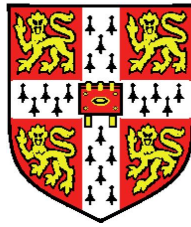# Tomographic-PIV Measurement of Coherent Dissipation Scale Structures

Nicholas Worth

Darwin College

University of Cambridge

A thesis submitted for the degree of

Doctor of Philosophy

April 2010

Declaration

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except where specifically indicated in the text. No part of this dissertation has already been, or is currently being reviewed for any other degree, diploma or qualification. This thesis contains approximately 64,900 words, and 122 figures.

Nicholas Worth
April 2010

Acknowledgements

Tomographic-PIV Measurement of Coherent Dissipation Scale Structures

Nicholas Worth

Further understanding the small scale coherent structures which occur in high Reynolds number turbulence would be of enormous benefit. Therefore, the aim of the current project was to make well resolved three-dimensional flow measurements of the mixing flow between counter rotating impellers, using Tomographic Particle Image Velocimetry (TPIV).

TPIV software was developed, with a novel approach permitting a significant reduction in processing time, and a series of numerical accuracy studies contributing to the fundamental understanding of this new technique. Basic flow characterisation determined the local isotropy, homogeneity and expected Reynolds number scaling. A favourable comparison between planar PIV and TPIV increased confidence in the latter, which was used to assess the dynamics and topology of the dissipation scale structures.

In support of previous investigations similar topology, strain rate alignment, scale-invariance, and clustering behaviours are demonstrated. Correlated high enstrophy and dissipation regions occur in the periphery of larger structures, resulting in intermittency. Geometry characterisation indicates a predominance of tube-like structures, which are observed to form from larger ribbon-like structures through unsteady breakdown and vortex roll-up. Significant correlation between intermittent fields of dissipation and enstrophy describe the fine scales effects. These relationships should pave the way for more accurate models, capable of relating small scales and large scales during the prediction of dynamically important quantities.

Contents

Nomenclature

Roman Symbols

a Decay rate

A Scaling argument dissipation rate variable

A_{ij} velocity gradient tensor

C_2 Second order Kolmogorov constant

C_{ij} Velocity Correlation

d_p Particle diameter

q Any velocity gradient term

D_0 Orifice diameter

E Energy

E_{global} Global Error

E_{local} Local Error

e_m Mean strain rate

f Longitudinal correlation function

F_k Calibration function

f_{ij}^k Coefficients in the calibration Taylor series expansion

$f/\#$ 'f-Number' aperture setting

g Transverse correlation function

H Shape parameter

k_{it} Number of iterations

L_0 Piston stroke length

L Characteristic length scale

$\overline{\overline{L}}$ Average line of sight through a volume

m, n Position of value in a 2D matrix

N_{cam} Number of cameras in setup

N_{ideal}, N_{PIV} Number of ideal and PIV data points

N_0 Random box count

N_{box} Box counts

N_{calcs} Estimated Tomo program FLOP count

N_{ppp} Seeding density (particles per pixel)

N_{total} Total number of structures

N_{vol} Number of volumes

N_{vox} Number of voxels

N_{Wvals} Number of values which require storage in the W matrix

n_p calibration polynomial order

p pressure

P_A, Q_A, R_A Topological invariants of the velocity gradient tensor

P_S, Q_S, R_S Topological invariants of the rate of strain tensor

P_W, Q_W, R_W Topological invariants of the rate of rotation tensor

Q Divergence correlation coefficient

$R_{1/e}$ Vortex core diameter

r Position vector connecting two points in a flow

Re Reynolds number

R_I Impellor radius

R_λ Taylors microscale Reynolds number

r_x, r_y, r_z Box counting dimensions

S_{ij} Strain rate tensor

t Time

T_0 Piston stroke time

u_0, v_0 Coordinates of optical center in image plane

u_θ Azimuthal velocity

u_{im}, v_{im} 2D camera coordinates

u Fluctuting velocity

U Characteristic large scale velocity

v_0 Measure of mean structure separation

V_x, V_y, V_z Reduced volume size in box counting method

W_{ij} Rate of rotation tensor

$W/p_r/f$ Reconstruction matricies: weighting / measured property / object

x_{cen}, y_{cen} Image centre co-ordinates

x_{im}, y_{im} Cartesian coordinates of a point in the image plane

x, y, z Cartesian Coordinates

Greek Symbols

α, β, γ Principle rates of strain in ascending order of magnitude

χ Local VGT normalised local divergence

γ_0 Vortex core circulation

Δ Discriminant of matrix invariants

Δh Interrogation window width

Δu Structure Function

Δx PIV vector spacing

ϵ Dissipation rate

Φ Enstrophy rate of change

Φ_{eq} Equivalent Enstrophy rate of change

η, υ, τ Kolmogorov length, velocity and time scales

Γ_0 Vortex ring circulation

θ, ϕ, ψ roll, pitch, and yaw angles

κ Wavenumber

λ Taylors microscale

λ_{ci} Swirling strength

Λ Integral length scale

$\lambda_1, \lambda_2, \lambda_3$ Eigenvalues of matrix invariants in ascending order of magnitude

ν Kinematic viscosity

Π Enstrophy stretching/compression

ρ Density

Σ Topological measure of vortex stretching

Υ Enstrophy advection

ξ Mean normalised divergence ratio

Ξ Enstrophy viscous dissipation

Ψ Enstrophy viscous diffusion

ω Vorticity

ω_0 Vortex core vorticity

Ω Enstrophy

Ω_I Rate of impeller rotation

Superscripts

p Structure function order

$*$ non dimensional

T Matrix transpose

Subscripts

$_D$ Domain size

$_{i,j,k}$ Tensor notation

$_m$ Mean $(= \overline{(\cdots)})$

Other Symbols

$\overline{(\cdots)}$ Mean $(= {}_m)$

$\langle(\cdots)\rangle$ Ensemble average (equal to time average for steady flows)

$sgn[\cdots]$ Signum function

$tr[\cdots]$ Matrix trace

Acronyms

$3C$ Three Velocity Components

amu atomic mass unit

AR Aspect Ratio

ART Algebraic Reconstruction Technique

CCD Charge-Coupled Device

CFD Computational Fluid Dynamics

$CMOS$ Complementary metal oxide semiconductor

$DDPIV$ Defocusing Digital Particle Image Velocimetry

$HPIV$ Digital Holographic Particle Image Velocimetry

DNS Direct Numerical Simulation

FOD depth of field

dpi Dots per Inch

$DPIV$ Digital Particle Image Velocimetry

$SDPIV$ Stereo Digital Particle Image Velocimetry

FFT Fast Fourier Transform

$FLOP$ Floating Point Operation

$FLOPS$ Floating Point Operations per Second

FOV Field of view

GS 2-pass Gaussian smoothing

HIT Homogeneous isotropic turbulence

$HPIV$ Holographic Particle Image Velocimetry

IW Interrogation Window

$JPDF$ Joint Probability Density Function

$K41$ Kolmogorov's 1941 paper

LDA Laser Doppler Anemometry

LES Large Eddy Simulation

$MART$ Multiplicative Algebraic Reconstruction Technique

MFG Multiplicative First Guess

PDF Probability Density Function

PIV Particle Image Velocimetry

ppm Parts per Million

PTV Particle Tracking Velocimetry

RAM Random Access Memory

$RMART$ Rescaled Multiplicative Algebraic Reconstruction Technique

SF Spatially filtered

SF/S Stable focus/stretching topology

SF/S Stable focus/stretching

SNR Signal to Noise Ratio

$SN/S/S$ Stable node/Saddle/Saddle topology

$SN/S/S$ Stable node/saddle/saddle

$TIFF$ Tagged image file format

$TPIV$ Tomographic Particle Tracking Velocimetry

UF Uniform field

UF/C Unstable focus/contracting topology

UF/C Unstable focus/contracting

$UN/S/S$ Unstable node/Saddle/Saddle topology

$UN/S/S$ Unstable node/saddle/saddle

VGT Velocity Gradient Tensor

Introduction

Due to its prevalence in a huge variety of practical engineering problems, understanding the complex phenomenon of turbulence is of great importance. The subject balances great interest stemming from its significance with a high level of complexity, the former fuelling it's intensive study while the latter has drawn this out for more than a century.

The definition of turbulence is not straight forward, and it is usually characterised by unsteady, three-dimensional motion, which contains a wide range of scales, shapes, and orientations of motions. Observation of the large range of eddying motions prompted Richardson ([101]) to develop an energy cascade concept; whereby flow instabilities result in the initial formation of large eddies from the mean flow. These large eddies tend to become unstable and break down into smaller eddies, which also contain instabilities, transferring energy down through coherent structures into increasingly fine-scale motions. This breakdown continues until viscosity begins to exert an effect on sufficiently small scale motions, resulting in energy dissipation into heat.

The large scale structures appear to scale with the flow width, and dominate momentum, mass, and heat transport. The small scales govern the dissipation of kinetic energy and also the behaviour of the inertial sub-range; a region where the scales of motion are much larger than the dissipative scales, but much smaller than the large scales. The universal equilibrium range and inertial sub-range were originally defined by Kolmogorov ([62]) using assumptions of local isotropy, statistical equilibrium and universality. According to these assumptions, motion in the universal equilibrium range will depend solely on the kinematic viscosity and average global dissipation, with the former dependence relaxed in the inertial sub-range.

Kolmogorov's theories have been tested extensively, and although there appears to be some support for these (80), the concept of universality is still debated. Sreenivasan & Antonia (116) detail a number of investigations which question the exactness of Kolmogorov's original hypotheses (119; 144), with discrepancies manifesting in higher order structure functions ($p \geq 4$). However, as noted by Frisch (35), due to the difficulty measuring these higher order structure functions, any departures from Kolmogorov's predictions are far from resolved.

Modifications to Kolmogorov's original theory suggest the dissipation at the finest scales does not have a uniform distribution; a phenomenon known as dissipative intermittency. Landau (64) was the first to object to Kolmogorov's assumptions, noting that large scale inhomogeneities will be passed to the small scale. Therefore, eddies of a particular size should be controlled by the locally averaged, rather than the globally averaged dissipation, as each local area will be subject to its own specific energy cascade. Although a large number of models have been generated to take account of this intermittency through extra scaling arguments, Sreenivasan & Antonia (116) note that due to the tenuous link between these models and physical understanding it is difficult to recommend one model over others. Additionally, model performance is often assessed with respect to experimentally determined high-order structure functions which, as already stated, may be subject to significant errors. Frisch (35) supports this view, noting that it is difficult to ascertain which intermittency model best describes the actual processes, or even if Kolmogorov's original hypothesis is accurate over a greater range than is currently assumed.

Coherent structures are thought to be related to this intermittency, and although the influences of certain structures as active or passive components are still not fully resolved, their connection with turbulent mixing phenomena has been demonstrated. It is the coherency of such structures that make them eligible for study, permitting greater understanding of the fluid processes at both large and small scales, which may lead to more accurate turbulence models based on physical observations rather than solely statistical data. The desire for a greater understanding of these small scale structures has motivated a large number of numerical and experimental investigations.

Numerical investigations are often focused on small-diameter, long, high vorticity *filaments* or *worms*, studying their scaling properties, formation processes and influence on the small scales in some depth. Experimental investigations have also been used in confirming the presence of these and other structures, identifying possible formation mechanisms, and assessing their significance in the flow. However both numerical and experimental approaches are somewhat restricted; the former limited to relatively low Reynolds numbers by computational costs, and the latter by accurate resolution of extremely small fast motions, typically allowing only one-dimensional data to be collected, making the study of these structures and their interactions difficult.

Therefore, the aim of the current investigation is to study high Reynolds number small scale turbulent motions experimentally using both conventional Digital Particle Image Velocimetry (DPIV) and Tomographic-PIV (TPIV), which is a new, fully three-dimensional approach. The use of a very large mixing tank will bring the smallest length and time scales within the reach of this technique, allowing full 3D experimental quantification of these structures. This should support other experimental, and more crucially, numerical findings, providing information for further statistical model development on the flow features which still retain the greatest universal potential. This is particularly relevant to improving computational simulations such as the increasingly affordable *Large Eddy Simulation* (LES), which depends on accurate modelling of these scales. Further knowledge of these structures may also contribute significantly to the understanding of turbulent mixing (59) and particle dispersion (86) at the small scales, each with a host of applications, and also allow accurate assessment of the local dissipation distribution. Additionally, through the use of this new tomographic technique, it is hoped that insight leading to its further development can be gained, and through implementation, the method's effectiveness can be increased.

Chapter 1 introduces some tools for the study of turbulence before covering some of the findings from previous studies regarding the nature of the small turbulent scales. Next a range of measurement techniques are reviewed before a TPIV is discussed in more detail, including its applicability to the current investigation. In Chapter 2 the current experimental set-up is introduced, and the design and implementation of an iterative tomographic reconstruction program

is detailed, before discussing data post-processing algorithms. The DPIV and TPIV simulations conducted in Chapter 3 allow quantification of the measurement uncertainty, allowing more definite conclusions to be drawn from the large mixing tank results, which are presented in Chapters 4 and 5. In the first of these the Reynolds number scaling and nature of the statistically steady flow is characterised, demonstrating the consistency of DPIV and TPIV data. In Chapter 5 the coherent flow structures are characterised using a combination of visualisation and statistical methods, before finally drawing some conclusions in Chapter 6.

Chapter 1

Turbulence Measurements at the Dissipation Scale

This chapter begins by introducing a number of tools for studying turbulence, ranging from statistical to topological, with the former useful in broadly classifying the range of scales present in fully developed turbulent flows, and the latter for identifying specific coherent structures. An understanding of these methods and their shortcomings is necessary before previous work can be critically analysed.

Next, a review of literature concerning coherent small scale turbulent motions is attempted, in order to identify areas which require further study and provide the objectives for the current investigation. A review of experimental techniques is then conducted, focusing on TPIV. The chapter ends with a summary, followed by a statement of aims and objectives.

1.1 Turbulence Analysis Tools

1.1.1 Statistical Tools

Although turbulence is intrinsically unsteady, useful predictions often relate to the time-averaged statistical flow behaviour. Statistical descriptions can be presented through probability density functions (PDFs), which denote the probability of a certain event occurring. Comparison of these functions with Gaussian fields can provide a useful characterisation of the flow field (8). Other higher-order

correlations such as skewness and flatness may also prove useful with regard to quantifying any deviation from Gaussian behaviour.

It is also useful to extend single point distributions to the product of velocities at separate points to capture structural flow features, with positive correlation suggesting coherency. For example, a two point double velocity correlation tensor is denoted by, $C_{ij}(x_i, x_i + r_i, t) = \langle u_i(x_i,t)u_j(x_i + r_i,t)\rangle$, where r_i is the position vector connecting the two points. Through the variation of this vector, r, the value of C_{ij} indicates the correlation of the velocity field around the point x_i. Homogeneity and statistical steadiness will remove this function's dependence on a fixed position, x_i, and time, t, respectively.

To determine the distribution of kinetic energy, this function can be extended to form longitudinal structure functions (moments of velocity increments) or the energy spectrum. A structure function, Δu, of order p can be defined by $\langle [\Delta u_r(x_i, x_i + r, t)]^p\rangle = \langle [u_x(x_i + r, t) - u_x(x_i, t)]^p\rangle$, where Δu_r is the difference between two velocities components, u_x, in the x direction, separated by a distance, r in the x direction. Thus, when $p = 2$, the structure function gives an indication of the energy of contribution of different sized eddies, and varying the distance, r, allows the contribution of different sized structures to be assessed, and hence the eddy population distribution to be inferred.

An alternative way of viewing the energy distribution is by wavenumber through the energy spectrum function. A Fourier transform is used to convert the velocity correlation function into wavenumber space, thereby decomposing the local fluid motion into different scales according to wavenumber, κ. The energy at a particular wavenumber, $E(\kappa)$, can be considered the energy contribution from eddying motions of size π/κ.

However, it should be noted, although the structure functions take account of the energy contained in eddies of size $\sim r$ and under, at the dissipation scales the enstrophy from larger eddies also contributes significantly. Also with regard to spectral methods, it should be stressed that although eddies of a particular size will contribute predominantly to a band of wavenumbers, they will contribute to a lesser extent across a broad range of wavenumbers. Therefore, these techniques only permit the range of scales and the energy they contain to be broadly categorised.

1.1.2 Turbulent Scales

It is useful to at this point to quantify what is meant by large and small scale motions. As previously mentioned the largest turbulent eddies extract energy from the mean flow, and therefore can be defined in terms of characteristic mean flow length and velocity scale, Λ and U respectively. The largest turbulent motions are observed to transfer their energy on a time-scale of, Λ/U, resulting in a rate of energy transfer, which in statistically steady turbulence can be equated to the mean dissipation rate, $\langle \epsilon \rangle$.

Expressed in terms of RMS velocity and integral length scale the mean dissipation rate can be calculated using Equation 1.1. From high Reynolds number homogeneous isotropic turbulence DNS ($R_\lambda > 200$) the value of A has been shown to be ≈ 0.5 (23; 115).

$$\epsilon = A\frac{u^3}{\Lambda} \tag{1.1}$$

According to Kolmogorov's analysis (62) at the large scale $U\Lambda/\nu \gg 1$, and the cascade is dominated by inertial forces. However, as the Reynolds number based on eddy size and mean velocity approaches unity, viscous forces begin to become important, and the smallest length, velocity and time scales, η, υ and τ respectively can be defined as functions of ϵ and ν alone, and derived from dimensional analysis.

$$\eta \sim \left(\frac{\nu^3}{\epsilon}\right)^{\frac{1}{4}}$$
$$\upsilon \sim (\nu\epsilon)^{\frac{1}{4}} \tag{1.2}$$
$$\tau \sim \left(\frac{\nu}{\epsilon}\right)^{\frac{1}{2}}$$

Another important small scale is the Taylor microscale, which gives a reasonable estimate of where the majority of dissipation occurs (8). As detailed in Hinze (43) it is possible to define the correlation function shape at $r = 0$ in terms of velocity derivatives at that point with a Taylor series expansion. At small

distances the shape of the correlation function approaches a parabola, which can be defined by introducing a length, λ, as in Equation 1.3.

$$f(r) \approx 1 - \frac{r^2}{2\lambda^2} \tag{1.3}$$

Alternatively this length scale can be directly related to the local velocity gradients through the dissipation rate, as shown in Equation 1.4 (23).

$$\lambda^2 = \frac{15\nu u^2}{\epsilon} \tag{1.4}$$

1.1.3 Role of Coherent Structures and their Identification

As previously mentioned, intermittency is thought to be responsible for the observed departures from the universality of Kolmogorov's small scale theories. Although a number of models have taken account of this through additional scaling arguments (35), further knowledge of the origins of this phenomenon may be useful in improving these. Understanding the evolution and interaction of dynamically significant structures with the background turbulence, other large and small scale structures, and the distribution of dissipation should provide valuable insight into dissipative intermittency, potentially permitting characterisation. Even more fundamentally, if such structures are responsible for transferring energy to each other as in the Richardson cascade, studying the evolution of such structures should enhance the understanding of this process, allowing the connection between large scale and small scale intermittency to be investigated.

One way to study the structural properties of the small scales of turbulence is through the vorticity and strain fields. By considering eddying motions as regions of concentrated vorticity the local structure of these regions can be studied (47). This view is shared by Cadot et al. (20) who note that the time-dependent features of a turbulent flow can be characterised by the squares of both vorticity and strain rate, with the former known as the enstrophy whereas the latter can be directly related to the dissipation rate through the kinematic viscosity. If the most significant contributions to these fields are governed by coherent structures down to the smallest scales, it should be possible to characterise the flow depending on these structures.

The interaction of these fields however is non-trivial, and although the enstrophy integrated over the domain can be related to mean dissipation, due to dissipative intermittency in general the local distributions of these quantities are quite separate. For example, a Burgers vortex type structure will have a core of enstrophy wreathed in dissipation (16).

The correlation between these regions can be examined using a several approaches: firstly through direct examination of local strain and vorticity distributions (or the balance of these quantities), popularised by DNS studies (4; 56; 103) but also reproduced experimentally (37); and secondly through single point statistical correlations of these properties or other quantities relating to local flow topology. The various properties required for coherent structure identification and statistical analysis will now be introduced and compared.

1.1.4 Velocity Field Decomposition

In practical terms the velocity field can be decomposed into three more fundamental effects of translation, rotation and distortion. These effects are clearly demonstrated by considering the change of a fluid element in a flow field. The average rate of rotation can be determined using Equation 1.5. The vorticity is most clearly written in vector notation as the curl of the velocity field (Equation 1.6).

$$W_{ij} = \frac{1}{2} \left(\frac{\partial u_i}{\partial x_j} - \frac{\partial u_j}{\partial x_i} \right) \tag{1.5}$$

$$\omega_i = \nabla \times u_i \tag{1.6}$$

The enstrophy is a scalar measure of vorticity magnitude, playing a similar role to the kinetic energy with respect to velocity, and can be defined by Equation 1.7.

$$\Omega = \frac{1}{2} \omega_i \omega_i \tag{1.7}$$

The enstrophy transport equation can be derived from the vorticity transport equation (see Davidson (23)), and is defined symbolically and in tensor notation in Equation 1.8. Here the terms from left to right represent the enstrophy rate of

change with respect to time (Φ), advection (Υ), stretching or compression (Π), viscous diffusion (Ψ) and viscous dissipation (Ξ), with the balance of these terms determining the evolution of this quantity.

$$\frac{\partial \left(\frac{1}{2} \omega_i \omega_i \right)}{\partial t} + u_j \frac{\partial \left(\frac{1}{2} \omega_i \omega_i \right)}{\partial x_j} = \omega_i \omega_j \frac{\partial u_i}{\partial x_j} + \nu \frac{\partial^2 \left(\frac{1}{2} \omega_i \omega_i \right)}{\partial x_j \partial x_j} - \nu \frac{\partial \omega_i}{\partial x_j} \frac{\partial \omega_i}{\partial x_j} \tag{1.8}$$
$$\Phi + \Upsilon = \Pi + \Psi - \Xi$$

In the absence of sufficient temporal resolution, the time rate of change of enstrophy can be calculated from the other terms, with this quantity denoted as the equivalent rate of change, Φ_{eq}, shown in Equation 1.9.

$$\Phi_{eq} = \Pi + \Psi - \Xi - \Upsilon \tag{1.9}$$

The distortion of an element is given by the angular strain rate, which can be defined through the strain rate tensor, S_{ij}.

$$S_{ij} = \frac{1}{2} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) \tag{1.10}$$

The local rate of energy dissipation is defined in terms of this strain rate (equation 1.11), where ν is the kinematic viscosity.

$$\epsilon = 2\nu S_{ij} S_{ij} \tag{1.11}$$

Thus, calculation of the spatial distribution of these quantities in numerical and some experimental studies permits both qualitative and quantitative study of their correlation and interaction, with single point analysis providing useful statistical clarification of these observations.

However, it should be noted that the velocity gradient tensor (VGT) $A_{ij} = \partial u_i / \partial x_j$ can be decomposed into these *symmetric* and *anti-symmetric* components respectively, $A_{ij} = S_{ij} + W_{ij}$. Strain rate and vorticity are not independent quantities (23) which complicates the study of these, prompting some investigators to use identification parameters which express the balance of these component quantities.

1.1.5 Pressure Minima

As noted by Cadot et al. (20) there exists a simple relationship between vorticity, strain and pressure, p (equation 1.12), and therefore the pressure difference can be used to determine the relative strength of the vorticity and strain fields. It is thought that this pressure difference is established through the balancing of pressure and centrifugal forces (the cyclostrophic balance), although this is only strictly true in steady inviscid planar flow (50).

$$\frac{1}{\rho}\nabla^2 p = (W_{ij}W_{ij} - S_{ij}S_{ij}) \tag{1.12}$$

Iso-surfaces of pressure can be used to identify coherent vortical motion, with some authors claiming advantage over vorticity schemes (102). However, it is noted that in densely populated regions pressure iso-surfaces may fail to capture all flow structures (27).

