immediately above a bubbling fluidised bed is extremely turbulent. The bed surface rises and falls as a function of time, according to the holdup of bubbles in the bed. The bed surface in these experiments could easily rise dozens of cm above its average position, just before a particularly large bubble reached the surface. For this bed depth, the minimum viewing height was kept at about 40 cm above the level of the unfluidised bed surface. Image-shifting was used, in order to compress the dynamic range of the measured vectors.

These aims were achieved to a certain extent. Difficulties arose because the horizontal velocities in the flow field were somewhat larger than had been expected, which means that the runs at the higher air flowrates could have been repeated with the rotating mirror revolving at a greater angular speed. This would have to be done with caution, however, or else the measured velocities (after subtracting the image-shifting displacement) will have a large relative error due to the difference between two large numbers. For example, suppose that a flow with a typical vertical velocity of $0.10\,\mathrm{m\,s^{-1}}$ is photographed using an image-shifting displacement of $1.0\,\mathrm{m\,s^{-1}}$. A 1% error in each displacement measured by autocorrelation corresponds to an absolute error of $0.01\,\mathrm{m\,s^{-1}}$, which is about 10% of the true vertical velocity. On the other hand, perhaps velocity data with a 90% yield of correctly-aligned vectors at an accuracy of only $\pm 10\%$ is preferable to results accurate to $\pm 1\%$ but at a yield of only 50%. It may depend on the flow...

Film 140 was used up in three parts. At first, the air flowrate was about $495\,\mathrm{l\,min^{-1}}$, so the superficial velocity was $0.0917\,\mathrm{m\,s^{-1}}$ ($1.31\,U_{mf}$). After that, 11 frames were taken at a flowrate of $660\,\mathrm{l\,min^{-1}}$ ($U = 0.122\,\mathrm{m\,s^{-1}} = 1.75\,U_{mf}$). Both of these runs used image-shifting, at a rotational speed of $15^\circ\,\mathrm{s^{-1}}$. Finally, some photographs were taken without image-shifting, at a higher air flowrate.

The PIV parameters included a camera exposure time of 1/30 s and an illumination time interval of 6.2 ms, so five images of each particle were recorded. The mean image-shifting displacement was about 20.1 pixels. This displacement corresponded to an object-plane velocity of $0.28\,\mathrm{m\,s^{-1}}$, which seemed large enough.

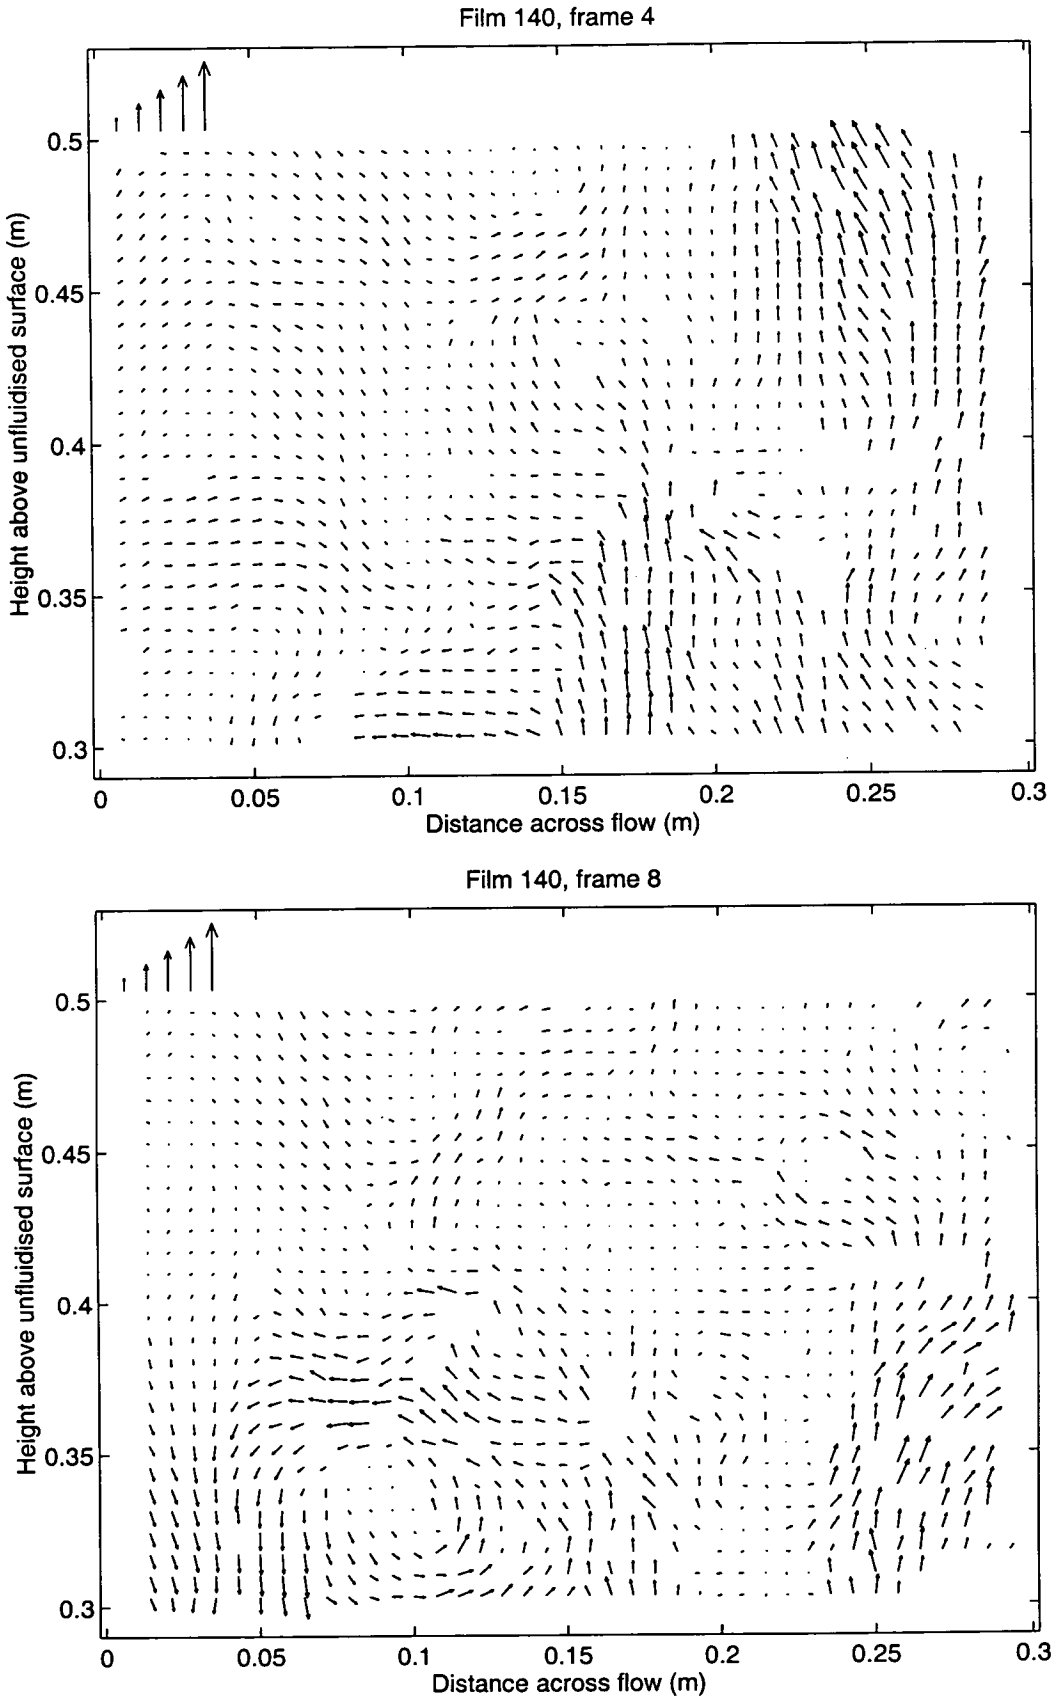


Figure 9.37: PIV vector maps for the flow about 40 cm above a 60 cm-deep bed of particles, after subtracting the image-shifting velocity. $U \approx 1.31\, U_{mf}$.

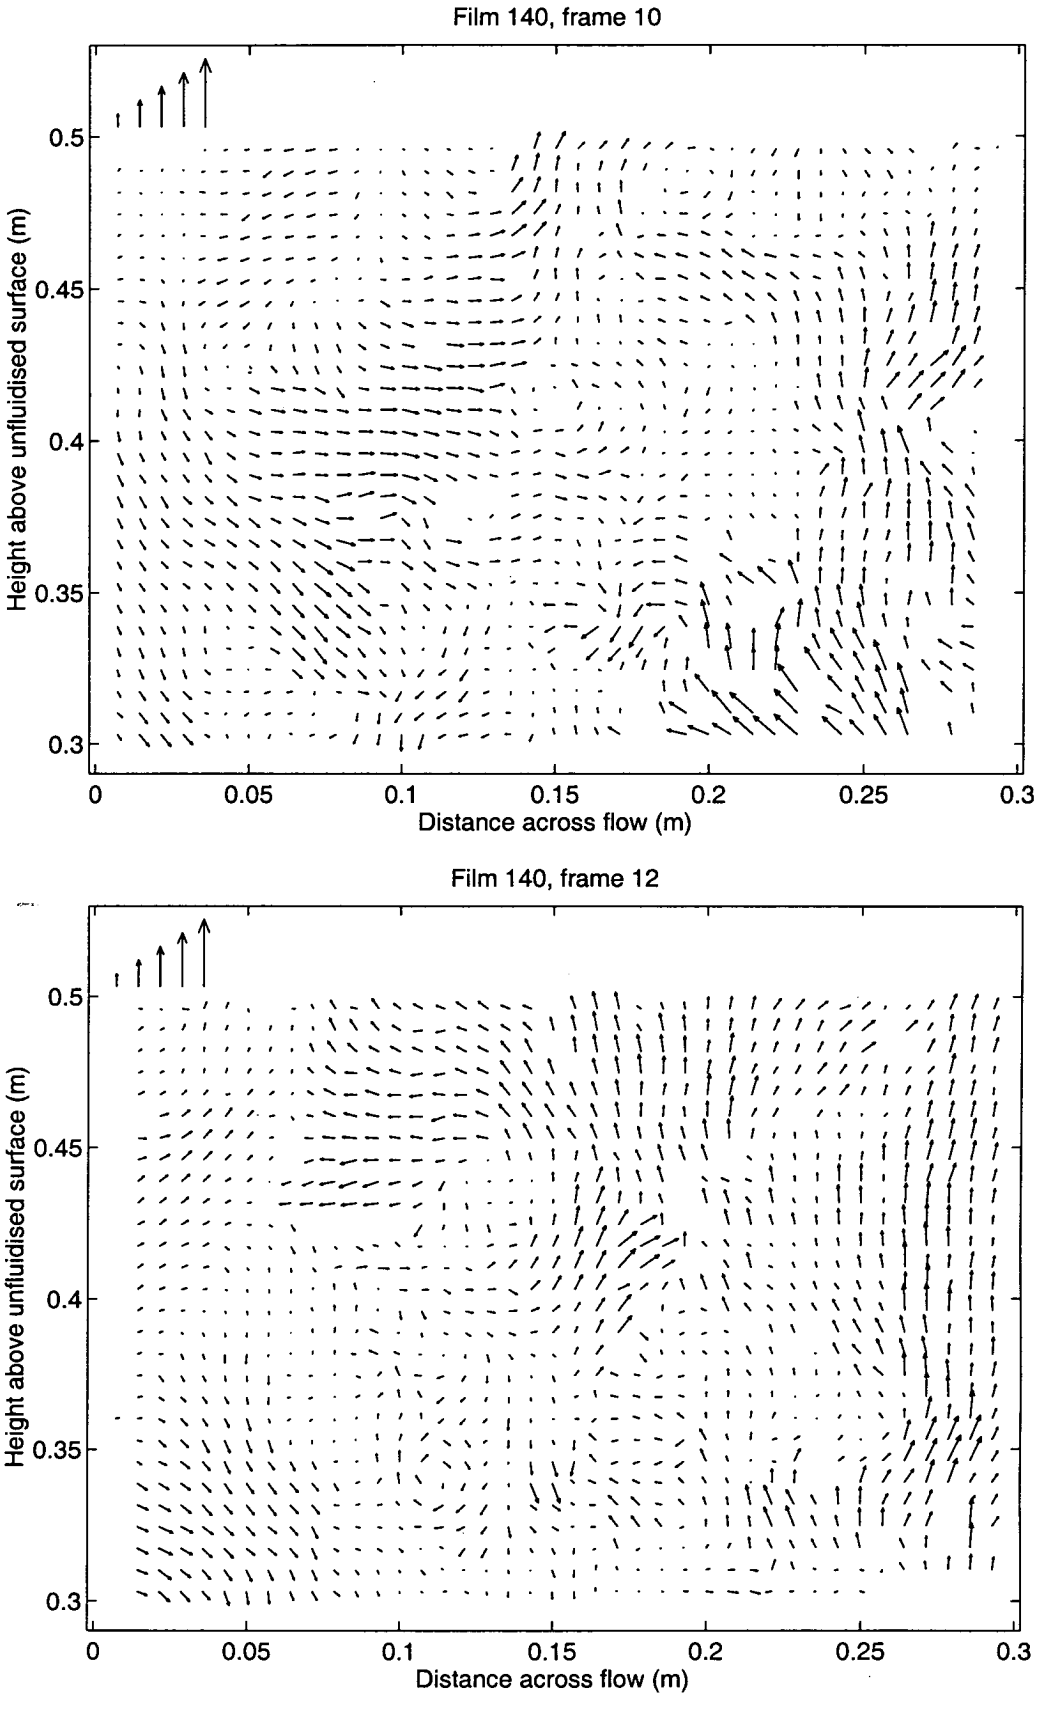


Figure 9.38: PIV vector maps for the flow about 40 cm above a 60 cm-deep bed of particles, after subtracting the image-shifting velocity. $U \approx 1.31\, U_{mf}$.

The run at $U = 1.31\,U_{mf}$ gave consistently good results. The yields of valid vectors were 85–94%. The flow was predominantly upward in frame 4 (Figure 9.37, top) although the left-hand third of the flow appeared to be mainly horizontal or downward. A large recirculating region was observed in the left-hand corner of frame 8 (Figure 9.37, bottom). In frame 10, (Figure 9.38, top), the flow was barely upward overall, with a section-average vertical velocity of only $0.0096\,\mathrm{m\,s^{-1}}$ (about $0.1\,U_{mf}$). A pair of velocity histograms for frame 12 are given in Figure 9.39, which can be compared with the vector map in the lower half of Figure 9.38.

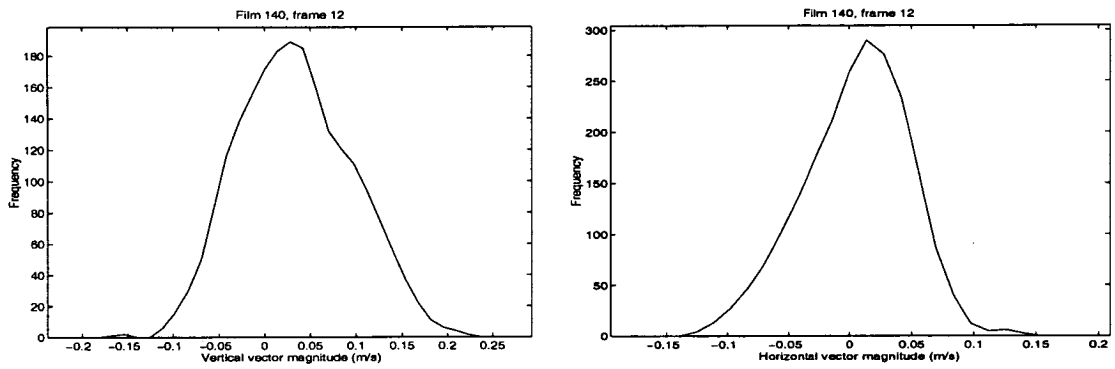


Figure 9.39: Histograms of the vertical (left) and horizontal (right) distributions of velocity in frame 12 of film 140, where $U = 0.0917\,\mathrm{m\,s^{-1}} \approx 1.31\,U_{mf}$.

The statistics given in Table 9.16 show the range of velocities under these conditions. The greatest velocity components in the horizontal and downward directions were usually rather large. Typical values were an average upward velocity of $0.0285\,\mathrm{m\,s^{-1}}$ (averaging across the section), a greatest upward velocity of about $0.22\,\mathrm{m\,s^{-1}}$ and a greatest downward velocity of about $0.13\,\mathrm{m\,s^{-1}}$.

A problem with using image-shifting is that a large horizontal displacement needs to be subtracted from the results obtained by autocorrelation analysis, in order to obtain the object-plane velocities. However, if the velocity is calculated as the difference of two large velocities, the relative error in each measurement may be large. The magnitude of the flow displacements can be increased by increasing the illumination time interval, but this will also increase the image-shifting displacement. As the magnitude of the overall displacement increases, the yield of valid vectors is reduced if the same analysis parameters are used. Large interrogation areas can be used, but only if the effects of velocity gradients

Frame	Yield (%)	Mean v_y ($m\,s^{-1}$)	Max v_y ($m\,s^{-1}$)	Min v_y ($m\,s^{-1}$)	Mean v_x ($m\,s^{-1}$)	Max v_x ($m\,s^{-1}$)
f140f4	91	0.0353	0.21	−0.07	−0.0007	−0.14
f140f5	90	0.0400	0.15	−0.07	−0.0124	−0.15
f140f6	87	0.0205	0.20	−0.10	+0.0032	+0.08
f140f7	85	0.0234	0.24	−0.13	−0.0100	−0.17
f140f8	88	0.0129	0.20	−0.18	−0.0017	−0.17
f140f9	92	0.0297	0.20	−0.15	−0.0155	−0.16
f140f10	92	0.0096	0.23	−0.15	+0.0091	−0.18
f140f11	94	0.0417	0.22	−0.14	−0.0078	−0.17
f140f12	94	0.0334	0.23	−0.16	+0.0041	+0.15
f140f13	94	0.0478	0.26	−0.16	+0.0088	+0.18
f140f14	93	0.0196	0.24	−0.15	+0.0249	−0.15
mean:	91	0.0285	0.22	−0.13	+0.0002	~ 0.15

Table 9.16: Summary of velocity values for the first part of film 140, where the superficial velocity U was $0.0917\,m\,s^{-1}$ ($1.31\,U_{mf}$). The bed depth was 60 cm and the mean viewing height was 40 cm above the bed. Image-shifting was used.

are not too severe. In fact, 128×128 pixel interrogation areas were used to analyse the frames from film 140, without too many side-effects being observed.

In some cases at this viewing height, *e.g.* for the low-flowrate results in the first part of film 140, better results might have been obtained if the illumination time interval had been increased to about 7.8 ms and the rotational speed of the mirror had been reduced. A mirror speed reduction of 25%, when combined with an increase of 25% in the illumination time interval, would have resulted in no change to the image-shifting displacement, whilst the longer illumination time interval would have helped to increase the displacements due to the flow. However, the image-shifting apparatus could only produce rotational speeds of $10°\,s^{-1}$, $15°\,s^{-1}$, $20°\,s^{-1}$, *etc.*, so a rotational speed of 11–$12°\,s^{-1}$ was not possible.

At this low air flowrate, the main alternative was to drop down to $10°\,s^{-1}$ using the increased illumination time interval, in the hope that the resulting image-shifting displacements of about 17 pixels to the *left* would still be large enough to allow measurements of the largest displacements to the *right* (*e.g.* a horizontal displacement of 16 pixels to the right would appear to be 1 pixel to the left if recorded using image-shifting, and would not be measured if the vertical displacement was also low). The mirror rotational speed of $15°\,s^{-1}$ was a safer choice,

resulting in all image displacements being clearly to the left. Overall, the PIV parameters used here were probably the best choice for the flowrate of 495 l min^{-1}.

If photographs were to be taken at a somewhat lower flowrate, the results presented in Table 9.16 suggest it would be best to rotate the image-shifting mirror at 10° s^{-1} whilst using an illumination time interval of about 7.8 ms. At higher flowrates than 495 l min^{-1}, however, it would be worth considering a mirror rotational speed of 20° s^{-1}, an illumination time interval of about 4.1 ms or 5.5 ms and a camera exposure time of 1/60 s, so that either exactly three or exactly four images of each particle would be recorded (providing a shorter-duration snapshot of the rapidly-moving seeding particles).

The image-shifted run at the greater air flowrate (*i.e.* $U = 0.122\,m\,s^{-1}$) revealed an extremely turbulent flow. A sequence of four consecutive frames is given in Figures 9.40 and 9.41. It can be seen from these vector maps that the greatest velocities in the flow were comparable to the largest marker vectors in the top left-hand corner of each vector map, whereas the flow vectors appeared small in comparison in some of other runs in the centre of the bed (*e.g.* Figure 9.32, at a freeboard height of 2.2 m and a greater superficial velocity of $0.157\,m\,s^{-1}$).

In frame 18 (Figure 9.40, top), there were very large vertical velocities in the left-hand half of the vector map, whilst the flow in the right-hand half was strongly downward. The greatest vertical velocity was $0.50\,m\,s^{-1}$, which was more than $4\,U$. Large downward velocities were also obtained, reaching as high as $0.39\,m\,s^{-1}$. After interpolating to fill in missing vectors, the section-average velocity was found to be $+0.0031\,m\,s^{-1}$, *i.e.* the net flow was upward but only very mildly so.

In contrast, the flow in other frames was actually downward. The section-average vertical velocity was fairly small in frame 19 ($-0.0159\,m\,s^{-1}$). A section-average velocity of $-0.1130\,m\,s^{-1}$ was found in frame 20, where most of the vectors were downward. These and other velocity statistics are given in Table 9.17. Some histograms of the horizontal and vertical velocity are shown in Figure 9.42.

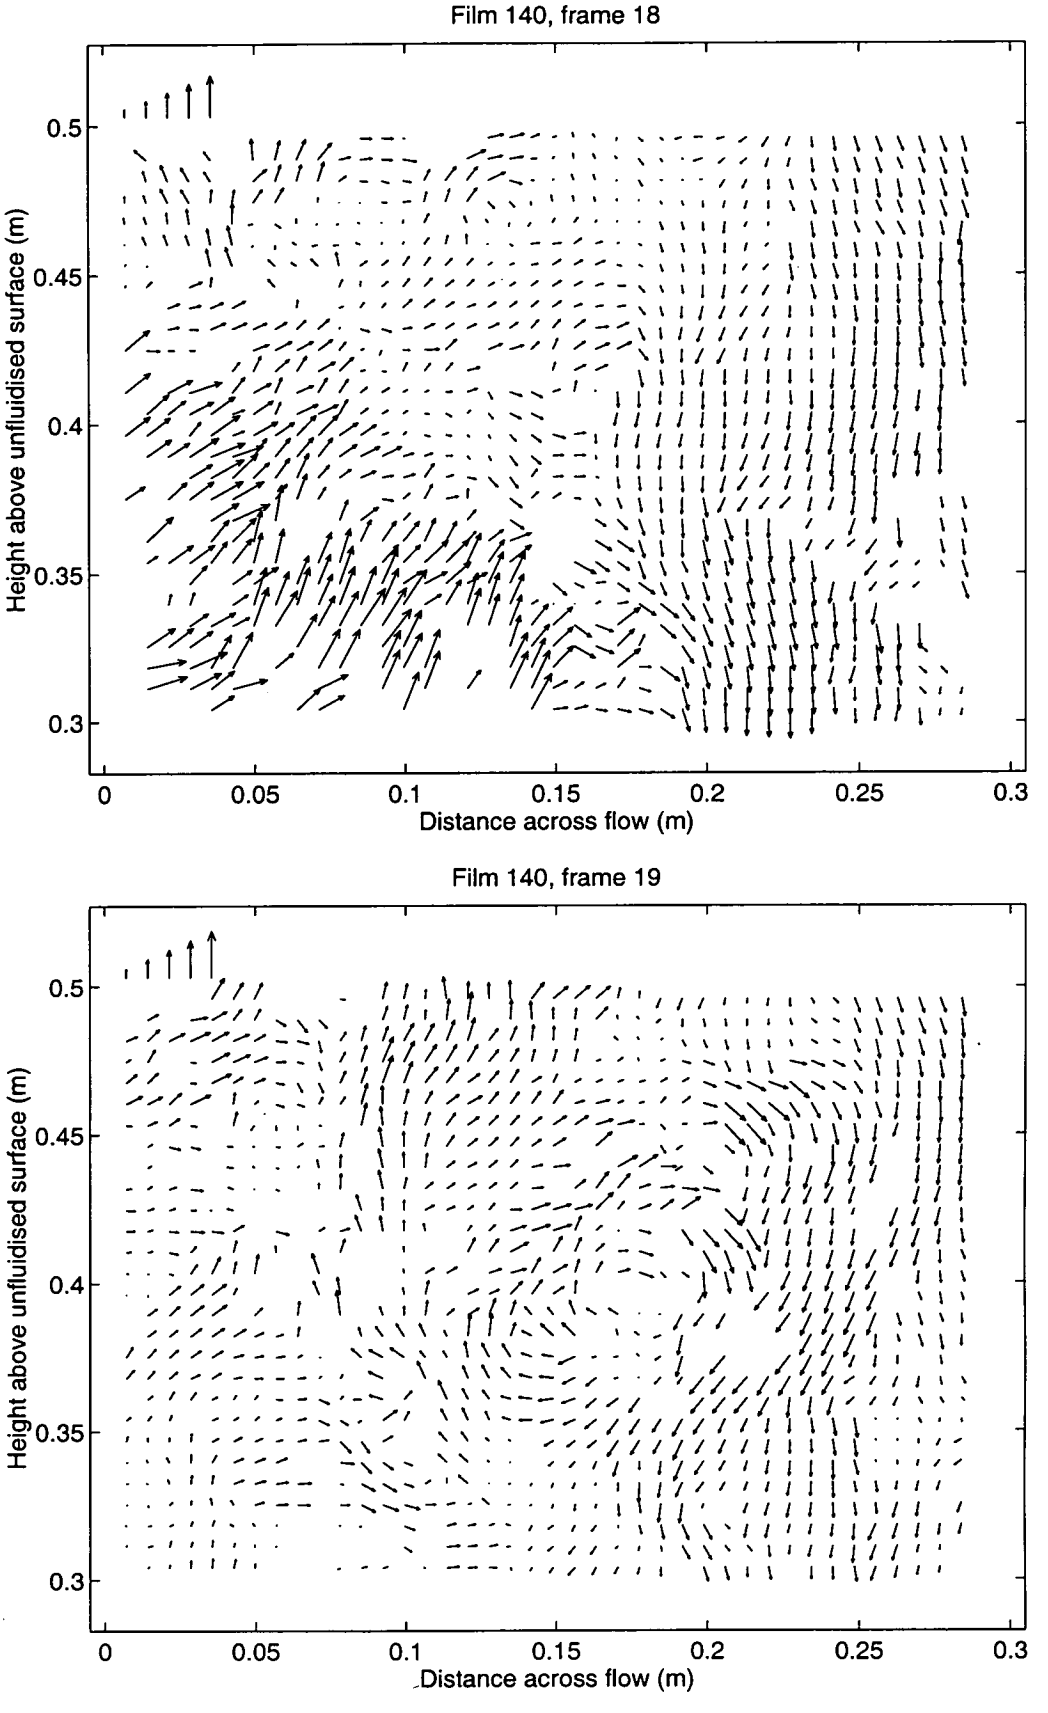


Figure 9.40: PIV vector maps for the flow about 40 cm above a 60 cm bed of particles, after subtracting the image-shifting velocity. $U \approx 1.75\, U_{mf}$.