1.1.6 Topological Methods

Introduced by Chong et al. (22) topological methods are theoretically able to identify vortical motion without threshold or reference-frame constraints, and thus, may provide a more general and widely acceptable definition of vortical motion. Following the description in Ooi et al. (85) the solution to the characteristic equation of the VGT is given in Equation 1.13, where P_A, Q_A and R_A are the first, second, and third tensor invariants respectively, which can be calculated as in Equations 1.14-1.16, where $tr[\cdots]$ is the trace of the matrix.

$$\lambda_i^3 + P_A\lambda_i^2 + Q_A\lambda_i + R_A = 0 \tag{1.13}$$

$$P_A = -A_{ii} = -tr\,[A] \tag{1.14}$$

$$Q_A = P_A^2 - \frac{1}{2}A_{ij}A_{ji} = \frac{1}{2}\left(P_A^2 - tr\,[A^2]\right) \tag{1.15}$$

$$R_A = -\frac{1}{3}A_{ij}A_{jk}A_{ki} = \frac{1}{3}\left(-P_A^3 + 3P_AQ_A - tr\,[A^3]\right) \tag{1.16}$$
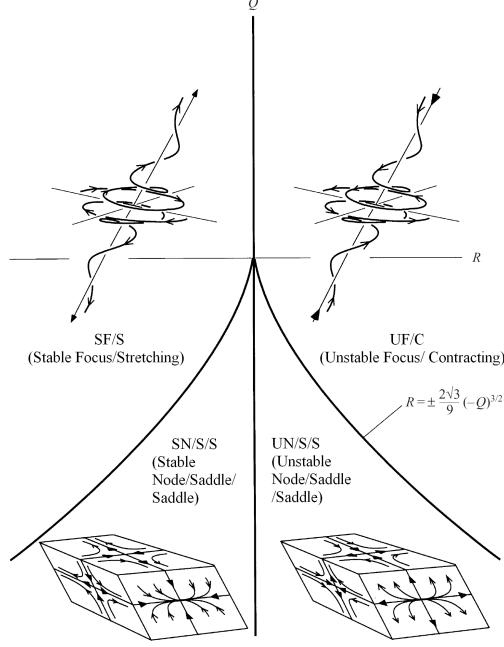
Figure 1.1: $Q_A - R_A$ plane flow topologies (85)

For incompressible flow, $P_A = 0$, and the local flow topology can be characterised solely in terms of Q_A and R_A. Therefore, the position of any flow point on the two-dimensional $Q_A - R_A$ plane defines it in terms of one of four non-degenerate topologies (see Figure 1.1). JPDFs of these quantities tend to have the same generic shape for a range of flows, with a predominance of vortex stretching and bi-axial strain[1] (23), for example see Ooi et al. (85) and O'Neill & Soria (84).

Similarly topological invariants can also be calculated for the symmetric rate of strain tensor S_{ij}, (P_S, Q_S, R_S) and the anti-symmetric rate of rotation tensor W_{ij}, (P_W, Q_W, R_W). Noting that again in incompressible flow $P_S = P_W = 0$, $R_W = 0$, and Q_S and Q_W are negative and positive definite respectively. The rate at which vorticity is stretched and contracted can also be defined in terms

---

[1]Bi-axial strain is defined here in terms of the principal strain field as one large compressive and two smaller extensional strains, whereas vortex stretching is defined as one large extensional and two smaller compressive strains.

of topological invariants, as in Equation 1.17.

$$\Sigma = \frac{\omega_i S_{ij} \omega_j}{\omega_k \omega_k} = \frac{R_S - R_A}{Q_W} \tag{1.17}$$

1.1.6.1 Q_A-Criterion

Hunt et al. (46) used the second invariant, Q_A, to define vortical motion (Equation 1.15). $Q_A > 0$ indicates a dominance of vorticity over strain and suggests vortical motion, and as vorticity increases towards the centre of a vortex Q_A is expected to remain positive inside a vortex core region. It is noted that there is an obvious relationship between this and the pressure criterion in Equation 1.12, as $2Q_A = \nabla^2 p / \rho$.

1.1.6.2 Δ_A-Criterion

Chong et al. (22) propose a critical point scheme, where a vortex core is defined at points where complex eigenvalues imply a closed or spiral streamline pattern around the point, and this occurs at points where the VGT discriminant is positive (Equation 1.18).

$$\Delta_A = \frac{27}{4} R_A^2 + Q_A^3 > 0 \tag{1.18}$$

This discriminant definition is also related to the Q_A-criterion by noting that positive values indicate a dominance of rotation over strain, and while it is thought to be less restrictive than the Q_A-criterion it is not not entirely clear which is more appropriate (50).

1.1.6.3 λ_2-Criterion

Jeong & Hussain (50) discard unsteady straining and viscous effects in the Navier-Stokes equations to obtain their definition of vortical motion, which is only valid for incompressible flow. Therefore, the two conditions of Hunt et al. (46) are combined into a single definition: that vortex cores exist in connected regions with negative eigenvalues of $S_{ik}S_{kj} + W_{ik}W_{kj}$. If the eigenvalues are ordered $\lambda_1 \geq \lambda_2 \geq \lambda_3$, this definition therefore requires that $\lambda_2 < 0$ in the vortex core. As

such this method is shown be closely related to the Q_A-criterion, although slight variations are observed (50).

1.1.6.4 Swirl Criterion

Zhou et al. (143) define swirling strength, λ_{ci}, as the imaginary part of the complex eigenvalue of the velocity gradient tensor. In addition to removing reference frame constraints, this method will not highlight regions of vorticity which have no local spiralling motion (for example shear layers), and therefore is widely applied in boundary layer investigations.

1.1.7 A Comparison of Identification Schemes

Jeong & Hussain (50) argue that the common methods of identifying vortices such as iso-vorticity surfaces, pressure minima, and closed or spiral stream lines, albeit intuitive, are ultimately fallible. Due to the lack of agreement between these definitions and the problems that beset them, they recommend using more general topological approaches.

Pressure minima do not necessarily indicate net vorticity, and can be created through straining or viscous effects and as such may be inadequate for a general definition of such structures. Additionally, considering a Burgers vortex, although the length scale of vortex core is fixed, the position of the local pressure maximum ($\partial p/\partial r = 0$, where p is the pressure and r is the radius from the centre) is not, and will increase with vorticity, and as such pressure-isosurfaces may fail to identify some motions (50).

Closed or spiral streamlines are not Galilean invariant and therefore changing the frame of reference may change the appearance of these. For example, in a moving frame of reference streamlines may no longer appear closed (or for a fixed frame of reference with rapid vortex advection), which could possibly alter their classification. Additionally, vortical motions which do not undergo a complete revolution will also not contain closed pathlines and may be omitted. Furthermore, Melander & Hussain (72) show that in reconnecting regions, particularly for high Reynolds numbers, pathlines and streamlines can become highly distorted and may become difficult to identify reliably.

Iso-vorticity surfaces suffer from thresholding problems, where variation of vorticity in a single structure may result in the prediction of several separate vortices. Additionally, under this approach the magnitude of the structures must exceed the general background level, and therefore it is impossible to identify all but the strongest structures (20).

To illustrate the differences between various vortex identification techniques, Dubief & Delcayre (27) compared schemes in a variety of flows, with their study of homogeneous isotropic freely decaying turbulence particularly relevant to the current study. As shown in Figure 1.2 all schemes identify filament structures, although the pressure criteria finds fewer but larger diameter structures, failing to capture some finer details. The vorticity criterion is the only scheme which identifies layered sheet structures in addition to the filaments. The Q_A-criterion and λ_2-criterion produce similar fields without these sheet-like structures, with the latter subject to additional noise.

The large differences between pressure and vorticity criteria illustrates the difficulty with comparing results analysed using different methods. After comparing all methods over a range of flows Dubief & Delcayre argue that the Q_A-criterion is more suitable than other methods, in that it is insensitive to mean-shear and large scale pressure variations, demonstrating an ability to pick out streamwise vortices close to Kelvin-Helmholtz instabilities in a mixing layer (27).

Jeong & Hussain (50) note the similarity between Q_A and λ_2-criteria in a number of cases, arguing that their scheme is superior, with the Q_A-criterion misrepresenting structures in the presence of strong external strain. However, Dubief & Delcayre (27) note that the flow configurations in this investigation, such as the Bödewadt vortex[1] are unlikely in common turbulent flows, and therefore, due to the intrinsic link between them, these methods largely predict similar results. Jeong & Hussain (50) also note that the discriminant criterion takes low positive values outside vortex cores in DNS mixing layer and vortex ring data, resulting in noisier visualisations with slightly larger core sizes and a large number of additional small scale structures.

Despite these comparisons, it is extremely difficult to determine if one method offers clear advantage over the others, as the performance of each is relative to the

---

[1]A vortex normal to a stationary wall.

(a) Pressure iso-surfaces, $P = -0.15\rho U^2$ (b) Vorticity iso-surfaces, $\omega = 0.36\omega_I$

(c) λ_2 Iso-surfaces, $\lambda_2 = -0.25$ (d) Q iso-surfaces, $Q = 0.5$
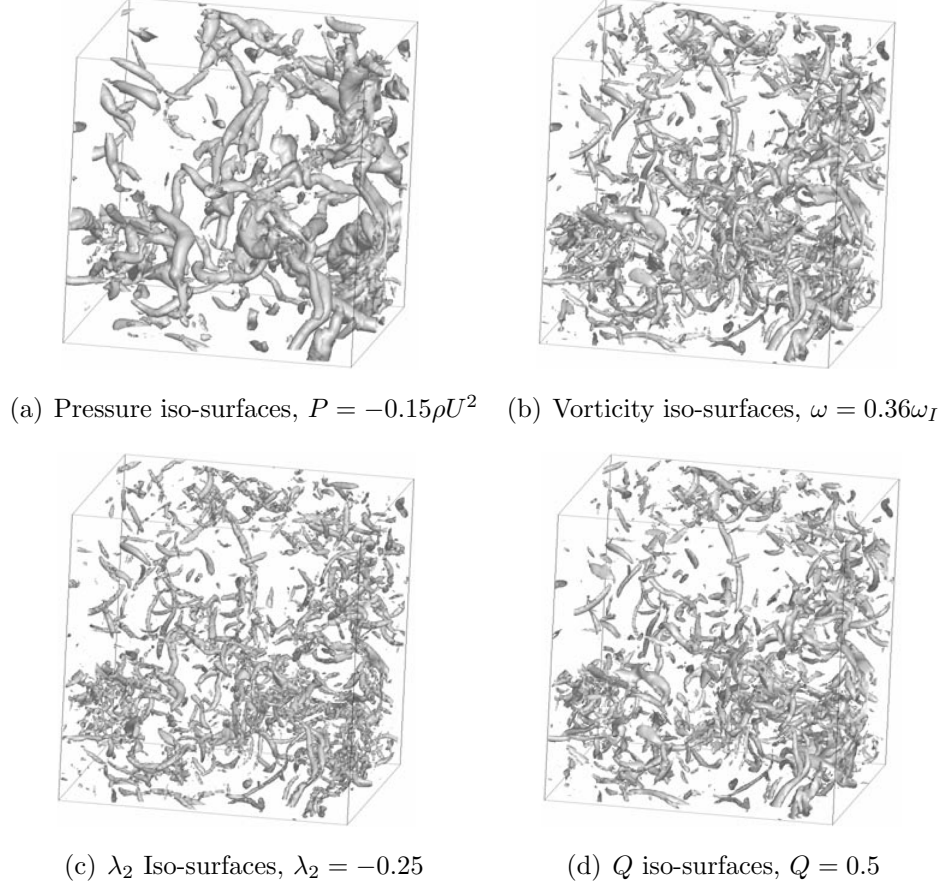
Figure 1.2: Comparison of vortex identification techniques in isotropic homogeneous turbulence (27) $U = \sqrt{3}$, $\omega_I = 2\pi/\sqrt{3}$

particular flow in question, and without a definitive answer there is no rigourous way to determine relative scheme accuracy.

1.2   Dissipation Scale Coherent Structures

After their discovery by Siggia (110), a large number of numerical and a smaller number of experimental studies have provided support for the idea that the smallest scales of high Reynolds number turbulence contain coherent turbulent structures. Further elucidation of these structures should contribute towards an understanding of dissipative intermittency, providing valuable material for con-

structing statistical models based on real flow behaviour, and also giving insight into the stretching of material lines and planes by these structures, which is of great importance with regards to the study of mixing phenomena (59).

Typical numerical and experimental methods will now be introduced, to permit assessment of their relative advantages and shortcomings, before discussing the findings of various studies, and identifying areas where further work would be useful.

1.2.1   DNS Approach

Direct integration of the Navier-Stokes equations is known as Direct Numerical Simulation (DNS); a technique which has gained popularity with the availability of increasingly powerful computing resources. However, due to large costs, DNS cannot currently be used as a practical engineering tool, and is limited to providing flow solutions to geometrically simple problems at relatively low Reynolds numbers. Typical simulations are currently of $O(R_\lambda) \sim 100 - 500$, although some have reached $R_\lambda = 1200$ (56). However, despite this constraint these simulations are very useful for studying small scale coherent structures, contributing significantly to the current understanding.

To maximise Reynolds number it is common practice to use the simplistic geometry of the periodic cube, which allows for application of very efficient spectral algorithms. The periodic cube is characterised by unphysical forced symmetry between opposite faces, with the assumption that these will have little influence on the bulk flow provided the largest eddies, $O(\Lambda)$, are much smaller than the domain size, L_D. However, the lure of high Reynolds numbers has resulted in rather high Λ/L_D ratios ($\sim \pi^{-1}$), meaning the large scales may be influenced by the artificial boundary conditions. Therefore, if the small scales remain dependent on the large scales at these Reynolds numbers (as shown by Shen & Warhaft (109) and Sreenivasan (113)), then the predictions pertaining to these scales may also be affected.

The small scale resolution of these studies is set by the need to accurately capture the dissipation spectra, typically around $\kappa_{max}\eta \approx 1.5$ based on the maximum wave number, which equates to a physical grid spacing of $\Delta x/\eta \approx 2.1$ (93).

Therefore, although sufficient for resolving the majority of the dissipation spectra, other resolution effects may persist. Caution is therefore advised concerning the results of these numerical investigations, as although filament-like structures are present in the large number of DNS simulations performed with different forcing schemes (20), the structure characteristics and Reynolds number dependence are yet to be definitively supported experimentally.

1.2.2 Experimental Approach

In experimental investigations adequate spatial and temporal resolution of high Reynolds number small scale turbulence is extremely challenging. For example in the laboratory setup of Belin et al. (7) using $L = 100$mm, the smallest length and time scales are of order $\eta \sim 12\mu$m and $\tau \sim 1.3$ms respectively at $R_\lambda \approx 800$. Therefore, both small diameter probes or measurement volumes with extremely high-frequency response must be used to capture these motions.

These difficulties have led to a number of novel approaches, for example using extremely low viscosity fluids such as cryogenic helium (7; 120; 144) to permit a large range of Reynolds numbers in small apparatus, or very rapid response equipment-such as position sensitive silicon photo diodes (63; 127), and sonar-doppler techniques (81) to increase measurement resolution. Approaches often include other more standard rapid response sensors such as hot-wire, hot-film, Laser Doppler Anemometers, high-frequency pressure transducers, and direct visualisation using very high frame-rate video. However, despite these measures, full small-scale resolution is either rarely achieved, or is so at the cost of Reynolds number or measurement dimensionality.

For example, in a helium gas flow experiment ($R_\lambda = 151 \rightarrow 2300$) Belin et al. (6; 7) measured velocities using a 7μm diameter carbon fibre hot wire, giving a spatial resolution between $0.2 \rightarrow 3\eta$ (although the longer wire length will result in a lower resolution).

Resolution constraints also typically limit measurements to a single dimension, through choice of experimental technique. In these cases, measurements must be combined with assumptions of local isotropy and Taylor's frozen turbulence hypothesis in order to infer the three-dimensional structure, and determine statis-

tical quantities such as dissipation; an approach with proven shortcomings (37). Furthermore, considering the possible persistence of anisotropic effects even at the high Reynolds numbers typical of these experiments (109) this approach may be considered unsatisfactory.

Although full three-dimensional measurements have been made, these too are subject to a number of restrictions. Given very limited information from single particle traces far short of full field resolution, previous 3D particle tracking experiments (63; 127) are limited to speculation regarding particle/flow structure interactions. Although Dual-plane Stereo-PIV measurements (82) permit volumetric measurements, the very thin volume restricts structural observations. Other Stereo-PIV investigations are forced to rely on Taylor's hypothesis to produce quasi-instantaneous volumes from time-resolved measurements (37). Furthermore, PIV methods such as these are also typically limited to moderate Reynolds number ($R_\lambda \lesssim 150$).

The use of pressure transducers limited measurements by Cadot et al. (20) to the near wall region to avoid interference effects. Here secondary flow stabilisation and attached filament stretching gave rise to extended lifetimes and broadening in these regions, biasing these measurements. Additionally, the measured pressure drop magnitude was strongly dependent on filament proximity to the transducers, making reliable strength measurement difficult. Bubble visualisations (20; 26) also suffer the same problem as iso-vorticity thresholding, and as such coherent structures with weaker pressure fluctuations may not be identified, potentially biasing results towards tubular structures (63). Additionally, due to their size the bubbles may not accurately follow the small scale structures, for example the 0.05mm bubbles used by Douady (26) to measure structures of approximately 0.1mm diameter.

1.2.3 Vorticity Filaments

Despite various numerical and experimental limitations these approaches have led to a rough picture of the small scale structures; where some details regarding the local vorticity and strain field interaction and structure scaling properties and Reynolds number dependence are known.

Simulations generally identify persistent tubes, sheets and blobs of small scale vorticity (116). As noted by She et al. (108) the highest vorticity tends to be contained in tubular structures, with moderate values in sheets, becoming pattern-less at lower amplitudes. Therefore, typically investigations focus on the most coherent of these structures: thin threads of concentrated vorticity known as vortex *tubes*, *filaments* or *worms*. Some remarks now follow on the speculated dynamic life cycle of these structures, beginning with their formation, discussing some of their typical properties, before ending with their breakdown.

1.2.3.1   Formation of Filaments

The creation of vortex filament structures has been the focus of much interest, and in general it can be stated that these are formed in the presence of axial or (via the Kelvin-Helmholtz instability) bi-axial strain fields (78). By expressing the rate of enstrophy generation in terms of the strain field, Betchov (9) demonstrated that given that enstrophy generation is positive, a pre-dominance of bi-axial strain is expected, which may favour the latter mechanism. Brachet et al. (13) support this after observing the formation and subsequent thinning of pancake-like sheets of vorticity subject to constant bi-axial strains. At some point during this vortex stretching, the sheet is subject to an instability and tubes of concentrated vorticity result from the decomposition of the layer.

Passot et al. (89) used a modulational perturbation analysis technique to kaneda06further the work of Lin & Corcos (67), demonstrating that when thin strained vortex sheets become unstable, vorticity tends to concentrate in tubular filament structures. From a 256^3 grid point DNS study of decaying turbulence, visualisation techniques were used to deduce that the filaments form under the combined influences of sheet compression and self-induced rotation. It was demonstrated that when regions of strong vorticity are aligned with a constant intermediate principal rate of strain, the vorticity becomes concentrated. As noted by She et al. (108), at any given time concentrated vorticity is present in the form of vortex tubes and more moderate values in the form of sheets. Therefore, it is assumed by some authors that the tubes form via the sheet instability, and increase in strength by continuously rolling up the sheets vorticity.
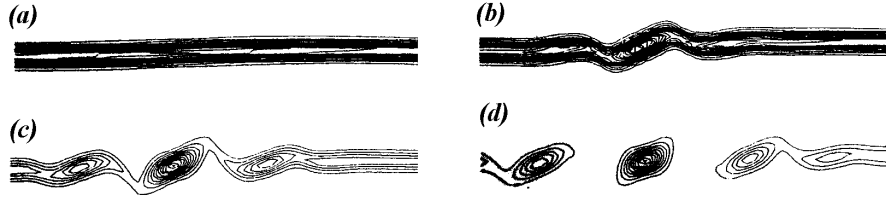
Figure 1.3: Visualisation of vorticity contours from asymptotic vortex filament formation model (Passot et al. (89))

To a certain extent these mechanisms are supported experimentally in studies of the flow between counter rotating impellers. Bonn et al. (10) and Douady et al. (26) suggest that tubes form from the roll-up of thin shear layers in the presence of shearing and stretching. Cadot et al. (20) also argue that filaments attached to the tank wall form as a consequence of mean shear layer roll-up, whereas other smaller weaker filaments may be generated in a process similar to classical mixing layer vortex generation by other large eddies, thereby inheriting length scale properties from their parent vortices.

In the case of larger scale filament structure formation, sheet decomposition takes place before the sheet has become sufficiently thin, and as such the time-scales are too slow to be observed using DNS (89). Cadot et al. (20) argue that both larger scale *low pressure filaments* and smaller scale *vorticity filaments* may form as previously described. However, further to the numerical investigations, they also suggest that *vorticity filaments* may also form as a *low pressure filament* bursts into several thinner intertwined filaments (see Figure 1.4). The interaction and stretching of these by one another lead to an increase in vorticity, which is accompanied by an increase in dissipation between separate strands, or sections of the same contorted strand.

In conjunction with the basic formation mechanism Vincent & Meneguzzi (125) suggest that the longest vortex filaments are formed by the amalgamation of several shorter parallel filaments originating from the same shear region, and it is this behaviour that allows these structures to develop to such extreme lengths. This view is supported to a certain degree experimentally by Cadot et al. (20) who note that although filaments exist in all orientations, the largest of these are aligned with mean shear layer.

(a) Low pressure filament   (b) Intertwined vortex filaments
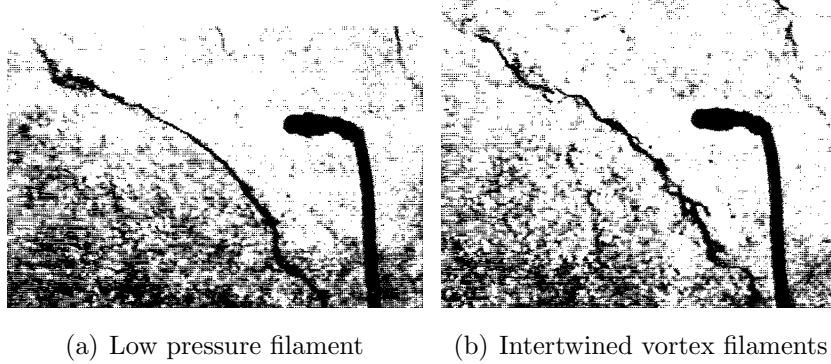
Figure 1.4: Visualisation of a *low pressure filament* bursting into *vortex filaments* (Cadot et al. (20))

1.2.3.2   Strain and Vorticity Interaction

The relatively low Reynolds number ($R_\lambda \leq 83$) results of Kerr (58) show concentrations of vorticity in tubular and sheet like structures, with scalar gradients wrapped around the tubes. The strain field appears to be well correlated with these structures, with the following arrangement of the principal rates of strain from largest to smallest with respect to the tubular structures: radial compressive, radial stretching, and axial stretching. It was noted that the intermediate strain was usually therefore aligned with the vorticity, and the compressive strain with the scalar gradient.