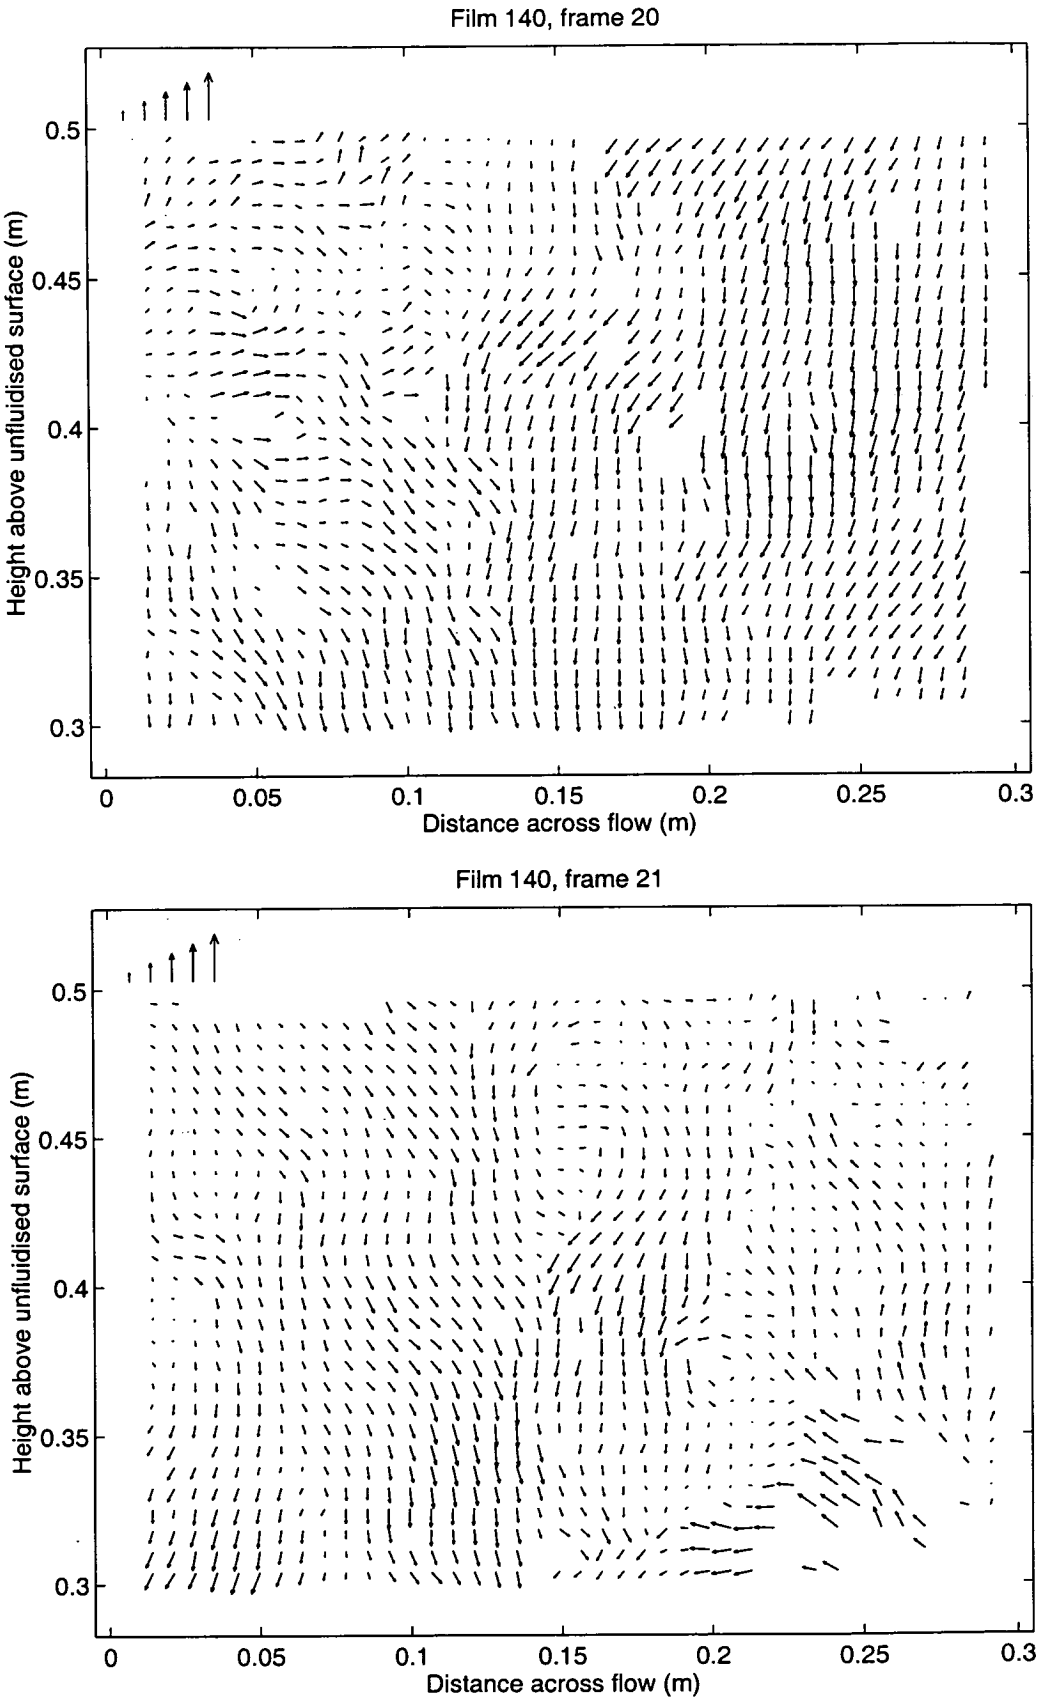


Figure 9.41: PIV vector maps for the flow about 40 cm above a 60 cm bed of particles, after subtracting the image-shifting velocity. $U \approx 1.75\, U_{mf}$.

Frame	Yield (%)	Mean v_y ($\mathrm{m\,s^{-1}}$)	Max v_y ($\mathrm{m\,s^{-1}}$)	Min v_y ($\mathrm{m\,s^{-1}}$)	Mean v_x ($\mathrm{m\,s^{-1}}$)	Max v_x ($\mathrm{m\,s^{-1}}$)
f140f17	83	−0.0067	0.31	−0.29	+0.0597	+0.29
f140f18	80	+0.0031	0.50	−0.39	+0.0816	+0.49
f140f19	84	−0.0159	0.31	−0.27	+0.0231	+0.23
f140f20	92	−0.1130	0.19	−0.29	+0.0035	+0.20
f140f21	89	−0.0559	0.21	−0.26	−0.0031	−0.23
f140f22	88	−0.1027	0.24	−0.38	+0.0665	+0.32
f140f23	76	−0.0175	0.32	−0.34	+0.0627	+0.32
f140f24	57	−0.0404	0.50	−0.50	+0.1230	+0.45
f140f25	83	−0.0182	0.26	−0.24	+0.0582	+0.30
f140f26	87	−0.0281	0.18	−0.19	+0.0519	+0.20
f140f27	73	−0.0089	0.32	−0.27	+0.0909	+0.40
mean:	81	−0.0367	0.30	−0.31	+0.0562	~ 0.31

Table 9.17: Summary of velocities for the second part of film 140, where the superficial velocity $U = 0.122\,\mathrm{m\,s^{-1}} = 1.75\,U_{mf}$. The bed depth was 60 cm and the mean viewing height was 40 cm above the bed surface. Image-shifting was used.

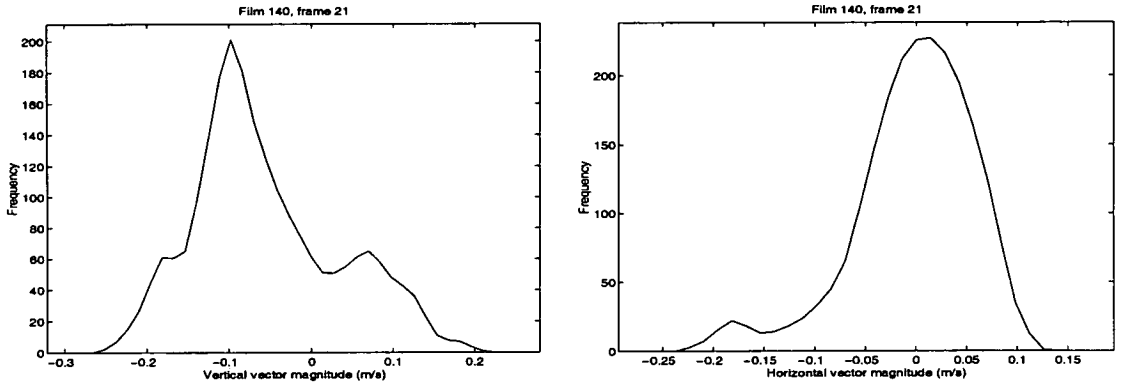


Figure 9.42: Histograms of the distribution of the vertical (left) and horizontal (right) velocities from film 140 frame 21, where $U = 0.122\,\mathrm{m\,s^{-1}} \approx 1.75\,U_{mf}$.

From the 11 frames taken at this air flowrate, the mean section-average vertical velocity was $-0.0367\,\mathrm{m\,s^{-1}}$. The magnitudes of the greatest upward, downward and horizontal speeds were comparable, with similar ranges (between $0.18\,\mathrm{m\,s^{-1}}$ and $0.50\,\mathrm{m\,s^{-1}}$) and similar mean values ($\pm 0.31\,\mathrm{m\,s^{-1}}$). It was curious that the increase in air flowrate (compared to the first part of film 140) led to a considerable *decrease* in the section-average vertical velocities measured by PIV.

However, it is not certain how valid these maximum and minimum values are. The vector maps showed high levels of turbulence, so it was difficult to validate the vector maps completely. Interpolation was used to fill in the missing vectors so

a number of extremely large vectors could have been generated artificially, either by interpolating close to spurious large vectors which had not previously been deleted or by extrapolating from a region of high velocity gradient into a large region where no valid vectors had been found. Indeed, the frames with the most extreme vectors often had a poor yield of valid vectors (*e.g.* frames 24 and 27).

Nevertheless, it is clear that some very large velocities are present, in all directions. Also, frame 20 demonstrated that it was not even valid to assume that most of the vectors in the flow would be upward (Figure 9.41, top). The implication must be that large upward velocities would have been found at this bed depth and at this viewing height: (a) at the same central position in the bed, if rather more PIV photographs had been taken under the same conditions, or (b) at different positions in the bed at the same time as the above photographs were taken (*e.g.* if there was a lot of bubbling activity near the four corners of the bed).

As mentioned earlier, these results might have been improved if they had been repeated using a greater image-shifting mirror rotational speed. Alternatively, it would have been better to have made the rotating mirror turn in the opposite direction: because the horizontal velocities at this flowrate appeared to be mainly to the right, it might be better to image-shift to the right rather than to the left.

9.4 60 cm bed, quarter of way in

Using the 60 cm-deep fluidised bed, PIV photographs were also taken of the flow one quarter of the way into the freeboard (as listed in Table 9.18). However, no vector maps or tables of velocity statistics from the greater freeboard heights will be presented here, because these films have not yet been re-analysed using 128×128 interrogation areas. Based on old analyses using 63×63 interrogation areas, the results seem qualitatively similar to those in the centre of the freeboard.

Bed depth (cm)	Viewing height (cm)	U/U_{mf}	Film number	Image-shifting used?
60	190	1.31	116	Yes
60	190	1.75	115	No
60	190	2.25	111	No
60	130	1.18	130	Yes
60	130	1.75	128	Yes
60	130	2.25	131	No
60	40	1.31	136	Yes
60	40	1.75	136	Yes

Table 9.18: Details of films taken using PIV to record the flow one quarter of the way into the freeboard space, when the bed depth was 60 cm.

9.4.1 Viewing heights of about 190 cm

At the greatest viewing height, the run with a superficial velocity of $1.31\,U_{mf}$ (film 116) produced a flow with the greatest vertical velocities near the walls. The velocities in the centre of the section were usually upward, but were often almost zero. The downward velocities observed in the centre of the bed (film 108; $U = 1.22\,U_{mf}$) were not found here.

The run with $U = 1.75\,U_{mf}$ (film 115) was very similar to the corresponding run in the centre (film 107). The largest velocities were usually found near the walls and the velocities elsewhere were very clearly upward (although smaller). Slow-rising low-velocity regions were a problem in some frames, so a very small horizontal image-shifting displacement (*e.g.* 3–4 pixels) would have been a useful addition when measuring this flow, in an ideal world.

Some problems were encountered at the highest flowrate, where the image displacements were large and some bed particle images were found in the flow records. Perhaps 128 × 128 interrogation areas would have given better results.

9.4.2 Viewing heights of about 130 cm

At the intermediate viewing height, the lowest flowrate gave $U = 1.18\,U_{mf}$. The largest velocities were always found near the walls and some downward velocities were measured near the centre of the section in all frames (film 130).

Image-shifting was also used when $U = 1.75\,U_{mf}$ (film 128). The largest velocities were usually found near the walls and no downward velocities were observed.

At the largest flowrate, the problems encountered were similar to those at the greatest viewing height. Larger interrogation areas would probably help (by analogy to the corresponding results in the centre of the freeboard, in film 119).

9.4.3 Viewing heights of about 40 cm

Some image-shifted photographs were taken of the chaotic flow 40 cm above the bed surface, on film 136. In the first part of this film, 11 frames were taken using a superficial velocity of $0.0917\,\mathrm{m\,s^{-1}}$ ($1.31\,U_{mf}$). Following that, a further 10 pictures were recorded using $U = 0.122\,\mathrm{m\,s^{-1}} = 1.75\,U_{mf}$.

In both cases, the rotating mirror revolved at $10^\circ\,\mathrm{s^{-1}}$ and the illumination time interval was 5 ms. Up to six images of each particle were recorded in the 1/30 s exposure time. As would be expected following the corresponding results in the centre of the freeboard (film 140), these parameters were highly suitable for the lower flowrate run but were not really adequate when the higher flowrate was used. With the benefit of hindsight, the higher-flowrate run would have yielded better results if it had used a mirror rotational speed of $20^\circ\,\mathrm{s^{-1}}$, an illumination time interval of about 3 or 4 ms and an exposure time of 1/60 of a second.

At the lower flowrate, an illumination time interval of 6.2 or 7.8 ms would have been better, to allow larger displacements to be recorded. Nevertheless, the results obtained were extremely good. For example, the average yield of valid vectors was 98% (as indicated in Table 9.19). It was found that the section-average velocity was always upward, but its magnitude was substantially less than the superficial

velocity. Comparing these results to the data from the centre of the freeboard (film 140, first part), it appears that the section-average vertical and horizontal velocities are similar, but wider ranges of velocities were found in the centre. For example, the magnitudes of the greatest upward, downward and horizontal velocities here were only half those in the centre of the freeboard.

Frame	Yield (%)	Mean v_y ($m\,s^{-1}$)	Max v_y ($m\,s^{-1}$)	Min v_y ($m\,s^{-1}$)	Mean v_x ($m\,s^{-1}$)	Max v_x ($m\,s^{-1}$)
f136f5	97	0.0177	0.09	−0.05	−0.0031	±0.06
f136f6	97	0.0245	0.12	−0.05	−0.0005	±0.06
f136f7	100	0.0253	0.11	−0.07	+0.0037	+0.06
f136f8	97	0.0239	0.10	−0.07	+0.0061	+0.08
f136f9	97	0.0270	0.10	−0.07	+0.0039	+0.08
f136f10	99	0.0208	0.10	−0.06	+0.0002	+0.06
f136f11	99	0.0206	0.12	−0.07	−0.0024	±0.06
f136f12	100	0.0273	0.13	−0.08	+0.0055	+0.09
f136f13	99	0.0324	0.14	−0.07	+0.0117	+0.10
f136f14	99	0.0217	0.12	−0.07	−0.0054	−0.08
f136f15	99	0.0257	0.10	−0.07	−0.0054	−0.08
mean:	98	0.0243	0.11	−0.07	+0.0013	~ 0.07

Table 9.19: Summary of velocities for the first part of film 136, where the superficial velocity $U = 0.0917\,m\,s^{-1}$ ($1.31\,U_{mf}$). The bed depth was 60 cm. The mean viewing height was 40 cm above the bed surface, one quarter of the way into the freeboard. Image-shifting was used.

The vector maps from the first part of film 136 revealed a considerable amount of variation from frame to frame. At least one recirculating region could be seen in every vector map. A pair of counter-rotating vortices was found on the right-hand side of frame 5 (Figure 9.43, top), but it is thought that this is the result of a jet of rapidly-rising air rather than a ghost bubble (because the flow speeds are reasonably constant below, in between and above the vortex pair, rather than greatest in between the vortices). Frame 9 captured a large, single vortex near the right-hand wall and also included a number of slow-moving regions (Figure 9.43, bottom). Frame 11 was rather messy, with most of the upflow occurring in jets near the left-hand wall and near the centre of the section (Figure 9.44, top). A reasonable amount of downward flow took place near the right-hand wall in frame 14, which also had a vortex near the centre of the field of view (Figure 9.44, bottom). Analysis with 128 × 128 interrogation areas gave good results.

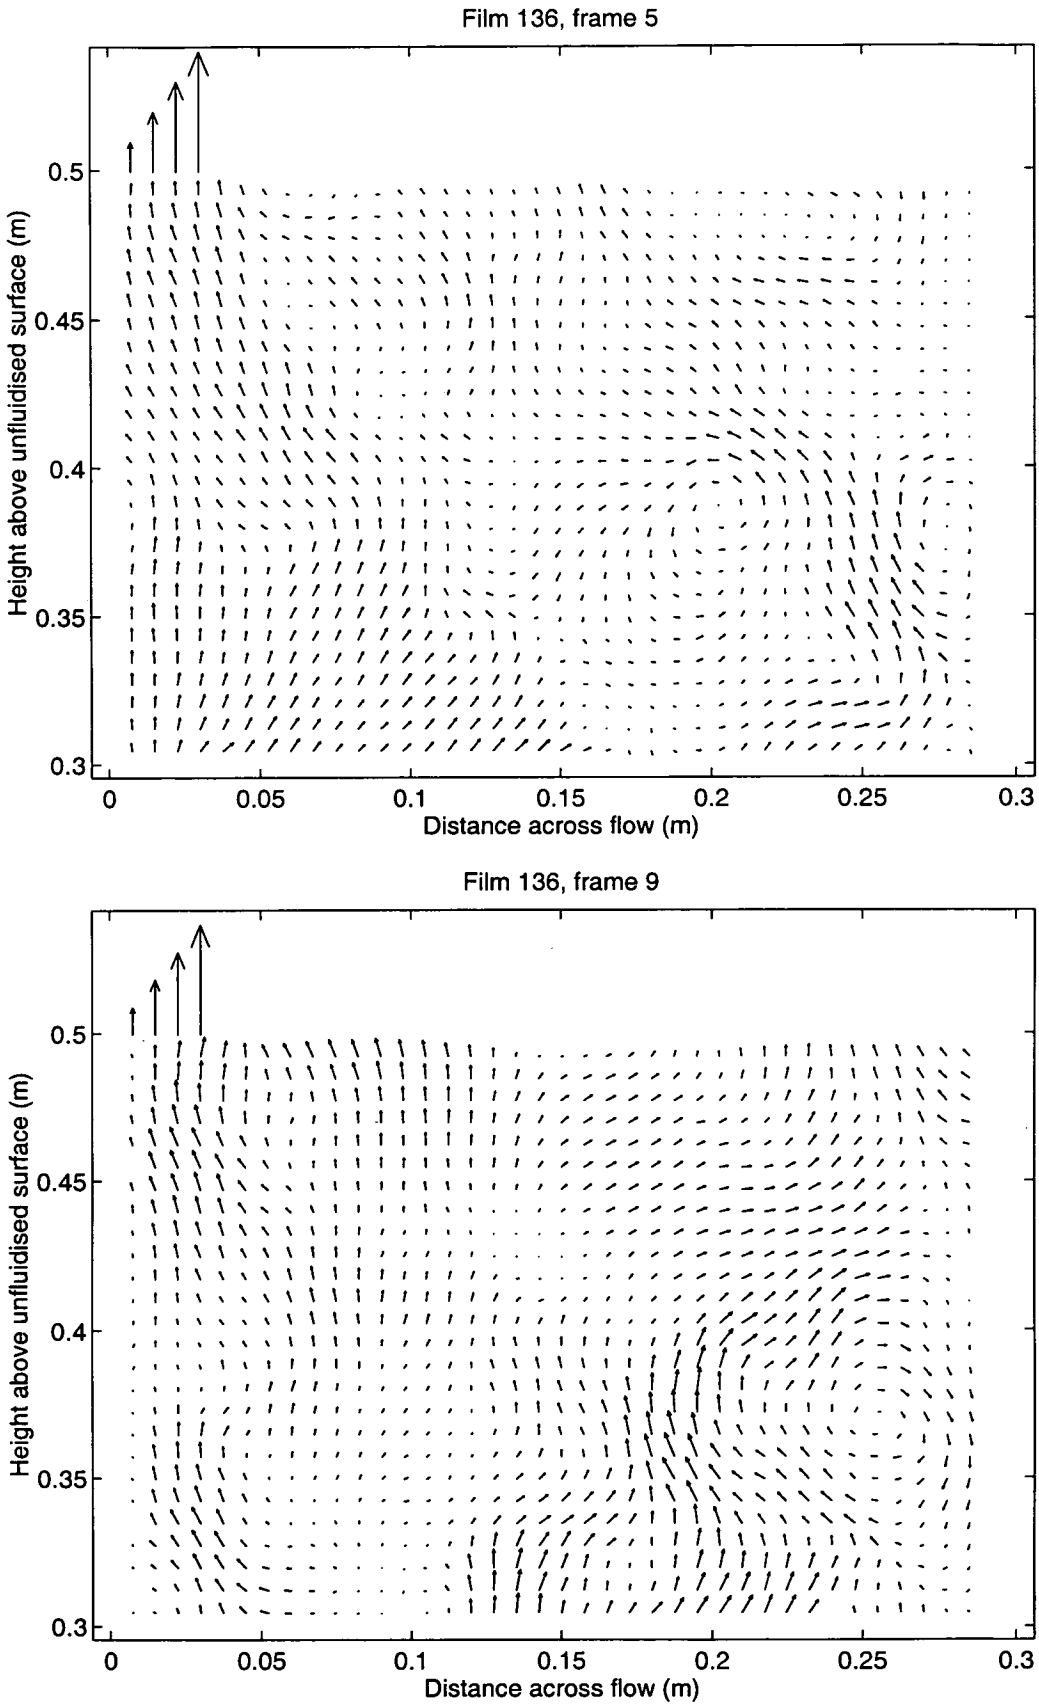


Figure 9.43: PIV vector maps for the flow about 40 cm above a 60 cm bed of particles, after subtracting the image-shifting velocity. $U \approx 1.31\, U_{mf}$.

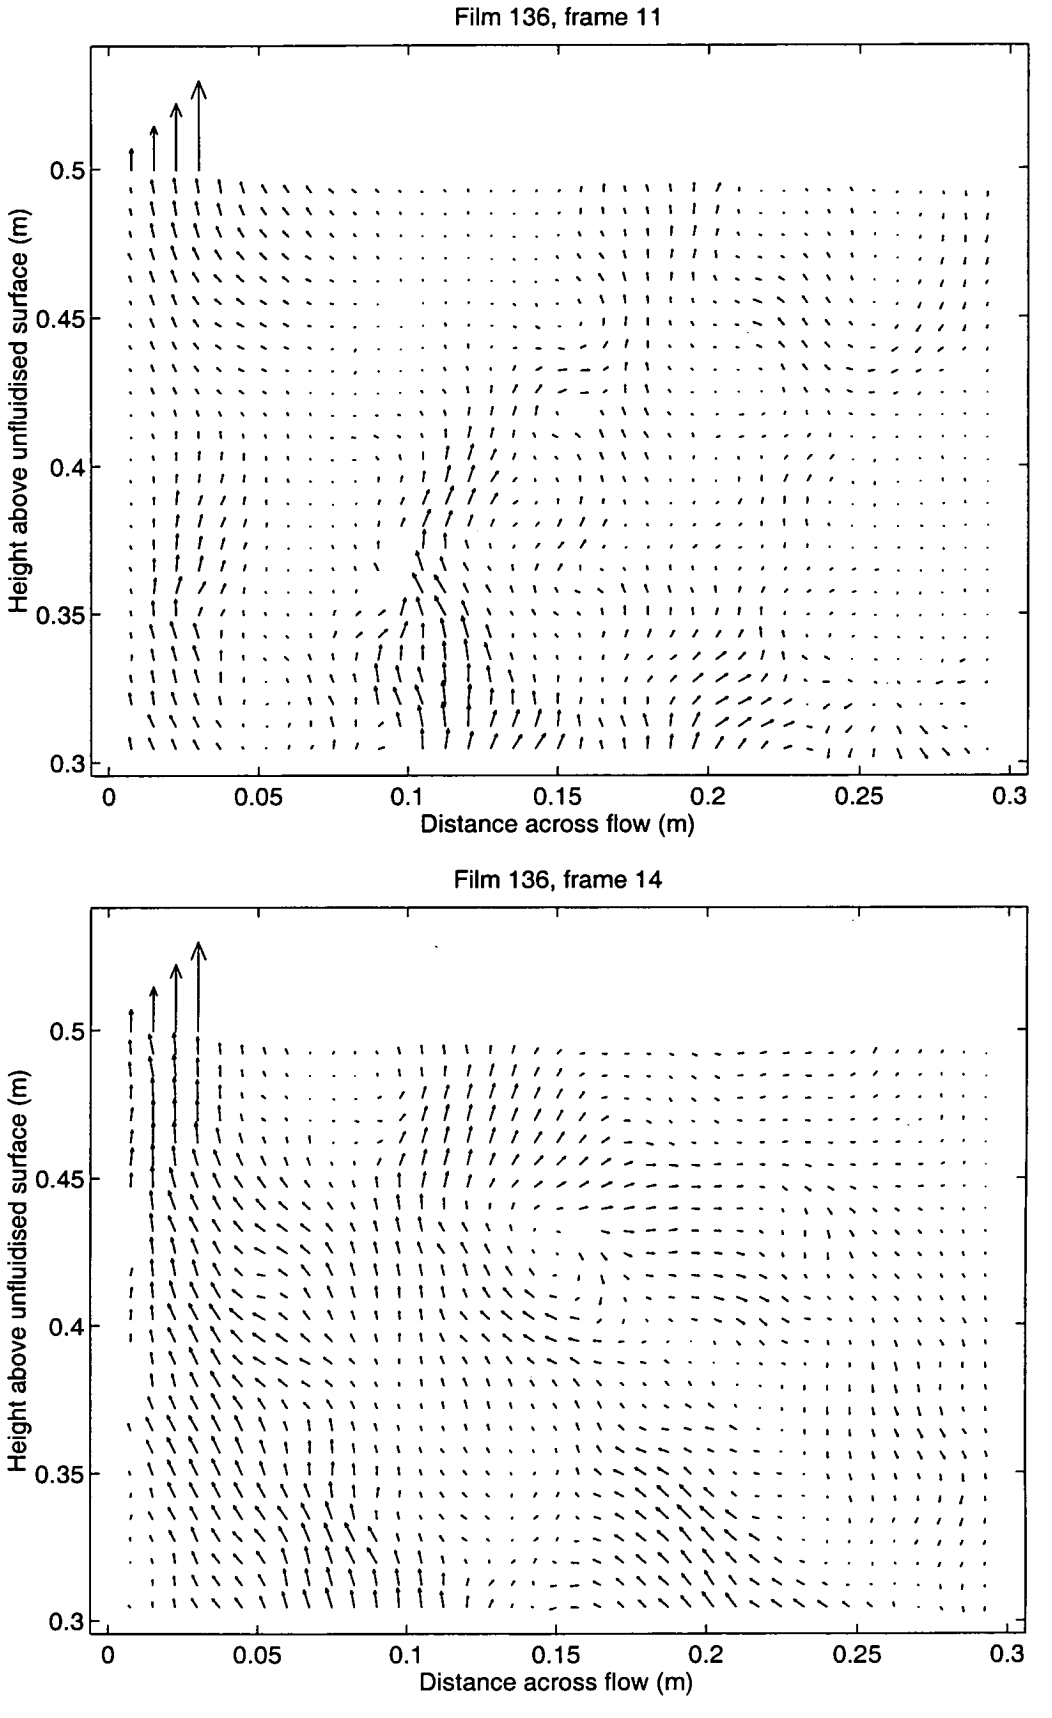


Figure 9.44: PIV vector maps for the flow about 40 cm above a 60 cm bed of particles, after subtracting the image-shifting velocity. $U \approx 1.31\, U_{mf}$.

Frame	Yield (%)	Mean v_y ($m\,s^{-1}$)	Max v_y ($m\,s^{-1}$)	Min v_y ($m\,s^{-1}$)	Mean v_x ($m\,s^{-1}$)	Max v_x ($m\,s^{-1}$)
f136f17	90	+0.0545	0.40	−0.28	−0.0238	−0.28
f136f18	79	+0.1139	0.60	−0.25	−0.0226	±0.30
f136f19	76	+0.0623	0.43	−0.25	+0.0222	+0.21
f136f20	71	−0.0032	0.30	−0.35	+0.0484	+0.48
f136f21	49	+0.0515	0.48	−0.30	+0.0737	+0.45
f136f22	83	+0.0369	0.55	−0.35	−0.0140	±0.27
f136f23	90	+0.0384	0.35	−0.30	−0.0355	+0.35
f136f24	89	+0.0428	0.50	−0.42	−0.0006	−0.27
f136f25	82	+0.0161	0.40	−0.23	−0.0016	±0.24
f136f26	86	+0.0305	0.60	−0.29	+0.0039	+0.30
mean:	83	+0.0444	0.46	−0.30	+0.0050	~ 0.31

Table 9.20: Summary of velocities for the second part of film 136, where the superficial velocity $U = 0.122\,m\,s^{-1} = 1.75\,U_{mf}$. The bed depth was 60 cm and the mean viewing height was 40 cm above the bed surface, one quarter of the way into the freeboard. Image-shifting was used.

The results in the second part of film 136 were less impressive, because the image-shifting parameters were not optimal. One practical difficulty was that the horizontal image-shift was not large enough, so some manual vector alignment was required before the image-shifting displacements could be subtracted. Some of the results obtained are summarised in Table 9.20. The average yield of valid vectors was 83%; this figure ignores the poor yield of frame 21, however.

In line with the corresponding results in the centre of the freeboard, it was found that the increase in air flowrate caused the magnitude of the greatest horizontal, upward and downward vectors to increase considerably. One interesting difference was that the increase in flowrate caused the section-average vertical velocities in the centre of the freeboard to become downward, whereas the corresponding change here, nearer the front wall, was an increase to a more upward value. In all four cases, however, the section-average vertical velocity was substantially lower than the superficial velocity. These might be freak results made possible by small sample sizes, or they might indicate that the most upward velocities are in a region not yet investigated (*e.g.* very close to the walls and corners, or in the rear half of the freeboard), especially at the higher flowrate.