Ashurst et al. (4) attempt to quantify this relationship through intermediate strain rate *PDFs*. For low magnitude intermediate strains these appear symmetric, however as strain rate is increased their distribution becomes increasingly asymmetric, with most probable principal strain rate distribution from highest to lowest suggested as 3:1:-4. The agreement of these predictions and previous analysis by Betchov (9) is noted by a number of authors (23; 116). Betchov predicted that on average one strong compressive strain is accompanied by two weaker stretching ones, with this bi-axial strain resulting predominantly in the formation of vortex sheets.

Kida (59) identified filaments using the characteristically low pressure cores in conjunction with the local peak in u_θ^2/r, where u_θ is the azimuthal velocity, to study steady simulations from $R_\lambda = 80 \rightarrow 170$. This study attempted to

quantify the location of these important quantities, suggesting that the majority of the enstrophy is located at the vortex cores ($\approx 60\%$ of the total), whereas the dissipation ($\approx 30\%$ of the total) as in the Burgers vortex surrounds the core. Although it is noted that these results remain approximately independent of Reynolds number, the range of Reynolds numbers is perhaps too small to support this conclusively. The separation of high magnitude vorticity and dissipation may account for vortex filament stability, explaining their typical existence of several turnover times. A consequence of this coherence is their observability in DNS simulations, although a more important feature may be their sustained contribution to the flow field.

Kida also suggests that the strong swirling motion around the vortex core may trap small stray filaments, resulting in the normal component of vorticity having a complex distribution around the filament. However, the strongest dissipation in the annulus region cannot be necessarily attributed to this cross-axial vorticity, as the Burgers vortex also contains a similar dissipation distribution (23). The trapping of stray filaments gives the filament a spiral nature (visible in Figure 1.5), although this is different to the spiral vortex in Lundgren's model (70).

(a) Vorticity mag.   (b) Axial vorticity   (c) Cross-axial vort.   (d) Dissipation
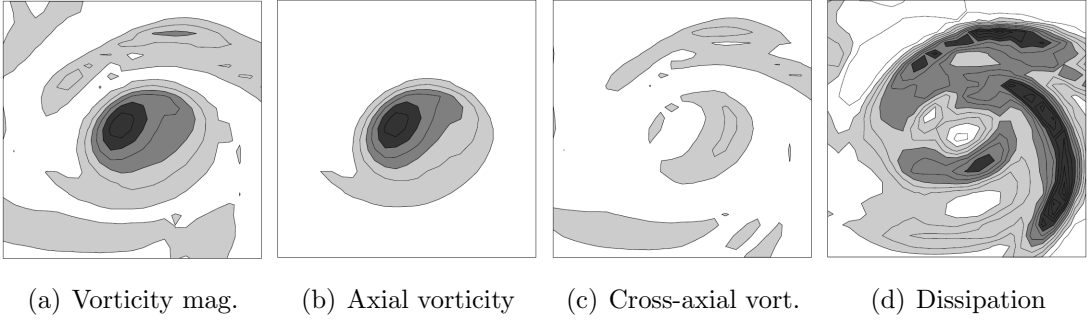
Figure 1.5: Cross-section ($41^2\eta$) of a spiral vortex filament, diameter $\approx 10\eta$, CCW rotation, $R_\lambda = 170$, dark indicates high values (Kida et al. (59))

Recently Kaneda & Ishihara (56) performed very large scale 4096^3, $R_\lambda = 1200$ simulations. Simultaneous plotting of high vorticity and high dissipation fields (see Figure 1.6) shows a strong correlation, similar to the findings of Kerr et al. (58). However, as stated rather disappointingly, no attempt is made to

understand or further quantify the relationship between the coherent structures and dissipation, or its variation with Reynolds number in this study.
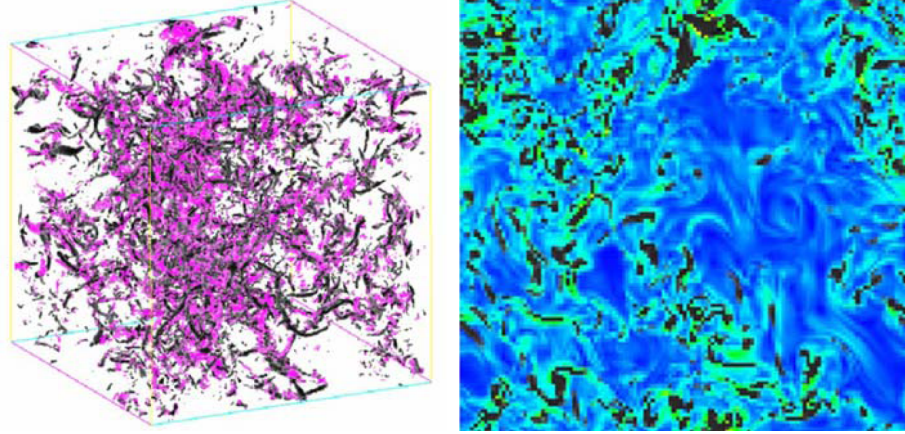
(a) 3D distributions, vorticity, $\omega >$ $\langle\omega\rangle + 3\sigma_\omega$ shown in purple, dissipation, $\epsilon > \langle\epsilon\rangle + 3\sigma_\epsilon$ shown in black

(b) Contour plot of dissipation, $\epsilon >$ $\langle\epsilon\rangle + 3\sigma_\epsilon$ (vorticity $\omega > \langle\omega\rangle + 3\sigma_\omega$ shown in black)

Figure 1.6: Comparison between regions of intense dissipation and intense vorticity $R_\lambda = 167$ (56)

A number of these findings have also been repeated experimentally. Ganapathisubramani et al. (37) made far field measurements of an axisymmetric co-flow jet using time-resolved stereo PIV. Taylor's hypothesis was used to reconstruct a quasi-instantaneous 3D volume, in which similar structural observations were made. The study also identified similar principal strain rate distributions and alignments, which have also been found in previous hot wire (123) and Dual-Plane SDPIV investigations (82).

1.2.3.3 Typical Scaling Properties

Some insight into typical vortex filament scaling has also been gained, including details such as: diameter, length, shape, population density, and the dependence of these on Reynolds number. Filaments are usually described as having approximately circular cross sections, with diameters ranging between λ and η, and lengths between λ and Λ, approximately resembling Burgers vortices. The diameter and length are however extremely difficult to predict accurately, and slight

differences in the findings of various authors may also be attributed to the use of different analysis techniques. Additionally, as noted by Sreenivasan & Antonia (116), inconsistencies may also arise from the poor scale separation in these low Reynolds number results, making an accurate assessment difficult.

Numerical investigations (53; 56; 58; 110; 125) often suggest dissipation scale filament diameters $\sim 6 \rightarrow 10\eta$, although larger diameter filament structures are not precluded. Jiménez & Wray (52) consider that less intense, larger diameter vortical motions may also become coherent given higher vorticity than the background level (ω_{rms}), and if these took azimuthal velocities of u_{rms}, then their radii could be $\sim \lambda$. However, the high thresholds required for visualisation show only the highest vorticity structures.
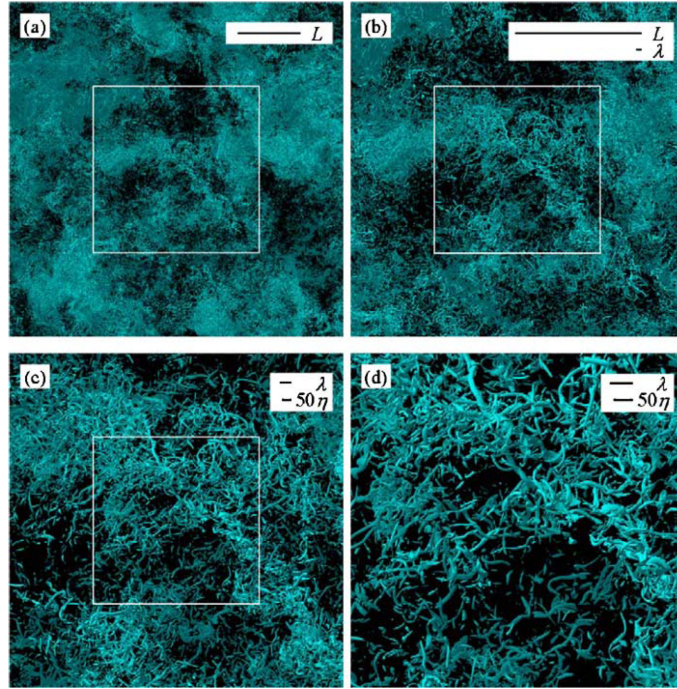
Figure 1.7: High Vorticity regions $\omega > \langle\omega\rangle + 4\sigma_\omega$ from a high Reynolds number $R_\lambda = 732$ simulation (56)

Experimental investigations tend to support the numerical findings, with similar scale structures first observed by Douady et al. (26) using bubble visualisations. Ganapathisubramani et al. (37) also find intense worm-like vortex structures with characteristic diameters and lengths $\sim 10\eta$ and $60 \rightarrow 100\eta$ respectively.

As previously mentioned Cadot et al. (20) identified a larger range of radii from $\sim \lambda \to \eta$. Furthermore, although only relatively long structures up to size L are observed, these authors claim that lengths down to the dissipation scale may be possible, assuming strain must be coherent over the length of the structure. However, this claim is disputed by Jiménez & Wray (52) on the basis that strong vortices will be de-coupled from the background, and even in the case of locally compressive strains the coherent vortex structure can be maintained (124).

Jiménez et al. (52; 53) performed a sequence of forced steady simulations from $R_\lambda = 40 \to 170$, proposing additional scaling relationships for the maximum azimuthal velocity of u_{rms}, and maximum axial vorticity of $\omega_{rms} R_\lambda^{1/2}$. However, due to the reasonably high threshold $\Omega \approx \Omega_{rms}$ employed in these studies it is clear that these results only relate to the most intense vorticity structures. Studies by Belin et al. (6; 7) support these length and velocity magnitude scalings up to Reynolds numbers of $R_\lambda \approx 700$. Additionally Cadot et al. (20) found the magnitude of the filament low pressure regions scales only with the square of the injection velocity, $\sim U^2$, and although the circulation was found to be approximately constant, the pressure drop depended on filament length. At much higher Reynolds numbers it is postulated that filaments may blend into the background (6), however much higher Reynolds number results may be required to confirm these Reynolds number scalings.

Both numerically and experimentally filament population density has been shown to remain approximately constant with Reynolds number (20; 52; 59), and therefore their fractional volume decreases as $R_\lambda^{-1/2}$ (52). These findings are however contrary to those by Abry et al. (1) who found that population density increased with Reynolds number, although the use of different stirring mechanisms does not permit a more conclusive analysis.

Belin et al. (6; 7) also observed that beyond $R_\lambda \approx 700$ filaments tend to widen, filling more space and reducing flow intermittency; a change confirmed by the plateau of skewness and flatness at this threshold. The transition is linked to the idea that the internal Reynolds number of the filaments increases with Reynolds number (52), and therefore the filaments will become unstable, with the $R_\lambda \approx 700$ threshold possibly setting the transition point. Although this transition is also observed by Tabeling et al. (120), it requires further experimental support (116)

as it is not observed in some other experiments (90), and therefore may simply be caused by a particular facet of the experimental set-up or resolution.

1.2.3.4   Filament Breakdown

As observed by Douady et al. (26) although initially straight, filaments rapidly become unstable, decaying in a process of vortex breakdown and forming larger, less intense eddies. This observation is supported by Cadot et al. (20) who note that after becoming unstable, roll up is shown to continue; characterised through an increasing circulation in the larger diameter eddy left in a filament's wake.

The decay of the larger *low pressure filaments* to smaller *vortex filaments* discussed by Cadot et al. (20) also signifies the breakdown of a coherent structure. It is interesting to note that this breakdown could at one stage bear resemblance to the spiral structures observed by Kida (59). Cadot et al. also note that the majority of the current DNS results are at insufficient Reynolds number to capture this form of decay, and therefore the structural evolution of this form of breakdown may benefit from further study.

1.2.3.5   Significance of Vortex Filaments

The influence of various structures on the statistical properties of turbulence is of great importance, with continuing debate over the relative dominance of sheets, tubes and ribbon like structures. Brachet (12) and Bonn et al. (10) suggest that filaments will not govern the dynamics of the small scales as they contain only a small fraction of the dissipation and vorticity. Jiménez et al. (53) added support to this argument by artificially removing these structures without significantly affecting dissipation. On the other hand, Kida (59) suggests rather more influence, through significant correlation between the most intense vorticity and dissipation fields. However, the variation of this fraction with Reynolds number is still not yet entirely certain (35), and the different definitions of these filament structures clouds this debate.

Bonn et al. (10) used Polyethylene oxide ($\approx 3 \times 10^6$ amu) and Polyox coagulant ($\approx 5 \times 10^6$ amu) polymers to inhibit vortex filament formation in the flow between counter-rotating disks. Bubble visualisations indicate the introduction

of polymers significantly reduces filament formation, and also inhibits large scale motions. For example, the smaller particles at a concentration of 40ppm (by weight) approximately halved the number of filaments produced. The small sized polymers are advected by the large scales, only interacting dynamically with the smaller scale motions, and therefore it is the large scale motions resulting from the decay of vortex filaments which are suppressed. However, it is concluded in this investigation that by disrupting the process by which the filaments form, then the process that controls dissipation may also be disrupted, and the absence of filaments are merely a symptom of this disruption. Further to this, Cadot et al. (19) use a variety of forcing schemes to show the suppression of coherent structures by the polymers can take place without any reduction in drag, and therefore the structures may not be directly related to the energy dissipation.

1.2.3.6 Inertial Range Structures

Although the structure of the finest scales has been widely identified through DNS results, whether similar structures are present at all levels in the energy cascade is still unclear. Vincent & Meneguzzi (125) support the idea that the cascade is characterised by the formation of sheets and their subsequent breakdown into tubes. Boratav & Peltz (11) found that ribbon-like structures are dominant in the inertial range, not tubes. Furthermore, by comparing the effects of both enstrophy and strain dominated structures, they suggest that an intermittent strain field as opposed to the vorticity field may result in the shortcomings of Kolmogorov's original model. In the numerical simulations of Porter et al. (94) Favre averaging was used to separate dissipation and inertial range structures, demonstrating that although larger scale structures were filament like, they were distinctly shorter and showed increased curvature compared to the smaller scales.

Lundgren (70) and Pullin & Saffman (98) attempted to use random distributions of elementary structures, such as those uncovered by DNS results to model the inertial range. However, as illustrated by Kaneda & Ishihara (56) (see Figures 1.6 and 1.7) and supported by other authors (53; 94), the distribution of these structures may not be random, with filament structures apparently clumped together in distinct regions. Porter et al. (94) suggest there is a correlation between

the large scale eddies, with the filaments filling the background between sparsely populated larger scale motions.

Moisy & Jiménez ([79]) attempted to quantify these essentially qualitative observations, albeit at the relatively low Reynolds number of $R_\lambda = 168$, applying box counting methods to an existing DNS dataset ([53]). An interesting approach is used to characterise the geometry of coherent vorticity and dissipative structures; each structure is assigned three characteristic inner, outer and intermediate lengths, which are used to classify structures over a range of thresholds between four extremes: sheet-like; spherical; tubular; and ribbon-like. The moderate and intense dissipative structures tend to be flat sheet-like or ribbon-like, and the moderate and strong vortical structures tend to be ribbon-like and tubular respectively. Moreover, individual structure centres appear to be grouped in hierarchical inertial scale clusters between $20 \rightarrow 400\eta$, and although the clusters are shown to be distributed self-similarly, the simulation Reynolds number is insufficient to effectively isolate a characteristic fractal dimension.

More knowledge of the inertial range structures would be extremely useful to further quantify the strain and vorticity field interaction at higher Reynolds numbers, the intermittency of these, and any clustering that can be related to larger scale structures. More accurate knowledge of this small scale clustering could improve statistical models through more accurate fundamental structure distributions.

1.3 Flow Measurement Techniques

The current investigation will follow an experimental approach, and therefore a discussion regarding method selection is now presented.

Dissipation scale measurements have been performed using a variety of instruments, such as hot-wire, hot-film, Laser Doppler Anemometry (LDA), pressure transducers, and high speed PTV through use of silicon accelerometers. These techniques are largely used in response to the high-frequency demands made by many of the experimental set-ups (see Section 1.2.2). However, the larger length and time-scales of the present experimental set-up (Section 2.1.1) relax these constraints somewhat, allowing consideration of a broader range of techniques.

As discussed by Jiménez & Wray ([52]), it is difficult to accurately assess the evolution of coherent flow structures using single-point measurements, which must typically invoke assumptions such as local flow isotropy, and Taylor's frozen turbulence hypothesis in order to describe statistical properties, such as the dissipation rate. Furthermore, the influence of invasive detectors placed in the flow cannot be altogether discounted ([20]), and therefore non-invasive techniques are preferable. Particle Image Velocimetry (PIV) therefore appears to be an appropriate choice, adequately satisfying the measurement frequency requirements, and providing good spatial resolution.

Additionally, in order to remove dependence on the assumptions of isotropy and permit a more intuitive study of coherent flow structures, full 3D measurements are highly desirable. Therefore, after a brief introduction to PIV, a number of 3D extensions to the basic technique are reviewed before selecting one of these methods for further study.

1.3.1 Digital Particle Image Velocimetry

Digital Particle Image Velocimetry is a non-invasive planar measurement technique (Figure [1.8]). The flow is seeded with small particles and illuminated twice in rapid succession typically using a laser light sheet ([2]). The scattered light is recorded on either a doubly exposed image or more commonly a pair of images. The particle displacement is obtained by dividing each image pair into small interrogation regions and correlating corresponding windows. Therefore, this process does not take individual particle motion into account, instead calculating the mean motion of a group of particles. The displacements are converted from image to global co-ordinates through calibration.

Spatial resolution is a function of the optical set-up and can reach $\sim 10\mu m$ in microfluid applications ([2]), producing 2D fields with $\sim 10^4$ vectors. However, sampling frequency is limited by laser repetition rate, which for high powered lasers can be as low as 10Hz.

Recent advances have enabled extension of the basic PIV technique into three dimensions, through a variety of different methods. An overview of the most popular of these techniques, and their relative merits follows.
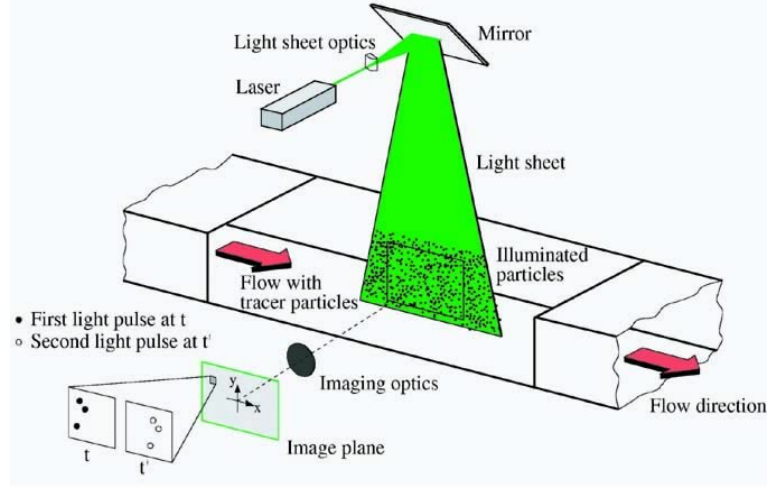
Figure 1.8: PIV Schematic

1.3.2 Stereoscopic PIV

Stereoscopic-PIV resolves all three velocity components in a planar region of interest (3; 95; 135). The illuminated plane is imaged from two different directions, with camera orientation determined through calibration with a known 3D object. The particle velocity projection on each image plane can be evaluated, correcting for perspective error, and the vectors can combined to reconstruct all 3 velocity components (100; 135).

The accuracy of the out-of-plane velocity component is $1/\tan(\theta)$ times the in-plane component, where θ is the camera angle with respect to the light sheet (95). Therefore, although an angle of $\theta = 45°$ is optimal this may be impractical due to other conflicting requirements such as: increased lighting power; aperture settings; or camera positioning limitations. Accuracy is also highly dependent on precise calibration plate positioning at the laser sheet location, although this can be improved somewhat through advanced self calibration methods (131).

Stereoscopic-PIV increases the number of gradient terms which can be directly measured to 6, which is a slight improvement over the 4 which can measured using DPIV. However, in order to calculate the entire velocity gradient tensor, it is possible to extend this technique by using multiple closely spaced parallel light sheets, which is known as Dual-Plane Stereo-DPIV (54). The two light

31

sheets are polarised differently, and therefore the light from each can be recorded by two cameras fitted with the appropriate filters. However, this increases the complexity of the set-up considerably, and accuracy is dependent on the precise alignment of the two laser sheets.

1.3.3 Holographic PIV

In Holographic-PIV (HPIV) light scattered from particles in a volume is recorded on a hologram in superposition to a reference light beam (42; 74). Real space particle positions are recovered by illuminating the hologram with the conjugate of the reference beam (illuminating the plate from the reverse direction). HPIV set-ups can be sub-classified as 'in-line' or 'off-axis' based reference light beam position. The simpler in-line technique suffers from excessive speckle noise, which limits the maximum seeding density, and the small numerical aperture formed by the forward particle scattering gives poor depth resolution. The more complicated off-axis technique allows higher seeding, better SNR and higher accuracy, although at the expense of higher laser power and more stringent coherence requirements (97).

In contrast to PIV, which is now almost exclusively digital, HPIV is still also performed on photographic film. Therefore in addition to the complexity of the set-up, standard HPIV requires the time consuming wet-processing of film and scanning of the hologram with a CCD, which also introducing errors from positioning, re-sampling, and image distortion (31).

In digital-HPIV (DHPIV) images are recorded directly on CCD chips, and through calculation of the diffraction integral using the Fresnel approximation the 3D particle information can be reconstructed (87). CCD chips are currently much lower resolution that photographic film, which limits the fine scale interference fringe detail which can be recorded, and therefore requires an in-line approach to be used; limiting accuracy and resolution (42; 87). Currently DHPIV is only capable of capturing 2 to 3 times fewer vectors than standard HPIV (31).

Off-axis HPIV allows capture of large data sets (up to 10^6 vectors), with accuracy comparable to DPIV, however the technique is complicated, expensive, and

capturing enough holograms to compute adequately converged turbulent statistics would be extremely difficult (2). DHPIV is limited in terms of its accuracy and resolution, and therefore more suitable for very small volumes; such as those used in micro-PIV (2).

1.3.4 Scanning-PIV

Another way to extend DPIV or Stereo-PIV into 3D is to rapidly scan the light sheet through the volume, sequentially capturing slices of data which can be combined to create a 3D volume of data (15; 44). The scanning rate must be high in comparison with flow time-scales (44) and therefore extremely high repetition rate lasers are employed ($\sim$kHz), which may limit illumination power. High speed cameras are also required, which maybe unable to operate at full resolution at the repetition rates required, reducing the maximum resolvable volume. Additionally, extra complexity will be introduced through the requirement to rapidly translate or rotate the light sheet in small increments.

1.3.5 Defocusing-DPIV

Defocusing-DPIV (DDPIV) modifies the optical aperture to produce 3D information from a single optical axis setup (91; 136). A mask with 3 off-axis pin-holes is placed before the lens, which is slightly defocused, causing each particle to create 3 image exposures. Then, through knowledge of the pin-hole arrangement, the global particle position can be determined by the relative particle image positions, with their separation proportional to the out-of-plane plane position. 3D particle positions are obtained in 2 volumes of interest at a small temporal displacement, and through cross-correlation of these volumes, the velocity field can be determined.

Complex set-ups can be imaged using this approach as optical access is only required from a single direction. However, the small angles between each of the 3 ray paths for this in-line method limits out-of-plane component accuracy. Additionally, particle seeding density is limited due to the increased image plane size of out-of-focus particles. Capturing all particle images on a single image sensor further limits the seeding density by a factor of three. Although this problem can

be alleviated, at the expense of additional complexity, by using 3 separate image sensors to capture the particle images (55), the maximum permissible seeding density, and therefore measurement resolution remain low.

1.3.6   Particle Tracking Velocimetry

In similarity with PIV, in particle tracking velocimetry (PTV) the flow is seeded with small particles, which are illuminated in a volume of interest (71). In PTV however, scattered light is recorded simultaneously by three or more cameras arranged around the volume of interest. After camera calibration, lines of sight through the volume are calculated from each particle image position, and groups of these are matched, with their closest intersection defining the 3D particle position. Global particle positions are matched over either frame pairs or more commonly short frame sequences to determine velocity and particle trajectories. Therefore, the technique permits Lagrangian measurements of particle trajectories, which are particularly useful for studying particle dispersion (86). These principles have been applied using extremely high response silicon strip photodiodes to make extremely high frequency measurements (up to 70kHz (126)). Eularian measurements require interpolation of the results onto a regular grid.

Due to finite calibration accuracy, scattering intensity variations from different camera angles, and illumination inhomogeneity, particle lines of sight will not converge perfectly to a single point instead requiring some relaxation criterion (86). Therefore, to prevent matching incorrect lines of sight the seeding density is severely limited to around 0.005 particles per pixel (31), limiting spatial resolution. Although a number of different algorithms are used, matching particle pairs often uses a nearest neighbour approach; requiring closely spaced sequential frames. Additionally, validation of matched lines of sight is often made by following a particle over a number of frames (typically 3 or 4) (137). This requires the laser pulse rate and camera frame rate to match the time interval, δt, which may restrict the choice of laser to a higher repetition rate lower power design.

1.3.7 Tomographic-PIV

Similar to PTV, in TPIV the flow is seeded with small particles and a volume of fluid is illuminated. Light scattered from the illuminated particles is recorded from a number of different angles, with typically between 3 and 5 cameras used. The camera positions are defined through camera calibration (imaging a known 3D object), which allows the derivation of a mapping function for each individual camera, to tranform image coordinates into global (real world) lines of sight. It is then possible to use these mapping functions and the recorded images to reconstruct the particle positions in the volume of interest (as a volume of light intensity) through use of a tomographic reconstruction algorithm (31). These light intensity volumes can then be used directly for PIV cross-correlation; distinguishing the technique from PTV.

Without the need to individually match particles, the seeding density restrictions can be relaxed somewhat, allowing higher resolution than PTV. The technique set-up is also relatively simple. However, both Scanning and HPIV allow higher spatial resolution due to limitations still imposed on seeding density by the tomographic technique. Accuracy is limited by imperfect reconstructions caused by optical defects (134) and relies heavily on high precision calibration (30). Furthermore the reconstruction problem is computationally intensive (107), and well resolved statistical data requiring hundreds of vector fields require significant resources. The need to illuminate a volume of fluid will also require an extremely high power light source.

1.3.8 Comparison of 3D Techniques

Despite its excellent resolution, the application of HPIV is complex and unable to provide sufficient data for turbulent statistics, or make time-resolved measurements to follow the evolution of flow structures. Relatively long flow time scales make Scanning-SDPIV possible, however the quasi-instantaneous nature of this technique may introduce inaccuracies at high Reynolds numbers as the small scale turnover time becomes increasingly fast. Furthermore, the high repetition light source requirement would limit illumination power, and use of a scanning mirror introduces additional complexity. The in-line approach of both DHPIV and

DDPIV result in limited in-plane measurement accuracy, and in addition to their complexity, both techniques suffer from seeding density and therefore resolution limitations. As the current set-up allows optical access from a range of angles (see Section 2.1.1), multiple camera set-ups are favoured.

Although both PTV and TPIV would be suitable, the former suffers from limited seeding density and therefore resolution, and if particle positions are to be validated over a sequence of frames, the choice of illumination is also limited. Although TPIV also imposes some restrictions on the seeding density, these are less severe than PTV, and allow higher measurement resolution. The technique uses a relatively simple set-up, and despite requiring accurate calibration and significant data processing time, offers a good balance between spatial resolution, measurement accuracy, and temporal resolution.

Therefore, a combination of DPIV and TPIV measurements will be made in the current investigation, with technique comparison allowing TPIV validation against the more widely used DPIV, and assessment of 2D isotropy assumptions.

1.4 Fundamentals of TPIV

Due to its relatively recent introduction, there now follows a more in depth review of TPIV and its application.

1.4.1 The Image Reconstruction Problem

Image reconstruction from projections is the process of producing an image distribution of some physical property from estimates of its line integrals along a finite number of lines at known locations (41). The reconstruction problem was initially solved by Radon, who showed that the distribution of some property along an infinitely thin projection slice can be determined exactly using an infinite set of line integrals around the object. However, in practice it is only possible to use a finite number of projections, which are unlikely to be exact due to inherent experimental measurements inaccuracies. Therefore the reconstruction problem is generally under-defined, and direct solution to Radon's Inverse Formula, which is extremely sensitive to any inaccuracies or finiteness is unattainable (41). A

finite set of projections could therefore define a number of possible reconstructed objects, and a reconstruction algorithm should be applied in order to determine the object which best matches the projection data.

Therefore, solutions to more practical problems require a discretised approach, such as algebraic series expansion. This method defines a set of basis functions, which can be combined linearly to give an adequate approximation of the object (41), reducing the reconstruction problem to Equation 1.19, where p_r represents a data matrix of some measurable projected property; f represents the object matrix; and W represents a two-dimensional weighting matrix, which provides a relationship between every discretised projected data point and every discretised point in the object volume. Therefore, the discrete reconstruction problem can be stated as: given the data, p_r, estimate the object, f.

$$Wf = p_r \qquad (1.19)$$

This apparently straight-forward problem of solving a set of linear-equations is complicated by its size, as the object f is a vectorisation of a discrete approximation to a function of typically two or three spatial variables (17). Therefore, iterative methods are often the tools of choice for image reconstruction due to their simple, versatile nature, ability to handle constraints and noise, and process large data sets efficiently (83). Despite the efficiency of an iterative approach the requirement of processing huge numbers of voxels (volume elements) in the discretised volume still gives rise to substantial computational costs (see Appendix A).

A large number of image reconstruction algorithms are available, based on the algebraic series expansion of a function; these can be categorised as iterative, block-iterative, or simultaneous. Broadly speaking, sequential schemes such as the Algebraic Reconstruction Technique (ART) and multiplicative version (MART) tend to converge more quickly than simultaneous schemes, and although block-iterative methods may provide small improvements they tend to be more unstable (18). The sequential order, under-relaxation parameters, and rescaling (dividing every row in the W matrix by its maximum value - RMART) must be carefully applied to ensure rapid convergence. Although testing of different reconstruction algorithms for TPIV may prove fruitful, simple, fast sequential

schemes appear to be well suited for this application. Moreover the relatively low number of iterations required to achieve adequate convergence (typically 5 to 6) might render any improvements as marginal.

1.4.2 Practical Application of TPIV

The methodology of Elsinga et al. (30) is followed to illustrate the practical application of TPIV. The object is modelled using simple 3D cubic voxel (volumetric pixel) elements as basis functions. These take a positive value inside the voxel, but are zero outside it, with these simple elements arrayed to form a volume. Hence, each weighting array element is defined simply as the intersecting volume between a sphere representing a voxel, and a cylinder representing the projection line of sight, taking a value between zero and one. After testing both additive (ART) and multiplicative (MART) reconstruction algorithms, the latter was shown to offer advantage through more distinct particle identification (Equation 1.20).
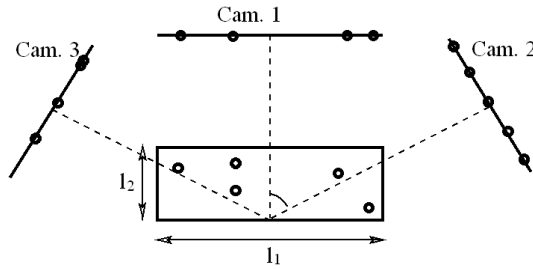
Figure 1.9: Schematic of numerical TPIV Setup

$$f_j^{k_{it}+1} = f_j^{k_{it}} \left(\frac{p_i}{W_i f^{k_{it}}} \right)^{\mu W_{ij}} \qquad (1.20)$$

In a series of 2D parametric studies it is suggested that reconstruction accuracy can be maximised by adding cameras to the system, and by using an optimum angle of 30°. To evaluate reconstruction quality a correlation coefficient between artificial and reconstructed object arrays was used, with a cut-off value of 0.75 above which reconstruction is considered successful. High seeding densities and camera angles where $\theta >> 30°$ or $\theta << 30°$ were shown to decrease reconstruction accuracy through an increase in *Ghost Particles*, which arise where

intensity from all lines of sight create a particle that does not exist. Seeding density is defined in terms of the number of particles per pixel (N_{ppp}), relatable to volumetric seeding density through camera resolution and domain volume. The importance of sub-pixel calibration accuracy was also demonstrated, with errors over 0.4 pixels exceeding the assigned accuracy cut-off limit. However, although *Ghost Particles* are noted to decrease accuracy, in cross-correlation these essentially random events may not significantly affect the vector field results. Therefore there is scope for further quantification of set-up parameters with respect to cross-correlation, rather than correlation coefficient based accuracy analysis (discussed further in Section 2.3.3).

Degradation due to image noise was shown to be controllable through image pre-processing. Subtracting the minimum pixel intensity from all pixels helps treat Hot Pixel, Window Dirt, and laser reflections (134). Removing the image background intensity reduces noise from erroneous *Ghost Particles* (29). Intensity normalisation takes account of forward and backscattering intensity differences in addition to laser profile variations and areas of weak seeding, and the 3×3 Gaussian smoothing, reduces noise and improves the small particle (≤ 2 pixels) reconstruction.

Additionally, a number of other important considerations were highlighted including: small camera apertures to keep the whole volume in focus; use of a reconstruction volume thicker than the illuminated volume in order to correctly place all recorded light; and removal of reconstruction edges, where placement of a finite volume of interest results in light from the entire thickness being erroneously concentrated in thin regions.

1.4.3 Previous TPIV Experiments

Due to its recent introduction there have been relatively few experiments performed using TPIV (28; 29; 61; 75; 107; 133; 134), which are almost exclusively based on the principles derived by Elsinga et al. (30). A number of set-up parameters for full-volume TPIV have been compared in Table B.1 in Appendix B.

Elsinga et al. (29) studied the turbulent wake velocity distribution behind two different circular cylinders in a 10ms^{-1} free stream velocity flow, demonstrating

the potential of TPIV. First a thin laser sheet was used to compare the technique to the more well established Stereo-PIV, after which a thicker light sheet was used to analyse the flow pattern. Using the thin-sheet (2.3mm) it was found that the mean particle displacement differences between Stereo-PIV and TPIV is 0.03 pixels with a standard deviation of 0.2 pixels and a maximum difference of approximately 1 pixel. Although this indicates the potential of TPIV, the accuracy of the thin-sheet may not reflect the full 3D accuracy, and requires further quantification.

(a) Stereo PIV   (b) TPIV
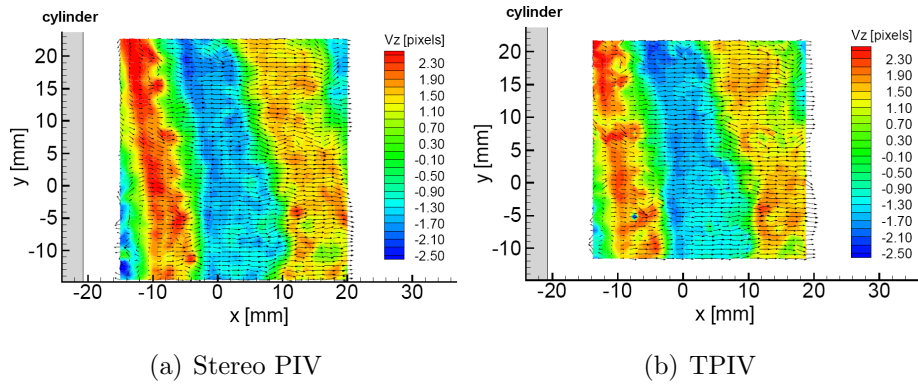
Figure 1.10: Comparison of Stereo and TPIV for a thin light sheet (29)

Through a series of water jet experiments Wieneke et al. (133) also demonstrate a favourable comparison between thin-sheet TPIV and Stereo-PIV. Additionally light from a single source was split into 4 thin sheets, and although this increased accuracy in comparison with full volume TPIV and 3D-PTV, it was concluded that this is still less accurate than Stereo or Dual-Plane PIV with a velocity error from $0.1 \rightarrow 0.2$ pixels.

Comparing Stereo-PIV and thin-sheet TPIV numerically, Wieneke & Taylor (134) found that the latter was more accurate, however, with real data Stereo-PIV proved to be more effective, especially in low gradient regions. Given very accurate calibration (< 0.1 pixels), the RMS displacement difference in the $[x,y,z]$ planes between the techniques was $[0.22,0.22,0.30]$ pixels (16×16 interrogation window, over 20 images) corresponding to 3% of the maximum velocity, with error based on local velocity variation estimated as 0.13 and 0.20 pixels for Stereo-PIV and TPIV respectively.

40

(a) Mean Velocity Profile

(b) RMS Velocity Components
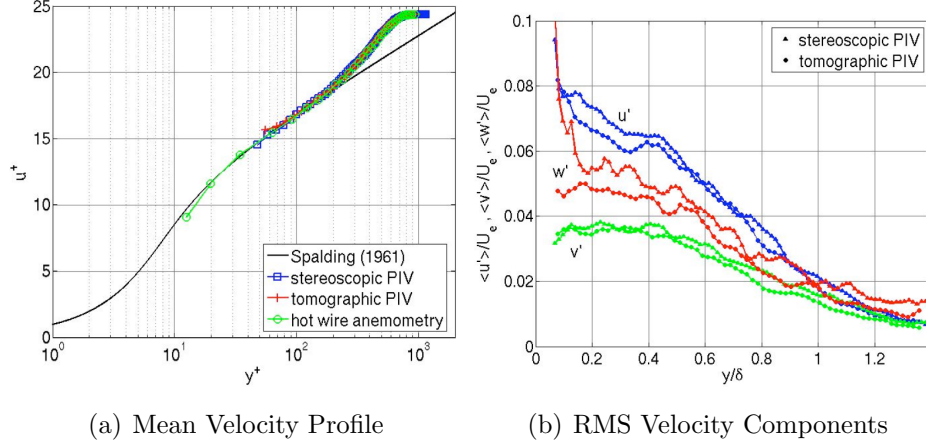
Figure 1.11: Comparison of Hot Wire, Stereo-PIV and TPIV (Elsinga et al. (28))

A Dual-Plane Stereo-PIV comparison was also conducted using a slightly thicker 'fat-sheet'. Dual-plane PIV accuracy is limited by light sheet alignment, instantaneous particle positioning, and large gradients due to non-overlapping light interrogation windows, in addition to standard errors from Stereo-PIV (with gradients containing at least a 10% error), producing typical in-plane and out-of-plane errors of at least $5 \rightarrow 10\%$ and $10 \rightarrow 20\%$ respectively. TPIV can avoid particle positioning problems through deforming interrogation volumes, and given accurate calibration compares favourably to Dual-Plane Stereo-PIV, with out-of-plane gradient accuracy between $5 \rightarrow 10\%$ and $10 \rightarrow 20\%$ reported for numerical and real data respectively.

A number of experiments demonstrating the technique's ability to identify turbulent flow structures have been conducted. Michaelis et al. (75) studied the wake flow behind a cylinder in a water flow channel, obtaining time-resolved measurements of the small scale 3D flow features. Although visualisation of both spanwise vortices, and secondary connected counter-rotating structures is enabled, the accuracy of these measurements is not indicated.

Schröder et al. (107) analysed the sub-structures within an artificially tripped turbulent spot in a laminar boundary layer, and flow structures in a tripped turbulent boundary layer. The measurements yielded full 3D spatial and temporal resolution of the spot, including elucidation of rapidly forming hairpin vortex structures. Cross-correlation peak errors of order ~ 0.2 voxels related to a 4%

uncertainty in the vector results.

Elsinga et al. ([28](#)) used TPIV to study the coherent structures in a turbulent boundary layer at $Re_\theta = 1900$, reporting a dynamic range (ratio of the highest to lowest velocities which can be resolved) of 260, which yielded a measurement accuracy 0.1 voxel. The mean velocity profile and the three RMS components of velocity through the boundary layer compare well with both Stereo-PIV and hot wire results (shown in Figure 1.11).

1.5 Summary and Aims

An introduction to various turbulence analysis techniques has been given, followed by an overview of the small scale turbulent structures and the investigations which uncovered these. Additionally a range of experimental techniques have been reviewed with emphasis on TPIV.

Although the scaling of the highest vorticity magnitude/lowest pressure structures has been loosely defined, there are some discrepancies between findings, with the use of different identification schemes making comparisons difficult. Glimpses of a hierarchical clustering system have been shown, although evidence for this is almost exclusively qualitative. Additionally, the majority of results are severely limited in terms of Reynolds number, and in both numerical and experimental approaches typically only the strongest structures are considered, which may not contribute significantly to the distribution of dissipation. Therefore, it is clear from this review that a comprehensive unified understanding of the small scales is proving elusive.

Previous investigations have raised a number of important questions, many of which remain under debate: How is dissipation distributed with regard to the coherent structures, and how does this changes with Reynolds number? Does a vortex instability transition region exist, and if so how does this relates to the scaling of the coherent structures? Why are vortex tubes more prevalent than sheets despite the predominance of bi-axial strain? Furthermore, Richardson might be perplexed to find that although there is wide acceptance of his original cascade concept, the mechanisms through which it propagates are still under discussion,

and according to some authors even the fundamental dissipation scaling has not been satisfactorily proven for high Reynolds number flow (115; 144).

Therefore, there is a need to establish these processes over a broader range of scales, to give a more comprehensive picture of the kinematic life cycles of these filaments and other significant small scale structures (35). The distribution of enstrophy and dissipation should also be examined at higher Reynolds numbers, where increased intermittency may permit more conclusive dissipation distribution visualisation. There is also a need to further understand the time-dependent evolution of structures, to bolster knowledge of the cascade process.

With regard to the experimental technique, although the potential of TPIV has been demonstrated a number of aspects are not yet well understood. Currently there is little evidence of the technique's full 3D accuracy, with studies instead focussing on relatively thin-light sheets rather than a full volume. Furthermore, previous parametric studies were limited to 2D reconstructions (30), and used a correlation coefficient to assess the reconstruction; a measure which does not fully represent velocity field calculation accuracy. Hence, it would be beneficial to extend these accuracy studies into 3D using a more applicable measure of accuracy.

Therefore, the aims of the current investigation are:

- To develop fast TPIV software in order to gain better technique understanding.

- To conduct a series of parametric TPIV tests and accuracy simulations.

- To make DPIV and TPIV measurements over a range of Reynolds numbers in the large mixing tank.

- To statistically characterise the flow, and compare DPIV and TPIV results.

- To characterise small scale turbulent structures in order to increase our fundamental understanding of the complete life cycle of these events, their significance, and their inertial scale clustering.

Repetition of previous TPIV parametric tests will permit a validation of the current code, with the extension of these tests into 3D and the use of a more applicable accuracy assessment addressing the current lack of literature regarding the accuracy of this method. The deduction of experimental measurement uncertainty through a series of simulations will also permit a more complete analysis of the results.

Due to the bias in terms of their extensive numerical study, further experimental study of coherent dissipation-scale structures will be extremely useful. These measurements will first be used to quantify measurement accuracy and resolution and statistically characterise the flow, with comparison of 2D and 3D results providing support for the less accurate TPIV approach. Next the 3D data will be used to investigate the dissipation scale coherent flow structures using a variety of methods.

It is hoped that this will provide support for the myriad of DNS and other experimental studies, contributing towards a better understanding of the fundamental physical flow interactions at these scales, and thus, support the derivation of physics-based statistical models.

Chapter 2

Experimental Set-up and Code Development and Optimisation

Having introduced the choice of experimental techniques in Chapter 1, implementation details will now be covered. The chapter begins by introducing the experimental procedure and hardware set-up, before giving an overview of the software development, optimisation, and validation through numerical and experimental tests, and post-processing details.

2.1 Experimental Set-up

This section begins by outlining the experimental facility, detailing the procedure and discussing limitations imposed by the current geometry. Following this individual set-up elements are examined, justifying the reasoning behind set-up choices. The choice of seeding particles, timing set-up, and camera configuration and calibration are discussed, in addition to water temperature and purity control.

2.1.1 Large Mixing Tank Facility

A large number of experimental studies investigating the small scale turbulent structures have measured the flow between between two counter-rotating disks or impellers (1; 6; 10; 20; 21; 26; 33; 63; 127; 144).

Rotating the disks in opposite directions at equal velocities causes the fluid nearest the impellers to spin in opposite directions, establishing a strong shear plane in the centre of the tank, which creates the majority of the turbulence. Two secondary toroidal recirculation regions are formed, stretching the flow in the central region. These component motions are shown in Figure 2.1. Due to the axial symmetry of the tank, the shear should be zero at the axial and vertical centre of the tank, where the measurements are made. Douady (26) observed that the size of the largest eddies are of the same order as the impeller radius, R_I, and the characteristic velocity, U, is that of the mean rotation velocity. The impellers are typically but not always rimmed and set at a separation distance, $\sim R_I$, with radial baffles used on the tank walls to inhibit global rotation, and further increase turbulence production. While anisotropy may persist at the large scale, the flow is known to approximate homogeneous isotropic turbulence at least locally.

(a) Central Shear   (b) Centrifugal Pumping
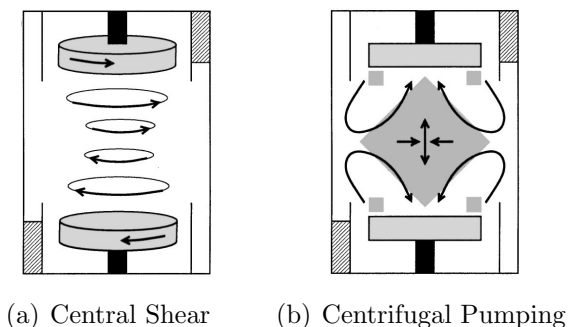
Figure 2.1: Streamline Representation of the flow between counter-rotating impellers (La Porta et al. (63))

This flow set-up is useful for generating high Reynolds numbers with constant energy injection at the largest scales, with the closed geometry permitting the spatial evolution of flow structures to be studied, without these being transported away too quickly by the mean flow. In fact the flow is similar to the Taylor-Green vortex, although instead of having periodic no-slip boundaries the current flow is contained by the tank geometry, therefore permitting direct comparison with DNS results (20; 26). A statistical investigation using this geometry by Zocchi et

al. (144) suggests the non-dimensionalised nature of the flow is comparable with other experiments based on energy spectrum results.

The flow can be characterised through the geometry of the tank using Equation 2.1, where Ω_I is the impeller rotation rate.

$$Re = \frac{UR_I}{\nu} = \frac{\Omega_I R_I^2}{\nu} \qquad (2.1)$$

Due to turbulent energy dissipation into heat, previous experiments have had to consider the effects of temperature and hence viscosity drift over time during long experimental runs, which required the use of temperature regulating pumps. This control is only required for long run times and, as will be shown later (see Section 2.1.11), the current experiments did not require such measures.