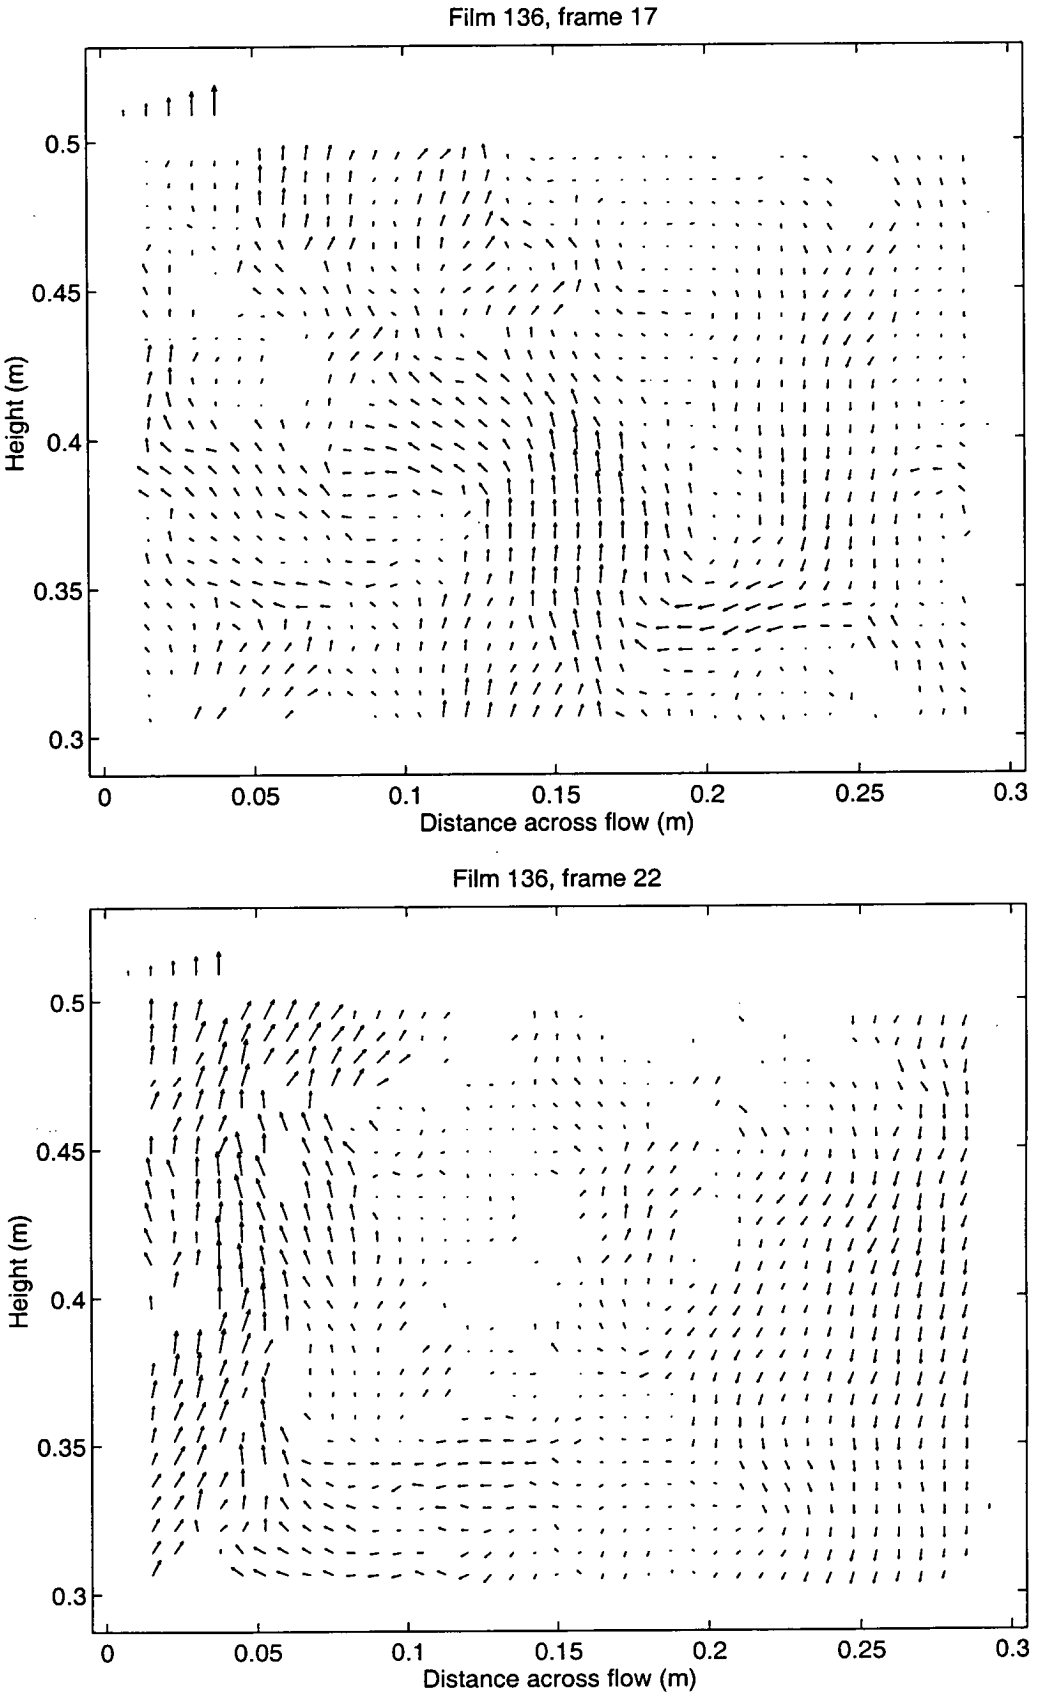


Figure 9.45: PIV vector maps about 40 cm above a 60 cm bed of particles using image-shifting, one quarter of the way into the freeboard. $U \approx 1.75\,U_{mf}$.

Frame 17 is given as a sample vector map in Figure 9.45 (top). The flow was rather messy, but a good yield of vectors was obtained. Frame 21 (not shown) included a very bright region in its bottom right corner, as a result of a bubble eruption. (This was one of the reasons that lower freeboard heights were not investigated.) The high particle density, possibly helped by a small amount of damage to the negative, meant that the image-separation program PIVSPLIT crashed after processing the top 60% of the bitmap image (hence the low yield of valid vectors). It was curious that the flow was strongly *downward* near the right-hand wall in frame 22, taken a short time later (Figure 9.45, bottom). Bubble eruptions were expected to eject particles and fluidisation gas up into the freeboard, at speeds greater than U_{mf}. However, recent work by Yórquez & Duursma (1999) has also found large downward velocities at low freeboard heights, immediately after bubbles burst.

9.5 Ensemble averages

A summary of the vertical velocity results from the most successful experiments in the larger apparatus is given in Table 9.21.

Ensemble-average vector maps have been calculated for each run, producing time-averaged velocities at each point. Interpolation was used first, to estimate values for the missing vectors.

The averaged vector maps fell into two main classes. Looking through the centre of the freeboard space about 1.3 m above a 31 cm bed, the ensemble-averaged velocities when $U = 1.27\,U_{mf}$ (film 133; Figure 9.46, top) revealed peaks near the walls and lowest velocities near the centre of the section. In the equivalent run using $U = 1.69\,U_{mf}$ (film 80), the vertical velocity was fairly constant across the section (Figure 9.46, bottom). (A similar picture was seen in the ensemble-averaged run when $U = 2.67\,U_{mf}$ (films 77 and 78).)

Viewing height (cm)	Bed depth (cm)	Superficial velocity ($m\,s^{-1}$)	Mean v_y ($m\,s^{-1}$)	Typical max v_y ($m\,s^{-1}$)	Typical min v_y ($m\,s^{-1}$)	Image-shifting used?	Film no.
220	31 C	0.089	0.036	0.15	−0.038	Yes	105
220	31 C	0.082	—	—	—	No	90
220	31 C	0.118	0.086	0.13	+0.044	No	88
220	31 C	0.157	0.130	0.19	+0.066	No	86
130	31 C	0.089	0.045	0.14	+0.003	Yes	133
150	31 C	0.119	0.090	0.17	+0.030	No	80
150	31 C	0.187	0.152	0.26	+0.061	No	77, 78
40	31 C	0.108	0.061	0.40	−0.24	No	71
220	31 Q	0.089	0.061	0.14	−0.019	Yes	102
220	31 Q	0.108	0.079	0.14	+0.028	No	93
220	31 Q	0.136	0.087	0.14	+0.041	No	91
130	31 Q	0.079	0.050	0.15	−0.007	Yes	132
140	31 Q	0.125	0.108	0.18	+0.040	No	84
140	31 Q	0.157	0.140	0.24	+0.050	No	85
190	60 C	0.085	0.020	0.09	−0.032	Yes	108
190	60 C	0.122	0.076	0.13	+0.039	No	107
190	60 C	0.165	0.116	0.17	+0.062	No	118
130	60 C	0.093	0.002	0.13	−0.057	Yes	121
130	60 C	0.130	0.095	0.16	+0.040	Yes	124
130	60 C	0.172	0.133	0.27	+0.036	No	119
40	60 C	0.092	0.029	0.22	−0.13	Yes	140
40	60 C	0.122	−0.037	0.30	−0.31	Yes	140
190	60 Q	0.092	—	—	—	Yes	116
190	60 Q	0.122	—	—	—	No	115
190	60 Q	0.158	—	—	—	No	111
130	60 Q	0.083	—	—	—	Yes	130
130	60 Q	0.122	—	—	—	Yes	128
130	60 Q	0.158	—	—	—	No	131
40	60 Q	0.092	0.024	0.11	−0.07	Yes	136
40	60 Q	0.122	0.044	0.46	−0.30	Yes	136

Table 9.21: Summary of velocity statistics for the main runs on the larger fluidised bed. The letter 'C' after the bed depth indicates a photograph of the centre of the test section, whilst 'Q' denotes a run one quarter of the way into the test section.

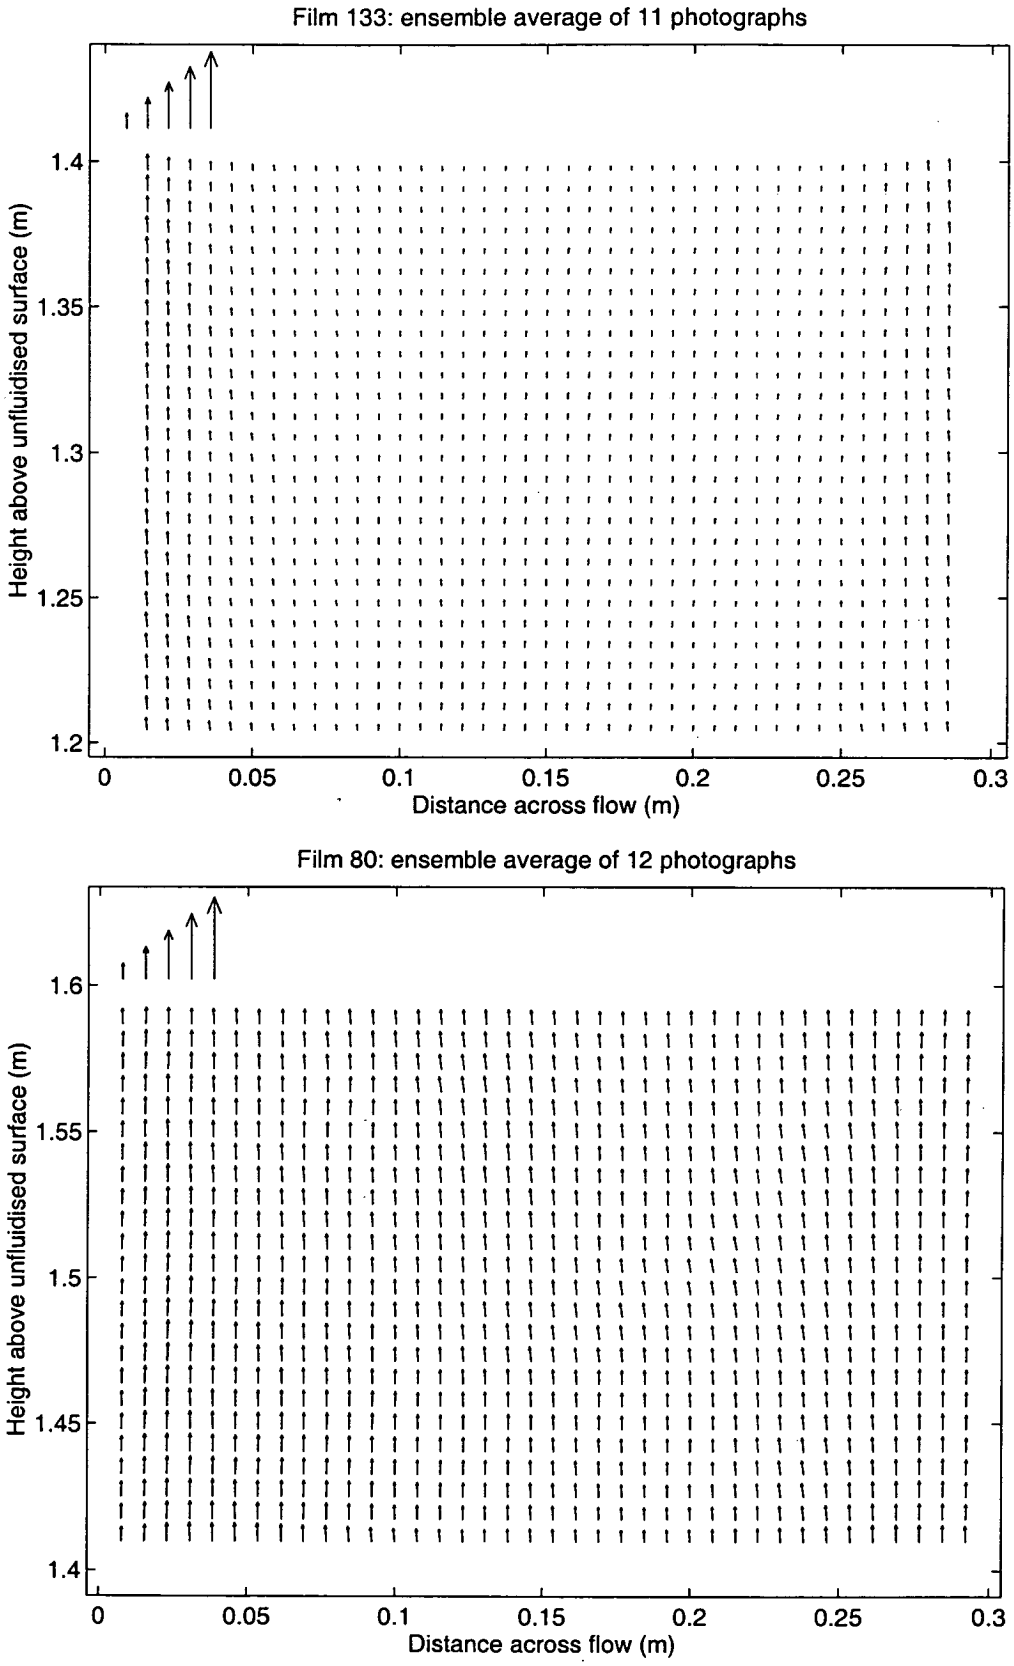


Figure 9.46: Ensemble-averaged vector maps for the flow 1.3 m above a 31 cm fluidised bed, looking through the centre of the freeboard. $U = 0.0889\,\mathrm{m\,s^{-1}} = 1.27\,U_{mf}$ in the upper vector map and $0.118\,\mathrm{m\,s^{-1}}$ $(1.69\,U_{mf})$ in the lower one.

It was found earlier that most of the runs at low superficial velocity ($U < 0.1\,\mathrm{m\,s^{-1}}$ or so) produced flows with some downward-moving regions in every frame (*e.g.* film 105, taken 2.2 m above a 31 cm bed, in the centre of the freeboard). When the frames from these films were ensemble-averaged, the time-averaged velocities were upward throughout the flow. This means that the downward-moving regions could not have been stationary; some points might have had downward velocities in 30% of the frames (say) and upward velocities in the remaining 70% of the frames, so that the time-average velocity was upward.

Velocity profiles could be calculated from this ensemble-averaged data. The graphs given in Figure 9.47 contain data from the centre of the freeboard, above a 31 cm bed. In each graph, the vertical component of the mean velocity was plotted against distance across the test section, for six evenly-spaced heights within each average vector map. This method produces a reasonable scatter of points, giving an indication of the variations in vertical velocity without cluttering up the graphs with over 1000 data points. Alternative methods were (a) to plot the vertical velocity profile for every row of averaged vectors and (b) to use the velocities obtained by averaging every *column* of ensemble-averaged vectors.

In the lower graphs in Figure 9.47, the 'dip' velocity profile was clearly apparent in the run at the lowest flowrate used under these conditions (film 133). The profile for film 80 was much flatter, but signs of a very slight dip were still present. In the run at the highest air flowrate (not shown), the vertical velocities were once again rather constant across the section, but slightly reduced values were found near the walls. In all cases, the velocities in the centre of the section were significantly lower than those predicted by a one-seventh power law model for turbulent flow in a cylindrical pipe (plotted using continuous curves).

For the flow in the centre of the freeboard 2.2 m above the surface of a 31 cm bed, the velocity profiles told a similar story (Figure 9.47, top). At the lowest flowrate tested at this height, the average vertical velocities near the centre of the section were close to zero (or even slightly downward). The dip profile had a lower minimum than was found at the same flowrate at the lower viewing position.

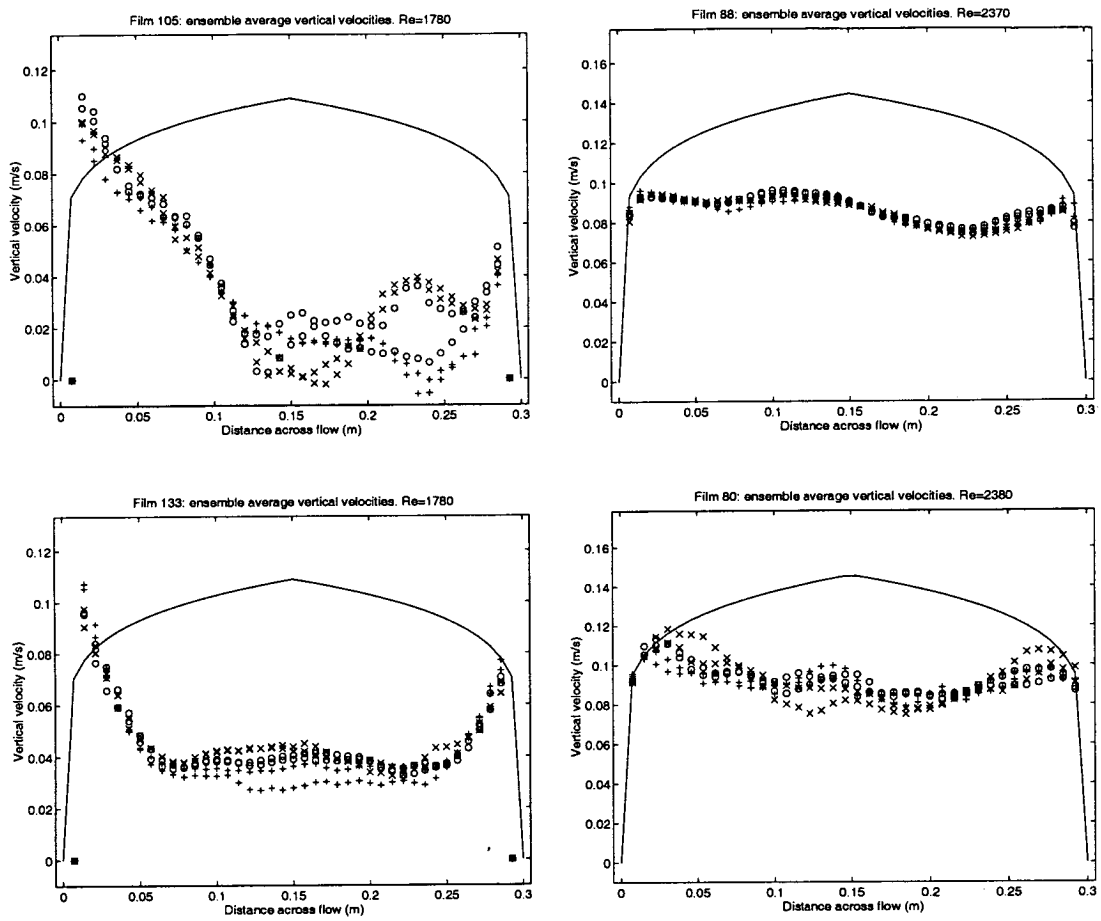


Figure 9.47: Profiles of vertical velocities along horizontal sections through the centre of the freeboard above a 31 cm deep fluidised bed, at heights of 2.2 m (top) and 1.3 m (bottom) above the bed surface. The superficial gas velocity was 0.0889 m s^{-1} in the left-hand figures and 0.118 m s^{-1} on the right.

Similar observations can be made for the velocity profiles at the greater viewing heights in the centre of the freeboard above a 60 cm bed (Figure 9.48) and also for the velocity profiles one quarter of the way into the freeboard space above a 31 cm bed (Figure 9.49). However, the flow 0.4 m above the 60 cm bed included an unexplainably-large number of downward velocities, as had been suggested by the individual vector maps and the velocity statistics for this run (Figures 9.40 and 9.41; Table 9.17).

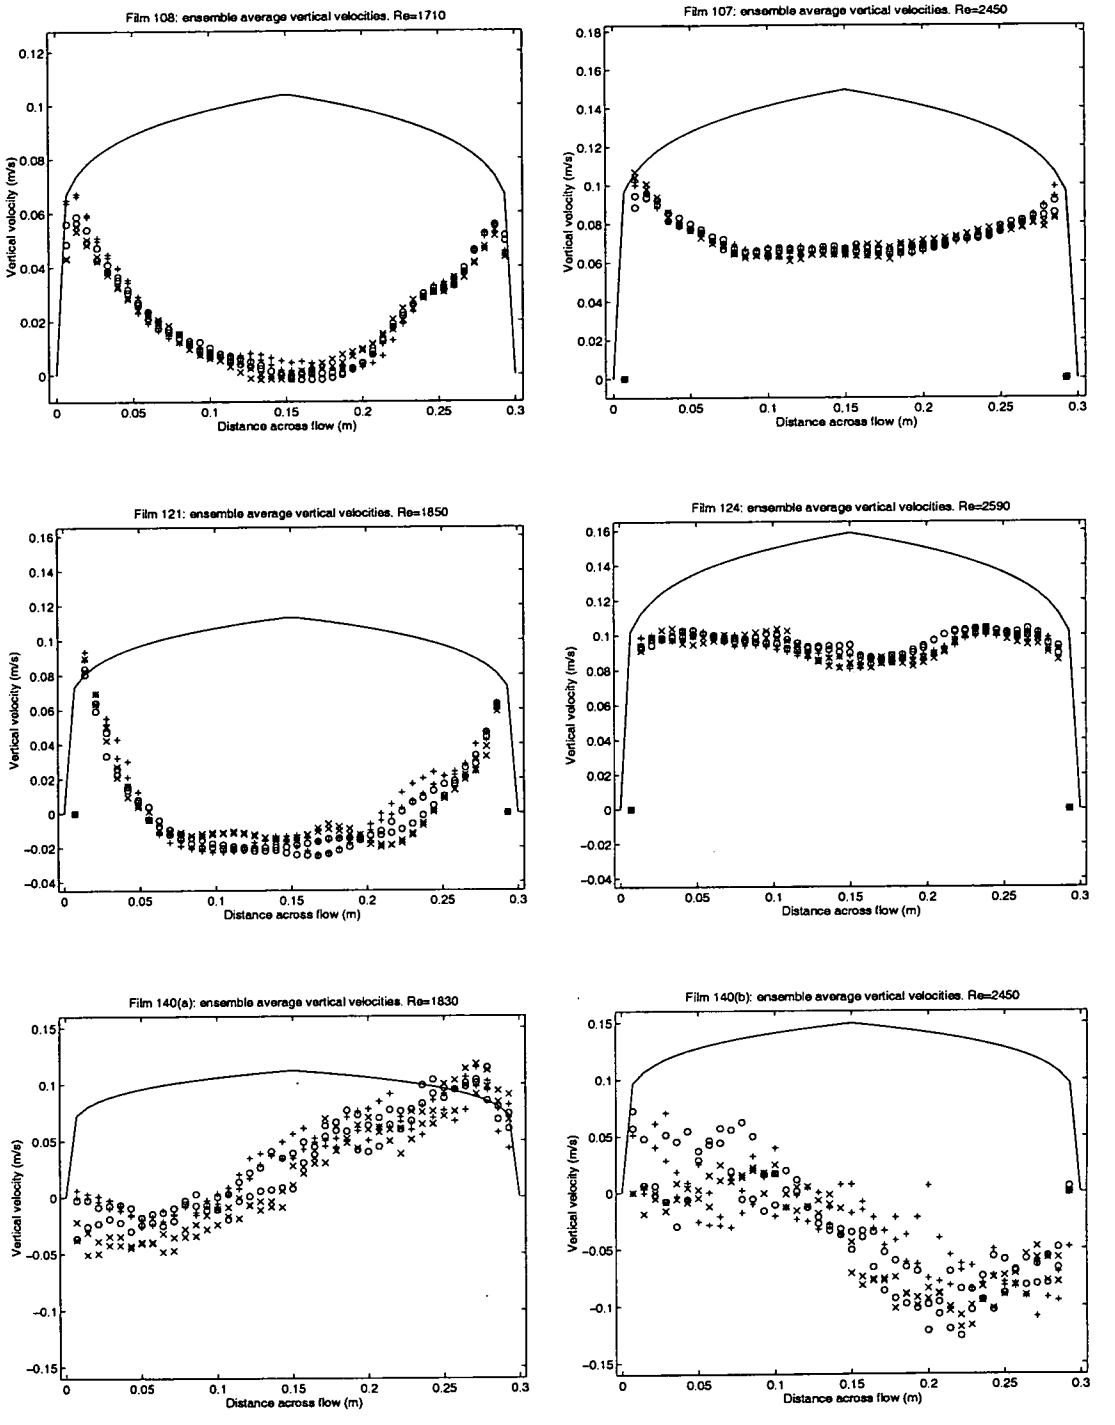


Figure 9.48: Profiles of vertical velocities along horizontal lines through the centre of the freeboard above a 60 cm deep fluidised bed, at heights of 1.9 m (top), 1.3 m (middle) and 0.4 m (bottom) above the bed surface. The superficial velocity was roughly 0.09 m s^{-1} (left) or 0.125 m s^{-1} (right). The continuous curves are profiles of the velocity under fully-developed turbulent flow, based on a one-seventh power law for cylindrical pipes; hopefully these will be reasonable estimates of turbulent flow in a square duct.

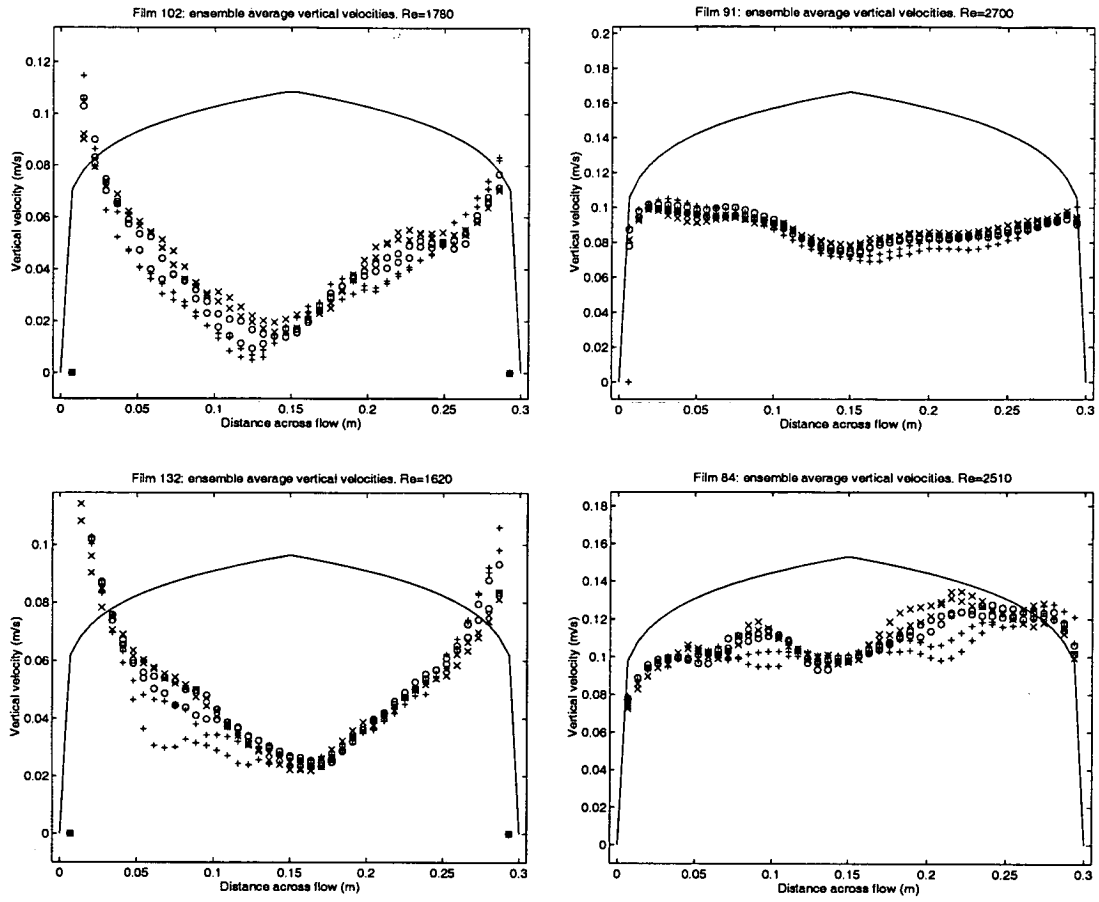


Figure 9.49: Vertical velocity profiles along horizontal lines one quarter of the way into the freeboard space above a 31 cm deep fluidised bed, at heights of 2.2 m (top) and about 1.35 m (bottom) above the bed surface. The superficial velocity was about $0.083\,\mathrm{m\,s^{-1}}$ (left) or about $0.130\,\mathrm{m\,s^{-1}}$ (right). The expected profile for turbulent flow is plotted using a continuous curve.

A possible way to categorise the flow patterns was seen when Reynolds numbers were calculated, based on the density and viscosity of air, the superficial velocity in the freeboard and the side length of the apparatus (0.3 m). These values were given in Figures 9.47–9.49, in the title of each graph. The flows with a Reynolds number below about 2400 all exhibited the dip profile reported by Levy & Lockwood (1983), with maxima near the walls (often greater than the superficial velocity) and minima near the centre (often small in magnitude, or even negative). As the Reynolds number increased, the minimum vertical velocity became more upward and the central dip was less pronounced. When Re $>$ 2500 or so, the velocity profiles were reasonably flat and dips were not always apparent. These observations apply both in the centre of the freeboard and one quarter of the way in, at all viewing heights more than 1.2 m above both depths of bed.

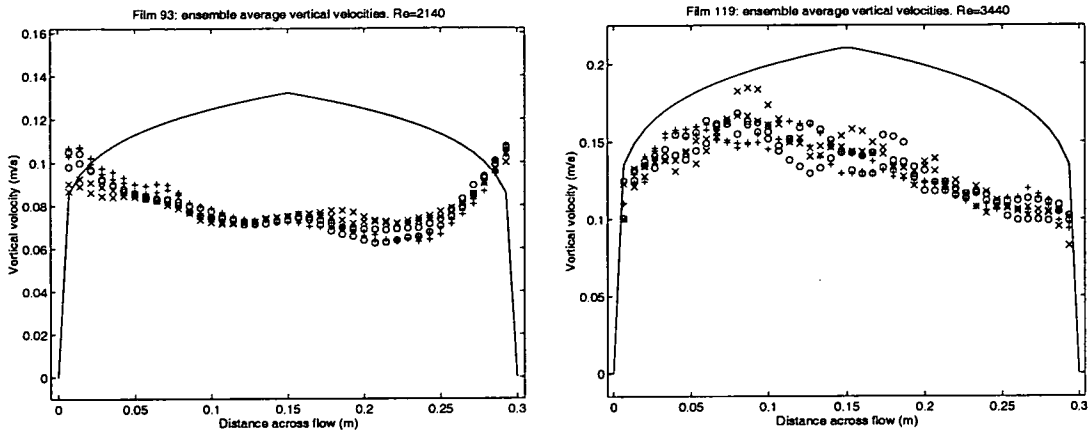


Figure 9.50: Further vertical velocity profiles along horizontal lines through the freeboard. Left: one quarter of the way into the section and 2.2 m above a 31 cm bed; Right: in the centre of the freeboard, 1.3 m above a 60 cm bed.

Figure 9.50 illustrates this further, showing two more ensemble-averaged velocity profiles. The left-hand figure, for film 93, shows the profile at the intermediate flowrate 2.2 m above a 31 cm bed. The dip profile was still clear at Re = 2140, but the velocity minimum at the centre had risen to about 75% of the maxima. This graph can be compared to those in the top row of Figure 9.49.

At the high-velocity end of the scale, no dip profile could be seen in film 119, which was the highest-flowrate run 1.3 m above the centre of the 60 cm bed (Figure 9.50, right; Re = 3440). The vertical velocity had a maximum about 0.1 m across the flow section and appeared to decrease as the walls were approached (as would be expected for normal pipe flow). The companion graphs to this were given in the middle row of Figure 9.48 (Re = 1850 and 2590).

For comparison, Levy & Lockwood (1983) used superficial velocities of at least $0.40\,\mathrm{m\,s^{-1}}$ in a $0.3\,\mathrm{m} \times 0.6\,\mathrm{m}$ fluidised bed. Their Reynolds numbers were at least 10,000. In contrast, Lang (1986) used superficial velocities of up to $0.04\,\mathrm{m\,s^{-1}}$, so his Reynolds numbers were not much more than 1000. In both cases, dip profiles were found. It appears that the Reynolds number cannot be the only factor to consider, when predicting whether the velocity profile in the freeboard region will show a dip profile or a more normal domed profile (*cf* laminar flow in a pipe).

Chapter 10

Conclusions

10.1 PIV and image-shifting

The principles behind direct multiple-image PIV have been reviewed. A digital analysis system was tested by means of Monte Carlo simulations. It was important that PIV recordings were simulated at displacement intervals of less than 1 pixel.

The analysis system was found to be very accurate most of the time, if Gaussian interpolation was used to calculate the image separations to sub-pixel accuracy (Chapter 5). Whittaker's reconstruction and the method of volume moments were also tried, but were found to be much less reliable (Figure 5.16).

The optimum e^{-2} image diameter (d_τ) for multi-image digital PIV was shown to be about 2 pixels, if the variable errors in the measurements are to be minimized (*e.g.* Figure 5.20). The seeding particles and the image:object magnification should be chosen with this in mind. Using greater magnification than this (*e.g.* a scanning resolution which produces images 4 pixels in diameter) can increase the relative error in each measurement and may thus be counter-productive.

From Figure 5.27, a new correlation was proposed for the minimum image separation to give an accurate measurement. Keane & Adrian (1990, 1991) had

recommended that consecutive images should be separated by at least 1.5 times their diameter. However, it was found here that, for all analysis parameters, digital images need to have a two-pixel 'halo' of clear space around their edges: $|\boldsymbol{d}|_{min} = d_\tau + 2$ pixels.

As the image displacement increases, the absolute accuracy of PIV remains very good (*e.g.* each reading might remain at ± 0.05 pixels on average); consequently, the relative error decreases. At the same time, however, the *probability* of actually obtaining these highly-accurate measurements begins to decrease (Figure 5.11); an apparently-random reading might be made instead.

If it is possible to use vector validation and interpolation to correct these spurious vectors, the high-displacement limit for an ideal PIV recording can be about half the width of the interrogation area. The yields of valid vectors were over 80% when measuring image separations of up to 28 pixels using 63×63 interrogation areas and no zero-padding (or up to about 40 pixels when using full zero-padding). For square interrogation areas measuring $d_I \times d_I$, the recommendation by Keane & Adrian (1990, 1991) that $|\boldsymbol{d}|_{max} = d_I/4$ might seem to be an overcautious rule of thumb. However, real PIV recordings will include background noise in the image plane, velocity gradients across each interrogation area and out-of-plane motion. It will not be possible to measure large displacements in practice as consistently as is suggested in this work.

If contours of the yield of valid vectors can be plotted (along the lines of Figure 5.31), one idea is to choose the illumination time interval so that, for example, 90% of the largest vectors will be determined correctly. These contours of vector yield depend on the analysis method: analyses using fully-padded FFTs give better yields than those using unpadded FFTs.

Elsewhere in Chapter 5, it was noted that autocorrelation wrap-around can cause large errors at certain image separations, if the digital analysis is carried out using unpadded FFTs. If N images are recorded of each particle and M is an integer satisfying $2 \leq M \leq N$, then these problematic displacements were found to

be $(\pm d_I/M, 0)$, $(0, \pm d_I/M)$ and $(\pm d_I/M, \pm d_I/M)$. Some effects of autocorrelation wrap-around were shown in Figure 5.6.

Wrap-around errors should probably be avoided on principle, but they could be ignored for some flows if it is known that only a few points will be affected. For example, a particular flow might contain thousands of velocities with many different speeds and directions, only a few of which will be measured incorrectly. On the other hand, very poor results will be obtained if the flow is predominantly unidirectional and the most commonly-occurring image displacement is close to one of the above values. This could due to the nature of the flow or the result of using image-shifting with an unfortunate choice of parameters. Further simulations demonstrated that wrap-around errors cannot arise if fully-padded FFTs are used instead (Figure 5.8) and that problems at certain displacements can be avoided if larger interrogation areas can be chosen (Figure 5.28).

A general relationship was derived for the ratio of 'good' correlations between neighbouring images in each particle track to spurious wrap-around correlations (equation 6.2). This became part of a theoretical model, based on the addition of Gaussian autocorrelation signals with known magnitudes and at known locations. This model successfully explained the magnitudes of the wrap-around errors found in Chapter 5 (as shown by Figure 6.17).

Interrogation area sizes need to be specified. The negative scanner used in this work produced bitmap images which were about 3450 × 2350 pixels in size (±3% or so from film to film, as the camera position varied slightly from run to run). Velocity gradients were reasonably small in most flows, so 128 × 128 interrogation areas produced the best results. The large side length allowed image displacements of up to 40 pixels to be measured reasonably consistently, whilst also avoiding wrap-around problems in the runs using image-shifting. For example, the re-analysed results from frame 23 of film 39 (Figure 8.6) are considerably better than they were in the two papers appended to this thesis. Initially 64 × 64 interrogation areas were used, but these produced many more vectors than could conveniently be processed or presented on a half-page (*cf* Figure 8.7). The repeat analyses

using 128×128 interrogation areas had a higher percentage yield of valid vectors and the vector arrays were more manageable.

A start was made on calculating image-pair dropout, as a function of the image separation and the number of images of each particle. A general formula was found for the average fraction of paired images in each interrogation area (Table 6.1). Probability distribution functions now need to be introduced, to describe the variations of the numbers of paired and unpaired images about their mean values. It is hoped that, in due course, the average yield of valid vectors can be calculated theoretically. This, however, has been left for future work.

Square interrogation areas were used in the digital analyses. The best yields of valid vectors were obtained when the image displacements were parallel to two of the interrogation area sides. This slight biasing means that diagonal displacements will be measured less consistently than 'horizontal' or 'vertical' displacements. In contrast, circular interrogation spots perform equally well in all directions.

Square interrogation areas will also give axi-symmetric yield contours, if used with an axi-symmetric windowing function. However, it is best to consider the nature of the flow before adding such a windowing function. For example, some photographs will record mainly-horizontal displacements (*e.g.* because image-shifting was used). Adding axi-symmetric windowing to PIVANAL will artificially reduce its performance when measuring horizontal displacements, so might not be worthwhile. Table 5.4 suggests that the 90% yield mark might be passed at horizontal displacements of about $18\sqrt{2} \approx 25$ pixels if axi-symmetric windowing is applied, rather than at about 35 pixels if windowing is not used.

PIV was used to measure the flows in the freeboard region above two different square-sided fluidised beds. It was found that image-shifting could be a very productive refinement to the technique, for two main reasons. In many experiments, especially when the flowrate was relatively low, there were substantial regions in which the measured velocities were too small to measure using direct PIV. By adding an extra horizontal velocity component to all flow velocities, image-shifting

could ensure that every displacement presented to the analysis program would be resolvable. The yields of valid vectors could be improved significantly by using image-shifting in this way.

Another problem was that the greatest flow velocities were considerably larger than the superficial velocity, with these vectors often pointing in all manner of directions. Image-shifting could also ensure that all analysed vectors pointed to the left, thus avoiding alignment difficulties: direct PIV has a directional ambiguity. In some of the higher-flowrate runs (films 136 and 140), image-shifting was only partly successful; a faster mirror rotational speed should have been used, because the flow velocities were faster than expected.

10.2 Fluidisation results

Leva (1951) demonstrated that the fine-particle concentration in a fluidised bed decayed exponentially with time. The particle concentrations in the freeboard could well have followed the same decay. This property was exploited experimentally: variations in seeding particle concentration meant that conditions for PIV were very good at some stage of each (batch-process) experiment.

Particle concentrations in the freeboard were not measured in this work, although later programs by Hind might have made this possible. No consistent method was used when adding seeding particles to the beds between runs; the only consideration was that there should be at least enough to seed the next run. If instead the beds had been fluidised for an hour before adding a known mass of fines, it might have been possible to investigate the particle concentrations above the bed as a function of time and freeboard height (*cf* Figure 2.3) in a reproducible manner. This was not an aim of these experiments, however.

In the smaller, 0.076 m × 0.076 m fluidised bed, all runs used a superficial velocity only slightly greater than U_{mf} (Chapter 8). The Reynolds number was

about 350. Two different viewing heights were used in the centre of the freeboard (about 17 cm and about 33 cm above the bed surface) along with two different bed depths (3 cm and 5 cm). The flow was chaotic in all cases, with consecutive vector maps usually looking rather different. Image-shifting would have been helpful in most runs because there were many low-velocity regions, but the technique was not available at the time.

The individual vector maps had a number of features. Jets of rapid upflow were often found, in which the vertical velocity was 3–4 times greater than the superficial velocity. Single vortices were often observed at the lower viewing height (film 41, frame 11; Figure 8.10, bottom) but were rarely seen at the greater height (films 33 and 37). It could be argued that paired, counter-rotating vortices were seen at the end of film 39 (Figure 8.6), but these are believed to be the result of a high-speed jet near the centre of the bed rather than evidence of the ghost bubbles proposed by Pemberton & Davidson (1984).

In this smaller bed, the mean vertical velocity was greater than the superficial velocity in almost every frame. This section-average vertical velocity was calculated by averaging the vertical velocities from every point in a given vector map. It was shown in Section 2.4.6 that these mean velocities would have been about $1.44\,U$, if laminar flow had been able to develop in the square freeboard; this value is greater than U, because the velocity terms were averaged directly and not weighted by the relevant fraction of the bed area. The average of these experimental section-average vertical velocities was slightly less than $1.44\,U$ at the lower viewing height but somewhat greater at the upper position.

Ensemble-average vector maps were calculated, by combining the analysed frames from each film into a single vector map. A laminar velocity profile gave a surprisingly good fit to the cross-bed profiles of most of these time-averaged velocities, despite the rather different appearance of the individual vector maps. This was consistent with the low Reynolds number in the freeboard. There would have been more confidence in these ensemble-average velocities, however, if they had been obtained by averaging hundreds of frames, rather than just 12–17.

The observed velocity profiles were similar to those measured by Horio *et al.* (1980) and Lang (1986), as the greatest velocities were found near the bed axis. Horio *et al.* also found descending zones near their walls, in which particles returned towards the bed (as sketched in Figure 2.5). Descending zones were not found here when using the Grade 10 ballotini, probably because the fluidisation conditions were so mild that few bed particles were ejected into the freeboard. The photographs in Figures 8.1–8.3, however, suggest that descending regions can exist when $U/U_{mf} \sim 20$.

Zenz & Weil (1958) attempted to calculate the TDH from the drag forces on ejected particles. The vertical velocities obtained in these experiments in the smaller bed suggested that the average vertical velocity past a particle with a long freeboard residence time could be about 2–2.5 U. This might be a good value for elutriation calculations if the particles drift upwards slowly but steadily, with a significant slip velocity between the particles and the turbulent air stream. On the other hand, some particles were seen to rise at roughly double this speed in most frames. If bed particles are ejected at high velocity into a jet of rapidly-rising gas and then carried through the freeboard very quickly, maximum vertical velocities of about 4–5 U seem more relevant than the average values.

The experiments in the larger fluidised bed used various combinations of air flowrate (low, intermediate, or high), bed depth (31 cm or 60 cm), section position (one quarter or one half of the distance into the freeboard) and viewing heights (about 40, 130 or 200 cm above the bed surface). At the two greater viewing heights, however, it was found that the bed depth and the viewing position had little effect on the nature of the flow.

Low-velocity regions caused big problems in all of the low flowrate runs, as well as in some of the intermediate flowrate experiments. Fortunately, image-shifting hardware had become available and was used when necessary. At the highest air flowrates, however, image-shifting was not usually required.

A key experiment was carried out in film 104, taken 2.20 m above a 31 cm bed, in the centre of the freeboard. Seeding particles were driven into the freeboard, then the air was switched off and the particles were allowed to fall back down. It was found that particles in the centre of the bed fell at up to $0.09\,\mathrm{m\,s^{-1}}$ (much faster than their terminal velocity). Particles nearer the wall were either stationary or moving upwards (Figure 9.15). The flow resembled a natural convection pattern, with downward flow in the centre and upward flow near the walls.

At low air flowrates, similar downward velocities were found in the central part of the test section. In film 105, the superficial velocity was $0.0889\,\mathrm{m\,s^{-1}}$ ($1.27\,U_{mf}$) and the other parameters were as for film 104. The largest downward velocities were about $0.038\,\mathrm{m\,s^{-1}}$ and many downward vectors were found in every frame. The flow could be described, to some extent, as the superposition of a uniform upward velocity (due to the air flow) and the above natural convection pattern.

This pattern could not be seen at the intermediate or higher flowrates. Film 88 used similar conditions to films 104 and 105, but for a larger superficial velocity of $0.118\,\mathrm{m\,s^{-1}}$ ($1.69\,U_{mf}$). The smallest upward velocities were $+0.044\,\mathrm{m\,s^{-1}}$ and absolutely no downward velocities were found in any of the frames. This increase in superficial velocity from about $0.09\,\mathrm{m\,s^{-1}}$ to about $0.11\,\mathrm{m\,s^{-1}}$ seemed to cause a significant change in the nature of the flow at the two larger viewing heights, resulting in no downward velocities. This transition might be caused by a change in the bubbling properties of the fluidised bed, governed by the ratio U/U_{mf}.

The values of Reynolds number corresponding to these flowrates were 1800 and 2200, respectively. Is it a coincidence that these are close to the values of Re found during the transition from laminar to turbulent flow in a cylindrical pipe?[1] Industrial fluidised beds can be an order of magnitude wider than the 'larger' fluidised bed used here, but they might well use smaller particles with a much reduced U_{mf} (*e.g.* FCC catalyst) and gases at high temperatures; perhaps Re = 2000 is approximately the case in some beds. Further tests are required (*e.g.* using smaller particles in the larger fluidised bed) in order to assess whether

[1]The corresponding value in a square-sided duct should be similar.

values of Re have more influence on the flow pattern in the freeboard than the ratio U/U_{mf} down in the bubbling fluidised bed.

The ensemble-average velocity profiles also showed this property (Figure 9.46). At the lowest flowrates, the 'dip' profiles reported by Levy & Lockwood (1983) and Lang (1986) were always found, following the pattern shown in Figure 2.6. Reasonable upward velocities were found near the walls, but the average velocities in the centre were usually very low (and possibly downward). At intermediate flowrates, the velocity profiles were much more uniform across the test section, due to the elimination of all downward velocities. However, the largest velocities were still found near the walls (Figure 9.47). Increases in the Reynolds number caused the ensemble-averaged velocities near the centre to become more strongly upward, further flattening the velocity profiles. Some of these intermediate and high flowrate results resembled the profiles found by Horio *et al.* (1980); however, descending regions were never observed near the walls.

The highest Reynolds number used was 3450, at which value the velocity profile showed signs of bulging up into a central dome (Figure 9.50, right). If Reynolds numbers of 4000–4500 had been achievable, it would have been interesting to see if this bulge grew further, showing increased resemblance to the velocity profile expected under fully-developed turbulent flow. The compressor could only deliver a limited range of flowrates, however, so this was not possible.

At the intermediate and highest viewing heights, all velocities were upward when the intermediate and highest flowrates were used (as noted above). No vortices were found in any frames from these runs. Downward velocities and single vortices were usually found in the runs at the lowest flowrate tested, but there was no evidence to support Pemberton & Davidson's (1984) theory of ghost bubbles rising up the freeboard. In all cases, the greatest vertical velocity was typically about $1–1.5\,U$. The section-average vertical velocities were always less than U, with the lowest deviations found in the runs with the highest flowrates (Table 9.21). It would have been interesting to perform experiments in more positions at each freeboard height (*e.g.* 1–2 cm beyond the front wall) in order

to see whether much larger upflows were found near the corners of the square apparatus (as continuity would suggest). However, there was a limit to what could be achieved in the time available.

The easiest PIV experiments to analyse were taken 2.20 m above the 31 cm bed, using the intermediate air flowrate and the centre of the freeboard (film 88). The yield of valid vectors was exactly 100% in several frames: every vector was resolved correctly, including those nearest the walls. (In contrast, image-shifting limited the yield to about 95%, because it was not usually possible to obtain valid vectors in the left-hand-most and right-hand-most columns of each vector map.) This intermediate flowrate seemed to be large enough to ensure that all measured velocities were upward (thus the dynamic range was low), but not so large as to carry many bed particles into the test region. The photograph in Figure 9.2 possibly shows that a few bed particles were present, even though correlations suggested that this viewing height was above the TDH (Table 9.2).

Comparing this run to the results 1.50 m above the same bed at the same flowrate (film 80), the vertical velocities were much more uniform at the 2.20 m height. At the intermediate height, the greatest vertical velocities were larger and the smallest vertical velocities were closer to zero (thus the dynamic range was much larger). This is not inconsistent with the expected decay pattern for isotropic turbulence; RMS velocity fluctuations have not been calculated, however, so a more categorical statement cannot be made. Plenty of bed particles were seen at the height of 1.50 m, which was close to the TDH.

It was considerably harder to obtain good results at the lowest viewing height, which was only 0.40 m above the bed surface when the air was off (comfortably below the TDH). Results without image-shifting were problematic (film 71). It was difficult to align all the vectors correctly. Vectors could be aligned in groups, but there is a chance that all of the vectors so aligned were in the wrong direction. Vector alignment and validation were both difficult, so some of the vectors presented in Figure 9.23 could well be spurious.

Using image-shifting, some encouraging results were obtained above the 60 cm bed using the lowest flowrate (films 136 and 140). Vector maps such as those in Figures 9.37 and 9.43 show that the flow was extremely turbulent. The vertical velocities were about double those at the intermediate and highest viewing positions, when using comparable air flowrates: speeds of 3–4 U were found.

At the lowest viewing height, the velocities increased considerably when moving from the lowest to the intermediate flowrate. The increase in superficial velocity was only 30%, but the most extreme velocities (*e.g.* the magnitudes of the greatest vertically-upward and horizontal velocities) doubled in the centre of the freeboard and quadrupled at the 'quarter-in' position (Tables 9.19 and 9.20).

The section-average vertical velocities were interesting: the increase in air flowrate meant that the average velocity in sections through the centre of the freeboard *fell* from $+0.029\,\mathrm{m\,s^{-1}}$ to $-0.037\,\mathrm{m\,s^{-1}}$, *i.e.* it became negative. One quarter of the way into the freeboard, the same increase in flowrate produced an increase in the section-average vertical velocity. This suggests that the air flow immediately above the bed surface followed a 'down at centre, up at walls' pattern, which was reinforced by the increase in flowrate. Further runs 1–2 cm into the freeboard would be required to confirm or refute this, however. The videos taken by Yórquez & Duursma (1999) show similar downward flows in the centre, both above an incipiently-fluidised bed (with bubble injection) and above a freely-bubbling bed ($U = 2\,U_{mf}$).

This 'down at centre, up at walls' profile above the bed surface affects a wide range of freeboard heights, when the air flowrate is low: the greatest upward velocities remain near the walls, even at heights of 2.2 m.

At the intermediate flowrate, however, the dip velocity profile soon becomes much flatter: compare the second part of film 140 ($z = 0.4\,\mathrm{m}$; Figures 9.40 and 9.41) to film 124 ($z = 1.3\,\mathrm{m}$; Figure 9.35). There certainly seems to be a fundamental change between the lowest viewing position (where a lot of downward flow was found) and the intermediate position (where upward flow was universal). It

would be interesting to investigate other freeboard heights (*e.g.* 0.9 m above the bed surface), to see how and where this transition takes place.

Above this 0.3 m × 0.3 m fluidised bed, the velocities appear to undergo a considerable amount of change in the first 1 m or so of the freeboard; the changes between 1.3 m and 2.2 m are fairly slight. As a result, it does not seem worth investigating even greater freeboard heights.

Chapter 11

Recommendations for future work

11.1 PIV work

Some further PIV simulations/analyses should be carried out, in order to determine the optimum seeding concentration. Four different sets of runs will be needed, to find the best values when 2, 3, 4 or 5 images of each particle are simulated.

Also concerning the PIV simulations, the work in Chapters 5 and 6 considered only the errors in each *average* measurement (from up to 2209 valid vectors). Further simulations are needed, to determine the image diameter which minimizes the standard deviation of these 2209 measurements about their mean value. Initial work suggests that this is about 2 pixels. Figure 5.24 (right) reveals that this optimum image diameter minimizes not only the standard deviation but also the variable error (which measures the bias in the mean values, from pixel to pixel).

The simple PIV technique used in this work has been found to be very useful, measuring a wide range of flow velocities. However, technological advances should be exploited in any future work, in order to facilitate the PIV process. A high-

resolution digital camera would be an advantage, saving the time required in this work to process 'wet' photographs. A digital cross-correlation camera would be a more significant improvement, avoiding directional ambiguities and also allowing even the smallest velocities to be measured (*cf* Figure 9.6).

Commercial PIV packages should also be investigated, such as Dantec's (1998) Flowmap 3D PIV system. This allows three-dimensional velocity measurements to be made in real time, using a stereoscopic digital cross-correlation method.

11.2 General fluidisation considerations

Considering the experimental results in the larger fluidised bed, it is felt that the 'quarter-in' viewing position was not sufficiently different from the central position. Further photographs should therefore be taken at the various viewing heights, at a position about 1–2 cm into the freeboard (beyond the front wall).

Before these results could be used with complete confidence to predict entrainment rates, it would also be necessary to test the symmetry of the flow in the freeboard. Similar experiments should therefore be carried out 22.5 cm and 28–29 cm into the freeboard.

A further improvement would be to record the bed surface and the freeboard flow on video, whilst taking experimental PIV photographs. A simple electric circuit could allow several LEDs to be illuminated simultaneously whenever a photograph is triggered (perhaps emitting a double flash after every six exposures). This would allow the recordings to be compared afterwards, so that events at the bed surface could be related to simultaneous (or later) events in the freeboard. The uniformity of the bubbling in the fluidised bed could also be assessed.

11.3 The natural convection pattern

The 'natural convection' pattern observed in film 104 (*e.g.* Figure 9.15) should also be investigated. It is possible to test whether this pattern is a result of using too many seeding particles and then finding an uneven distribution of fine particles in the freeboard (thus varying the bulk density across the freeboard).

The existing photographs from film 104 should therefore be analysed in a new way, to measure the number of particle images in various vertical regions within each frame. This would reveal whether there actually was a density difference across the freeboard during the experiments. If so, further experiments could then be carried out using smaller seeding particles at a lower concentration in the fluidised bed (hopefully producing a much weaker natural convection pattern). For comparison, seeding particle concentrations could also be measured in conditions of more uniform flow (*e.g.* film 88, shown in Figure 9.1).

The initial conditions should also be tested, because the results in film 104 might simply be due to the low gas flowrate used to drive particles into the freeboard in the first place (*cf* Figures 9.7–9.9). A repeat run should be carried out using a greater initial flowrate, then a sudden reduction to zero flowrate.

11.4 Dip profiles and 'ghost bubbles'

It was found here that flows in the freeboard above the larger fluidised bed could exhibit either a 'dip' or a 'domed' profile. It is not clear whether this fundamental difference in the nature of the flow is caused by the Reynolds number in the freeboard (Re = 2000 appearing transitional) or the bubbling properties in the fluidised bed (with a transition at about $U/U_{mf} = 1.5$), or both of the above.

The conditions under which dip profiles are obtained should therefore be investigated. The larger fluidised bed and the existing air supply could be used, but

with different grades of glass beads (*e.g.* with mean sizes of about 150 μm and about 350 μm) forming the fluidised bed. The aim would be to assess the nature of the flow, as a function of Re and of U/U_{mf}. This work could initially be largely qualitative, *e.g.* using the light sheet to visualise the flow.

It appears that ghost bubbles are unlikely to be found above fluidised beds of Geldart Class B material. Erupting bubbles throw bed material into the freeboard, which then falls back down towards the bed surface. If this mass is too large, any circulation pattern within the ejected gas will be destroyed (Figure 2.7).

Conditions which minimize the mass of ejected particles have been investigated by Yórquez & Duursma (1999): $U \approx U_{mf}$, singly-injected bubbles and large bed depths. No ghost bubbles were found. Similar work should be carried out using Class A materials at $U \approx U_{mb}$. Class A particles are smaller and lighter than those in Class B, so they should cause less disruption to any vortex ring in the gas phase. Ghost bubbles might survive into the freeboard region only when certain grades of bed particle are used.

11.5 Analysing existing results

As well as further experiments, future work should also attempt to obtain more information from the flow maps which have already been obtained (only some of which have been presented here). Various post-processing techniques can be used to analyse the data or to present the data more clearly.

In particular, vorticity maps should be obtained. These reveal vortex patterns more clearly than velocity vector maps, because the vorticity is independent of any superimposed uniform rise velocity. If paired vortices are found on either side of the freeboard, it would be interesting to know whether these rotate downwards in the centre and upwards near the walls (Levy & Lockwood, 1983) or upwards in the centre and downwards near the walls (Pemberton & Davidson, 1984).