In order to increase the measurability of the dissipation scale structures the current investigation made use of a more simplistic strategy than those previous discussed (see Section 1.2.2); significantly increasing the impeller radius, R_I, which is the characteristic flow length, which increases the size and scale of these structures. This also alleviates the problem of using relatively large particles to study extremely small scales (127).

To this end, a very large 2m diameter dodecahedral perspex mixing tank with two eight vane 0.8m diameter impellers, set at a separation distance of 1.25m was used in the current investigation (shown in Figure 2.2). Water was used as the working fluid and radial baffles were fitted to remove any net rotation. Two micro-stepping motors controlled impeller rotation, permitting accurate speed matching. In comparison with other experimental set-ups the predicted length and time scales in this set-up are one and two orders of magnitude larger respectively, permitting study of dissipation scale structures using the chosen measurement technique.

2.1.2 Overview of Experimental Procedure

The mixing tank (Section 2.1.1) was seeded with small neutrally buoyant particles (Section 2.1.4), and two stepper motor controlled impellers were set to counter-rotate at a range of speeds to create a high Reynolds number mixing flow. Double pulsed lasers (Section 2.1.6) were used to illuminate a thin volume (Section 2.1.7)

(a) Mixing Tank Schematic   (b) Mixing Tank during TPIV Experiment
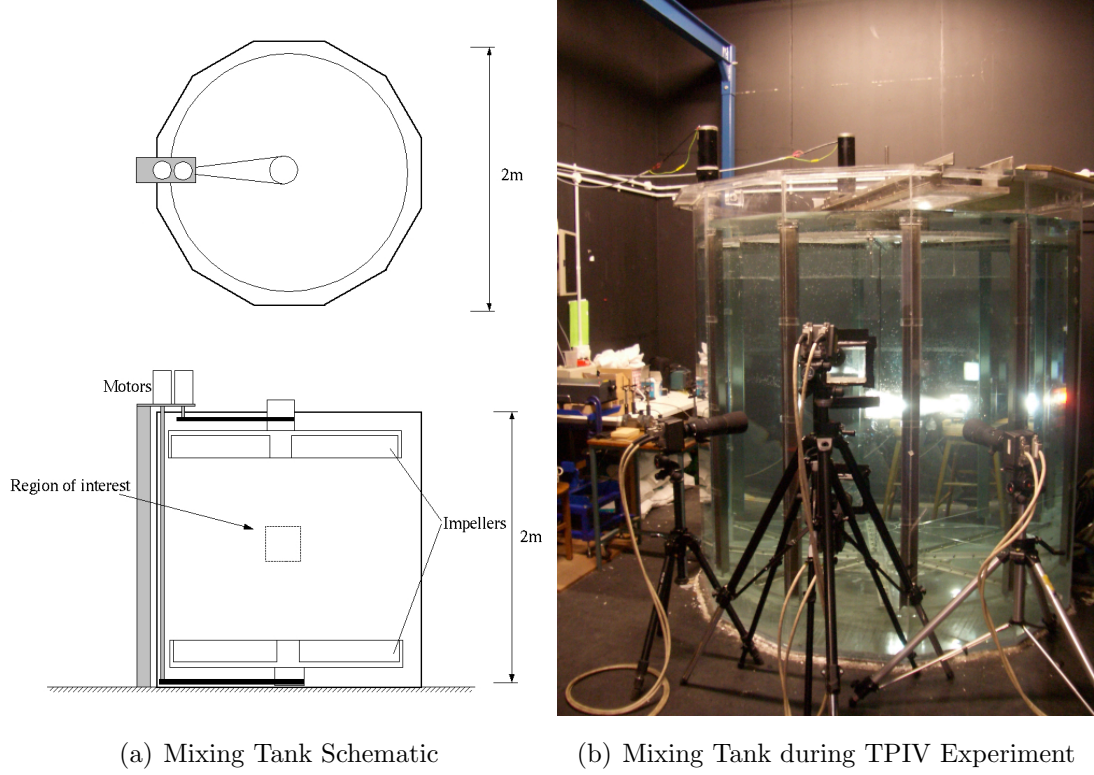
Figure 2.2: Large Mixing Tank Facility

at the centre of the tank, using signal generators (Section 2.1.8) to control laser pulse separation and camera synchronisation.

For the DPIV measurements, a single camera perpendicular to the light sheet (Section 2.1.9.1) was calibrated (Section 2.1.10) and used to record the light scattered from the seeding particles. Images were recorded at a small temporal displacement (Section 2.1.8) allowing these to be cross-correlated (Section 2.2.6) to produce a 2D velocity vector field.

For the TPIV measurements a number of cameras were positioned around the volume of interest (Section 2.1.9.2), calibrated using a 3D calibration object (Section 2.1.10), and used to record the light scattering from particles at two instances separated by a small temporal displacement (Section 2.1.8). The images were pre-processed (Section 2.2.3) before being used to tomographically reconstruct the light intensity within the discretised volume of interest (Section 2.2.4). These light intensity volumes were directly cross-correlated to obtain 3D velocity

vector fields (Section 2.2.6).

Data was be collected at two sampling speeds: the first over ~ 80 impeller rotations in order to achieve statistical convergence (Section 2.1.8); and the second at the maximum frequency available (Section 2.1.6), in order to achieve quasi-time resolved results for analysis of coherent structure evolution. A total of four DPIV experiments were conducted at different resolutions, which from lowest to highest resolution are denoted as cases $D1$, $D2$, $D3$ and $D4$. A single TPIV case was conducted, denoted as case $T1$.

2.1.3 PIV Limitations in the Current Geometry

In order to study coherent dissipation scale structures it is desirable to study a large volume of interest at high resolution. However, the experimental technique and facility both impose limitations on these measurements, which are now discussed.

In all PIV the resolution and domain size are linked (through CMOS resolution and optical set-up magnification), and therefore it was decided to make measurements at a range of resolutions; allowing a broader range of structures to be measured and resolution effects to be quantified. The four different DPIV resolutions are introduced in Section 2.1.9.

The application of TPIV is subject to a number of additional restrictions stemming from hardware limitations, which limit the maximum domain size and resolution through SNR and image density requirements. The scattered light intensity is a function of particle seeding density, particle size, laser power, laser sheet size, camera sensitivity, and camera aperture setting, while the image density is a function of the seeding density and laser sheet thickness.

TPIV is particularly sensitive to these limitations for the following reasons: 1. A thick light sheet is used to illuminate a volume of fluid, which spreads the beam energy over a greater area; 2. To maintain a sufficient depth of field to keep the entire volume in focus, particularly from cameras set at an angle to the volume, smaller aperture settings must be used; 3. At the high seeding densities required the effect of multiple scattering increases light diffusion (140). These all serve to reduce the SNR and impose strict limits on domain size and resolution.

The maximum measurement frequency is also limited by the SNR through the choice of laser system; making use of lower powered but higher frequency lasers difficult.

During preliminary testing the current set-up geometry and PIV hardware were found to impose strict limits on the attainable volume and resolution of measurements. Increasing the seeding beyond 10g/6000 litres (based on the seeding particles described in Section 2.1.4) resulted in a rapid drop in SNR due to multiple scattering generated diffusion, limiting the maximum attainable resolution. At the other end, increasing the illuminated volume dimensions over $60 \times 60 \times 10$ mm, resulted in an unworkable reduction in the SNR. Therefore, although DPIV measurements were made at a range of resolutions, the more demanding TPIV requirements only permitted accurate measurements at a single resolution.

A range of Reynolds numbers were measured for each case, with only 5 points chosen to minimise the considerable processing costs (see Section 2.2.4). The Reynolds number range was chosen to maintain good resolution at the low end while covering a sufficiently large range. Various set-up details for the different cases can be found later in Tables 2.2 and C.2.

2.1.4 Seeding Particles

PIV is an indirect measurement technique, in that flow velocity is deduced from the movement of small suspended particles. Therefore particle choice is extremely important, with conflicting requirements demanding particles large enough to scatter a reasonable quantity of light, and small enough to faithfully follow the smallest flow structures (73). The image plane particle diameter affects both the PIV accuracy and the image density, with a diameter of 2-3 pixels recommended (100). Due to diffraction limited imaging, this size is a function of image magnification, particle size, aperture setting, and wavelength of the illuminating light. Furthermore, in the current set-up particle size is also limited by optical attenuation through the volume, due to strong multiple scattering effects (140).

After testing a number of potential seeding particles, Dantec Dynamics silver coated hollow glass spheres offered the best compromise. These are spherical hol-

Case	$D1$	$D2$	$D3$	$D4$	$T1$
Seeding density (N_{ppp})	0.015	0.010	0.010	0.010	0.020

Table 2.1: PIV seeding densities (approximate)

low Borosilicate Glass spheres with a density of 1.4g/cm³ and a diameter ranging from $2 \to 20\mu$m, with a mean size of $d_p = 10\mu$m. The highly reflective metallic coating gives a more intense and even scattering profile across a broad range of scattering angles, known to increase accuracy through intensity fluctuation removal (106). This may be particularly advantageous in TPIV, increasing image uniformity between cameras at different scattering angles. However, due to increased scattering and absorption isotropy these spheres are likely to increase optical attenuation.

Particle diameter is an order of magnitude less than the smallest flow structures, $\eta/d_p \approx 10$ (see Table 4.1), and due to their similar density the particles are therefore expected to follow the flow faithfully. Furthermore, these mono-dispersive particles should limit the variation in scattering intensity between particles, and ensure similar response to fluid acceleration.

Seeding density estimates for each case were measured after each experiment using a particle tracking algorithm to identify particle centres, averaged over ten randomly selected images, and expressed in Table 2.1 in particles per pixel.

2.1.5 Water Purity Control

To ensure the water contains no large particulate matter other than the intended seeding, and to prevent lime-scale deposits, a series of water control measures are used. Supply water is passed through a water softener, before being filtered twice using 100μm and 1μm mesh diameter filters. The tank is kept in darkness, and a weak Copper Sulphate II solution is used (1g/1000 litres) as a herbicide to prevent organic growth.

2.1.6 Laser Systems

The seeding requirements require the use of very small particles (Section 2.1.4). Therefore, to produce images with a reasonable SNR and meet stringent pulse

duration and pulse separation restrictions (95) a high powered double pulse laser light source is required (2).

A *New Wave Gemini PIV 120-15* laser system was used for the majority of the results. This is a water-cooled, dual-head, flashlamp pumped, Nd:YAG system. At the maximum repetition rate of 15Hz, the energy per pulse is 120mJ at a wavelength of 532nm, with a pulse width of 5ns. Due to the relatively slow nature of the flow (See Table. 4.1) it was possible to achieve quasi-time resolved results at this repetition rate for the majority of cases. That is, it is possible to follow the evolution, and interaction of flow structures as they pass through the volume of interest. The sacrifice of full time resolution at all Reynolds numbers is necessary due to the lighting requirements, and extension of these results may have to await further development in high-power, high-repetition rate laser systems (See Section 6.5).

The less demanding DPIV power requirements allowed a *New Wave Pegasus PIV* laser system to be used for the *D3* case. This is a water-cooled, dual-head, high repetition rate, diode-pumped Nd:YLF system. At 2kHz, the energy per pulse is 10mJ at a wavelength of 527nm, with a pulse width of 135ns.

Despite the advantages offered by Nd:YAG lasers in terms of high energy and short pulse duration, Raffel et al. (100) note that due to thermal lensing issues the beam quality tends to be low in comparison with continuous wave lasers. High power beams tend to have very poor beam profiles with hot spots and different ring modes, producing uneven illumination between pulses and within the same light sheet, and increasing measurement uncertainty.

2.1.7 Sheet Expansion

Both DPIV and TPIV experiments require laser beam expansion from a point source into sheets of varying height; a thin sheet for the DPIV, and a thick sheet or 'slab' of light for TPIV.

In the DPIV experiment the beam was expanded and then re-collimated by passing it through two cylindrical lenses (see Figure 2.3). The re-collimation is imperfect by design and the resulting sheet is weakly expanding, with relative

lens positioning allowing a variety of sheet sizes. The measurement point sheet thickness is 1.5mm.

To produce the thick sheet the laser beam was first passed through a Galilean beam expander (consisting of a plano-convex and an achromic lens series, as shown in Figure 2.3) to create the desired beam width, before being expanded into a sheet as before. This produces a ovoid shaped beam cross-section which could then be clipped to produce the required rectangular cross-section, with a thickness of 10mm.
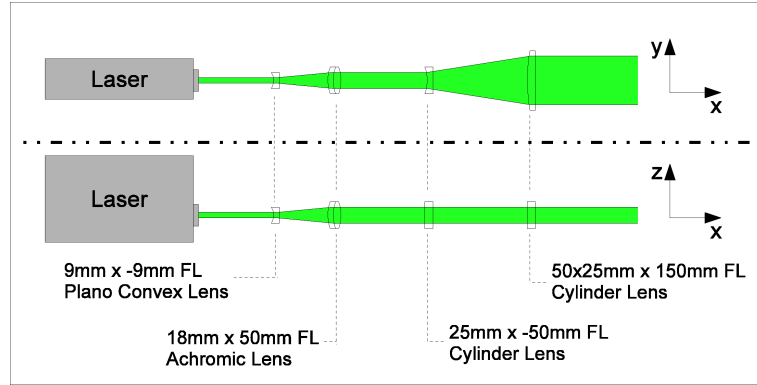
Figure 2.3: Beam Expansion

2.1.8   Timing Set-up

DPIV and TPIV timing set-ups differ only in the control of multiple cameras. Three Thurlby TGP110 pulse generators were used to create the correct input signals for the cameras and laser, with signal timing measured using a Tektronix TDS 2002 digital oscilloscope, and cameras controlled through Photron FAST-CAM Viewer V2.4.5.2 software. Exact set-up details, and configuration diagrams can be found in Appendix C.

The time increment, δt, must be carefully chosen to allow the particles to displace by a distance which is both large enough to give a reasonable dynamic range, and small enough to allow accurate displacement evaluation and minimise out-of-plane particle loss. This time increment is adjusted for each Reynolds number to maintain a reasonable mean image particle displacement ~ 5 pixels/voxels.

This rule of thumb was derived from a comparison of different time increments, and is similar to previous investigations (82)

The shear plane location is likely to be unsteady with a frequency proportional to the impeller rotation rate, and therefore the measurement frequency should be much longer or shorter than this. For the statistical study 1000 fields were collected over 80 impeller rotations to ensure good statistical convergence (see Section 4.1.2), with these timings and camera software set-up details again included in Appendix C. 500 fields were also captured at the maximum laser repetition rate of 15Hz.

2.1.9 Camera Configurations

In PIV the appropriate lens, distance between the object plane and lens, and aperture setting must be carefully chosen to provide the desired field of view (FOV) and depth of field (DOF), while allowing enough light to give a reasonable SNR. Narrow apertures produce darker images with larger DOF, and wider apertures produce brighter images with a smaller DOF. Additional choices must be made to accommodate multiple cameras in the TPIV set-up, with reconstruction accuracy consideration. Images were written to disk as loss-less 16 bit TIFF files.

2.1.9.1 DPIV Camera Configuration

A single 1024×1024 pixel CMOS Photron APX was positioned perpendicular to the light sheet (through alignment with the relevant tank face) using a Manfrotto 055XPROB tripod and Manfrotto 410 geared head, at a distance of 1100mm from the measurement point. Details of the lenses, aperture settings, FOV, and resolution can be found in Table 2.2 for the different resolution cases. The camera was focused manually, initially using the calibration plate, and then using the illuminated particles themselves. Correct particle focus is extremely important, and therefore this extra focusing stage is crucial to ensure high quality images. After the particle focus, the calibration plate is repositioned, and imaged.

Case	lens(es)	$f/\#$	FOV (mm^2)	resolution (pix/mm)
$D1$	60mm Nikon	2.8	212×212	4.8
$D2$	105mm Sigma	2.8	117×117	8.8
$D3$	180mm Sigma	3.5	77.4×77.4	13.2
$D4$	180mm Sigma & 2$\times$ Komura teleconverter	3.5	32.0×32.0	32.0

Table 2.2: PIV camera configuration parameters

2.1.9.2 TPIV Camera Configuration

In TPIV the volume of interest must be imaged from different angles, with the choice of these directly related to reconstruction accuracy and illumination level (through depth of field requirements and corresponding aperture restrictions, as discussed in Section 2.1.3). Previous numerical investigations recommend an angle of $40-50°$ (see Section 2.3) based solely on reconstruction accuracy. However, illumination difficulties required the following compromise to be reached.

Four 1024×1024 pixel CMOS Photron APX cameras were positioned using Manfrotto 055XPROB tripods and Manfrotto 410 geared heads. Three cameras were fitted with Sigma 180mm lenses and the fourth with a 200mm Nikon lens, using aperture settings of $f = 1/8$. These were angled in the yaw and pitch planes by $\theta = [-30, 0, 0, 30]°$ and $\phi = [0, 15, -15, 0]°$ respectively, contained therefore within a solid angle of $[60, 30]°$. The cameras are positioned 1100mm from the measurement volume, which gives a FOV of 70×70mm^2, with a resolution of ≈ 15pixels/mm.

As the cameras are now angled with respect to the object plane, the depth of field required to keep the region of interest in focus will increase with the tangent of the angle. Therefore, in order to maintain a low aperture settings and keep the light intensity high the lens plane has been angled with respect to the image plane; a practice described by the Scheimpflug condition. The technique results in a wedge shaped depth of field, bringing an angled object plane into focus, and allowing reasonably low aperture settings to be used. At these settings the lenses give a depth of field of ≈ 20mm, which is large enough to keep the entire volume of interest in-focus.

Water filled perspex prisms are positioned between the central cameras and the tank, to correct for the refractive index effects through the air/perspex/water interface. Prasad & Jensen (96) showed through use of liquid prisms the optical radial distortion can be significantly reduced.

2.1.10   Camera Calibration Hardware

In DPIV, to convert from a displacement in pixels measured on the image plane to an object plane displacement in SI units it is necessary to calibrate the camera. By carefully aligning the light sheet parallel to the image plane and assuming negligible image distortion, calibration can be performed by imaging a target of known dimensions. A calibration constant can be defined by measuring the distance between points (spanning the majority of the field of view) on the image plane and relating these to the actual distance between the two reference points.

TPIV requires a more involved calibration function to relate each image plane pixel position to a line of sight through the volume of interest for each camera. To calculate the necessary 3D calibration mapping function, an image or images of a calibration target with known 3D geometry should be acquired (112). A sufficient number of points for least squares regression should spread over the entire field of view, with spacing inversely proportional to image distortion (112).

Using a translation stage to move a single plane calibration object in the mixing tank would be difficult, due to the access point distance from the region of interest, adding unnecessary complexity. It was therefore decided to manufacture a single piece 3D calibration object, with a number of different designs created and tested before a printed design was chosen (see Figure 2.4). This method allows extremely small calibration dots to be produced, which are less affected by uneven illumination.

Two regular arrays of 23×23 7 pixel diameter black circles with a point spacing of 5.67mm were created using the GNU Image Manipulation Program 2.4.2, and printed at 600dpi onto acetate sheets. These were cut to size, and affixed to either side of a $75 \times 75\text{mm}^2$ piece of clear 10mm perspex using Tensol-70 cement, pressed flat to create a clear interface. The two arrays are misaligned to facilitate point identification. Due to imperfect pressing the z-direction variation

<table>
(a) Drilled perspex   (b) Anodised Aluminium   (c) Printed
</table>
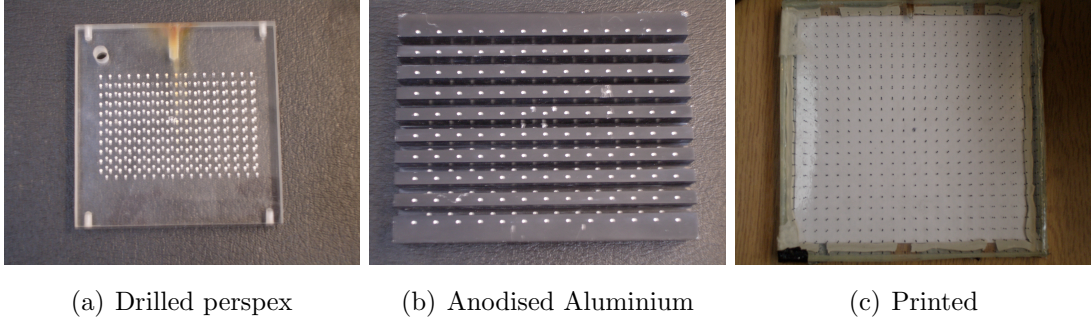
Figure 2.4: Different calibration plate designs

was measured as ± 0.05mm, and the accuracy in the x and y directions can be related to the printing resolution of 0.04mm. To take account of any misalignment between the planes, the misalignment is measured at the four corners of the plate using microscope and traverse table (accurate to within 0.01mm), and the global locations are adjusted according to a linear interpolation of these measurements. The printed calibration plate was attached to an adjustable stainless steel frame, allowing it to be precisely positioned at the tank centre, through careful use of adjustment screws. Correct positioning was confirmed by firing the laser on a low power setting and adjusting the frame as required. Following the sheet path through the perspex calibration plate allowed accurate parallel alignment. A translucent white plastic backing was affixed to the calibration plate, which could then be backlit using two 250W projector beams. The calibration plate images were used to derive image/global scaling parameter for DPIV, and camera calibration functions for the TPIV experiments (see Section 2.2.1).

2.1.11 Water Temperature

The water temperature is measured at the start of each experiment with a mercury thermometer, to an accuracy of $\pm 0.5°C$. Other similar experimental designs are required to regulate the temperature, as the dissipation of turbulent kinetic into heat raises the fluid temperature significantly (127). However, the large volume of water used in the current set-up results in a negligible temperature change. The accuracy of the temperature reading ($\pm 0.5°C$) will give rise to around a 1% uncertainty in the kinematic viscosity. Therefore, the dissipation estimates will

include an additional 1% uncertainty, although the estimates of other quantities such as the Kolmogorov length and time scales will only vary by 0.3% and 0.6% respectively. In comparison with other error sources this error is considered negligible.

2.2 Data Processing

In order to gain a better understanding of TPIV it was decided to develop a complete software suite, to process raw data in the form of calibration and particle field images, and produce velocity vector fields. Additionally, development of this software has allowed an independent study of the factors affecting reconstruction accuracy to be made (see Section 2.3).

The flow diagram shown in Figure 2.5 gives an overview of the data processing procedure. Each process will now be discussed in the following sections, including software design details and choice of key parameters.

2.2.1 Camera Calibration

Popular choices for the camera calibration function are the pin-hole (135) and polynomial model (112); the former creates a volumetric mapping by fitting the object and image plane points to a simple pin-hole camera model with added lens-distortion parameters; whereas the latter fits the data to a third-order polynomial in x and y at multiple z locations. Both approaches typically use least-squares regression to fit the data to the selected camera model.