Chapter 12

Symbols

Symbol	Description
A	The cross-sectional area of a fluidised bed
A	The distance from the camera lens to the film when using image-shifting)
A_0	The area of distributor plate per orifice ($A_0 \approx 0\,\mathrm{m}^2$ for a porous distributor)
A_2	The limit of penetration around a bubble in a two-dimensional fluidised bed
A_3	As A_2, but in a three-dimensional fluidised bed
$A(\Delta x, \Delta y)$	The autocorrelation function in PIV analysis
Ar	An Archimedes number
B	The effective optical path from the camera lens to the light sheet when image-shifting is used
C	The concentration of seeding particles
C	An offset distance between the optical axis of a camera and the image-shifting mirror surface's axis of rotation
C_D	A drag coefficient
D	A characteristic length for a Reynolds number
D_b	The diameter of a bubble in a liquid or a fluidised bed
D_{b0}	The initial diameter of bubbles in a fluidised bed, on formation immediately above the distributor
D_{be}	The equilibrium bubble diameter in a fluidised bed, if splitting and coalescence both occur
D_{bm}	The maximum obtainable bubble diameter in a fluidised bed, if only bubble coalescence occurs

Symbol	Description
D_e	The equivalent diameter of a non-spherical bubble in a fluidised bed, given by: bubble volume $= \pi D_e^3/6$
D_t	The diameter of a cylindrical fluidised bed
F	The focal length of a camera lens (image-shifting)
F_{45}	The mass fraction of fines smaller than 45 μm
F	The fraction of images in an average interrogation area which form part of at least one image pair
G_s	The mass flux of elutriated solids
H	The height or depth of particles in a fluidised bed
H_f	The height of freeboard region above a fluidised bed
I	The (general) intensity of a laser beam
I_0	The intensity of a laser beam at a reference point
$I(x, y)$	The image intensity at a point (x, y) in an interrogation area
J_1	A first-order Bessel's function, of the first kind
L	A length, *e.g.* of the parabolic mirror in a scanning-beam system
M	The image:object photographic magnification
M	A natural number used to describe wrap-around correlations, which can occur at displacements d_I/M $(N \geq M)$
N	A counter (natural number), *e.g.* for the number of images recorded of a seeding particle using PIV
N_{im}	The total number of particle images in an interrogation area
N_{ipt}	The average number of images of a track in an interrogation area
$N_{pairs/track}$	The average number of correct image pairs obtained from each particle track
$N_{pairs/IA}$	The average number of correct image pairs found in each interrogation area
$N_{harmonics/track}$	The average number of second harmonic correlations obtained from each particle track
$N_{harmonics/IA}$	The average number of second harmonic correlations in each interrogation area
Q_b	The volumetric flowrate of bubbles in a fluidised bed
P	Pressure, so ΔP is a pressure difference
R	The distance from the axis of rotation of an image-shifting mirror to the light sheet
R_a	The radius of peaks in the autocorrelation plane

Symbol	Description
Re	A Reynolds number
U	The superficial velocity of fluid in a fluidised bed
U_b	The rise velocity of a bubble in a freely-bubbling fluidised bed (in the laboratory frame)
$U_{b,\infty}$	The rise velocity of a single bubble in an incipiently-fluidised bed
U_{ej}	The initial rise velocity of particles ejected from a bubbling fluidised bed
U_{max}	The maximum value of U, if it varies across the flow
U_{mf}	The particular value of U in a fluidised bed at which fluidisation is just achieved
U_{mb}	The value of U in a gas-fluidised bed at which bubbling is first obtained
U_o	The interstitial fluid velocity in a fluidised bed
U_{slug}	The rise velocity of a single slug in an incipiently-fluidised bed
W	The mass of bed particles in a fluidised bed
a	The radius of a void in a fluidised bed; the rate of bubble growth is hence given by $\dot{a} = da/dt$
a	One side length of a rectangular duct
a	A convenient length for an image-shifting system. $a = 2\,M\,R\omega\,\Delta t$
b	The second side length in an $a \times b$ rectangular duct
b	A handy inverse length for an image-shifting system. $b = 2\omega\,\Delta t\,F^{-1}\,(1+M)^{-1}$
$\boldsymbol{d} = (d_x, d_y)$	The seeding particle displacement in the image plane
$\lvert\boldsymbol{d}\rvert_{min}$	The minimum image displacement needed to give a measurement of satisfactory accuracy
$\lvert\boldsymbol{d}\rvert_{max}$	The maximum image displacement which still gives an acceptable yield of correct vectors
d	A distance between the centres of successive images
d_I	The characteristic size of an interrogation area.
d_g	The grain size of a photographic system
d_p	The particle diameter in the object plane
d_{pix}	The pixel size in a digital recording
$\boldsymbol{d}_s$	The displacement vector added to a flow when using image-shifting (*e.g.* with a rotating mirror)

Symbol	Description
d_{sx}, d_{sy}	Components of the image-shifting displacement $\boldsymbol{d}_s$ in the x- and y-directions
d_1	The separation of successive digital images of each particle
d_2	The wrap-around distance when using unpadded autocorrelation
d^*	The position of the global maximum when wrap-around occurs
d_τ	The image diameter recorded on a photographic negative
$f_\sharp$	The camera aperture ('f-number')
g	The acceleration on earth due to gravity ($9.81\,\mathrm{m\,s^{-2}}$)
h	The separation of two circles
k	Various constants, *e.g.* for Beer's Law
m, n	Various constants, defined locally in the text
r	The radius of a circle, *e.g.* the distance from the origin in polar co-ordinates: (r, θ)
t	Time
Δt	The illumination time interval in a PIV system
u_0	An initial velocity
u_g	A mean gas velocity
u_p	The velocity of a seeding particle in a fluid flow
u_t	The Stokes Law terminal velocity
u_y	The streamwise fluid velocity in a vertical rectangular duct
v_r	The radial velocity in polar co-ordinates, measured away from the origin
v_θ	The tangential velocity in polar co-ordinates, measured in an anti-clockwise direction
$\boldsymbol{v} = (v_x, v_y)$	The velocity of seeding particles in the object plane
w	The mass of fines in a fluidised bed
x, y, z	Three mutually-orthogonal co-ordinate axes
$(\Delta x, \Delta y)$	The autocorrelation displacement
(X, Y)	The particular autocorrelation displacement which corresponds to the image separation
z	A height above the distributor in a fluidised bed, or a height above the bed surface in the freeboard region
z	The mathematical array calculated by PIVANAL, which stores the calculated image displacements
z_0	The height above the distributor at which bubbles form
z_{TDH}	The Transport Disengaging Height above a fluidised bed

Greek letter	Description
β	The dimensionless ratio U_b/U_{mf}
γ	A dimensionless bubble growth parameter, defined by $\gamma = (1-\varepsilon)\,\dot{a}/(\varepsilon\,U_o)$
ε	The void fraction in a packed bed
ε_{mf}	The void fraction in a incipiently-fluidised bed
θ	The angle of a point in polar co-ordinates, measured anti-clockwise from the x-axis
θ	Also, the angle of rotation of a rotating mirror; thus $\omega = d\theta/dt$ is its rotational speed
λ	The wavelength of a laser beam used in PIV
μ	The viscosity of the fluid present in the flow
$\overline{\rho}$	The general bulk density of particles in the freeboard
$\overline{\rho_R}$	The bulk density of particles in a very tall freeboard region, in the conditions of 'complete reflux'
$\overline{\rho_{R,0}}$	The bulk density of particles in a very tall freeboard region, measured just above the bed surface
ρ_1, ρ_2, ρ_3	The particle bulk densities in the freeboard region, for phases 1, 2 and 3 of Kunii & Levenspiel's model
ρ_f	The density of the fluid in the flow
ρ_s	The solid density of a seeding particle
τ	The velocity relaxation time: $\tau = \rho_s\,d_p^2/18\,\mu$
ϕ	The potential function used in potential flow calculations
ϕ_p	The sphericity of non-spherical particles
ψ	The stream function, measuring the volumetric flowrate between the origin and the point ψ is measured at
ω	The rotational speed of a rotating mirror

Subscript	Description
f	Property of a fluid
g	Property of a gas
i	A property found in potential flow inside a bubble
mf	Values at the condition of minimum fluidisation
o	A property found in potential flow outside a bubble
p	Property of a particle
s	Property of a solid

For example, $\mathrm{Re}_{p,mf}$ is a Reynolds number based on a particle diameter and the fluid properties found at the point of minimum fluidisation, whilst $v_{r,o}$ is a radial velocity calculated outside a fluidisation bubble.

Most of the above are to be used in SI units, *e.g.* densities such as ρ_s are in $\mathrm{kg\,m^{-3}}$, velocities such as U_{mf} are in $\mathrm{m\,s^{-1}}$, stream functions such as ψ_o are in $\mathrm{m^3\,s^{-1}}$ and the mass flux G_s is in $\mathrm{kg\,m^{-2}\,s^{-1}}$. Dimensionless groups (such as Re) obviously have no units.

Most distances relating to digitised images and analysis using autocorrelation should be in pixels, however. For example, an image of diameter d_τ (in pixels) will produce autocorrelation peaks with an e^{-2} radius of R_a (pixels), if the image separation d (pixels) is greater than $|\boldsymbol{d}|_{min}$ (pixels).

Some variables can be used with two different units. In an optical analysis, the magnification M is in metres of image plane per metre of object plane; in a digital analysis, it is more convenient to use M in pixels per metre of object plane.

Chapter 13

References

In alphabetical order of first-named author, then year, then second authors.

1. Adrian, RJ & Yao C-S (1985): 'Pulsed laser technique: application to liquid and gaseous Flows and the scattering power of seed materials', *Appl Opt*, vol. 24, 44–52

2. Adrian, RJ (1991): 'Particle-imaging techniques for experimental fluid mechanics', *Ann Rev Fluid Mech*, vol. 23, 261–304

3. Arroyo, MP & Greated, CA (1991): 'Stereoscopic particle image velocimetry', *Meas Sci Technol*, vol. 2, 1181–1186

4. Azevedo, JLT & Pereira, JCF (1991): 'Flow patterns around and inside a bubble in a gas-fluidised bed', *Chem Eng Sci*, vol. 46, 155–165

5. Baron, T, Briens, CL & Bergougnou, MA (1988): 'Study of the transport disengaging height', *Can J Chem Eng*, vol. 66, 749–760

6. Busignani, S, Cenedese, A, Querzoli, G & Romano, G (1994): 'Three colours trajectory recognition in PIV', IMechE Symp on *Optical Methods and Data Processing in Heat and Fluid Flow*, London, 75–79

7. Clift, R, Grace, JR, Cheung, L & Do, TH (1972): 'Gas and solids motion around deformed and interacting bubbles in fluidized beds', *J Fluid Mech*, vol. 51, 187–205

8. Collins, R (1965): 'The flow patterns within a void in a porous mass' — comments, *Trans Inst Chem Engrs*, vol. 43, T218–T219

9. Collins, R (1982): 'Cloud patterns around a bubble growing in a gas-fluidised bed', *J Fluid Mech*, vol. 122, 155–167

10. Collins, R (1983): 'Gas flow patterns associated with a growing spherical bubble', in *Fluidization IV* (Ed. Kunii, D & Toei, R), 1–8

11. Collins, R (1989): 'A model for the effects of the voidage distribution around a fluidization bubble', *Chem Eng Sci*, vol. 44, 1481–1487

12. Coulson, JM & Richardson, JF (1978): *Chemical Engineering Volume Two* (third edition), Pergamon Press

13. Dantec Measurement Technology (1998): 'Particle Image Velocimetry (PIV)', http://www.dantecmt.com/scientific/piv/pivmain.htm

14. Darton, RC, La Nauze, RD, Davidson, JF & Harrison, D (1977): 'Bubble growth due to coalescence in fluidised beds', *Trans Inst Chem Engrs*, vol. 55, 274–280

15. Davidson, JF, Paul, RC, Smith, MJS & Duxbury, HA (1959): 'The rise of bubbles in a fluidised bed', *Trans Inst Chem Engrs*, vol. 37, 323–328

16. Davidson, JF & Harrison, D (1963): *Fluidised Particles*, *Cambridge Univ Press*

17. Davidson, JF & Harrison, D (1966): 'The behaviour of a continuously-bubbling fluidised bed', *Chem Eng Sci*, vol. 21, 731–738

18. Davidson, JF, Harrison, D & Guedes de Carvalho, JRF (1977): 'On the liquid-like behaviour of fluidized beds', *Ann Rev Fluid Mech*, vol. 9, 55–86

19. Davies, RM & Taylor, GI (1950): 'The mechanics of large bubbles rising through extended liquids and through liquids in tubes', *Proc Roy Soc*, vol. A200, 375–390

20. Didden, N (1979): 'On the formation of vortex rings: rolling-up and production of circulation', *J Appl Mech Phys (ZAMP)*, vol. 30, 101–116

21. Di Felice, R (1995): 'Hydrodynamics of liquid fluidisation', *Chem Eng Sci*, vol. 50, 1213–1245

22. Dryden, HL, Murnaghan, FP & Bateman, H (1956): *Hydrodynamics*, Dover Press

23. Fournol, AB, Bergougnou, MA & Baker, CGJ (1973): 'Solids entrainment in a large gas fluidized bed', *Can J Chem Eng*, vol. 51, 401–404

24. Geldart, D (1973): 'Types of gas fluidisation', *Powder Technol*, vol. 7, 285–292

25. Geldart, D & Abrahamsen, AR (1980): 'The effect of fines on the behaviour of gas fluidized beds of small particles', in *Fluidization* (Ed. Grace, JR & Matson, JM), Plenum Press, 453–460

26. Glicksman, LR, Lord, WK & Sakagami, M (1987): 'Bubble properties in large-particle fluidized beds', *Chem Eng Sci*, vol. 42, 479–491

27. Grace, JR & Harrison, D (1969): 'The behaviour of freely-bubbling fluidised beds', *Chem Eng Sci*, vol. 24, 497–508

28. Grace, JR & Clift, R (1974): 'On the two-phase theory of fluidization', *Chem Eng Sci*, vol. 29, 327–334

29. Grace, JR (1990): 'High-velocity fluidized bed reactors', *Chem Eng Sci*, vol. 45, 1953–1966

30. Grant, I (1997): 'Particle image velocimetry: a review', *Proc Inst Mech Engrs*, vol. 211 Part C, 55–76

31. Gray, C, Greated, CA, McCluskey, DR & Easson, WJ (1991): 'An analysis of the scanning beam PIV illumination system', *Meas Sci Technol*, vol. 2, 717–724; Re-printed in *Eng Optics*, Vol 4 (4) (1991

32. Gugnoni, RJ & Zenz, FA (1980): 'Particle entrainment from bubbling fluidized beds', in *Fluidization* (Ed. Grace, JR & Matson, JM), Plenum Press, 501–508

33. Halloin, VL & Jottrand, B (1994): 'A new technique for the experimental study of 3D velocity fields', *Exp Fluids*, vol. 17, 115–118

34. Hamdullahpur, F & MacKay, GDM (1986): 'Two-phase flow behavior in the freeboard of a gas-fluidized bed', *AIChE J*, vol. 32, 2047–2055

35. Harrison, D & Leung, LS (1962): 'The rate of rise of bubbles in fluidised beds', *Trans Inst Chem Engrs*, vol. 40, 146–151

36. Hind, AK & Christy, JRE (1992): 'Digital PIV applied to flows around artificial heart valves: analysis by autocorrelation', in *Proc Euromech Colloquium 279*, Kluwer Academic Publishers, Netherlands

37. Hind, AK & Christy, JREC (1996): 'A method to correct for velocity gradients in the analysis of PIV recordings', IMechE Symp on *Optical Methods and Data Processing in Heat and Fluid Flow*, London, 201–210

38. Hind, AK (1997): 'An investigation into the flow fields around bluff bodies and artificial heart valves', *PhD thesis*, University of Edinburgh

39. Horio, M, Taki, A, Hsieh, YS & Muchi, I (1980): 'Elutriation and particle transport through the freeboard of a gas-solid fluidized bed', in *Fluidization* (Ed. Grace, JR & Matson, JM), Plenum Press, 509–518

40. Horio, M & Nonaka, A (1987): 'A generalized bubble diameter correlation for gas-solid fluidized beds', *AIChE J*, vol. 33, 1865–1872

41. Horio, M & Kuroki, H (1994): 'Three-dimensional flow visualisation of dilutely dispersed solids in bubbling and circulating fluidised beds', *Chem Eng Sci*, vol. 49, 2413–2421

42. Jackson, R (1963a): 'The mechanics of fluidised beds: Part 1: the stability of the state of uniform fluidisation', *Trans Inst Chem Engrs*, vol. 41, 13–21

43. Jackson, R (1963b): 'The mechanics of fluidised beds: Part 2: the motion of fully developed bubbles', *Trans Inst Chem Engrs*, vol. 41, 22–28

44. Jakobsen, ML, Hossack, WJ, Greated, CA & Easson, WJ (1993): 'PIV analysis using an optically-addressed spatial light modulator', *Optics and Lasers in Eng*, vol. 19, 253–260

45. Keane, RD & Adrian, RJ (1990): 'Optimisation of Particle Image Velocimeters, Part I: double-pulsed systems', *Meas Sci Technol*, vol. 1, 1202–1215

46. Keane, RD & Adrian, RJ (1991): 'Optimisation of Particle Image Velocimeters, Part II: multiple-pulsed systems', *Meas Sci Technol*, vol. 2, 963–974

47. de Kock, JW & Judd, MR (1965a): 'The flow patterns within a void in a porous mass', *Trans Inst Chem Engrs*, vol. 43, T78–T84

48. de Kock, JW & Judd, MR (1965b): 'The flow patterns within a void in a porous mass' — reply to comments, *Trans Inst Chem Engrs*, vol. 43, T218–T220

49. Khan, MJ & Ahmed, A (1995): 'On the modality of velocity histograms in the plane of symmetry of a wing-body juncture', *Exp Fluids*, vol. 18, 164–167

50. Knudsen, JG & Katz, DL (1958): *Fluid Dynamics and Heat Transfer*, McGraw–Hill

51. Kreysig, E (1983): *Advanced Engineering Mathematics*, John Wiley & Sons

52. Kunii, D & Levenspiel, O (1991): *Fluidization Engineering*, Butterworth-Heinemann

53. Kunii, D & Levenspiel, O (1995): 'Effect of exit geometry on the vertical distribution of solids in circulating fluidised beds. Part I: solution of fundamental equations; Part II: analysis of reported data and prediction', *Powder Technol*, vol. 84, 83–90

54. Landreth, CC (1986): 'Two-dimensional measurements of complex flows using Pulsed Laser Velocimetry with directional resolution', *MSc thesis*, Univ of Illinois

55. Lang, IWP (1986): 'The structure of flow in the freeboard region of gas/solid fluidised beds', *PhD thesis*, University of Edinburgh

56. Leva, M (1951): 'Elutriation of fines from fluidized systems', *Chem Eng Prog*, vol. 47, 39–45

57. Leva, M (1959): *Fluidization*, McGraw Hill

58. Leva, M & Wen, CY (1971): 'Elutriation', Ch. 14 of *Fluidization* (Ed. Davidson, JF & Harrison, D), Academic Press, 627–650

59. Levy, Y & Lockwood, FC (1983): 'Laser-Doppler measurements of flow in the freeboard of a fluidised bed', *AIChE J*, vol. 29, 889–895

60. Lewis, WK, Gilliland, ER & Lang, PM (1962): 'Entrainment from fluidized beds', *AIChE Symp Series*, vol. 58 (Fluidization), 65–78

61. Lin, JS, Chen, MM & Chao, BT (1985): 'A novel radioactive particle tracking facility for measurement of solids motion in gas fluidized beds', *AIChE J*, vol. 31(1), 465–473

62. Lockett, MJ, Davidson, JF & Harrison, D (1967): 'On the two-phase theory of fluidisation', *Chem Eng Sci*, vol. 22, 1059–1066

63. Lourenco, L & Krothapalli, A (1995): 'On the accuracy of velocity and vorticity measurements with PIV', *Exp Fluids*, vol. 18, 421–428

64. Manga, M & Stone, HA (1993): 'Interactions between bubbles at low Reynolds numbers', *Phys Fluids A*, vol. 5, S3

65. Maxworthy, T (1972): 'The structure and stability of vortex rings', *J Fluid Mech*, vol. 51, 15–32

66. Maxworthy, T (1974): 'Turbulent vortex rings', *J Fluid Mech*, vol. 64, 227–239

67. McCluskey, DR, Hind, AK, Rix, SJL, Jakobsen, ML, Christy, JRE, Easson, WJ, Glass, DH & Greated, CA (1994): 'Development and application of a simultaneous two phase PIV system for the analysis of air-particle flow fields', IMechE Symp on *Optical Methods and Data Processing in Heat and Fluid Flow*, London, 99–109 (**appended**)

68. Mori, S & Wen, CY (1975): 'Estimation of bubble diameter in gaseous fluidised beds', *AIChE J*, vol. 21(1), 109–115

69. Morrison, IG (1995): 'The hydrodynamic performance of an oscillating water column wave energy generator', *Ph.D. thesis*, University of Edinburgh

70. Murray, JD (1965): 'On the mathematics of fluidization Part 2', *J Fluid Mech*, vol. 22, 57–80

71. Oschwald, M, Bechle, S & Welke, S (1995): 'Systematic errors in PIV by realising velocity offsets with the rotating mirror method', *Exp Fluids*, vol. 18, 329–334

72. Osberg, GL & Charlesworth, DH (1951): 'Elutriation in a fluidized bed', *Chem Eng Prog*, vol. 47, 566–570

73. Pemberton, S.T & Davidson, J.F (1984): 'Turbulence in the freeboard region of a gas-fluidised bed', *Chem Eng Sci*, vol. 39, 829-840

74. Pemberton, S.T & Davidson, J.F (1986): 'Elutriation from fluidised beds - I. Particle ejection from the dense phase into the freeboard', *Chem Eng Sci*, vol. 41, 243–251

75. Prasad, AK & Adrian, RJ (1993): 'Stereoscopic Particle Image Velocimetry applied to liquid flows', *Exp Fluids*, vol. 15, 49–60

76. Pyle, DL & Rose, PL (1965): 'Chemical reaction in bubbling fluidized beds', *Chem Eng Sci*, vol. 20, 25–31

77. Pyle, DL & Harrison, D (1967b): 'An experimental investigation of the two-phase theory of fluidization', *Chem Eng Sci*, vol. 22, 1199–1207

78. Raffel, M & Kompenhans, J (1995): 'Theoretical and experimental aspects of image-shifting by means of a rotating mirror system for Particle Image Velocimetry', *Meas Sci Technol*, vol. 6, 795–808

79. Ramírez de Santiago, M (1990): 'Measurement of size and velocity distributions of droplets produced by bubbles bursting', 5th International Symp on *Applications of Laser Techniques to Fluid Mechanics*, Springer-Verlag, 203–226

80. Reuter,H (1966): 'On the nature of bubbles in gas and liquid fluidised beds', *Chem Eng Prog Symposium Series*, vol. 62, no. 62, 92–99

81. Reynolds, O (1876): 'On the resistance encountered by vortex rings, and the relation between the vortex rings and the streamlines of a disc', *Nature*, vol. 14, 477–479

82. Richardson, JM & Zaki, WN (1954): 'Sedimentation and fluidisation. Part 1', *Trans Inst Chem Engrs*, vol. 32, 35–53

83. Rix, SJL, Glass, DH & Greated, CA (1996): 'Preliminary studies of elutriation from gas-fluidised beds using Particle Image Velocimetry', *Chem Eng Sci*, vol. 51, 3479–3489 **(appended)**

84. Rott, N (1990): 'Note on the history of the Reynolds number', *Ann Rev Fluid Mech*, vol. 22, 1–11

85. Rowe, PN & Partridge, BA (1963): 'Gas flow through bubbles in an ideal fluidized bed — I: Flow through an ideal bubble', *Chem Eng Sci*, vol. 18, 511–524

86. Rowe, PN, Partridge, BA & Lyall, E (1964): 'Cloud formation around bubbles in gas-fluidised beds', *Chem Eng Sci*, vol. 19, 973–985

87. Rowe, PN & Partridge, BA (1965): 'An X-ray study of bubbles in fluidised beds', *Trans Inst Chem Engrs*, vol. 43, T157–T175

88. Rowe, PN, Partridge, BA, Cheney, AG, Henwood, GA & Lyall, E (1965): 'The mechanisms of solids mixing in fluidised beds', *Trans Inst Chem Engrs*, vol. 43, T271–T286

89. Rowe, PN (1971): 'Experimental properties of bubbles', Ch. 4 of *Fluidization* (Ed. Davidson, JF & Harrison, D), Academic Press, 121–191

90. Rowe, PN & Matsuno, R (1971): 'Single bubbles injected into a gas fluidised bed and observed by X-rays', *Chem Eng Sci*, vol. 26, 923–935

91. Rowe, PN & Nienow, AW (1975): 'Minimum fluidisation velocity of multi-component particle mistures', *Chem Eng Sci*, vol. 30, 1365–1369

92. Rowe, PN (1976): 'Prediction of bubble size in a gas fluidized bed', *Chem Eng Sci*, vol. 31, 285–288

93. Rowe, PN & Yacono, CXR (1976): 'The bubbling behaviour of fine powders when fluidized', *Chem Eng Sci*, vol. 31, 1179–1192

94. Rudinger, G (1980): *Fundamentals of Gas-particle Flow*, Elsevier Science Publishing

95. Shariff, K & Leonard, A (1992): 'Vortex rings', *Ann Rev Fluid Mech*, vol. 24, 235–279

96. Sridhar, G & Katz, J (1995): 'Drag and lift forces on microscopic bubbles entrained by a vortex', *Phys Fluids*, vol. 7, 389–399

97. Stewart, PSB (1968): 'Isolated bubbles in fluidised beds — theory and experiment', *Trans Inst Chem Engrs*, vol. 46, T60–T66

98. University of Edinburgh (1993): 'Health and Safety policy; Part 7: radiation protection', published by the University of Edinburgh

99. van Dyke, M (1982): *An Album of Fluid Motion*, The Parabolic Press, Stanford

100. Wace, PF & Burnett, SJ (1961): 'Flow patterns in gas-fluidised beds', *Trans Inst Chem Engrs*, vol. 39, 168–174

101. Wen, CY & Hashinger, RF (1960): 'Elutriation of solid particles from a dense-phase fluidized bed', *AIChE J*, vol. 6, 220–226

102. Wen, CY & Yu, YH (1966): 'A generalized method for predicting the minimum fluidizing velocity', *AIChE J*, vol. 12, 610–612

103. Werlé, H (1989): 'Liquids', Ch. 4 of *Handbook of flow visualization* (Ed. Yang, WJ), Hemisphere Publishing Corp

104. Westerweel, J (1994): 'Efficient detection of spurious vectors in Particle Image Velocimetry data', *Exp Fluids*, vol. 16, 236–247

105. Willert, CE & Gharib, M (1991): 'Digital Particle Image Velocimetry', *Exp Fluids*, vol. 10, 181–193

106. Yamada, H & Matsui, T (1978): 'Preliminary study of mutual slip-through of a pair of vortices', *Phys Fluids*, vol. 21, 292–294

107. Yerushalmi & Cankurt (1979): *Powder Technol*, vol. 24, 187–190

108. Yórquez, MI & Duursma, GR (1999): 'Study of the flow pattern above an erupting bubble in an incipiently fluidized bed using image shifting', in preparation

109. Zenz, FA & Weil, NA (1958): 'A theoretical-empirical approach to the mechanism of particle entrainment from fluidised beds', *AIChE J*, vol. 4, 472–480

110. Zenz, FA & Othmer, DF (1960): *Fluidization and Fluid-particle Systems*, Reinhold Chemical Engineering Series

111. Zhang, Z & Eisele, K (1995): 'The two-dimensional velocity shift caused by the use of a rotating mirror in PIV flow field measurements', *Exp Fluids*, vol. 20, 106–111

112. Various (1961): 'Symposium on fluidisation — Discussion', *Trans Inst Chem Engrs*, vol. 39, 223–240

Appendix

The following two papers are appended to this thesis:

- McCluskey, DR, Hind, AK, Rix, SJL, Jakobsen, ML, Christy, JRE, Easson, WJ, Glass, DH & Greated, CA (1994): 'Development and application of a simultaneous two phase PIV system for the analysis of air-particle flow fields', IMechE Symp on *Optical Methods and Data Processing in Heat and Fluid Flow*, London, 99–109.

  This material is taken from the IMechE Conference Transactions 'Optical Methods and Data Processing in Heat and Fluid Flow'. ISBN 1 86058 142 0. It is reproduced with the kind permission of Professional Engineering Publishing Limited, UK (formerly known as Mechanical Engineering Publications Limited), publishers to the IMechE.

  For more information on this or any other title produced by Professional Engineering Publishing Ltd, please telephone: +44 (0) 1284 724384, e-mail: sales@imeche.org.uk or fax: +44 (0) 1284 704006.

- Rix, SJL, Glass, DH & Greated, CA (1996): 'Preliminary studies of elutriation from gas-fluidised beds using Particle Image Velocimetry', *Chem Eng Sci*, vol. 51, 3479–3489.

  Copyright © 1996. Reproduced from the journal *Chemical Engineering Science* (vol. 51, no. 13) with the kind permission of Elsevier Science.

  For more information, please write to the Subsidiary Rights Department, Elsevier Science Ltd, PO Box 800, Oxford, OX5 1DX, UK, send e-mail to: permissions@elsevier.co.uk or fax: +44 (0) 1865 853333.

C485/048

Development and Application of a Simultaneous Two Phase PIV System for the Analysis of Air-Particle Flow Fields

by D McCluskey, PhD BSc MInstP CPhys, A K Hind, MEng MA, S Rix, MEng BA, M L Jakobsen, PhD MA, J R E Christy, PhD MA, W J Easson, PhD BSc, D H Glass, PhD MA and C A Greated, DSc PhD BSc FRSE MICE FInstP FIMA, *The University of Edinburgh, UK*

SYNOPSIS: This paper presents a PIV analysis method which produces an instantaneous velocity map of both phases of an air-particle flow field. In addition to the usual criteria for obtaining good PIV records of single-phase flows, further requirements need to be satisfied when studying two-phase flow fields. The main extra constraint is that the images of the air-phase seeding must be distinct from those of the particle phase, in terms of both image size and intensity. After discussion of the two-phase PIV analysis technique, results are presented from a particle-laden jet and a fluidised bed.

NOTATION

$\boldsymbol{d}$	Displacement vector of successive PIV images on the photographic negative
$\boldsymbol{u}$	Local flow velocity vector $\boldsymbol{u} = (u, v, w)$
M	Image:object magnification
Δt	Time interval between laser pulses
d_p	Diameter of particles in the flow
d_τ	Diameter of particle images on the negative
d_I	Diameter of interrogation area
d_g	Diameter of grains on photographic film
$f_\sharp$	F-number of the imaging lens system
Δz_0	Thickness of laser sheet
λ	Wavelength of laser light
U_g	Superficial gas velocity through a fluidised bed

All the above are in SI units, except for M and $f_\sharp$, which are dimensionless.

1. INTRODUCTION

The structure of air-particle flow fields is of relevance to many industrial processes, such as pulverised coal combustion and fluidised beds. Of the velocity measurement methods reviewed by Taylor (1), the most advanced non-intrusive optical technique is Phase-Doppler Anemometry (PDA), which is widely used as an experimental tool. A fundamental limitation of point measurement techniques such as PDA is that it is not possible to obtain the instantaneous structure of the flow field. This issue is particularly important when considering the effect on the nature of the flow field of large scale turbulent structures, such as secondary flows in bends and recirculation zones in coal burner combustion. Furthermore, the effect of turbulent eddies on the behaviour of the particle phase is of fundamental importance in fluidised beds and many other two-phase flow regimes, but as yet there is little detailed information to quantify the extent of these interactions. For

these reasons, the authors have chosen to move from point Eulerian measurement methods to instantaneous whole field techniques.

2. PARTICLE IMAGE VELOCIMETRY

Particle Image Velocimetry (PIV) has emerged as a powerful experimental tool which can provide an instantaneous velocity map of a flow field (2). The basic concept of conventional PIV is that a two-dimensional section of a seeded flow field is illuminated stroboscopically and a multiple exposure photograph is taken. If a region of the flow has a finite local velocity $\boldsymbol{u}$, each seeding particle will move a distance $\boldsymbol{d} = \boldsymbol{u}\Delta t$ in the time interval Δt between successive illuminations. By examining small regions of the processed photographic image, the local average particle displacement and hence the local velocity vector can be determined. By successively analysing small regions of the processed image, the velocity vectors throughout the flow field are determined.

Since PIV detects the movement of particles, whether naturally present in the flow or deliberately introduced as seeding, it is well suited to recording the particle phase of an air-particle system. Furthermore, when analysing particle flow fields, it is also possible to determine concentration boundaries (3). PIV can be extended to study both the air and the particle phases of a two-phase flow field. However, just as for single-phase measurements, it was found to be important to consider data acquisition constraints in order to obtain good quality experimental data.

2.1 Air-particle PIV data acquisition procedure

PIV is a highly accurate velocity measurement technique; reports have cited r.m.s. errors of 0.5–1.0% of the full-scale velocity (4). However, the actual velocity error is strongly influenced by the experimental procedure used to obtain the PIV recording. That is, the errors in a set of PIV measurements depend upon the illumination system and the recording parameters, as well as the nature of the flow field. Keane and Adrian (5, 6) have analysed the single-phase PIV technique statistically and have established a number of criteria which should be satisfied in order to minimise experimental error:

1. In order to obtain a statistically significant mean velocity vector, the number of effective particle image pairs in each interrogation region should be greater than 15.

2. The magnification (M) and the time interval between exposures (Δt) should be chosen together, so that the separation of successive particle images on the photographic film ($\boldsymbol{d}$) for a given diameter of the interrogation region (d_I) avoids an excessive loss of image pairs due to in-plane motions:

$$|\boldsymbol{d}|_{max} = M|\boldsymbol{u}|_{max}\Delta t < 0.25 d_I$$

3. If the flow field is 3-dimensional, Δt and the thickness of the illuminating light sheet (Δz_0) should satisfy

$$w\Delta t/\Delta z_0 < 0.25$$

where w is the out-of-plane velocity component of $\boldsymbol{u}$.

4. The magnification and the diameter of the interrogation region should be designed for the particular flow field under investigation. The velocity variations within each interrogation region should be less than 5%, to avoid loss of signal and the introduction of various bias effects.

5. For analysis systems which interrogate the PIV negative directly, the ratio between the size of particle images and the interrogation area should also be small, to avoid another bias effect due to the spatial weighting of light within the interrogation beam (which is usually Gaussian). This is illustrated in Table 1 for some relevant particle sizes; the effect may be compounded by intensity variations across the illumination plane.

d_p (μm)	M	$f_\sharp$	d_τ (μm)	Velocity Error (%)
1	1	5.6	15	0.02
1	0.5	5.6	13	0.01
1	0.25	5.6	11	0.01
75	1	5.6	78	0.6
75	0.5	5.6	40	0.2
75	0.25	5.6	22	0.04

Table 1: The effect of recording parameters on the size of particle images and the bias error; illumination by an Argon-Ion laser ($\lambda = 480nm$); $d_I = 1mm$; film has 125 lines/mm.

In the authors' experience, further requirements for good single-phase PIV results are:

6. A minimum image separation on the photographic film (related to a minimum velocity in the flow field) has been found, below which distinct peaks cannot be obtained in the auto-correlation function produced by the analysis. In the authors' experience, this limit is given by

$$|\boldsymbol{d}|_{min} = M|\boldsymbol{u}|_{min}\Delta t > 2d_\tau$$

where d_τ is the diameter of the particle images on the negative, given by

$$d_\tau = [(Md_p)^2 + (2.44(1+M)f_\sharp\lambda)^2 + {d_g}^2]^{1/2}$$

where $f_\sharp$ is the focal number of the camera and d_g is the average grain size on the photographic film. This equation assumes that the particles do not move significantly during each illumination pulse. The image size on the film is strongly influenced by the choice of recording parameters, as shown above in Table 1.

7. The image intensity on the photographic film (in this case, from both the particles and the air-phase seeding) should meet the detectability criterion discussed in (7), *i.e.* that 45% of a Gaussian image spot should exceed the gross fog level of the film.

The authors are proposing that the following further requirements should be satisfied in order to extend the PIV technique to the study of two-phase air-particle flows.

8. The images of the air-phase seeding should be distinct from those of the particles.

9. When photographing air-particle flow fields, the illumination and recording parameters should be selected so as to suit both phases of the flow field as far possible. The analysis parameters can be optimised independently for each phase.

2.2 PIV analysis of air-particle flows

The analysis of two-phase PIV recordings requires that the images of each phase can be separated from one another. This is difficult to achieve using optical interrogation methods, so the authors investigated ways of digitising a photographic image, to allow further processing before analysis.

Early work (8) used image intensity alone to separate the phases, based on the fact that, for particles much larger than the wavelength of the illuminating light, scattering efficiency is proportional to surface area. A particle phase will therefore normally produce a much more intense photographic image than the air-phase seeding[1]. The particle phase was extracted by applying a simple intensity threshold, whilst the air phase was obtained by removing all image data above a particular intensity and then removing the edges of the particle images (found by a combination of standard edge detection techniques and thresholding). However, this method is not robust where there are overlapping images or variations in image intensity across the PIV negative.

The method developed in this paper separates the phases in a more sophisticated way: images which satisfy a selectable combination of criteria for intensity, size and shape are associated with the generally larger

[1] Provided that the smaller phase does not produce saturated images on the film and that the scattering efficiencies of the materials are comparable. The particle images will also be larger: see Section 2.1, item 4.

and brighter particle phase. Techniques recently used for particle tracking (9), which can resolve overlapping images, have been used to detect particles in the image; the centre of each particle is then found by volume moments, and characteristic lengths in two directions are calculated using the second volume moment. The original image gives particle brightness information, whilst the two orthogonal measurements of diameter yield a size and an estimate of the roundness of the particle. Overlapping images are resolved using this eccentricity factor and the gradient detection part of the algorithm used to identify particle images.

A file, containing only air-phase image data, is produced from the original file, by deleting around each particulate image a region whose radius is a definable fraction greater than that calculated; this ensures that all the information associated with the particle phase is removed. This air-phase file can then be subtracted from the original image data to produce a second file, containing particle information, which can be used for conventional PIV analysis or for particle tracking.

The eccentricity of images may also be used to separate phases with appreciable velocity differences, where the images associated with one phase are circular and those with the other are mis-shapen.

Figure 1 shows a section of the air-particle PIV image obtained in Section 3.1, together with the two phases after separation.

3. EXPERIMENTAL RESULTS AND ANALYSIS

3.1 Particle-laden jets issuing into a co-current air stream

A particle-laden jet was injected, at a velocity of $25ms^{-1}$, via a $5mm$ bore tube, into the centre of a $52mm$ diameter duct, co-current with a turbulent air flow. The air flow was travelling at a mean velocity of $6ms^{-1}$ and its fluctuations were grid-generated. The particles in the jet were solid glass spheres with a density of $2500kgm^{-3}$, a mean diameter of $76\mu m$ and a loading of $2kgm^{-3}$. The turbulent air stream was seeded with $1\mu m$ corn oil droplets and was produced by the wind tunnel and grid arrangement described in (10).

The air-particle flow field was illuminated by a pulsing light sheet from two double-Q-switched Nd:YAG lasers. A cylindrical lens converted the train of 4 pulsing beams into a train of pulsing light sheets, which were then collimated and reflected into the flow field by a parabolic mirror, as shown in Figure 2a. The light sheet was approximately $0.5m$ wide and $1mm$ thick; the pulse duration was $8ns$, with the time interval between each pulse selected to be $0.10ms$.

PIV records of the air-particle flow field were recorded onto $35mm$ Kodak T-Max $100ASA$ film, using a Nikon camera fitted with an iris shutter and a $50mm$ lens. The image:object magnification was 1:4. Figure 3a shows a print made from one of these negatives. The experimental results are discussed in Section 4.

3.2. The flow field above a fluidised bed

Fluidised beds are commonly used in the chemical industry to carry out reactions between a fluid phase (often a gas) and a solid phase; examples include hydrocarbon cracking (where the solid phase consists of catalyst particles) and coal combustion (when the solid phase is burning coal particles). Very small bed particles can be elutriated, *i.e.* carried out of the fluidised bed by the gas flow. Problems can occur if these particles enter downstream processing equipment (*e.g.* the erosion of turbine blades), whilst if an intermediate cyclone or filter is used, it needs to be sized according to the solids loading. The purpose of this research project is to investigate air velocity profiles above bubbling fluidised beds, with a view to developing schemes to minimise elutriation.

In a two-phase flow such as this, it is useful to be able to identify particle images in the region above the bed (known as the freeboard), in order to determine the particle concentration. Separating these images from the air-phase data allows the air flow field to be measured more accurately.

A seeded air flow above a $75mm$ square-sided fluidised bed was illuminated using the scanning beam technique (11): a rotating octagonal mirror scanned a $3W$ argon-ion laser beam through the central cross-section of the bed (Figure 2b). PIV photographs were taken at a range of heights above the unfluidised surface and for two different bed depths. The bed consisted of solid glass spheres with diameters in the range $150-325\mu m$, which were fluidised by air at a superficial velocity $U_g = 0.07ms^{-1}$. The air phase was seeded by hollow glass spheres with a mean diameter of $8\mu m$ and a solid density of $1100kgm^{-3}$. Seeding was mixed in with the bed fill material; when the fluidising air flow was turned on, the seeding was gradually released into the freeboard region and was carried out with the flow. The resulting flow field was photographed onto $35mm$ Kodak T-Max $100ASA$ film, at an image:object magnification of 2:5, by a Nikon camera fitted with a $55mm$ lens.

Figure 5 is a photograph of the freeboard region about 180mm above the bed. This was analysed to yield the vector map shown in Figure 6.

4. DISCUSSION

Figure 1 shows that the phases of a two-phase PIV recording can be separated successfully using the algorithm outlined in section 2.2. The air-phase seeding clearly contains little residual particle information after the larger phase has been removed. However, the particle phase still contains a small number of images from the air phase, a result of removing from the original data file an image area slightly larger than that detected by the algorithm. This ensures that the air-phase data, whose analysis would be greatly affected by undersized deletions, retains its integrity, at the expense of an artificially-increased signal drop-out. In contrast, if the particle phase contains one or two air-phase seeding particles, which are small in comparison, there is little effect on the results of the analysis. This approach minimises the errors introduced as a result of the phase separation.

Figure 3a shows a photographic print of a PIV negative of a particle-laden jet injected into a seeded co-current airstream. The particle phase and air-phase seeding are readily distinguishable from one another. The air-phase seeding satisfies all the recording and analysis crtieria discussed in Section 2.1, except in the middle of the particle-laden jet, where the number density of the seeding is low (this could be avoided in future studies by seeding the jet air flow). The number density of the particle phase does not satisfy some of the constraints discussed above in regions where isolated particles have been dispersed from the jet. This situation is a consequence of the nature of the flow field and particle tracking algorithms may be a more suitable means of determining particle velocity in such cases. The air and particle phases were separated using the method outlined in Section 2.2. Figure 3b shows the vector map of the particle phase, superimposed on the vorticity contour map of the air phase. Plotting a contour map of the air phase seeding is possible because the velocity values of the air flow were accurately determined throughout the majority of the flow field. Presenting the data in this format readily displays the differing behaviour of the two phases of the flow: there is little noticeable interaction between the two phases of the flow field; the particle jet is gradually dispersing as it travels downstream from its injection nozzle; the few particles which have been dispersed from the jet are travelling with a velocity around half the values of those travelling in the jet.

A more complex air-particle flow field is shown in Figure 4a. In this photograph of the flow field above a fluidised bed, the air phase shows many interesting and varied features: there is a vortex on the left, rapid vertical flow in the central region and noticeable horizontal motion near the right-hand wall. There is also significant out-of-plane motion in some regions of the flow. Very few bed particles are found at this height, and those which are present have significantly different trajectories from one another. Furthermore, their velocities are lower than those of the air phase.

In this preliminary study, image shifting techniques were not deployed. The "reverse flow" regions could be identified with the aid

of simultaneous video tape footage. In future studies, the direction of the flow field will be determined using the rotating mirror image-shifting method, which will also avoid signal drop-out in regions of low velocity.

With regard to the recording criteria outlined in Section 2.1, in both of the investigations presented, the first constraint was satisfied by selecting interrogation areas for each phase which were large enough to contain at least 5 distinct sets of particle images. Using four illumination pulses per particle gives 15 or more effective particle image pairs in the each interrogation area and also gives a good signal-to-noise ratio in areas of the particle phase which are relatively sparsely-seeded. The values of M, Δt and Δz_0 were chosen to satisfy criteria 2, 3 and 4 for both phases. Some two-phase flows may exist where these conditions may not be met simultaneously; for these, a judicious compromise will be required when setting these parameters. A small size was chosen for d_I, to reduce velocity gradients (and the resulting bias) in the analysis area.

The analyses were performed digitally, *i.e.* without the spatial light intensity variation that would be introduced by an optical analysis technique. The criteria for producing good two-phase recordings are the same as for producing high contrast, evenly illuminated conventional PIV recordings.

Digital image recording, or subsequent processing of the PIV negative to give a digital image, allows two-phase images to be separated to a degree that would be difficult to achieve using optical techniques.

5. CONCLUSION

A PIV method has been presented, which can separate and analyse each phase of an instantaneous two-phase PIV image. The authors have discussed the general PIV criteria which should be met for both phases in order to obtain reliable and good quality data, and have included further criteria for good separation of the two phases.

Two relevant applications have been presented: a particle-laden jet issuing into a co-current air flow and the flow field above a fluidised bed. This research has demonstrated the principle of the analysis method and has indicated its value when applied to two-phase PIV in general. Further experimental research will involve other relevant topics, such as particle transport in pipe lines, swirling flows and the conveyance of aspherical particles. These will provide the PIV analysis technique with further challenges, due to significant out-of-plane motion and the increased difficulty of phase separation when scattering efficiency depends on the orientation of non-spherical particles.

6. ACKNOWLEDGEMENTS

The authors would like to thank the following for funding the work described in this paper: The Science & Engineering Research Council (Fellowship Section), Shell UK Exploration and Production, and the Science & Engineering Research Council (Process Engineering Committee).

The authors would also like to thank Mr D Anderson, for his workmanship in constructing the wind tunnel rig, and Mr F Morris, for maintaining the Argon Ion and Nd:YAG lasers.

7. REFERENCES

1. Taylor, AMKP, 'Optically-based measurement techniques for dispersed two-phase flows', Proc. NATO Advanced Study Institute on Combusting Flow Diagnostics, Kluwer Academic Publishers, 1992

2. Adrian, RJ, 'Particle-imaging techniques for experimental fluid mechanics', Ann. Rev. Fluid Mech., vol. 23, pp261–304, 1991

3. McCluskey, D R, Easson, W J, Greated, C A, Glass, D H, 'The characterisation of particle jet dispersal', presented at the 6th Int. Symp. on Applications of Laser Techniques to Fluid Mechanics, Portugal, July 1992

4. Adrian, RJ, 'Image shifting technique to resolve directional ambiguity in double-pulsed velocimetry', Appl. Opt., vol. 25, pp3855–3858, 1986.

5. Keane, RD, Adrian, RJ, 'Optimisation of Particle Image Velocimeters, Part I: Double-pulsed Systems', Meas. Sci. Technol., vol. 1, pp1202–1215, 1990

6. Keane, RD, Adrian, RJ, 'Optimisation of Particle Image Velocimeters, Part II: Multiple-pulsed Systems', Meas. Sci. Technol., vol. 2, pp963–974, 1991

7. Adrian, RJ, Yao C-S, 'Pulsed laser technique application to liquid and gaseous flows and the scattering power of seed materials', Appl. Opt., vol. 24, pp44–52, 1985

8. Hind, A K, McCluskey, D R, Christy, J R, Easson, W J, Greated, C A, Glass, D H, 'PIV analysis of two-phase air-particle flows', IMechE symposium on optical methods and data processing in heat and fluid flow, City University, London, 2–3 April 1992

9. Maas, HG, Gruen, A, Papantoniou, D, Particle Tracking Velocimetry in Three-Dimensional Flows, Expts in Fluids, vol. 15, pp133–146, 1993

10. McCluskey, DR, Elgaard, C, Easson, W J, Greated, CA, 'The application of PIV to turbulent two-phase flows' in Flow Visualization and Image Analysis, Ed FTM Nieuwstadt, Kluwer Academic Publishers, Dordrecht, The Netherlands, 1993

11. Gray, C, Greated, CA, McCluskey, DR, Easson, WJ, 'An analysis of the scanning beam PIV illumination system', Meas. Sci. Technol., vol. 2, pp717–724, Aug 1991, Re-printed in Engineering Optics, Vol 4 (4), Nov 1991

5. Keane, RD, Adrian, RJ, 'Optimisation of Particle Image Velocimeters, Part I: Double-pulsed Systems', Meas. Sci. Technol., vol. 1, pp1202–1215, 1990

6. Keane, RD, Adrian, RJ, 'Optimisation of Particle Image Velocimeters, Part II: Multiple-pulsed Systems', Meas. Sci. Technol., vol. 2, pp963–974, 1991

7. Adrian, RJ, Yao C-S, 'Pulsed laser technique application to liquid and gaseous flows and the scattering power of seed materials', Appl. Opt., vol. 24, pp44–52, 1985

8. Hind, A K, McCluskey, D R, Christy, J R, Easson, W J, Greated, C A, Glass, D H, 'PIV analysis of two-phase air-particle flows', IMechE symposium on optical methods and data processing in heat and fluid flow, City University, London, 2–3 April 1992

9. Maas, HG, Gruen, A, Papantoniou, D, Particle Tracking Velocimetry in Three-Dimensional Flows, Expts in Fluids, vol. 15, pp133–146, 1993

10. McCluskey, DR, Elgaard, C, Easson, W J, Greated, CA, 'The application of PIV to turbulent two-phase flows' in Flow Visualization and Image Analysis, Ed FTM Nieuwstadt, Kluwer Academic Publishers, Dordrecht, The Netherlands, 1993

11. Gray, C, Greated, CA, McCluskey, DR, Easson, WJ, 'An analysis of the scanning beam PIV illumination system', Meas. Sci. Technol., vol. 2, pp717–724, Aug 1991, Re-printed in Engineering Optics, Vol 4 (4), Nov 1991

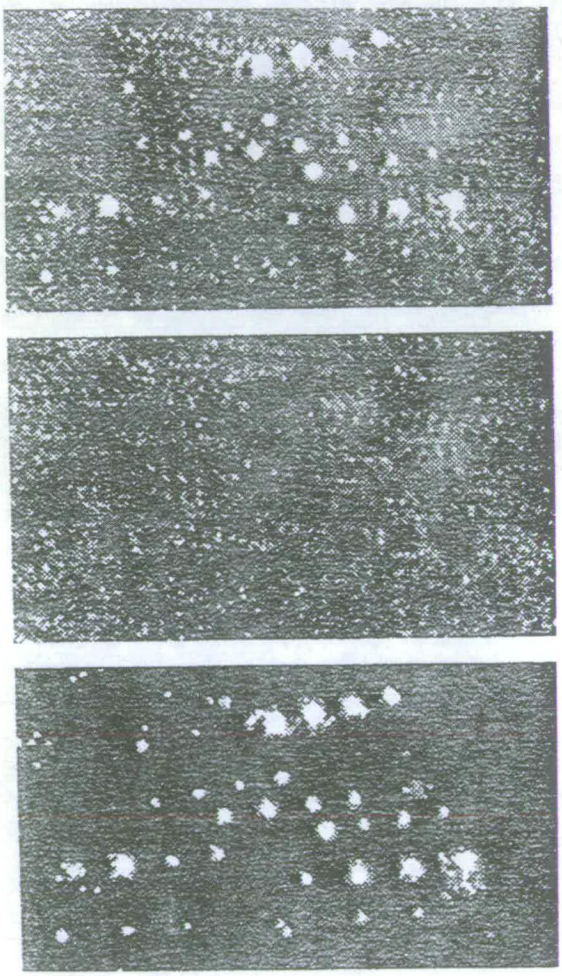


Figure 1: The upper picture shows both phases in a section of the raw PIV image. The middle and lower pictures show the corresponding air- and particle-phase respectively when they have been separated by the image process.

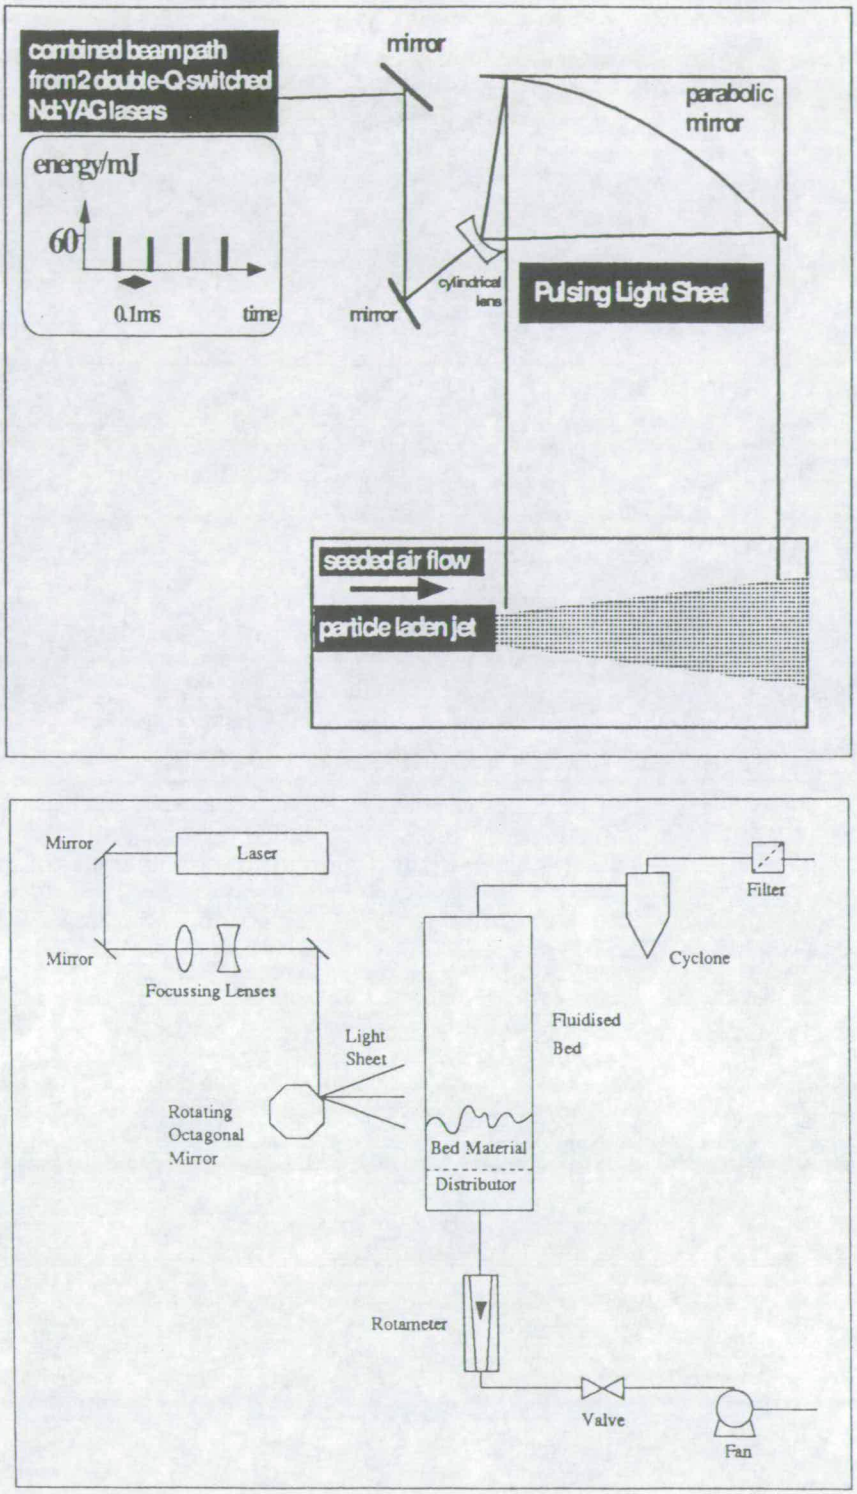


Figure 2: The experimental setups. (a): The upper figure shows setup of the jet. (b): The lower figure shows the setup of the fluidised bed.

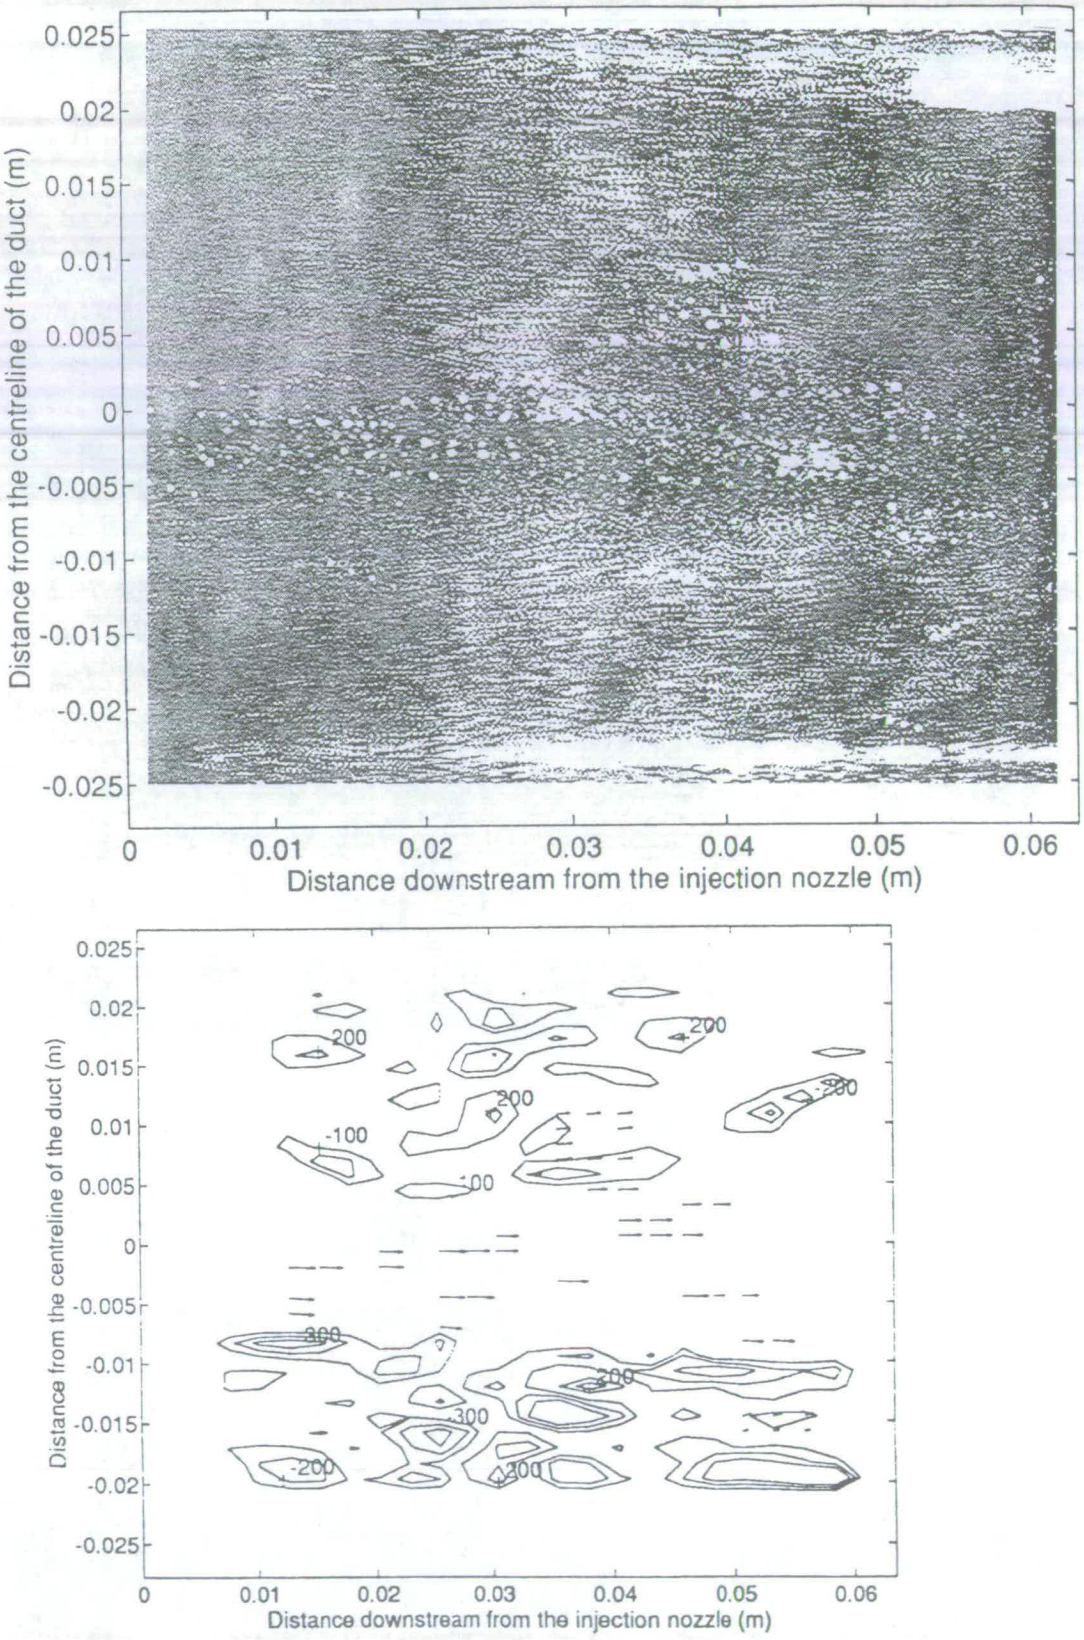


Figure 3: (a): PIV recording of the particle laden jets issued into a co-flowing airstream. (b): Corresponding vector and vorticity contour map (s^{-1}) of the particle and air phase respectively.