In addition to defining the image-object mapping, the polynomial model has the ability characterise and correct for optical distortion caused by optical misalignment, lens aberrations, and/or refraction from media interfaces (112). Therefore, for these reasons and its implementation simplicity, a 3rd order polynomial calibration function was used in the current investigation (similar to Mann et al. (86)). Sollof et al. (112) recommend a 3rd order model, and despite testing higher order functions up to 5th order in the present investigation, negligible accuracy benefits were obtained.
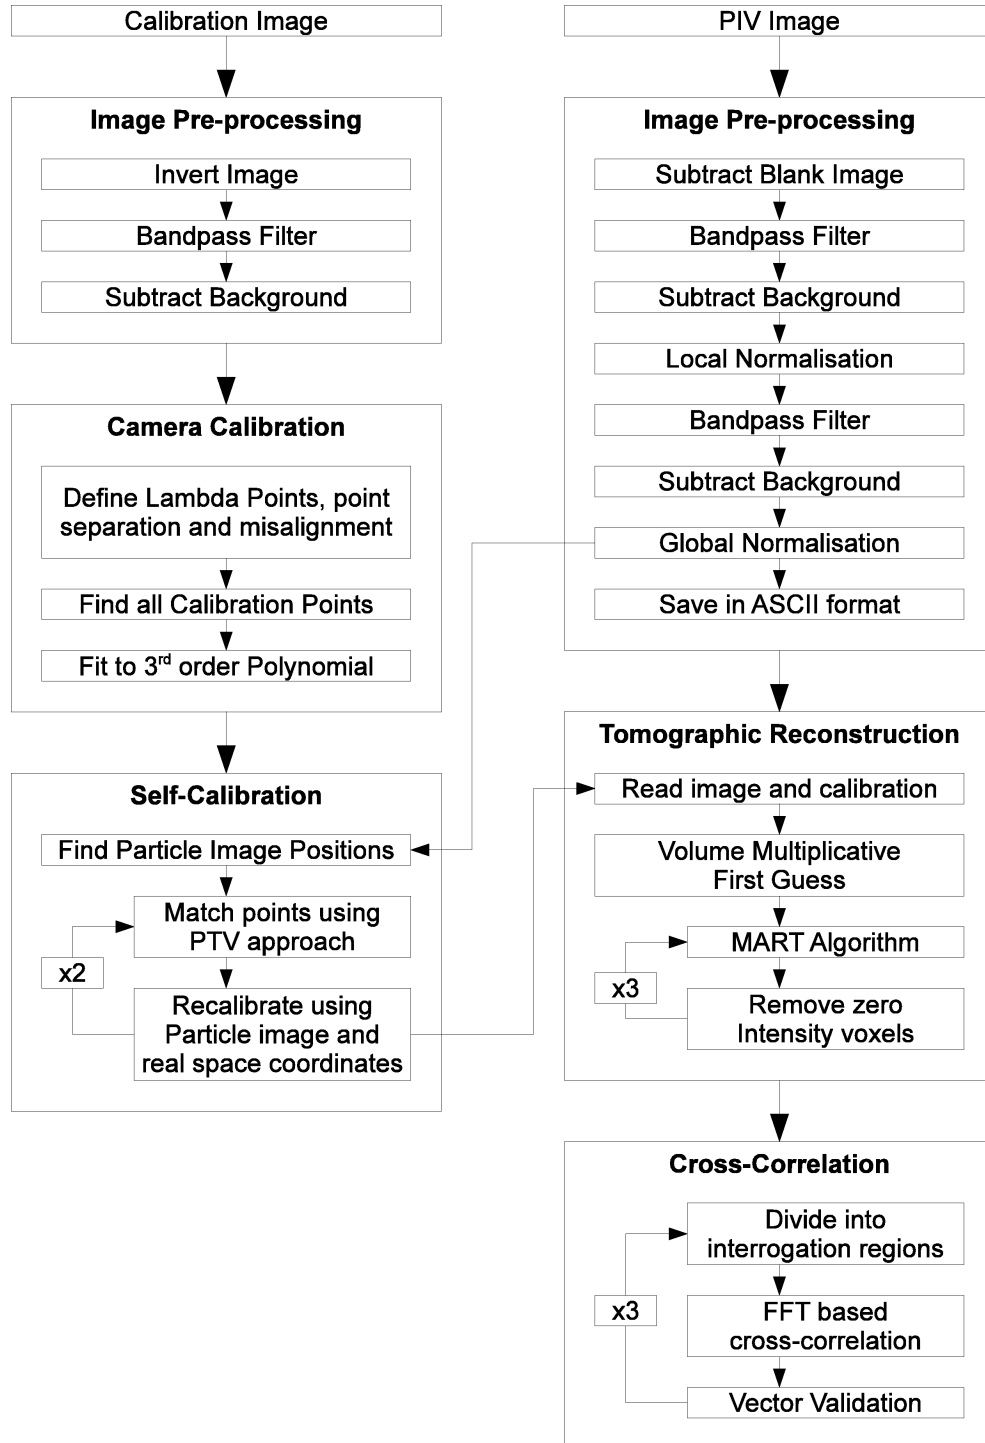
Figure 2.5: TPIV Flowchart

The calibration routine was written in MATLAB and requires the user to estimate the image plane position of three central points for each plane (referred to as the 'Lambda' points) and define the point spacing, the separation between planes, and the misalignment between planes at four locations. The calibration images are inverted, bandpass filtered, and the background noise is subtracted, after which the Lambda point estimates are refined using a centroid approach. The routine finds the remaining calibration points by estimating their position based on extrapolation from the Lambda point positions and spacing, using a robust least squares approach, which is refined to include new points as they are identified and validated. Examples of these stages are shown in Figure 2.6. The real space co-ordinates are also calculated, taking account of the misalignment between planes, and making a refractive index correction for the back plane points.

(a) Original Image   (b) Lambda Points   (c) All Points
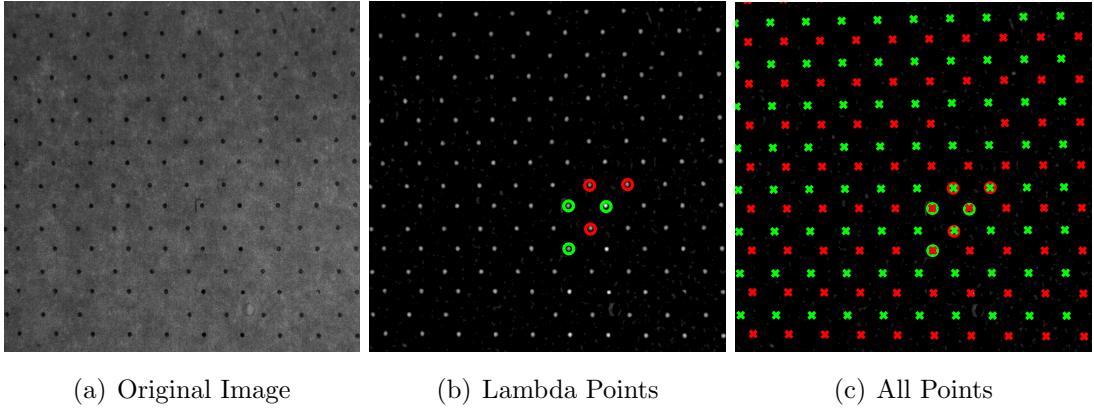
Figure 2.6: Example Calibration Point Detection

Calibration functions are defined by a series of coefficients, f_{ij}^k, in a 3rd order Taylor series expansion in x and y for each z plane and for each camera, centred at the image centre co-ordinates, x_{cen} and y_{cen} (see Equation 2.2). The coefficients are obtained by fitting the real and image space mark co-ordinates using a robust least squares approach. The reverse transformation is also calculated for later use in the self-calibration routine (Section 2.2.2) and PIV simulations (Chapter 3).

$$F_k(x_{im}, y_{im}) = \sum_{i=0}^{n_p} \sum_{j=0}^{n_p - i} f_{ij}^k (x_{im} - x_{cen})^i (y_{im} - y_{cen})^j \qquad (2.2)$$

Once the Calibration function has been derived, it can be used to assess the calibration accuracy. The real space coordinates of the calibration points are transformed (mapped) into image space coordinates using the calibration function. These coordinates are then compared with the original positions of the calibration marks on the image plane, with the mean difference representative of calibration accuracy. This accuracy value, expressed in image coordinates, is comparable with the measure used in previous investigations (31), and the accuracy studies conducted in the present report. A typical mean mapping accuracy of 0.25 pixels is achieved, which is sufficient for accurate TPIV results (see Section 2.3.3). However, it should be noted that this assessment fails to take account of real space co-ordinate errors, and as such may overestimate calibration accuracy.

2.2.2 Self-Calibration

Stringent TPIV calibration requirements can be addressed by 'self-calibration' (132), in which the calibration is refined based on particle images, thereby removing any residual error from calibration plate inaccuracies. In self-calibration real space particle positions are triangulated from their image projections as functions of the current calibration. Particle field images are used, with thresholding applied to retain only the strongest particle signals. The mean perpendicular distance between the triangulated particle centres and the converging lines of sight gives an estimate of the calibration error, which can then be minimised by updating the calibration to reduce this distance.

Due to the use of a pin-hole camera model, Wieneke details a method where this disparity calculation is performed at discrete points in the volume, with the update based on the mean displacement at these locations (132). In the current investigation, self-calibration software was written in MATLAB, with the current two plane 3rd order polynomial calibration system allowing a slightly different approach.

Again particle location through triangulation was used, with a similar method to Wieneke (132). Groups of possible matching lines of sight are defined as follows: after image pre-processing particle centres are located on each image using a

centroid approach and stored as lists. For each particle, the line of sight through the volume from the first image is mapped onto the second image. Particle centres close to this line of sight are shortlisted, and their the lines of sight together with that from the particle in the first image are mapped onto the third image. This process is repeated for the fourth image. All possible combinations of shortlisted lines of sight are tested, and the set which minimises distance between the lines of sight and particle centre location are grouped (86), and removed from the respective particle centre lists. An example of this process is shown in Figure 2.7, with the line of sight from a single particle identified in Figure 2.7(a), traced across all other images to identify possible matches.

(a) Image 1   (b) Image 2   (c) Image 3   (d) Image 4
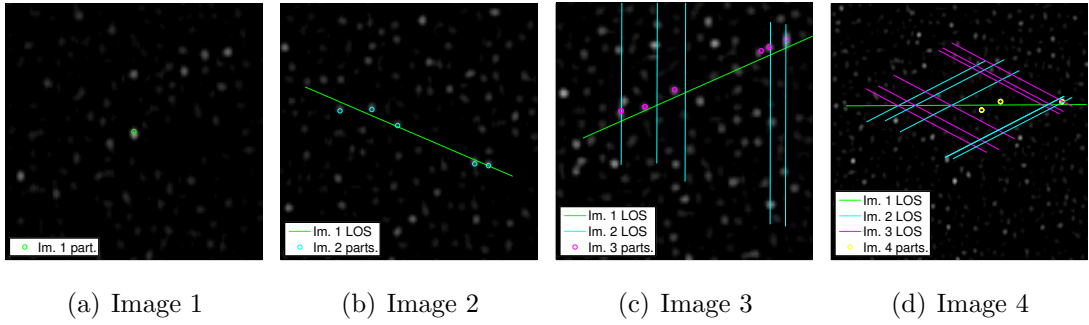
Figure 2.7: Example self-calibration point detection

Real space particle locations were then projected using the inverse transformation onto the front and back calibration planes, using an iterative refinement approach. Updated calibration coefficients are calculated from a robust least squares fit relating the updated calibration plane positions to the original particle image locations. These updated calibration coefficients are then used to refine the entire self-calibration process, which is iteratively executed over 3 passes.

This approach offers a number of potential advantages over the pin-hole camera volumetric approach, such as: use of all particle positions on each plane, ensuring an accurate fit and requiring fewer fields; and a simpler approach which removes any effects of arbitrary disparity map discretisation. For later use in the TPIV simulations (see Section 2.8) calibration disparity maps were still produced (shown in Figure 2.8) by identifying the mean local difference between the original and updated particle image co-ordinates.

(a) Image 1 (b) Image 2 (c) Image 3 (d) Image 4
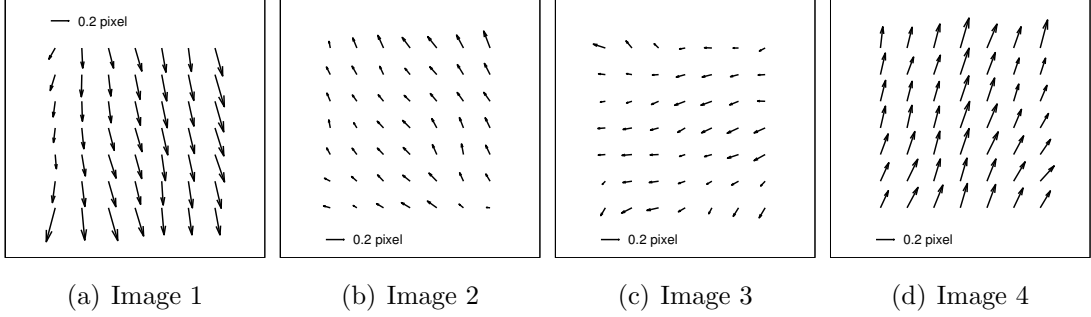
Figure 2.8: Self-calibration disparity maps; LOS=line of sight

The accuracy is assessed by transforming the real space coordinates of the particles into image space, and calculating the mean difference between these positions and the original image plane particle positions. The typical mean accuracy after self-calibration is reduced to around 0.15 pixels. Quantitatively, only a slight improvement in accuracy is realised, which may be due to the previous accuracy overestimation. It should also be noted that this measure of accuracy is dependent on user defined threshold for grouping particles. For example a high threshold will only allow very closely defined particle matches, which will result in a low calibration error, and therefore this measure should be treated with some caution. However, after the bundle adjustment implicit in the self-calibration method, there is a significant reduction in camera misalignment and corresponding improvement in the tomographic reconstruction (132).

2.2.3   Image Pre-processing

As discussed by Elsinga et al. in several papers (29; 31) image pre-processing significantly increases accuracy through a reduction in noise, with local intensity normalisation and Gaussian smoothing recommended (29). It is also necessary for a number of the cost reduction methods detailed in Section 2.2.4.

An image pre-processing routine was implemented in MATLAB. Initially a blank image is defined as the mean image intensity from all particle fields. Subtracting this from each image reduces camera noise due to light sheet diffusion and hot-pixels. After which bandpass filtering is performed using a 3×3 Gaussian kernel, subtracting background noise below a user defined threshold. Local

intensity normalisation is performed, before repeating the bandpass filtering and background noise removal operations (shown in Figure 2.9).

(a) Original   (b) After bandpass   (c) After local norm.
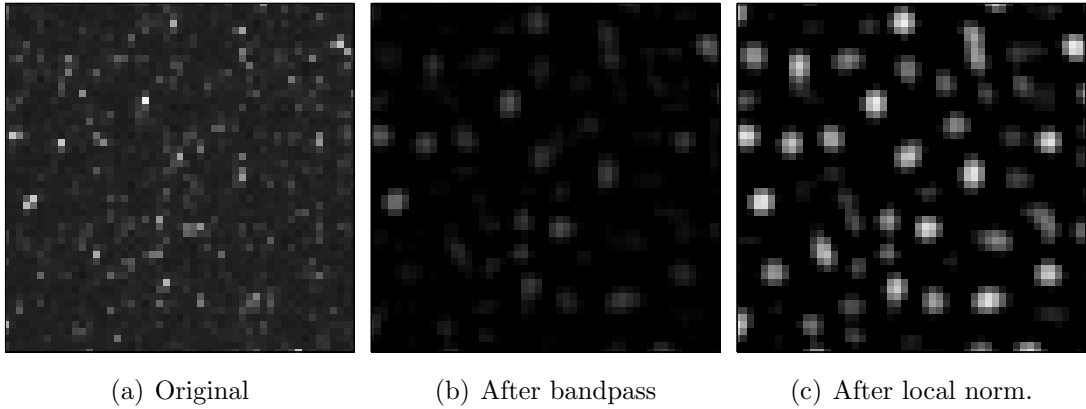
Figure 2.9: Example of image pre-processing in a 50×50 pixel window

2.2.4 Tomographic Reconstruction

In order to test the fundamental principles of TPIV and to make an independent study of the factors affecting reconstruction accuracy a 2D Tomographic program was written in MATLAB and a series of 2D studies performed (see sections 2.3.1). This also allowed some familiarity with iterative reconstruction techniques to be established before extending the code into three-dimensions, and translating it into FORTRAN 90.

The tomographic reconstruction software is based on an implementation of the MART algorithm (see Section 1.4.2), employing similar basis functions and pixel/voxel intersection relationships to Elsinga et al. (30). Thus, the volume is discretised into voxels which take a non-zero value inside and zero outside. The weighting values between pixels and voxels are determined by the distance between their centres and their relative sizes, based on the approximation of the intersecting area between a circle and rectangle, which are a simplified representation of the voxel volume and pixel line of sight respectively.

During code development the technique restrictions due to computational cost became clear; large spatial domains, high resolution images, and use of an iterative technique require both a huge number of calculations and a vast quantity

of RAM (random access memory) to store variables during processing. For example, resolution of a $700 \times 700 \times 140$ voxel volume imaged with 3 cameras set at $0°$ and $\pm 30°$ would require 4.8×10^{12} floating point operations (FLOP), and ignoring programming overheads, if run on a standard single core 3GHz processor would take approximately 25 minutes. Moreover, storing the W matrix for this calculation would require 13.2 GBytes of RAM, which far exceeds the current standard of approximately 1-2 GBytes (see Appendix A). Addressing these computational costs was necessary during the program development, and a number of programming solutions were employed.

MATLAB is originally an interpreted rather than a compiled language, and despite the ability to compile code, executing times are considerably slower than languages such as FORTRAN and C (25). Therefore, re-writing the code in FORTRAN 90 with the *Intel* Fortran compiler permitted a speed increase of approximately 50 times. Additionally, the code was parallelised using the OpenMP interface.

A new processing method was developed based on a 'Multiplicative First Guess' (MFG) of the intensity field (See Section 2.2.5), which cuts the number of voxels requiring the expense of MART algorithm iteration by approximately 95%, for typical seeding densities and domain aspect ratios (138). This first guess is so close to the final solution, that it contains sufficient information to directly obtain reasonable cross-correlation results (explored further in Section 2.3.3).

To reduce memory requirements W matrix values are only calculated when required. This 'on the fly' approach requires values to be re-calculated for every iteration and for every volume requiring reconstruction. However, after the first guess only the values relating to non-zero voxels are required. Testing demonstrated that this method only increases the reconstruction time by $\sim 15\%$ (assuming 5 iterations) and therefore although being slightly more computationally intensive, this method provides an acceptable solution.

As discussed in Section 1.4.1 voxel update order has a significant effect on convergence (18). Finding the optimum order would be computationally expensive, however even a random order will increase convergence (17) and is therefore employed. The RMART algorithm is also employed, with re-scaling affected through

the basis functions, although this does not generally result in significant time savings. Other measures include filtering insignificant updates, and use of ordered look-up tables to reduce weighting array value calculation times.

Specifically in the current investigation, the $T1$ case uses a $60 \times 60 \times 12$mm volume discretised at 15voxels/mm to create a $900 \times 900 \times 180$ voxel volume, with a pixel to voxel ratio of approximately 1. An under-relaxation parameter of 1 is used in the RMART algorithm, which after the multiplicative first guess is applied iteratively over 3 passes. Application of the first guess procedure accelerates convergence, and allows fewer iterations in comparison with previous investigations (see Section 2.3.3). Processing was performed on a 2×3GHz Quad-core Intel Xeon workstation (MAC OSX 10.4.11) with 16GB of RAM. The mean processing time per field (including cross-correlation) was 12.4 minutes, with 7500 fields processed in around 6 weeks.

2.2.5 Multiplicative First Guess

Using an iterative algebraic approach requires an initial volume intensity distribution to be set, which can then be updated using a reconstruction algorithm. The most basic first guess is uniform intensity throughout the volume; a solution commonly employed in other investigations (28; 31). As in CFD, a poor first guess will slow solution convergence, with a uniform intensity distribution forcing the algorithm to search every voxel during the initial iteration. A better approach would be to set the intensity field so that only regions that will contain blobs of intensity after the reconstruction are given an initial value. In other words, find regions within the volume to which non-zero intensity contributions are made by all cameras, and set any voxels outside these regions to zero intensity. This is justified by considering the update of a voxel within the volume: irrespective of updates from other high intensity pixels, if just one strongly correlated pixel contains zero intensity this will dominate the voxel update, and its intensity will tend towards zero.

Several new first guess methods were created by calculating the current projection $p^T W$ of each camera through the volume (Figure 2.10). This extends

(a) Artificial particle distribution

(b) Camera 1 current projection

(c) Camera 2 current projection

(d) Camera 3 current projection
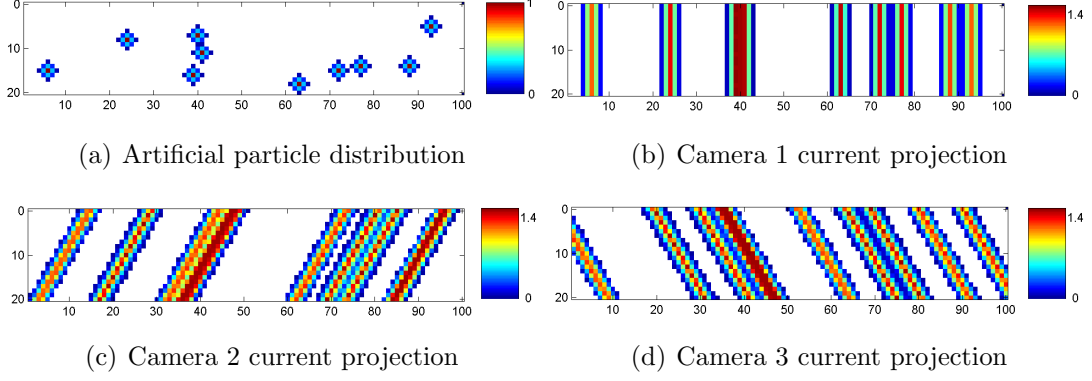
Figure 2.10: 2D projections (3 cams at 0° and ±30°)

(a) Artificial particle distribution

(b) Uniform first guess field

(c) Additive first guess field

(d) Multiplicative first guess field
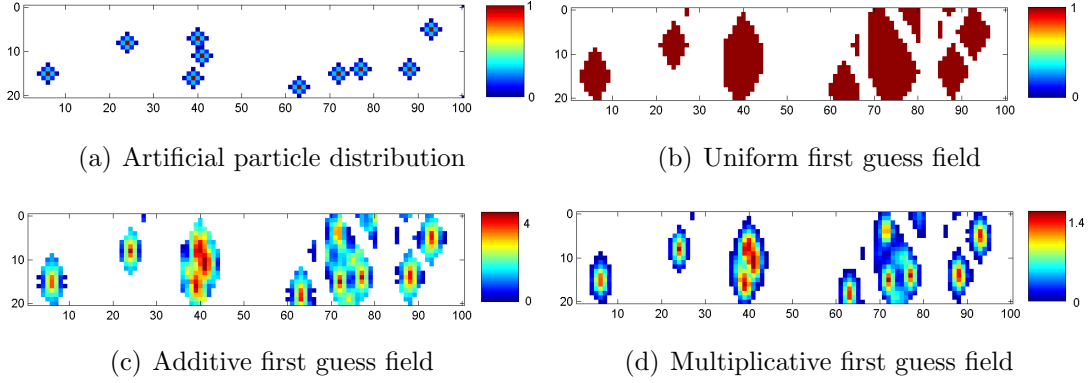
Figure 2.11: First guess scheme comparison (3 cams at 0° and ±30°)

intensity from each pixel along its line of sight, creating a series of constant intensity *streaks* through the volume for each camera. Overlapping regions were either set to a uniform value, summed, or multiplied (Figure 2.11), creating different initial fields. This places intensity only where all lines of sight converge, and in the latter two cases the magnitude of that intensity becomes dependent on all camera contributions. If any contribution to a particular voxel is zero, then the voxel intensity will be zeroed; accelerating the inevitable fate of the voxel under the MART scheme. In the case of the multiplicative first guess (MFG) scheme, the intensity magnitude is also normalised by raising it to the power $1/N_{cam}$, where N_{cam} is the number of cameras, to ensure relative field magnitude continuity. Thus, the elimination of a huge number of voxels from the iterative calculation is performed through simple matrix multiplication, without requiring

67

the expense of the MART algorithm (Equation 1.20), thereby reducing compu-
tational costs significantly (Detailed in Appendix A).