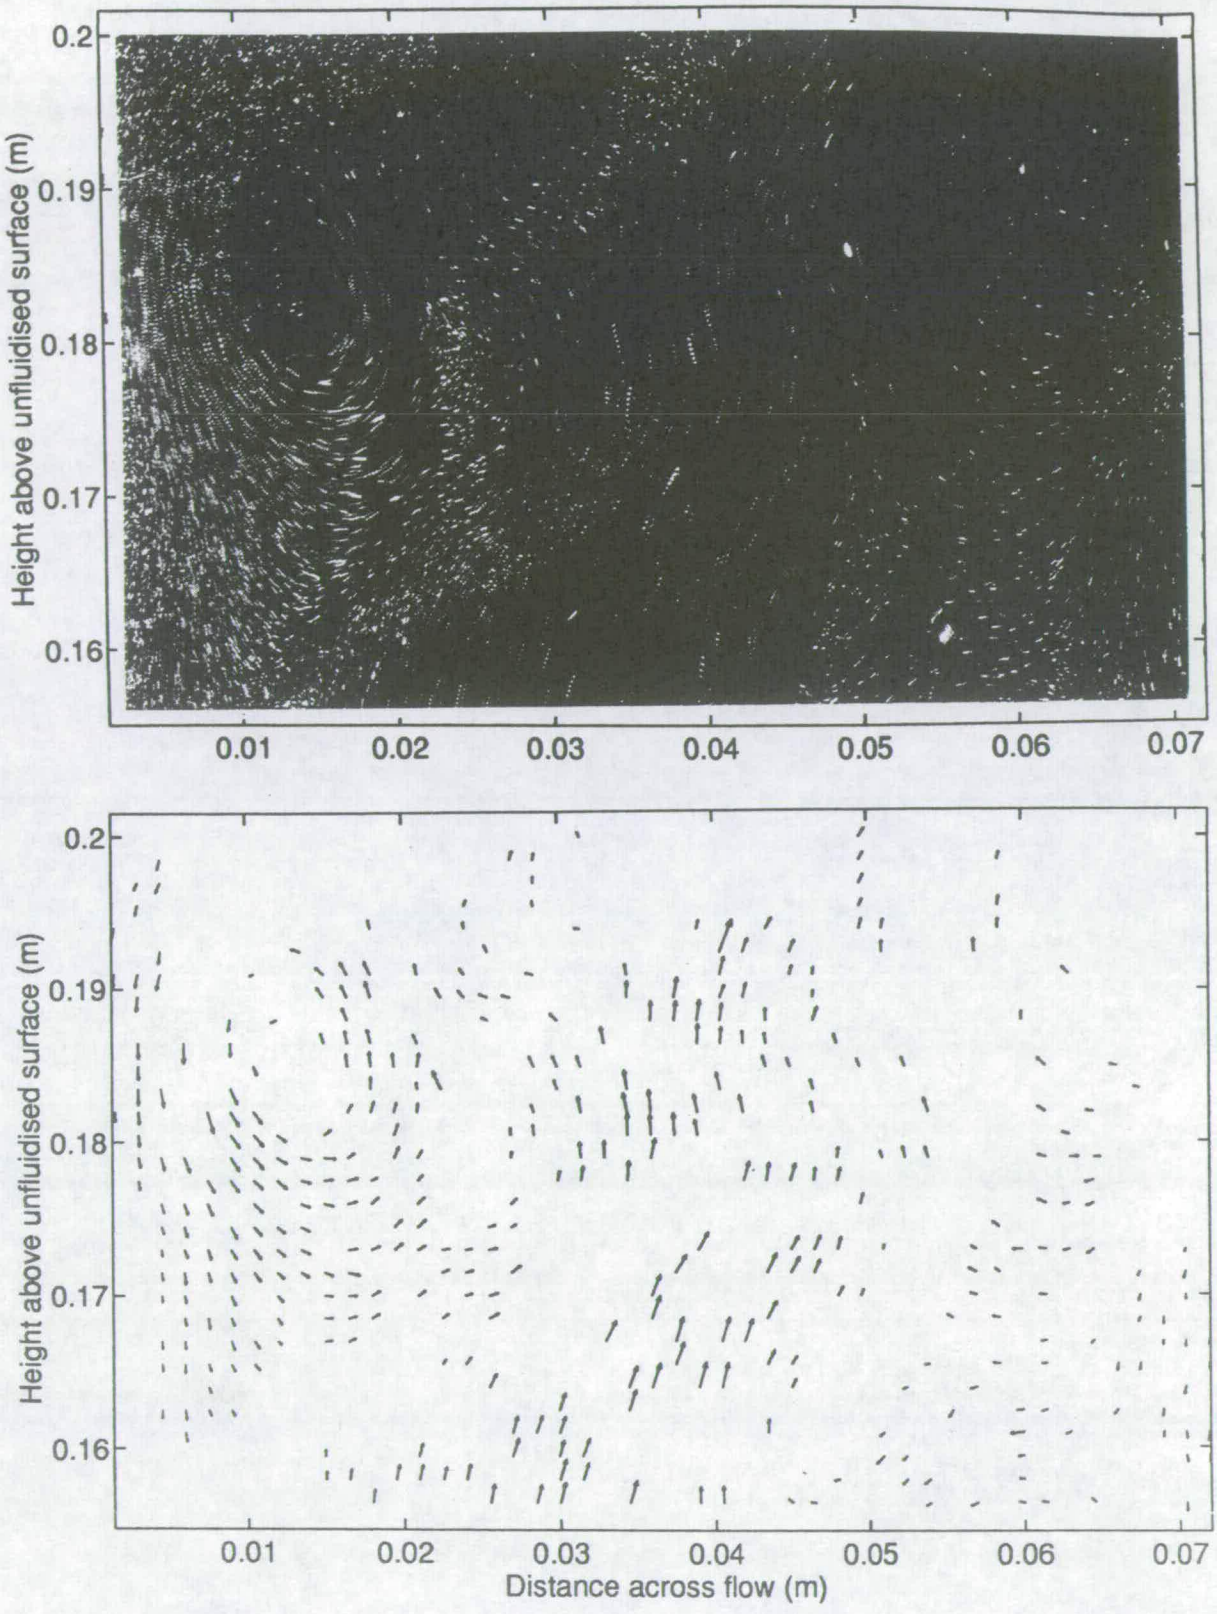


Figure 4: (a): PIV recording of the free-board region above the fluidised bed. (b): Corresponding vector map.

Pergamon

Chemical Engineering Science, Vol. 51, No. 13, pp. 3479–3489, 1996
Copyright © 1996 Elsevier Science Ltd
Printed in Great Britain. All rights reserved
0009–2509/96 $15.00 + 0.00

0009–2509(95)00405–X

PRELIMINARY STUDIES OF ELUTRIATION FROM GAS-FLUIDISED BEDS USING PARTICLE IMAGE VELOCIMETRY

S. J. L. RIX and D. H. GLASS*
Department of Chemical Engineering, The University of Edinburgh, King's Buildings, Mayfield Road, Edinburgh, EH9 3JL, U.K.

and

C. A. GREATED
Department of Physics and Astronomy, The University of Edinburgh, King's Buildings, Mayfield Road, Edinburgh, EH9 3JZ, U.K.

(*First received* 7 *July* 1995; *revised manuscript received* 15 *November* 1995; *accepted* 27 *November* 1995)

Abstract—In this work, the method of Particle Image Velocimetry has been used to produce velocity maps of the complex flow above a small, square-sided fluidised bed. The results presented illustrate the nature of the flow at two different heights above the bed surface. Copyright © 1996 Elsevier Science Ltd

1. INTRODUCTION

Previous studies of gas velocities in the freeboard region have used point measurement techniques, such as Laser Doppler Anemometry (LDA) and Hot Wire Anemometry [e.g. Lang (1986) and Pemberton and Davidson (1984)]. These yield a time history at a point in the flow, from which mean and RMS-fluctuation velocities can be calculated. However, a disadvantage is that readings at different points in the flow are not simultaneous.

If the flow being measured is reproducible, then this is not a problem, but if the velocity fluctuations are as large and as chaotic as they appear to be above a fluidised bed, then measurements of gas velocity at different positions in the freeboard, taken at different times, may not be fully comparable because the flow fields might not be identical. Further, the structure of the flow is not known; it is not clear whether the maximum velocity values at a point A in the flow coincide with maxima at some other point B, or whether there is a phase difference.

Intuitively, one might expect the gas flow immediately above a fluidised bed to be determined by the bubbling properties of the bed. At large heights above the bed surface, one would expect the turbulence to have decayed down to "normal" levels for flow in a pipe; in an established turbulent pipe flow, these fluctuations might be only $\pm 5\%$. Using Phase Doppler Anemometry, Ramírez De Santiago (1990) has reported gas velocity fluctuations of $\pm 1200\%$ of the mean for the flow of gas in the region immediately above a gas-sparged bubbling liquid. This situation is similar to fluidisation, except that the "liquid" has a non-zero surface tension.

Pemberton and Davidson (1984) measured the vertical gas velocity above a fluidised bed using a hot wire anemometer. They compared the decay in freeboard turbulence with height with that predicted by a rising "ghost bubble", which is a form of vortex ring. Photographed in air and in water by many workers [e.g. Maxworthy (1972)], vortex rings are structures which propagate through a medium without losing their character. Vortex rings can still be observed large distances downstream from their source [e.g. Reynolds (1876) tracked some in water moving over distances of up to 6 m].

In ghost bubble theory, recirculation zones in the flow in the freeboard are due to vortex rings or turbulent "puffs", which are formed by "real" bubbles bursting at the surface of the fluidised bed. Once released from the bed surface, ghost bubbles continue to rise as if they were still in the fluidised bed.

Usually, the word "bubble" implies a gas moving through a liquid; here, a ghost bubble consists of the same gas as the surrounding medium. After escaping from a fluidised bed, a ghost bubble is assumed to keep on rising, maintaining its vortex ring nature whilst entrapping gas from the freeboard region, with which it shares its momentum. Reynolds (1876) calculated that the momentum of his rings was conserved during their propagation, whilst their velocity fell from over $1.5\,m\,s^{-1}$ to about $0.07\,m\,s^{-1}$. Consequently, ghost bubbles should increase in size and decelerate as they rise (Fig. 1).

*Corresponding author. Tel.: 0131-650-4860; Fax: 0131-650-6551; E-mail: Don.Glass@ed.ac.uk.

In fluidisation, a bubble rising through a bed has an internal gas recirculation due to particle motion past the bubble surface and the throughflow of gas from the dense phase (the void space has the lower resistance to flow). Pyle and Rose (1965) derived stream-functions for the inviscid potential flow of gas around and through a small bubble in two- and three-dimensional fluidised beds, in a moving frame of reference which made the bubbles appear to be stationary ("the bubble frame"). The left-hand side of Fig. 2 is an air-velocity vector map obtained by differentiating their stream functions for the three-dimensional case and inserting values for the relevant constants.

The experimental work performed here has yielded flow patterns in a stationary frame of reference ("the laboratory frame"). The right-hand side of Fig. 2 differs from the left only in that the velocity vectors have been shifted from the bubble frame to the laboratory frame by the addition of an upward velocity of magnitude U_b to all vectors. After this transformation, the vortex ring appears to be much less pronounced.

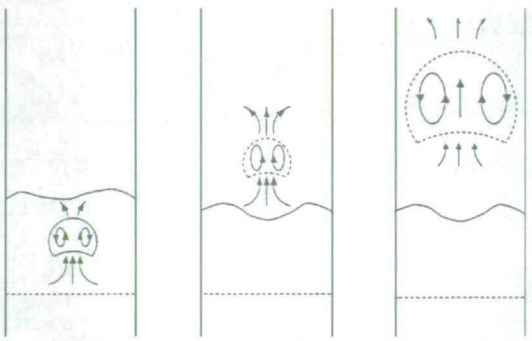


Fig. 1. Ghost bubbles from a fluidised bed and their influence on freeboard velocity.

The features of this figure are the recirculation of gas within the spherical "limit of penetration" (the outer circle in Fig. 2) and potential flow of the other gas around this same sphere. The inner circle represents the bubble boundary, around which particles flow as if round a sphere. There is a step change in values of void fraction across this boundary, but the tangential gas velocity is continuous, as is the stream function if suitably defined so as to take into account the sudden change of voidage; radial velocity is reduced by multiplication by the void fraction. The size of the limit of penetration relative to that of the bubble is a function of the bubble rise velocity (U_b) and the interstitial gas velocity ($U_o = U_{mf}/\varepsilon_{mf}$), as discussed by Pyle and Rose (1965).

Collins (1982) extended the streamlines of Pyle and Rose to the outside of a two-dimensional bubble to include extra source and sink terms; this allowed the particles to move outwards from the bubble surface whilst the gas moved inwards, modeling an expanding bubble. Collins (1983) later dealt with the three-dimensional case.

Azevedo and Pereira (1991) used slightly different boundary conditions when deriving stream functions inside the bubble. In their work and that of Collins

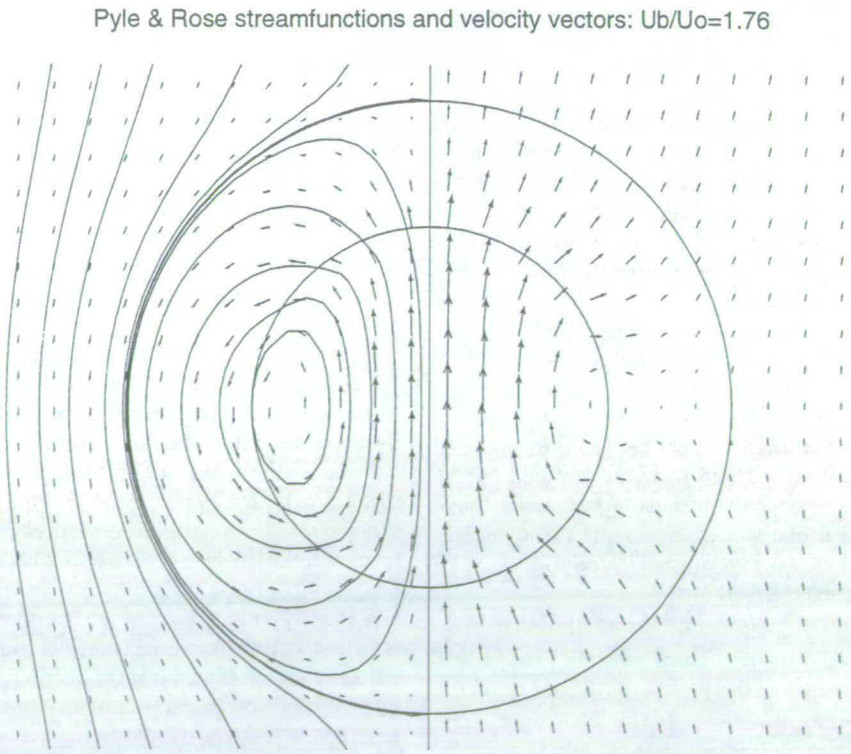


Fig. 2. Gas velocities after Pyle and Rose: bubble frame on the left, laboratory frame on the right.

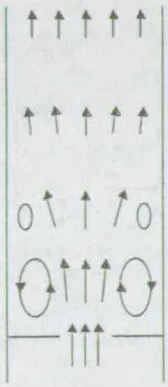


Fig. 3. Flow through an orifice.

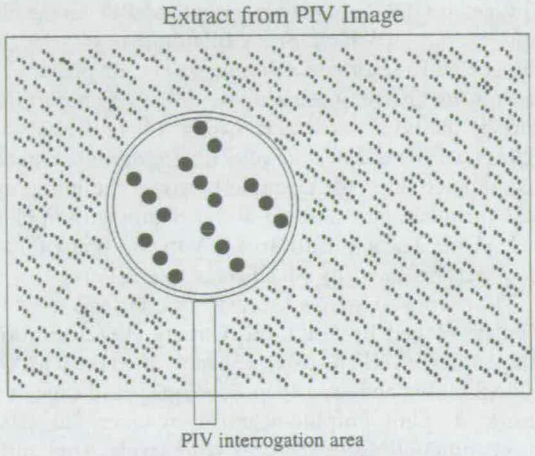


Fig. 4. Schematic PIV photograph.

(1982, 1983), the stream functions presented look similar to those of Pyle and Rose, which have been illustrated for reasons of simplicity.

Alternatively, the flow above a fluidised bed could be more influenced by the existence of a "jet" of bubbles than by the vorticity within each bubble; this was the model used by Zenz and Weil (1958) to model elutriation and the Transport Disengagement Height. A stream of bubbles could produce a flow resembling that of turbulent flow through an orifice. This would appear to generate a vortex ring (Fig. 3), albeit as secondary feature.

2. PARTICLE IMAGE VELOCIMETRY

In the work completed so far, the gas flow above a small fluidised bed has been investigated by directing a focused laser beam onto a multifaceted rotating mirror. The reflected beam behaves as a pulsing, diverging "sheet" of light, which is directed at the flow field. For large scale work, the light sheet can be collimated by reflecting it onto a parabolic mirror (Gray *et al.*, 1991), but this was not necessary for this work.

The flow is seeded with fine particles, which are small enough to follow local fluctuations in gas velocity but large enough to scatter a measurable amount of the incident light. Particles in the freeboard region appear to be illuminated stroboscopically by the light sheet, at an apparent flashing frequency controlled by the speed of the rotating mirror. Two or more successive images of each particle can be recorded photographically (Fig. 4).

Having obtained a record of the flow, velocities can be calculated by determining, for each element of a grid of interrogation areas, the mean displacement per illumination interval. Parameters such as the area photographed (i.e. magnification), exposure duration and rotating mirror speed need to be optimised for the particular flow and the method chosen for image analysis; this is often an iterative progress. The guidelines discussed by Keane and Adrian (1990, 1991) have been found to work well in practice. For a general discussion of PIV and other velocity-measuring techniques, the reader is referred to Adrian (1991).

The mean displacement per illumination interval referred to above can be determined in a number of ways. Direct visual examination of the PIV record would be very tedious, so researchers have sought more automated methods.

Some PIV researchers have used an optical Young's fringes analysis system, which works well for most single-phase flow fields (e.g. Greated *et al.*, 1992). The principle behind the system is that, when a coherent light source (e.g. a low-power laser beam) is passed through a portion of a PIV negative, each particle image acts as a source of scattered light; when any two such sources combine, Young's fringes are formed. The intensity of each possible fringe set depends on the number of source pairs having that particular relative displacement; the brightest set of fringes comes from the most frequently occurring displacements, which are those corresponding to the motion of seeding particles between two successive illuminations.

In the studies presented here, analysis was carried out in the digital domain. In this method, an image file is created on a computer, which records the brightness of each pixel of an array. The mean displacement per illumination interval can be determined by numerically calculating the autocorrelation function for each interrogation region, using Fourier transforms; this function has a maximum at a point one mean particle displacement away from the origin.

Digital analysis offers a number of advantages. The size of the PIV interrogation area can be varied more conveniently than in a Young's fringes method, because altering the input parameters of a computer program is easier than fitting a small iris to a laser beam used for optical interrogation, so as to limit the size of the interrogation beam.

Further, working digitally makes it possible to analyse both phases of a two-phase air–particle flow: pre-processing of a PIV record is possible, which can separate particle-phase images from images of the seeding (the former are generally significantly larger, brighter and less circular). Software described by

Hind *et al.* (1992) can read in a digitised PIV image file and output two files, one containing only particle images and the other containing only air-phase images. After this step, each phase of a two-phase flow can be analysed separately (using different analysis parameters if need be); the velocity vector maps can be recombined later, for example to assess the influence of a turbulent air phase on a higher-momentum particle phase. The application of PIV to two-phase flows has been reviewed by McCluskey *et al.* (1994).

The particular analysis software used in this work is that developed by Hind and Christy (1992). Photographs were taken using 35 mm Kodak TMAX 100ASA film; this was then developed and digitised using a photographic-negative scanner, at 100 pixels/mm and using 256 grey scale levels. After suitable image-processing, computer files were produced which could be analysed. The calculations were performed on the Department of Chemical Engineering's network of Sun Sparc stations.

Some researchers have started to use PIV to investigate flows of interest to the chemical engineer. Northrup *et al.* (1993) investigated the interstitial velocities of a liquid flowing through a packed bed, using fluorescent particles. Unfortunately, their technique is unlikely to be applicable to the flow within a three-dimensional gas-fluidised bed, because it relies on closely matching the refractive index of the fluid to that of the (static) bed material. Whilst some solids have a refractive index close to that of liquids, few have one close to that of air.

Although it is acknowledged that the technique used by Northrup *et al.* is unlikely to find immediate application in the field of gas-fluidisation, their paper also provides a description of the PIV technique. Being published in *Chemical Engineering Science*, this should be readily available to most readers of this journal.

Similar to PIV is the method of using light sheets to produce flow visualisation; using this technique, Horio and Kuroki (1994) videoed the freeboard of a fluidised bed. Light from a 10 mW He–Ne laser was expanded through a cylindrical lens to produce each light sheet; light sheets intersected in two 'T'-type arrangements, allowing some three-dimensional visualisation to take place around each intersection. Their work supported the ghost bubble theory of Pemberton and Davidson (1984).

3. EXPERIMENTAL APPARATUS

The apparatus used is depicted in Fig. 5 and has involved a number of elements.

Common to all the experiments was a small, square-sided fluidised bed, consisting of three 7.5 cm × 7.5 cm × 30 cm sections, stacked vertically in an interchangeable order. Two of these sections were made of perspex. The third section was made of glass, which is optically superior to perspex; this section was the one used for PIV measurements. On top of these sections sat a reducing duct, in which the area available for flow was gradually reduced until it met a 2.8 cm × 2.8 cm block, into which a 1.25 cm diameter exit nozzle was set.

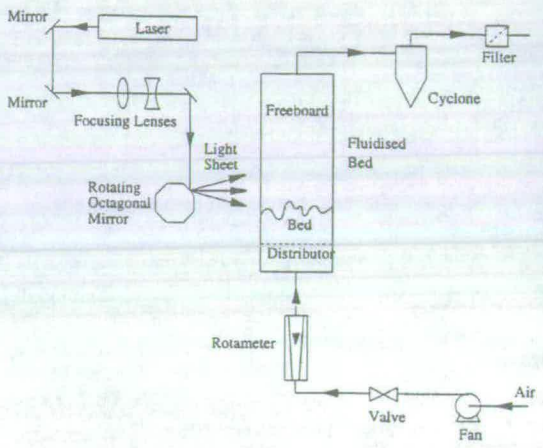


Fig. 5. Experimental apparatus.

Illumination was provided by a 15 W Argon-ion laser, which was focused and directed vertically downwards onto a rotating octagonal mirror. The axis of this mirror was horizontal, so that the reflected beam scanned through a thin vertical plane, approximately 0.5 mm wide. The apparatus was aligned so that the plane of this pseudo light-sheet passed through the centre of the square-sided bed and was parallel to two of the walls. Particles illuminated in this manner were photographed by a camera placed at right angles to the light sheet.

Fluidisation was carried out using air, compressed by a small fan. This was capable of delivering a superficial velocity of 0.03–0.25 $m\,s^{-1}$.

Bed fills were of glass ballotini (density $\sim 2500\ kg\,m^{-3}$), at depths of 3 and 5 cm. Two different particle sizes have been used, namely 53–106 μm and 150–325 μm. Based on drag coefficient calculations for spheres and a typical glass solid density of $2500\ kg\,m^{-3}$, these have terminal velocities $U_t \approx 0.4\ m\,s^{-1}$ and $\approx 4\ m\,s^{-1}$, respectively. Using the Ergun equation, minimum fluidising velocities were estimated to be $0.0067\ m\,s^{-1}$ and $0.070\ m\,s^{-1}$, respectively; the U_{mf} of the larger particles was then measured to be $0.0715\ m\,s^{-1}$, in good agreement. High U_g/U_{mf} ratios could be obtained using the smaller ballotini, resulting in very vigorous fluidisation; U_g/U_{mf} values closer to unity could be obtained with the larger ballotini, giving calmer conditions.

Small hollow glass spheres were used to seed the air flow in most experiments. These particles have a stated mean diameter of 8 μm, with 30% greater than 10 μm and 5% greater than 20 μm. The hollow nature of the seeding particles reduces their effective density to $1100\ kg\,m^{-3}$, which is less than half that of solid glass. Together with their smaller size, this means that the Stokes Law terminal velocities of the seeding were about 200 times lower than those of the solid 53–106 μm particles mentioned earlier. When

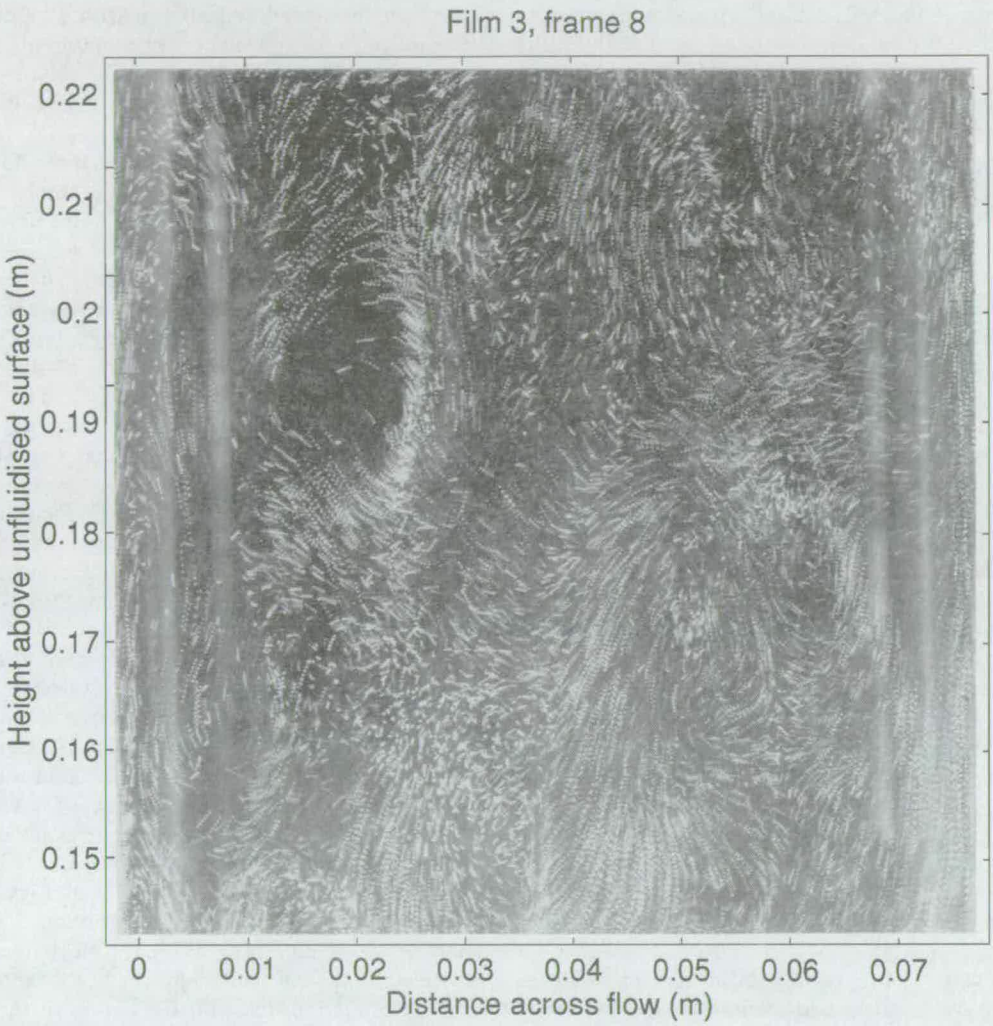


Fig. 6. PIV photograph (film 3, frame 8).

seeding a gas flow of 0.07 m s^{-1}, the terminal velocity of the seeding particles equates to about 4% of the flow velocity.

Methods of injecting these hollow spheres into a bubbling fluidised bed were considered, but it was simpler to mix some of these seeding particles into the bed with the gas feed off. By running each experiment as a batch process with the initial seeding concentration within the fluidised bed somewhat higher than the ideal, it was possible to ensure that the seeding concentration was satisfactory for an adequate period of time at some point during the course of each experimental run.

When a quantity of relatively-fine particles is mixed into a fluidised bed, the concentration of fines within (and above) the bed decreases as a function of time, due to elutriation. Leva (1951) showed that this decay of fines concentration in a mixed-size fluidised bed follows approximately first order kinetics; in other words, the fines concentration above the bed is expected to decay exponentially after start-up.

In PIV work, the concentration of seeding particles is quite important: too much seeding means that a cloud of particles between the camera and the light sheet acts as a fog obscuring the light sheet, whilst inadequate seeding in some regions of the flow means that the correlations calculated by the analysis method have a low signal-to-noise ratio and thus the velocities calculated are more prone to error. Keane and Adrian (1990, 1991) recommended that at least 15 particle image pairs should be present in each interrogation region, so as to allow a statistically-significant velocity vector to be calculated.

4. EXPERIMENTAL RESULTS

The first work performed was on unseeded 53–106 μm ballotini. Because of the vigorous fluidisation, many particles were ejected into the freeboard region, where they covered much of the illumination plane. A sample PIV photograph is shown in Fig. 6, which shows a complete cross-section through the

centreline of the bed, centred approximately 18 cm above the position of the bed surface. The fluidising velocity was 0.19 m s^{-1}, with $U_g/U_{mf} \sim 30$.