(a) Artificial particle distribution

(b) MFG field (3 cams at 0°, ±30°)

(c) MFG (4 cams at 0°, ±30° and 60°)

(d) MFG field (5 cams at 0°, ±30° and ±60°)
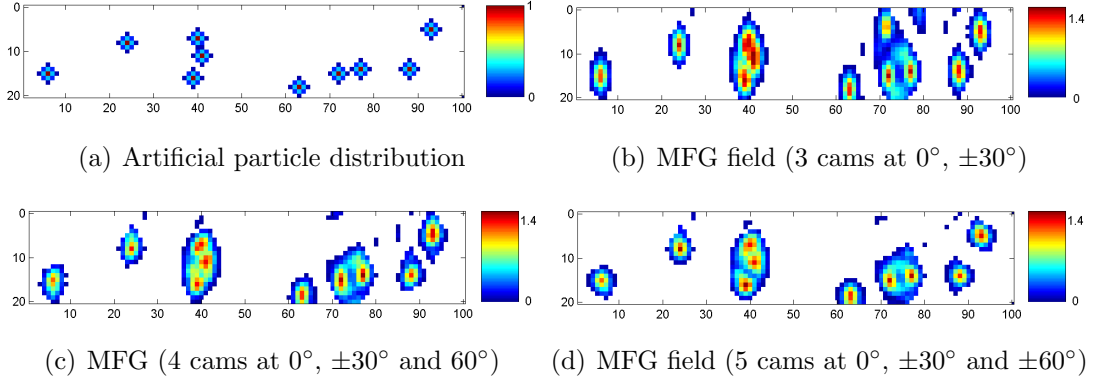
Figure 2.12: Effect of additional cameras on the MFG intensity field

The currently proposed MFG scheme appears to offer the most accurate es-
timate of the solution, and is therefore recommended over the other previously
mentioned schemes. The effect of adding cameras to the system can also be shown
to further increase the accuracy of this first guess (see Figure 2.12), offering more
well-defined intensity distribution in local regions of high particle density, and
more effective elimination of *Ghost Particles*.

(a) Artificial image

(b) Initialised intensity field
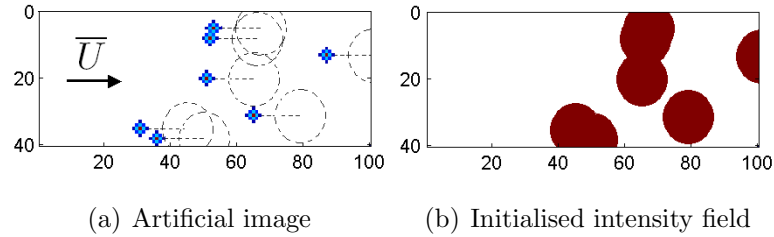
Figure 2.13: Predictive regions first guess procedure

Although not considered in this study, an additional PIV based first guess
mechanism could be used. If the maximum possible particle displacement is well
understood, the probable region in which a particle may be found in a subsequent
frame of a pair of images, or indeed subsequent pairs of images should be deter-
minable to within certain limits. Therefore a spherical radius of possible intensity
surrounding the likely particle location could be created, which could also take

account of the mean flow velocity (shown in Figure 2.13). Taking account of local flow velocity could permit a tighter radius of probable locations, although this would be more involved. Combined with the previous criterion, further computational savings may be possible. However, highly intermittent flow may require the use of large radii, and the implementation costs of this scheme may render it unworkable.

2.2.6 Displacement Evaluation

Currently a bewildering array of correlation techniques exist with no commonly agreed standard. In response to this the 'PIV challenge' initiative (117; 118) compared algorithms using synthetic and experimental data. Despite classification difficulties arising from various algorithm subtleties, the following general trends were observed: Gaussian sub-pixel estimation and multi-pass window shifting are both widely adopted; around 50% of the algorithms tested employed some form of window deformation; and local median or mean filters are popular choices for post-processing. In their analysis, Stanislas et al. (117; 118) conclude that the majority of algorithms produce similar results. Only modest improvements are seen for the more advanced algorithms (most results are within $0.1 \rightarrow 0.2$ pixels), which interestingly come at the expense of robustness (especially with regard to image quality (117)), in addition to computational cost and implementation complexity. Furthermore, despite evidence to suggest that these advanced algorithms improve results in high gradient regions, it should be noted that they also have a tendency to overestimate quantities such as the vorticity (117).

Initially both 2D and 3D cross-correlation software was developed in MAT-LAB, to gain a better understanding of different correlation techniques and reduce development time. The code was then translated to FORTRAN 90 and incorporated into the tomographic reconstruction algorithm. *Davis 7.2* software was used to process the DPIV results.

The present TPIV cross-correlation software employs a Fast Fourier Transform (FFT) based correlation process, using Gaussian correlation plane smoothing before the inverse transformation. The correlation plane is multiplied with a triangular weighting function to correct for the known velocity under-prediction

bias (detailed by a number of authors (100; 129)). A Gaussian curve fit method is used to achieve sub-pixel accuracy, which due to its effective minimisation of bias error and implementation simplicity, is recommended by a number of authors (100; 129).

A multi-pass recursive window-shifting approach is adopted, using vector validation and Gaussian smoothing of the velocity field between passes, with integer window displacements based on linear interpolation of the smoothed fields. This approach is relatively simple to implement, has a low computational overhead (105), and will reduce the detrimental effects of in-plane particle losses (57), allowing smaller interrogation regions to be used to obtain higher spatial resolution.

Due to increased algorithm complexity and computational cost, balanced against the potential gains in accuracy (2) it was decided not to implement window deformation.

Increasingly over-sampled data is subject to more severe coupling effects as neighbouring cells lose independence. Furthermore, the measurement noise remains constant even if the distance between neighbouring grid cells is reduced, increasing the uncertainty of the differential quantity estimation as the overlap in increased. Therefore, an overlap of 50% is recommended (100), providing a good compromise between resolution and accuracy. The size of the smallest interrogation windows was set by the requirement to have a minimum number of particles within the window (detailed in Table 2.3).

The vector validation subroutine follows the methodology of Westerweel (128), comparing vectors across a $3 \times 3 \times 3$ mask, rejecting vectors which exceed twice the median RMS velocity in any component. A multi-pass approach is again applied, retaining only valid vectors for the final comparison, and extending the mask to 5×5 points in the case of an insufficient number of valid data points. Spurious vectors are replaced with vectors derived from the second and third highest correlation plane peaks and revalidated as detailed by Raffel et al. (100), to be finally replaced by a mean value in the case of repeated failure.

Similar options were selected for the DPIV processing, with the only major difference being the use of window deformation in this code. A summary of the cross-correlation options is presented in Table 2.3.

Technique	Overlap	Passes	IW_1	IW_2	IW_3	IW shift	IW def.	Vectors	Vec. Val.
DPIV	50%	3	128^2	64^2	32^2	yes	yes	64×64	median
TPIV	50%	3	80^3	50^3	32^3	yes	no	$51 \times 51 \times 8$	median

Table 2.3: PIV processing options

2.3 Code Validation and Optimisation Studies

During the project the TPIV software was validated and optimised using a variety of numerical and experimental tests, which are over-viewed in this section.

Initial testing was by performed through a series of numerical studies. 2D and 3D parametric studies were conducted and compared with previous results (28), based on a similar measure of reconstruction accuracy. Following these, the new MFG method was analysed, with parametric studies based on an accuracy analysis technique more applicable to PIV.

Following this a series of experimental tests were carried out to first reconstruct a real object, then assess the displacement of 'frozen' particles on a traverse, and finally to resolve a simple laminar vortex ring flow field. These tests formed a valuable part of the code development and optimisation; permitting an informed choice of reconstruction parameters to be made in the main investigation.

2.3.1 2D Parametric Testing

A parametric reconstruction accuracy study was conducted using the MATLAB based 2D Tomographic reconstruction code, allowing comparisons with Elsinga et al. (31). A summary of the results are included below, and for further information the reader is referred to Worth & Nickels (140).

The 2D code effectively reconstructs a 1 voxel slice of the 3D volume; so the object can be considered as a 2D area and the camera projections as 1D images, which are related through simple geometric relationships as opposed to a more complex camera model. A random distribution of particles, modelled as Gaussian intensity blobs, are created in the discretised domain, and the intensity of each pixel is projected to create a set of artificial images. The images are processed using the MART algorithm to reconstruct the original particle distribution.

A 500×200 pixel domain was reconstructed using artificial 500 pixel images, with a correlation coefficient between the reconstructed and original artificial domain applied as a measure of reconstruction accuracy. Parametric tests show the effect of varying seeding density, number of cameras, camera angle, calibration error, image noise, and discretisation; performed with and without the interference and capping blocking models.

(a) Seeding density

(b) Camera number

(c) Camera angle

(d) Discretisation
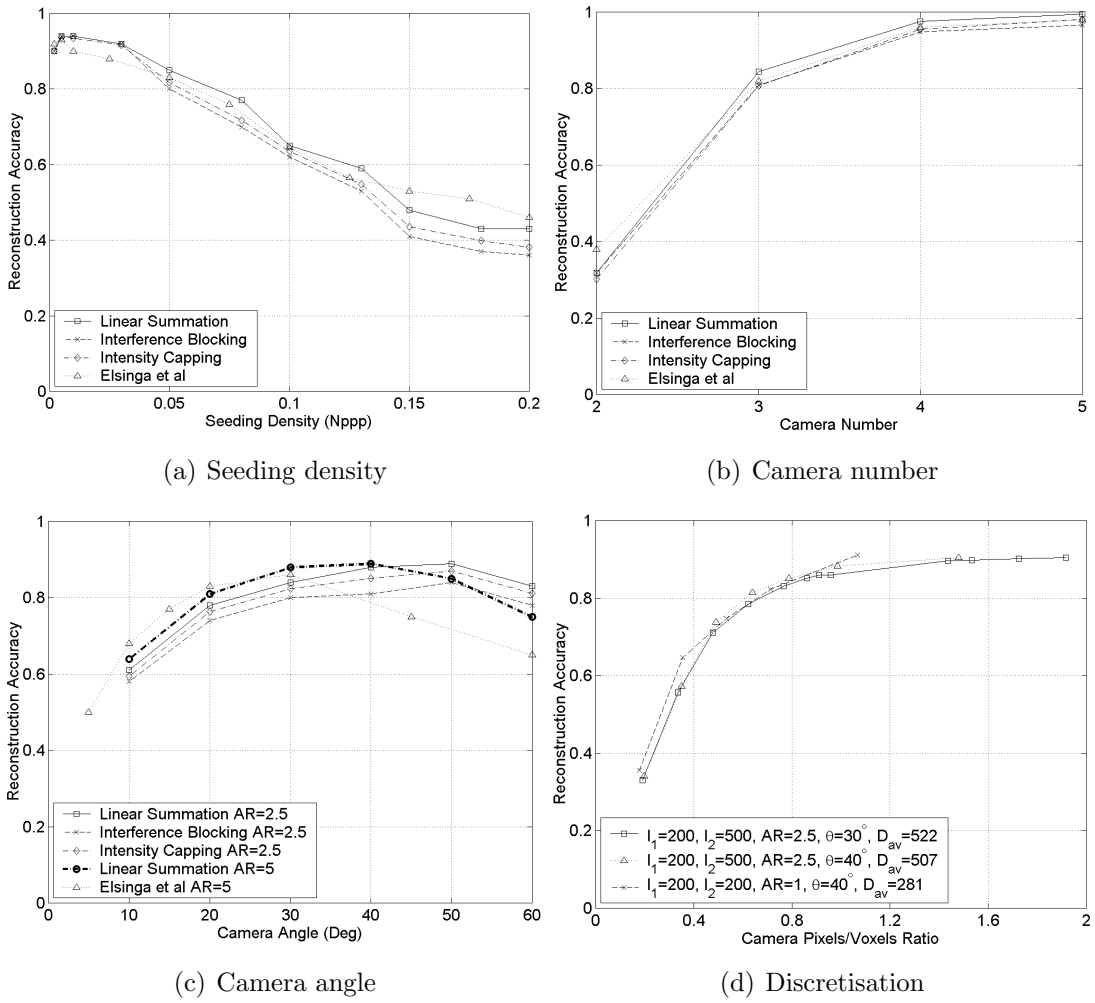
Figure 2.14: Example of 2D parametric study results

Figure 2.14 shows close agreement with the previous study (30) for all variables except camera angle, which are only slightly different. A useful addition was the quantification of volume discretisation effects shown in Figure 2.14(d),

which demonstrate that increasing the pixel to voxel ratio beyond a value of 1 yields only modest accuracy benefits. Therefore, through this ratio, a given camera resolution will dictate volume discretisation according to users' preference of accuracy against computational costs.

2.3.2 3D Parametric Study

A number of parametric studies were repeated using the 3D algorithm in order to determine additional reconstruction effects. Using a similar methodology to the 2D study, an artificial particle distribution was created in a $700 \times 700 \times 140$ voxel volume, which was then reconstructed from three artificial 700×700 pixel images.

(a) Seeding density   (b) Image noise
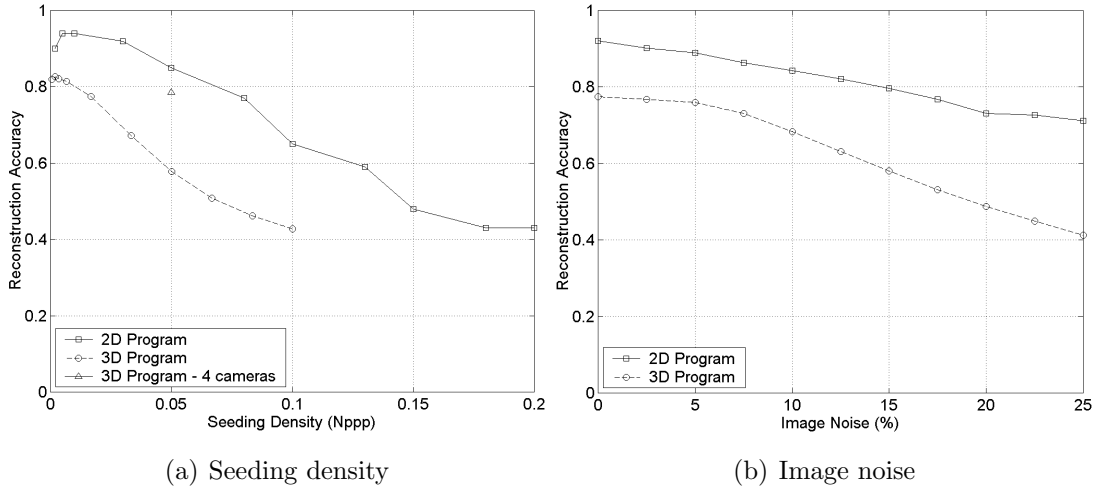
Figure 2.15: Example of 3D parametric studies

Comparing 2D and 3D results (Figures 2.15) highlights a sharp accuracy drop for similar set-ups. Although not explicitly stated, a similar drop in accuracy is shown in Elsinga et al. (31), where a 4-camera 3D reconstruction returns a correlation coefficient of 0.75; far lower than the value of 0.94 given for a similar 2D reconstruction.

This accuracy drop can be attributed to increased 'intensity smearing' due to the addition of an extra dimension. Image intensity from a single pixel will be distributed over 2 voxels (in each row of the object) in a 2D case, but 4 voxels in

3D. Therefore this extra 'smearing' may affect the correlation coefficient, but is unlikely to affect cross-correlation accuracy, which can take account of coarsely discretised intensity in the correlation plane using a sub-pixel estimator. Aside from this accuracy drop these results showed similar trends to the 2D study, with increases in seeding density, calibration error, and image noise all leading to a loss of reconstruction accuracy.

2.3.3 A Cross-Correlation Based Accuracy Analysis

In order to quantify the impact of the new MFG method on reconstruction accuracy and computational workload, a numerical study of an angled line vortex flow field was conducted; including a parametric investigation of iteration number, seeding density, camera angle, image noise, calibration error, and first guess scheme. In this study a cross-correlation based accuracy assessment was used in place of the previous correlation coefficient method, which is a more applicable measure for PIV investigations. A summary of the results follows, and for more detailed set-up information and discussion the reader is referred to Worth & Nickels (138).

(a) MFG, $k_{it} = 0$ (b) MFG, $k_{it} = 5$ (c) UF, $k_{it} = 5$
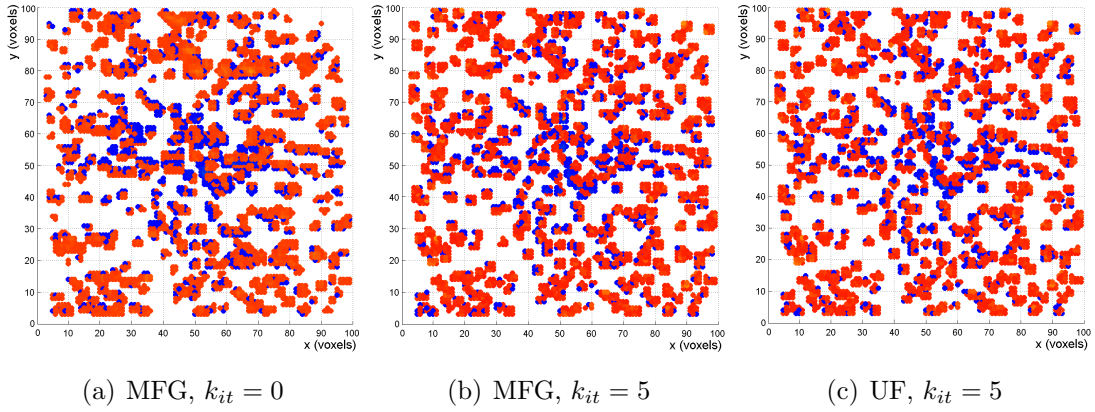
Figure 2.16: Comparison of reconstructed displaced particle fields X-Y plane view; MFG=multiplicative first guess; UF=uniform field; k_{it}=number of iterations

The new MFG method was shown to accelerate convergence, providing a reasonably accurate solution even before MART algorithm implementation. Recon-

(a) Iteration # - large volume

(b) Seeding density- large volume

(c) Camera Angle - small volume

(d) Calibration Error - small volume
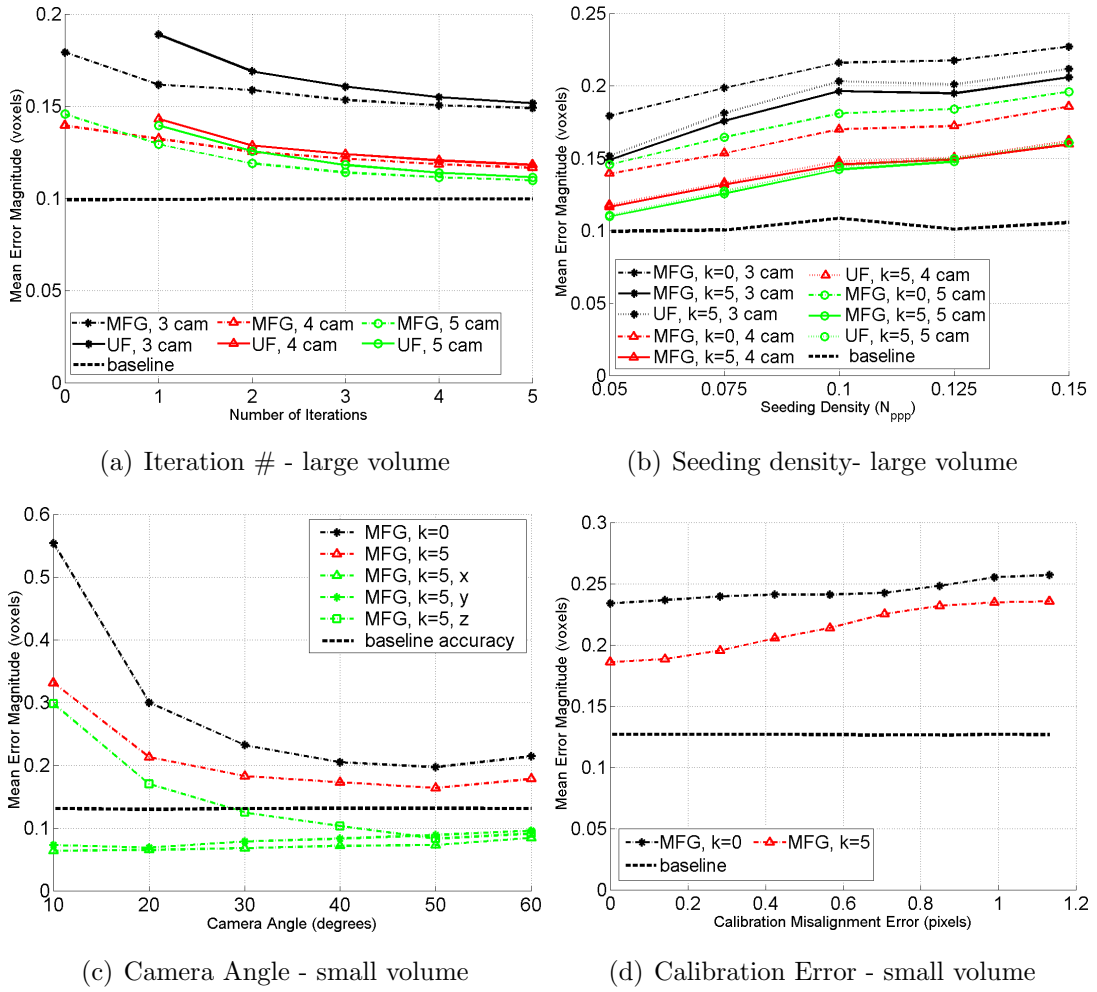
Figure 2.17: Error variation with number of iterations and seeding density

struction times were 175 and 35 times faster theoretically and during practical application respectively (See Appendix A). Solution robustness to degradation through camera misalignment was also enhanced, although after iteration this advantage appears not to hold. These results also indicate fewer iterations can be used in comparison with previous investigations (31) to reduce computational expense with only a fractional reduction in accuracy.

The parametric study results for seeding density and camera number appear to largely support the previous investigation (31), although the present study suggests that the benefit of a fifth camera is less clear, especially at high seeding densities. These results also show the solution is less noise sensitive than previously indicated, and slightly wider camera angles give optimum accuracy, although only yaw was considered.

2.3.4   Experimental Validation Testing

The first practical operability test was to reconstruct a translucent plastic frame with a number of regularly spaced cells (see Figure 2.18), reconstructing a $200 \times 150 \times 50$ voxel volume from four 360×258 pixel images, taken within a solid angle of $30° \times 40°$, with a calibration accuracy of 0.55 pixels.