As can be seen from this photograph, the particle trajectories are certainly not uniform. A particle vortex can be seen just to the left of centre, whilst there is a strong upward flow slightly below right of centre. However, particles cannot form a stable fluidised bed ($U_t > U_g$) whilst at the same time perfectly following the gas flow ($U_t \ll U_g$), so it is not possible to calculate the actual gas flow from these particle trajectories.

Further, whilst a large number of particles were thrown from the bubbling bed into the freeboard, the vast majority of these soon disengaged from the gas flow and fell back down into the bed. For this reason, the particle trajectories recorded by Fig. 6 appear to cross one another in a number of places; it did not appear to be the case that particles rose in one region within the flow and fell back down in another (near the walls, perhaps).

For this and other reasons (such as the non-uniform distribution of particles across the measurement area), conventional PIV analysis (based on average velocities across small sample areas) was not very successful. But even though quantitative results could not be obtained, the illumination technique at least allowed the flow to be visualised and monitored for long periods of time. In this way, a ten-minute video recording was made, which illustrates features of the flow similar to those reported by Horio and Kuroki (1994), in particular the existence of well-lit clusters of particles surrounding small regions devoid of particles.

One way of progressing from the above work would have been to add seeding to the same bed under similar fluidising conditions, so as to attempt a two-phase analysis. Instead, to keep things simpler, it was decided to use the larger particle size as a basic bed fill, fluidised at a lower air flow rate ($U_g/U_{mf} \sim 1$, but bubbles were present). This dramatically reduced the number of bed particles in the freeboard, producing an essentially single-phase flow above the bed. This air phase was then seeded by mixing hollow glass spheres in with the bed fill, batchwise.

Two or more rolls of 24-exposure photographic film were taken at various heights above a bed of 150–325 μm Glass Ballotini, for bed depths of 3 and 5 cm. Fluidisation was carried out using a constant air flow of $U_g = 0.07$ m s^{-1}, very close to incipient fluidisation. PIV photographs in these sets of films were taken using a motor-driven camera on continuous exposure mode. The time between successive photographs was approximately 0.3 s, so it is possible to observe the development of the flow with time.

Sample PIV photographs and their resulting vector maps are shown in Figs 7–9. For Figs 7 and 8, the bed depth was 3 cm and the photographs recorded a region centred approximately 16–20 cm above the bed surface. These photographs were consecutive. Figure 9 was taken under the same fluidising conditions as Figs 7 and 8, but used a mean viewing height of about 31–35 cm. In each photograph, an artificial vector has been added to the top left corner, representing a velocity of 0.1 m s^{-1} vertically upwards; this vector is enclosed in a small box to isolate it from the flow vectors.

From these figures, it is apparent that the flow field is by no means uniform. In Fig. 7, for example, there is an obvious vortex near the left-hand wall and a region of rapid upflow near the centre of the section. In the lower right-hand corner, the flow appears to be downward and to the left, suggesting the presence of a vortex outwith the illumination plane. Figure 8 shows similar features, but unfortunately this photograph is truncated due to the inopportune arrival of the end of the film; this is why the abscissa "Distance across flow" covers the full 7.5 cm in Figs 7 and 9 but stops at about 6.5 cm in Fig. 8.

In Fig. 9, further downstream, the velocity vectors are upward over almost all of the area photographed. There may be some vortex-like features near the left-hand wall, but these are smaller than those in Figs 7 and 8 and have lower velocities.

According to the theory of Pyle and Rose, the upward vertical component of gas velocity across a horizontal diameter through a three-dimensional bubble rises from $-3U_b/2$ at the bubble surface to $(6U_{mf} + 1.5U_b)$ at the bubble centre, these velocities being in the bubble frame. Where $U_b/U_o \approx 1.76$ and $\varepsilon_{mf} \approx 0.4$, as here, these velocities are equivalent to $-6.6U_{mf}$ to $12.6U_{mf}$.

The experimental work presented in Figs 7 and 8 found that the maximum velocity obtained by PIV analysis was about $2.5U_{mf} \approx U_o$, the mean interstitial gas velocity; this measurement was in the laboratory frame. Comparing this with the figures in the paragraph above, there has clearly been a considerable reduction in gas velocity within a hypothetical vortex ring during its rise from the fluidised bed to a position approximately 19 cm above the bed surface.

In terms of continuity, it is possible to evaluate the mean vertical component of velocity at any of the heights above the bed surface shown in Figs 7–9. Each figure shows velocities within one thin vertical slice of the flow, as explained in Section 3. The vertical components of all vectors in each horizontal row of a vector map can be averaged, using interpolation to fill in the missing vectors. The flow is clearly very three-dimensional, however, so it is not reasonable to expect a very close match to the mean velocity given the thinness of the illuminated section.

For Fig. 8, averaging across the flow (horizontally), the mean vertical velocities at each height have been plotted against the height above the bed surface (Fig. 10). The mean of all the mean values displayed is 0.073 m s^{-1}, which is very close to U_g. On this picture, the mean upward gas velocity seems to lie in the range $0.7U_g$ to $1.5U_g$. A full mass balance would need to take into account velocities perpendicular to the illumination plane, but the two-dimensional PIV technique used here does not yield this information.

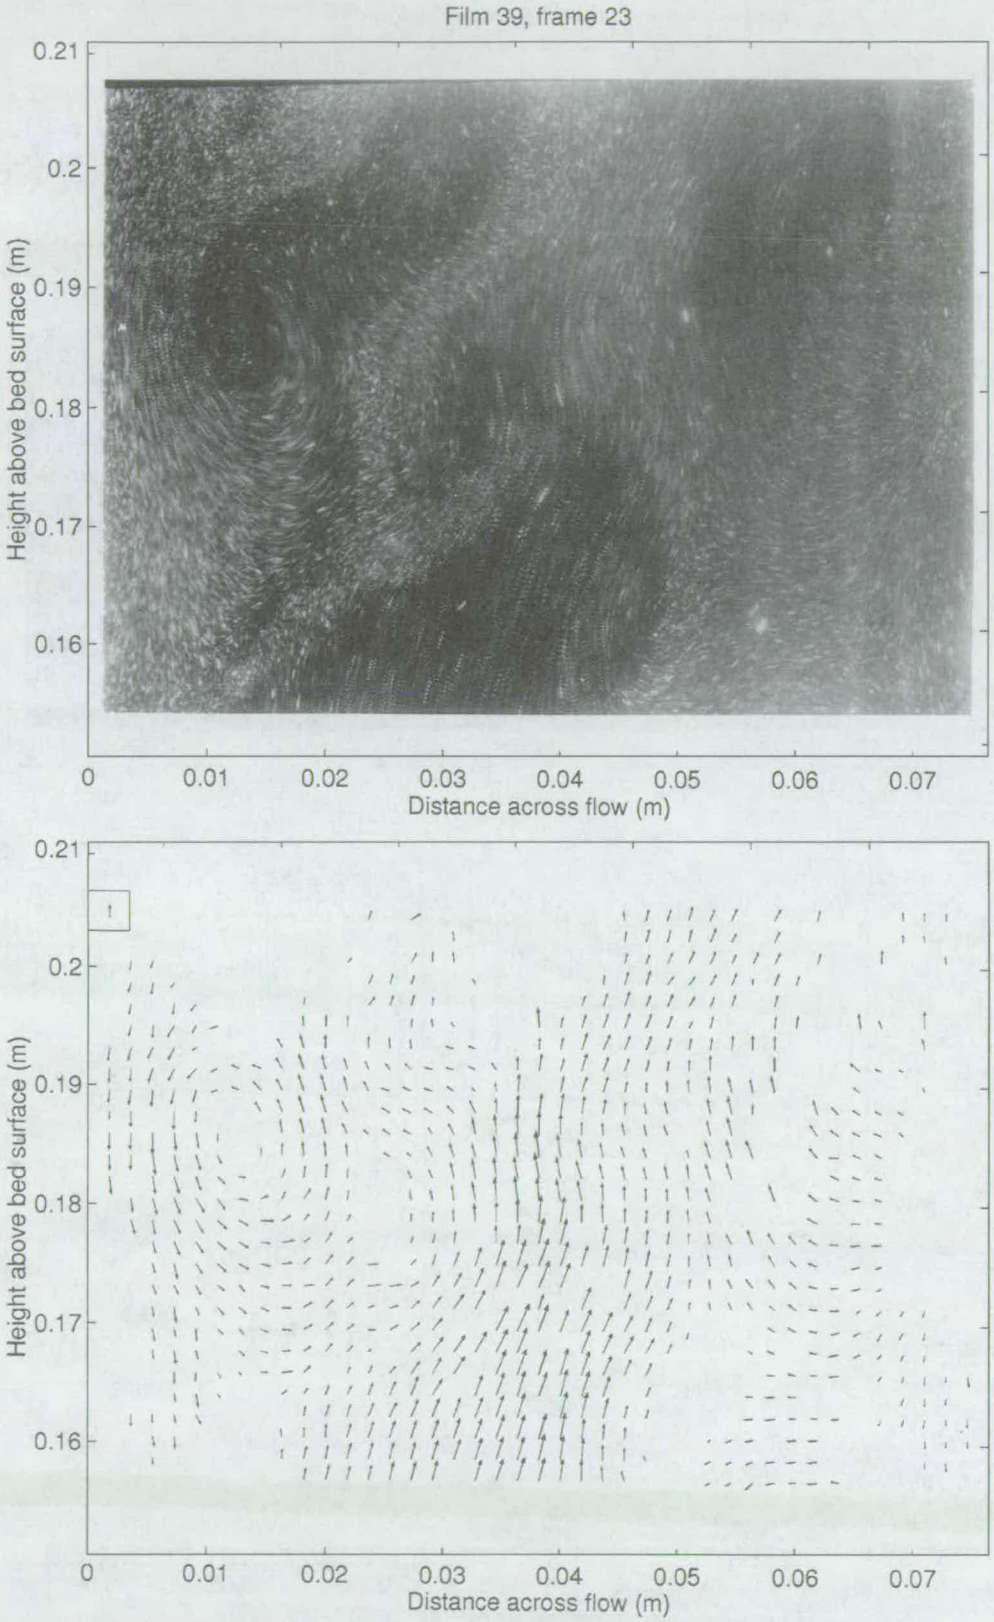


Fig. 7. PIV photograph and vector map: a 3 cm bed of 150–325 μm glass ballotini fluidised at $U_g = 0.07\ \mathrm{m\,s^{-1}}$.

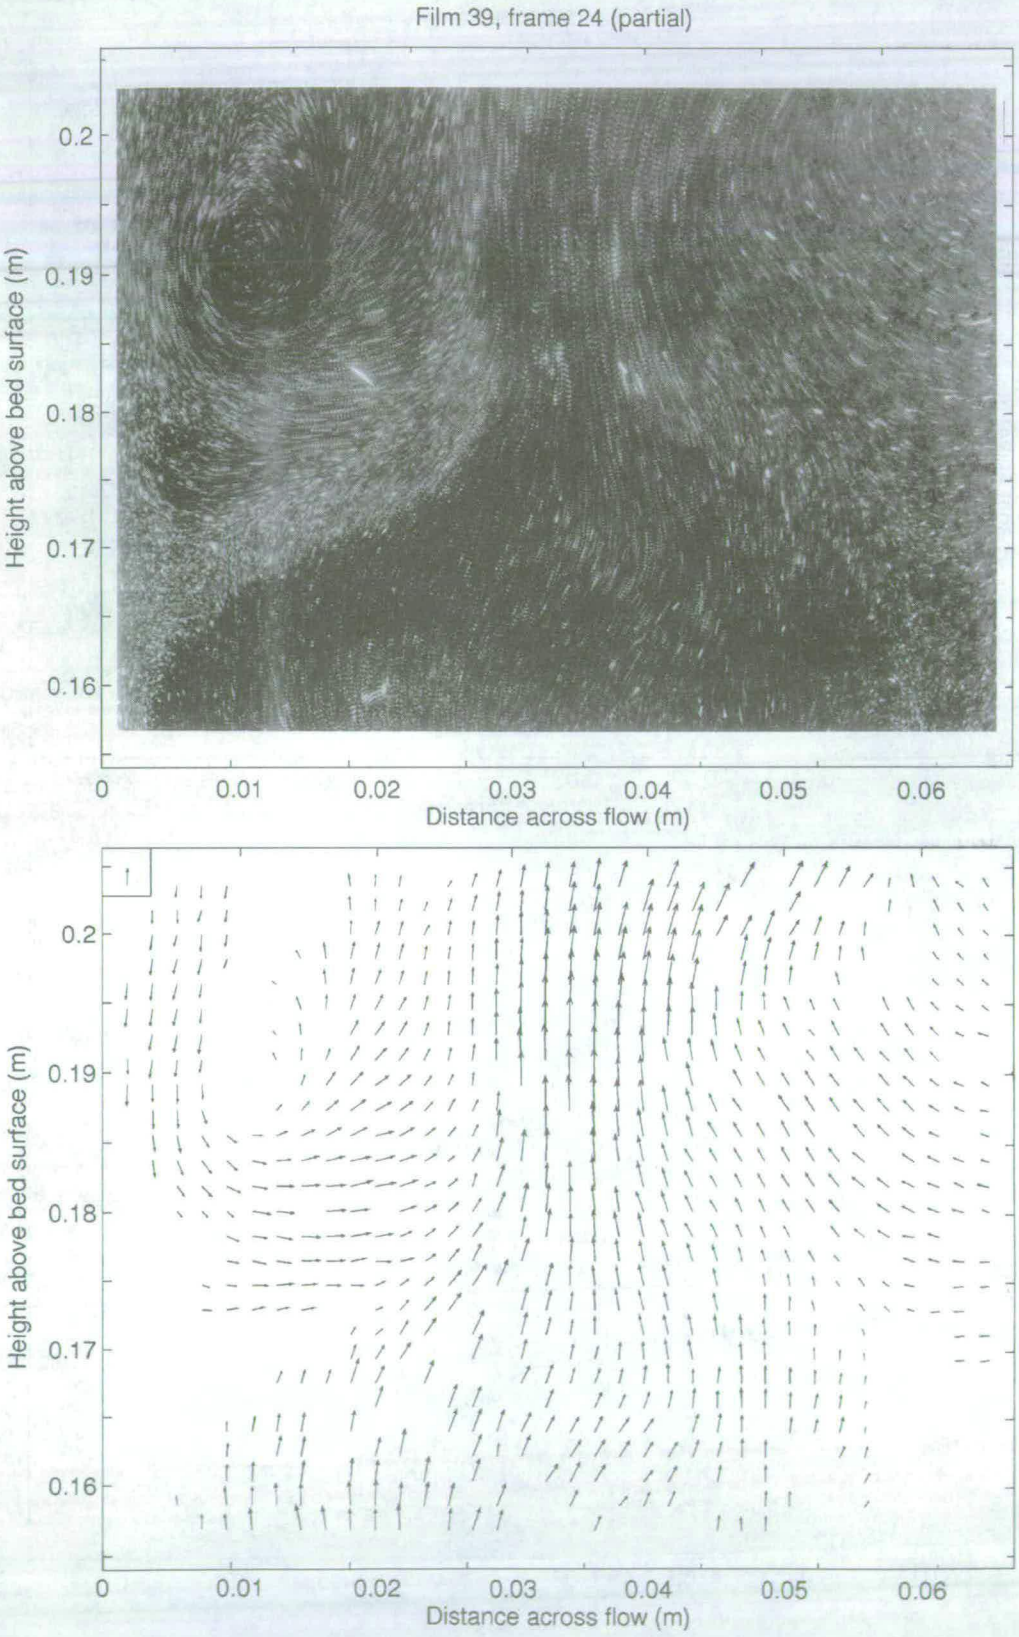


Fig. 8. PIV photograph and vector map: a 3 cm bed of 150–325 μm glass ballotini fluidised at $U_g = 0.07\ \mathrm{m\,s^{-1}}$.

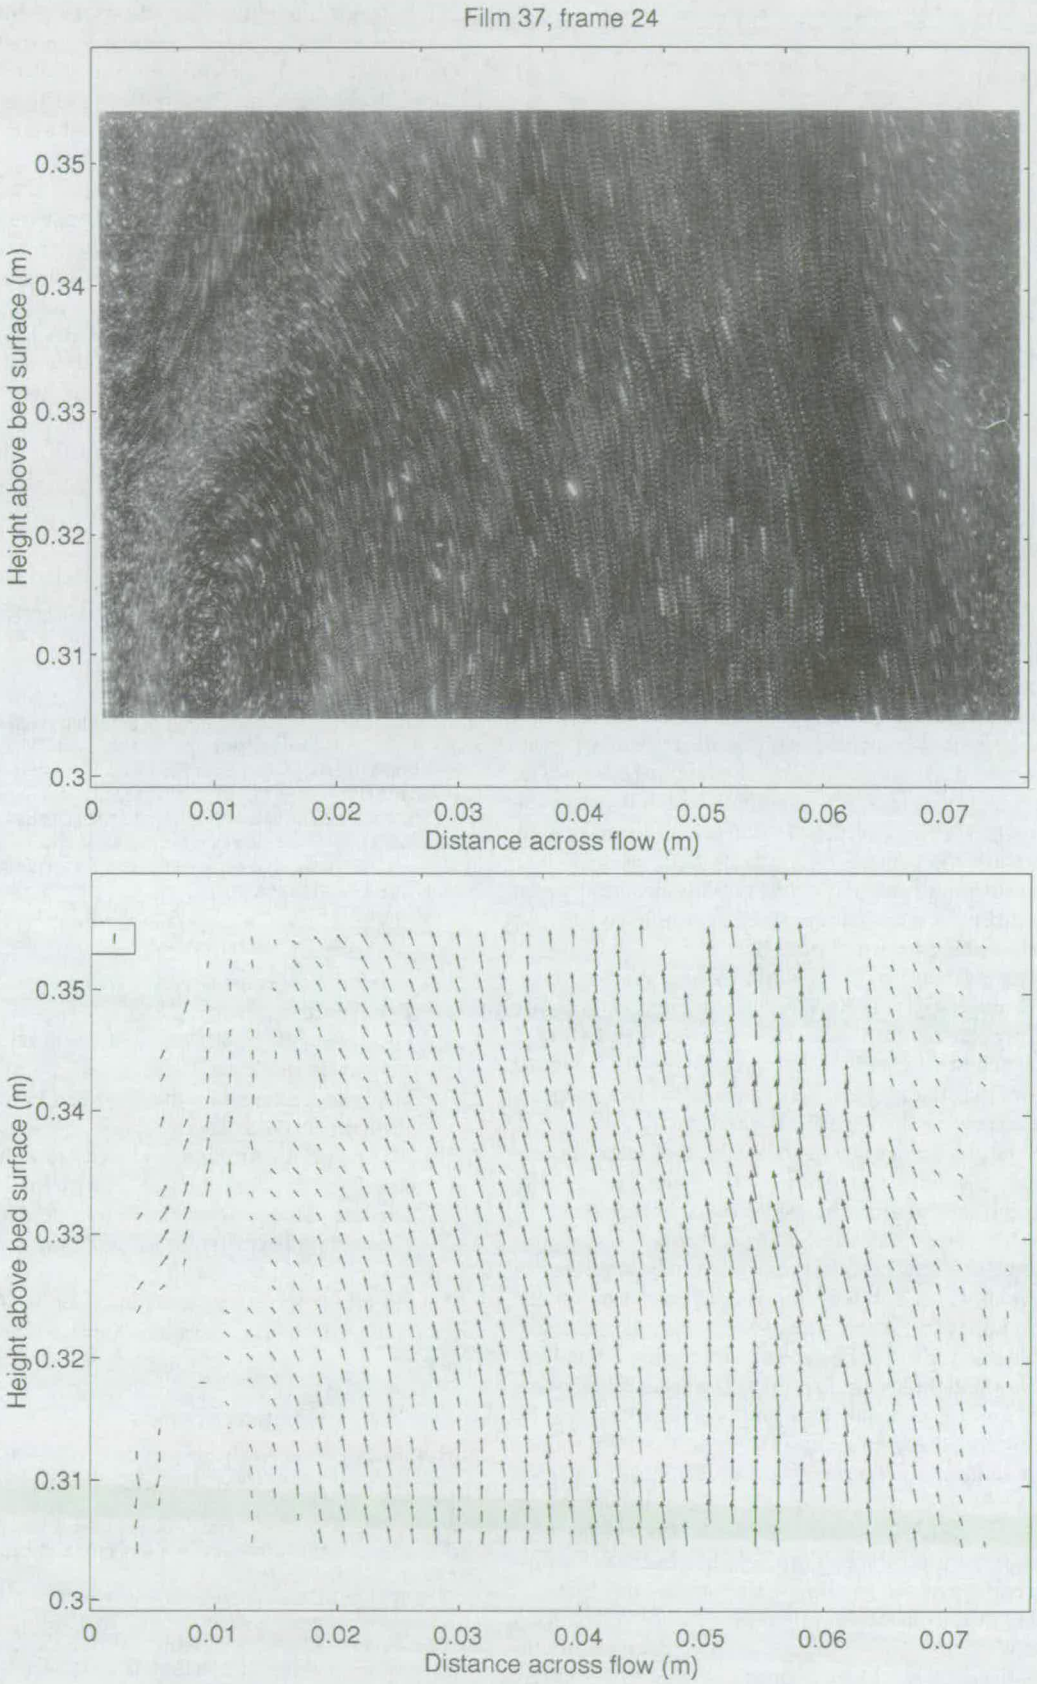


Fig. 9. PIV photograph and vector map: a 3 cm bed of 150–325 μm glass ballotini fluidised at $U_g = 0.07\ \mathrm{m\,s^{-1}}$.

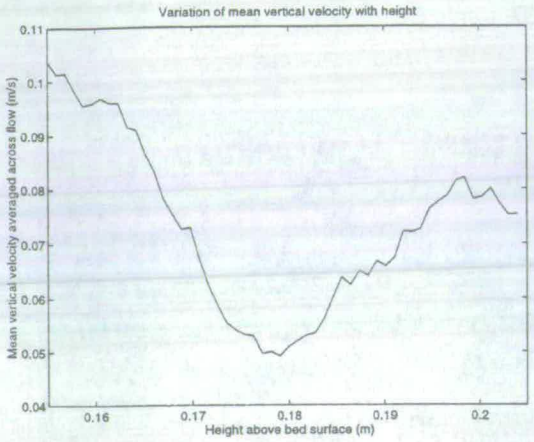


Fig. 10. Mean vertical velocities, averaging horizontally across the flow, calculated from the data presented in Fig. 8.

5. DISCUSSION

If representative, the results presented in Figs 7–9 show a clear decay in turbulence with height. At 18 cm above the bed surface, not all vectors are upward; some are predominantly horizontal or even downward. At about 33 cm above the bed surface, recirculation zones are less significant but are occasionally visible. The flow is not yet uniform, in that there are still regions of relatively high velocity near the centre and relatively low velocity near the walls; the turbulence has decayed but not sufficiently to be comparable to normal duct flow.

The existence of recirculation zones such as those shown in Figs 7 and 8 has interesting implications for the process of elutriation. Consider the flow in Fig. 8 at a height of 0.195 m. The graph shown in Fig. 10 shows that the mean upward velocity at this height is fairly close to the superficial gas velocity.

Concentrating now on the 0.025 m of section nearest to the left-hand wall: as the velocities on the descending and ascending sides of the vortex are comparable, the overall upward flow through this region is low. This suggests that approximately one third of the duct width at this height is not being used for the flow. Suppose, for the sake of argument, that this also applies to the other horizontal direction, perpendicular to the light sheet; if so, then the actual area being used for upflow is only four ninths of the total area. By continuity, the average upward velocity where significant upflow does occur must equal $9U_g/4$; particles from the bed whose size puts them in the range $U_g < U_t < 9U_g/4$ can still be carried up by the flow, even though they will ultimately fall back down into the bed when the gas flow becomes less turbulent.

The results as 18 cm might be qualitatively consistent with the ghost bubble theory of Pemberton and Davidson (1984). There seems to be a pair of counter-rotating vortices in Fig. 8. However, it must be remembered that these figures are in the laboratory reference frame. "Ghost bubble" theory holds that ghost bubbles maintain their vortex ring character when viewed in the bubble frame. A ghost bubble decelerates as it rises, so the two frames of reference become closer together; their initial condition, in the fluidised bed, is given by the left- and right-hand sides of Fig. 2.

There are many complicating factors in this analysis, however. For example, it is not reasonable to expect a potential flow solution to be fully representative of the flow inside a bubble in a fluidised bed because of vorticity, even though the technique has achieved many successes in the past. A fluidisation or ghost bubble will deviate from the Pyle and Rose model due to wall effects, an uneven gas distribution within the bed, the action of breaking the bed surface, and interactions with other bubbles both in the bed and in the freeboard.

6. CURRENT WORK

A larger-scale fluidised bed has been designed and built, having a 30 cm × 30 cm cross-section and a total height of over 4 m. PIV work on this ring is currently in progress.

Acknowledgements—The authors are indebted to an EPSRC grant, which provided the funding for this work. They would like to thank their colleagues in the Fluid Dynamics Unit, in particular Denise McCluskey, Callum Gray and Tom Bruce (for their patience in explaining the techniques involved in PIV), and Andy Hind (for developing and maintaining the analysis software used in this work and for explaining the eccentricities of Matlab).

NOTATION

U_g General superficial gas velocity through a packed bed
U_{mf} The value of U_g at which a fluidised bed is incipiently fluidised
ε_{mf} The void fraction of a fluidised bed in the state of incipient fluidisation
U_o The mean interstitial gas velocity in a fluidised bed; $U_o = U_{mf}/\varepsilon_{mf}$
U_b The rise velocity of a bubble in a fluidised bed
U_t The terminal velocity under gravity of a particle

All the above are to be used in SI units. All velocities referred to are measured in the laboratory frame.

REFERENCES

Adrian, R. J., 1991, Particle-imaging techniques for experimental fluid mechanics. *Ann. Rev. Fluid Mech.* **23**, 261–304.

Azevedo, J. L. T. and Pereira, J. C. F., 1991, Flow patterns around and inside a bubble in a gas-fluidised bed. *Chem. Engng Sci.* **46**, 155–165.

Collins, R., 1982, Cloud patterns around a bubble growing in a gas-fluidized bed. *J. Fluid Mech.* **122**, 155–167.

Collins, R., 1983, Gas flow patterns associated with a growing spherical bubble. In *Fluidization* IV (Edited by Kunii and Toei), pp. 1–8.

Gray, C., Greated, C. A., McCluskey, D. R. and Easson, W. J., 1991, An analysis of the scanning beam PIV illumination system. *Measure Sci. Technol.* **2**, 717–724.

Greated, C. A., Skyner, D. J. and Bruce, T., 1992, PIV in the coastal engineering laboratory. *Proc. 23rd Int. Conf. Coastal Eng.* **1**, 212–225.

Hind, A. K. and Christy, J. R. E., 1992, Digital PIV applied to flows around artificial heart values: analysis by auto-correlation. *Proc. Euromech. Colloq. 279*. Kluwer Academic Press, Netherlands.

Hind, A. K., McCluskey, D. R., Christy, J. R. E., Easson, W. J., Greated, C. A. and Glass, D. H., 1992, PIV analysis of two-phase air-particle flows. *IMechE Symp. on Optical Methods and Data-processing in Heat and Fluid Flow*, London.

Horio, M. and Kuroki, H., 1994, Three-dimensional flow visualisation of dilutely dispersed solids in bubbling and circulating fluidised beds. *Chem. Engng Sci*, **49**, 2413–2421.

Keane, R. D. and Adrian, R. J., 1990, Optimisation of particle image velocimeters, Part I: double-pulsed systems. *Meas. Sci. Technol.*, **1**, 1202–1215.

Keane, R. D., Adrian, R. J., 1991, Optimisation of particle image velocimeters, Part II: Multiple-pulsed systems, *Measure Sci. Technol.* **2**, 963–974.

Lang, I. W. P., 1986, The structure of flow in the freeboard region of gas/solid fluidised beds. PhD Thesis, University of Edinburgh.

Leva, M., 1951, Elutriation of fines from fluidised systems. *Chem. Engng Prog.* **47**, 39–45.

Maxworthy, T., 1972, The structure and stability of vortex rings. *J. Fluid Mech.* **51**, 15–32.

McCluskey, D. R., Hind, A. K., Rix, S. J. L., Jakobsen, M. L., Christy, J. R. E., Easson, W. J., Glass, D. H. and Greated, C. A., 1994, Development and application of a simultaneous two phase PIV system for the analysis of air–particle flow fields. *IMechE Symp. on Optical Methods and Data-processing in Heat and Fluid Flow*, London, pp. 99–109.

Northrup, M. A., Kulp, T. J. and Angel, S. M., 1993, Direct measurement of interstitial velocity field variations in a porous medium using fluorescent-particle image velocimetry. *Chem. Engng Sci.* **48**, 13–21.

Pemberton, S. T. and Davidson J. F., 1984, Turbulence in the freeboard region of a gas-fluidised bed. *Chem. Engng Sci.* **39**, 829–840.

Pyle, D. L. and Rose, P. L., 1965, Chemical reaction in bubbling fluidised beds. *Chem. Engng Sci.* **20**, 25–31.

Remírez de Santiago, M., 1990, Measurements of size and velocity distributions of droplets produced by bubbles bursting. *Proc. 5th Int. Symp on Applications of Laser Techniques to Fluid Mechanics* (Edited by Adrian, R. J., Durão, D. F. G., Durst, F., Maeda, M. and Whitelaw, J. H.), pp. 203–226. Springer, Berlin.

Reynolds, O., 1876, On the resistance encountered by vortex rings, and the relation between the vortex rings and the streamlines of a disc. *Nature* **14**, 477–479.

Zenz, F. A. and Weil, N. A., 1958, A theoretical-empirical approach to the mechanism of particle entrainment from fluidised beds. *A.I.Ch.E. J.* **4**, 472–480.