(a) Actual image   (b) Intensity plot   (c) Background removed
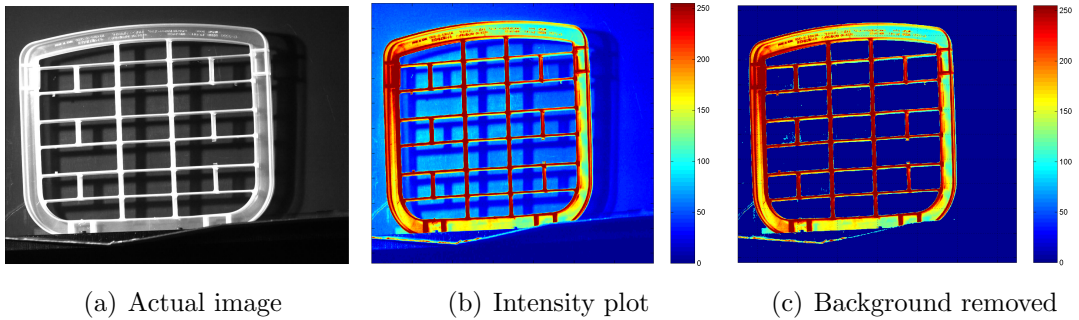
Figure 2.18: Object image before and after pre-processing

Reconstruction accuracy was assessed through frame cell spacing, depth, and wall thickness measurements. Despite imperfections due to insufficient views of the relatively complex object, poor calibration accuracy, and limited camera

angles, the object was reconstructed reasonably successfully. The real and reconstructed object's cell dimensions were measured a number of times (25 in each dimension), with the difference between the two used as a measure of accuracy. The RMS differences in the $[x, y, z]$ directions were calculated as $[1.16, 1.34, 1.66]$mm. Due to a pixel/voxel ratio of ≈ 1, accuracy to within a single voxel (1mm) is all that can be demanded of the reconstruction, and therefore this test was considered successful.

(a) Reconstruction - isometric view   (b) Reconstruction - $x - y$
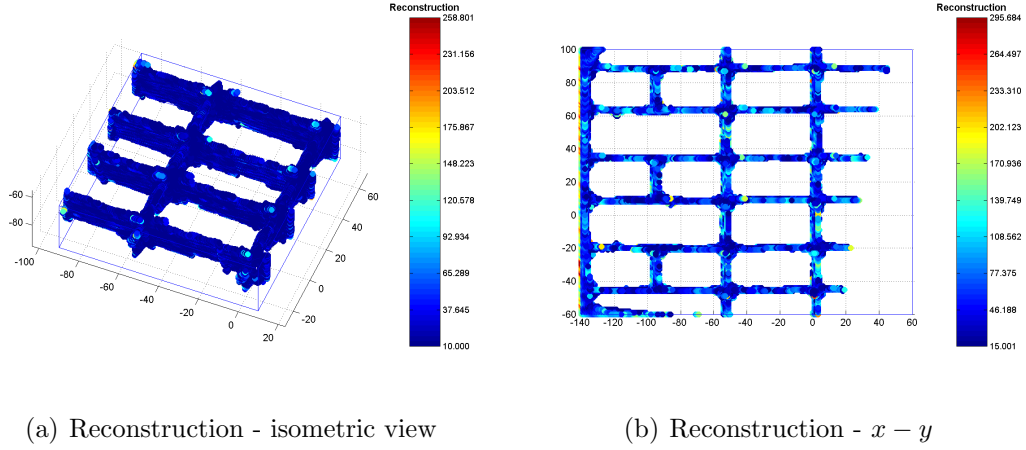
Figure 2.19: 4 camera plastic frame reconstruction - small inner section

2.3.5 Reconstructing Frozen Particles

A more applicable test was performed by imaging a random distribution of 100μm particles set into a cube of clear polyurethane resin. A known displacement was applied using a motorised 2-axis translation stage (shown in Figure 2.20). A $30 \times 30 \times 30$mm volume was reconstructed over a range of discretisations using three 360×258 pixel images, taken within a solid angle of $50°$, with a camera calibration accuracy of approximately 0.28 pixels.

An example of the volumetric reconstruction is shown in Figure 2.21(a) with blue and red colour maps representing the original and displaced particles respectively. The mean displacement was calculated using cross-correlation and a correction for the known backlash error is applied. Figure 2.21(b) demonstrates

Figure 2.20: Perspex block and calibration plate mounted on traverse

the software ability to produce good particle displacement estimates. The accuracy is assessed comparing the measured PIV displacement with the traverse displacement, with the percentage error expressed as a fraction of the latter, taking a value of approximately 4%.
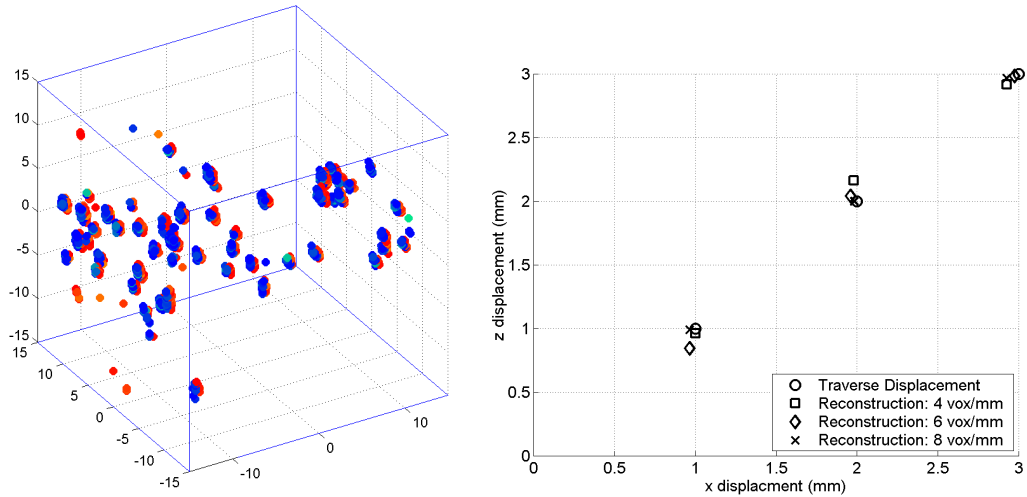
(a) Particle reconstruction - isometric view (b) Mean displacement $x - z$ plane

Figure 2.21: Reconstructed volume and Mean Displacement after Cross Correlation; (a) blue and red colour maps for original and displaced particle fields respectively

2.3.6 Simple Vortex Ring Flow Field

To test the software on a real 3D flow field and demonstrate the MFG method, a simple vortex ring flow was chosen. A summary of the results is included below, and for more information the reader is referred to Worth & Nickels (139).

A computer controlled piston and cylinder arrangement was used to generate vortex rings ($Re = 1 \times 10^3$ and $L_0/D_0 = 1$) in a small water tank. A $50 \times 50 \times 20$mm volume was discretised at 10 voxels/mm and reconstructed from three 512×512 pixels images imaged at $0°$ and $\pm 45°$ with a mean calibration error of 0.2 pixels. Three cases were compared: *1.* The MFG alone (MFG $k_{it} = 0$); *2.* A uniform field initial guess and 5 MART iterations (UF $k_{it} = 5$); *3.* The MFG and 5 MART iterations (MFG $k_{it} = 5$). Cross-correlation yielded a velocity field at a spatial resolution of 2.1mm which was smoothed using a 2-pass box filter.

(a) Vertical velocity profile

(b) v-vorticity profile
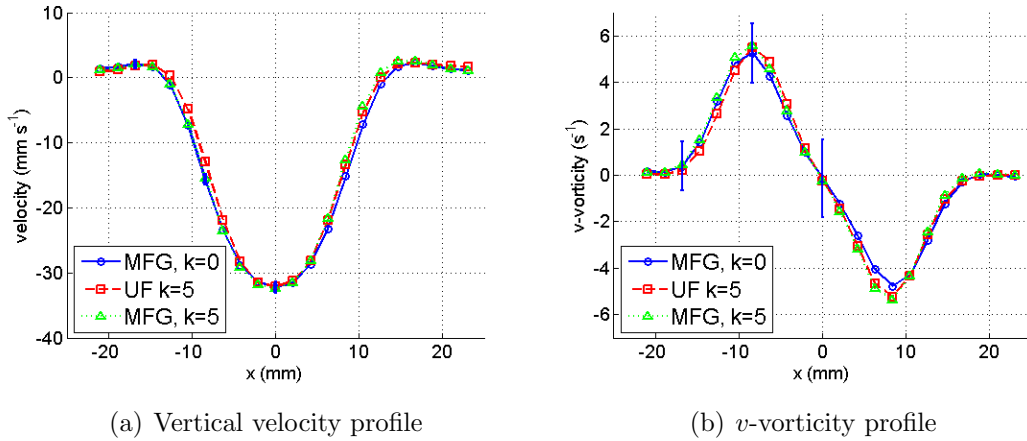
Figure 2.22: Velocity and vorticity profile comparison

Analysis of velocity profiles, circulation estimates, and an examination of the vorticity field demonstrate that the main flow details are captured well using all processing variations. A further comparison can be made using the video files included in Appendix D. The deviations between velocity profiles are within the bounds of the experimental error, and even the more sensitive vorticity distributions are very closely matched. The Slug Model assumes the vortex ring is formed from a cylindrical slug of fluid ejected from the cylinder (66), predicting the circulation, Γ_0, in terms of the piston velocity, U_p, and the piston stroke

(a) MFG $k_{it} = 0$ (b) UF $k_{it} = 5$ (c) MFG $k_{it} = 5$
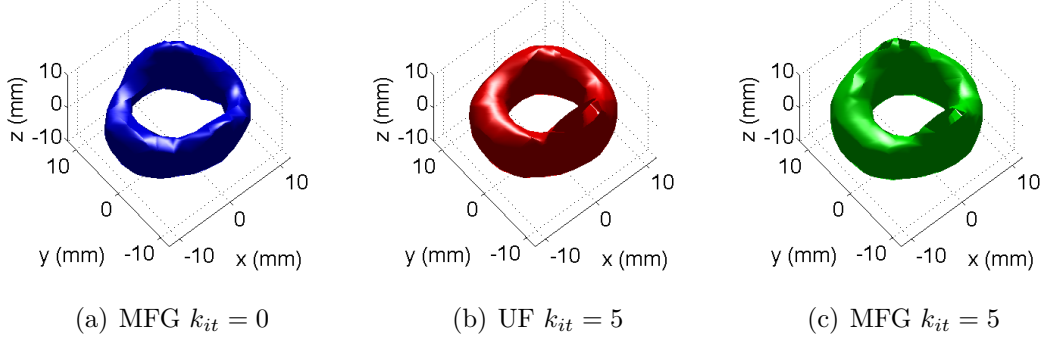
Figure 2.23: Vorticity magnitude isosurface comparison (65% maximum vorticity magnitude)

time, T_0 (Equation 2.3). Comparing these circulation predictions to line integral measurements confirms the similarity between schemes. The variation of the MFG method alone is consistent with the numerical study (see Section 2.3.3), and the advantage of this new method is demonstrated; reducing the computational workload by up to 70 times (139).

$$\Gamma_0 = \int_0^{T_0} \frac{1}{2} U_p^2(t) dt \tag{2.3}$$

2.4 Post-Processing

After velocity field processing the mean is calculated as an ensemble average, and subtracted from each field to obtain the fluctuating component of velocity, from which all other quantities are derived. The TPIV results are smoothed over 2-passes using a $3 \times 3 \times 3$ Gaussian kernel with a standard deviation of 0.65, to reduce the effects of noise (see Chapter 3).

This section contains details of the finite differencing scheme used to calculate the gradient fields, various direct dissipation rate calculation methods, the energy spectrum calculation method, and vortex identification algorithms.

2.4.1 Differential Field Calculation

Finite differencing must be applied to evenly-spaced PIV data in order to calculate the velocity gradient tensor. A first order central differencing scheme was adopted (Equation 2.4) to maintain an effective spatial filter size of three grid points, which at a 50% overlap is the same as the interrogation window size. This scheme cannot be applied to the edge points, and therefore either forward or backwards differencing must instead be used (Equations 2.5 and 2.6). Due to the higher error associated with these alternative differencing methods (100), these points are excluded from the majority of the statistical analysis, and only used during visualisation and vortex geometry characterisation (sections 5.1.3 and 5.2.3) where a large domain size is important.

$$\left(\frac{du}{dx}\right)_i = \frac{u_{i+1} - u_{i-1}}{2\Delta x} \tag{2.4}$$

$$\left(\frac{du}{dx}\right)_i = \frac{u_{i+1} - u_i}{\Delta x} \tag{2.5}$$

$$\left(\frac{du}{dx}\right)_i = \frac{u_i - u_{i-1}}{\Delta x} \tag{2.6}$$

2.4.2 Dissipation Estimation

The turbulent kinetic energy dissipation rate for a Newtonian fluid is defined in Equation 2.7, the calculation of which requires knowledge of the velocity gradients in all 3 dimensions. While TPIV provides sufficient information for this calculation, the DPIV measurements do not, and it is necessary to eliminate some terms.

$$\epsilon = 2\nu\langle S_{ij}S_{ij}\rangle = \nu\left\langle \frac{\partial u_i}{\partial x_j}\frac{\partial u_i}{\partial x_j} + \frac{\partial u_i}{\partial x_j}\frac{\partial u_j}{\partial x_i} \right\rangle \tag{2.7}$$

Using assumptions of isotropy and homogeneity this can be simplified to Equations 2.8 and 2.9 in one and two dimensions respectively (43).

$$\epsilon = 15\nu\left\langle \left(\frac{\partial u_1}{\partial x_1}\right)^2 \right\rangle \tag{2.8}$$

$$\epsilon = 6\nu \left[\left\langle \left(\frac{\partial u_1}{\partial x_1}\right)^2 \right\rangle + \left\langle \left(\frac{\partial u_1}{\partial x_2}\right)^2 \right\rangle + \left\langle \frac{\partial u_1}{\partial x_2}\frac{\partial u_2}{\partial x_1} \right\rangle \right] \tag{2.9}$$

The gradient terms in Equation 2.9 are replaced arbitrarily, which is justified given the assumption of isotropy. However, depending on measurement resolution and location within the energy cascade, local anisotropy may persist; in which case the assumptions of isotropy may fail. This is noted by George & Hussain (39), who propose an 2D axisymmetric dissipation estimate suitable for the current flow geometry (Equation 2.10), which should account for persistent anisotropy at the velocity gradient level. Although the persistence of anisotropy may indicate insufficient measurement resolution and therefore failure to accurately render gradient terms, given sufficient resolution it will be interesting to compare these estimates with that of Hinze (43).

$$\epsilon = \nu \left[-\left\langle \left(\frac{\partial u_2}{\partial x_2}\right)^2 \right\rangle + 2\left\langle \left(\frac{\partial u_2}{\partial x_1}\right)^2 \right\rangle + 2\left\langle \left(\frac{\partial u_1}{\partial x_2}\right)^2 \right\rangle + 8\left\langle \left(\frac{\partial u_1}{\partial x_1}\right)^2 \right\rangle \right] \tag{2.10}$$

Three other estimates are also assessed, the first by de Jong et al. (24) (Equation 2.11), the second by Michelet et al. (76) (Equation 2.12), and the third derived in the current investigation (Equation 2.13) based on the arbitrary replacement of gradient terms using all four available gradients, as opposed to the three used by Hinze (43).

$$\epsilon = 4\nu \left[\left\langle \left(\frac{\partial u_1}{\partial x_1}\right)^2 \right\rangle + \left\langle \left(\frac{\partial u_2}{\partial x_2}\right)^2 \right\rangle + \left\langle \left(\frac{\partial u_1}{\partial x_1}\frac{\partial u_2}{\partial x_2}\right) \right\rangle + \frac{3}{4}\left\langle \left(\frac{\partial u_2}{\partial x_1} + \frac{\partial u_1}{\partial x_2}\right)^2 \right\rangle \right] \tag{2.11}$$

$$\epsilon = \nu \left[2\left\langle \left(\frac{\partial u_1}{\partial x_1}\right)^2 \right\rangle + 5\left\langle \left(\frac{\partial u_2}{\partial x_2}\right)^2 \right\rangle + 2\left\langle \left(\frac{\partial u_1}{\partial x_2}\right)^2 \right\rangle + 2\left\langle \left(\frac{\partial u_2}{\partial x_1}\right)^2 \right\rangle \right] \tag{2.12}$$

$$\epsilon = \frac{10}{3}\nu \left[\frac{1}{2}\left\langle \left(\frac{\partial u_1}{\partial x_1}\right)^2 \right\rangle + \frac{1}{2}\left\langle \left(\frac{\partial u_2}{\partial x_2}\right)^2 \right\rangle + \left\langle \left(\frac{\partial u_1}{\partial x_2}\right)^2 \right\rangle \right. $$
$$\left. + \left\langle \left(\frac{\partial u_2}{\partial x_1}\right)^2 \right\rangle + \left\langle \left(\frac{\partial u_1}{\partial x_2}\frac{\partial u_2}{\partial x_1}\right) \right\rangle \right] \tag{2.13}$$

2.4.3 Energy Spectrum Calculation

The 1D energy spectrum is calculated directly from the PIV data using a similar approach to Foucaut et al. (34). 1D lines of the instantaneous velocity are transformed using the FFT and multiplied by their conjugate values, with the energy averaged over each field according to a homogeneity hypothesis and ensemble averaged over all fields. To compensate for non-periodic PIV data which will lead to truncation errors, a weighting function can be used (14). de Jong et al. (24) show that these can have a dramatic impact on energy spectrum calculation, and therefore to quantify this effect in the current investigation, results are shown with and without *Hann* windowing. It should be noted that applying the windowing function results in energy loss through signal appodisation, and therefore to take account of this, the integral of the windowed spectra energy is re-normalised by the integral of the energy spectra without windowing. The TPIV results are processed using Gaussian smoothing as before.

2.4.4 Vortex Identification Algorithms

In order to characterise coherent structure geometry and clustering it is necessary to algorithmically define the set of points which constitute a single structure. Two different techniques, following the vortex core tracking approach of Jiménez (53) and the box-counting methodology of Moisy & Jiménez (79), are introduced here.

2.4.4.1 Vortex Core Tracking

As in Jiménez et al. (53) the radius and circulation of the high vorticity structures are assessed using a vortex core tracking algorithm. An enstrophy threshold is applied, $\Omega > \Omega_{rms}$, with points above this marked as valid. Initially the tracking algorithm identifies highest enstrophy point within the valid point subset, and calculates the vorticity vector orientation at this location. Next the algorithm extrapolates lines from the vorticity vector in both directions until these intersect adjacent grid planes. The valid points surrounding the intersections in these planes are searched to find the local enstrophy maxima, which are stored as the next vortex core locations. The search for new core locations is repeated until

no valid points surround the new extrapolated intersection, or the domain edge is reached. Vortex cores with fewer than 6 points are rejected.

At each core point the vortex radius and circulation are assessed by slicing the flow perpendicular to the vorticity vector. The slice is discretised and the distance of each element, r, and vorticity value relative to the maximum, ω/ω_0 are plotted (as shown in Figure 2.24(b)). A Gaussian distribution is fitted to these points using a robust least squares approach, with the vortex radius, $R_{1/e}$, defined as the distance at which the vorticity drops to $1/e$ of the core value, ω_0 (see Figure 2.24(c)). The circulation is approximated through integration of the Gaussian as in Equation 2.14, and expressed in terms of its Reynolds number, $Re_\gamma = \gamma_0/\nu$.

$$\gamma_0 = \omega_0 \pi R_{1/e}^2 \tag{2.14}$$

After structure identification, the core points and any additional points within the calculated vortex radius are removed from the list of valid locations. The tracking algorithm is then re-applied until all valid points have been assessed.

2.4.4.2 Box-Counting Method: Geometry Classification

In order to characterise the shape of different flow structures a box counting approach was applied (79). Initially all points below the chosen enstrophy threshold of $\Omega > \Omega_{rms}$ are discarded. The algorithm then searches through the remaining points, grouping these with any neighbouring points. These connected regions are then grouped with other connected regions to build up larger structures, with this process repeated until no further connections can be made. Groups containing less than 10 points are rejected, as these smaller groups contain too few points for accurate classification. Structures are therefore defined as spatially connected high enstrophy regions.

The limited domain size makes structure classification difficult (see Section 5.1.3). So in order to classify structure geometry it is necessary to consider a subset which will not be significantly affected by the domain truncation. Therefore, the mean alignment between the vorticity vector and a unit vector in the out-of-plane direction is calculated for each structure ($\overline{\hat{\omega} \cdot \hat{z}} = cos\theta$), and only structures which are
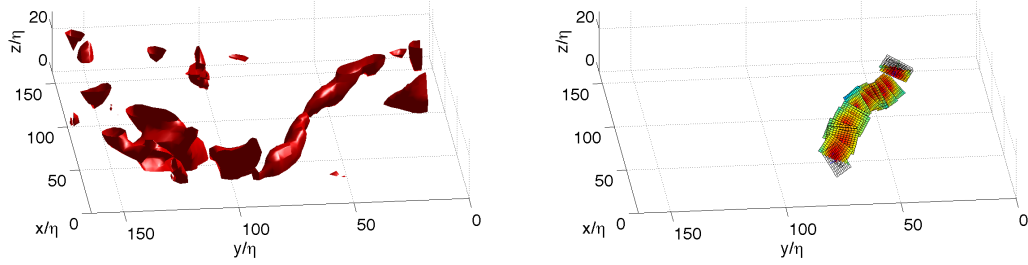
(a) Iso-surfaces of Enstrophy Magnitude

(b) Slices parallel to vorticity vector
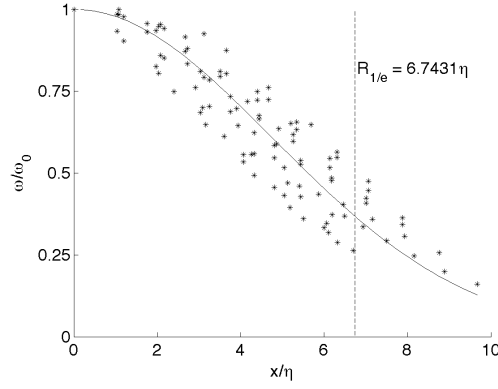
(c) Gaussian fit

Figure 2.24: Demonstration of vortex line tracking algorithm

predominantly aligned in this direction ($\cos\theta > 0.65$) are considered; removing the potential ambiguity associated with in-plane aligned tube-like structures and ribbon-like structures.

The remaining structures can then be classified as tube or ribbon-like by their in-plane dimensions. Due to the alignment condition, the remaining structures only vary significantly in the in-plane dimensions, and therefore the geometry classification can be reduced to a 2D approach. Each structure is divided into planes perpendicular to the out-of-plane direction. For each of these in-plane layers, major and minor axes of an ellipse with the same normalised second central moment as the structure can be calculated, with the ratio of these giving a

measure of the structure shape (roughly equivalent to the inverse of aspect ratio). Each layer value is weighted by the structure fraction contained within the layer, with the sum of these defining the structure shape in the range, $0 > H > 1$. Low and high values of this shape parameter represent ribbon and tube-like structures respectively.

An example of this classification is shown in Figure 2.25, where high enstrophy structures have been coloured according to shape parameter, allowing easy differentiation between tube and ribbon-like structures. Structures which have a mean alignment below the threshold are also included, but these are plotted using a black wire mesh.
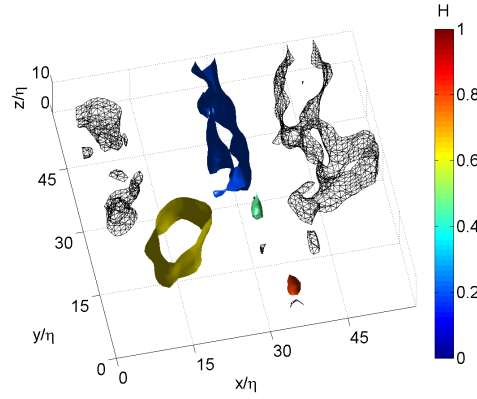
Figure 2.25: Example of box-counting structure classification

2.4.4.3 Box-Counting Method: Cluster Analysis

In order to quantify structure clustering a box-counting methodology similar to Moisy & Jiménez (79) was applied, although differences in the present data set require a slightly different approach to be taken.

The method is based on the clustering of structure mass centres, which were calculated during the geometry classification (see Section 2.4.4.2). Similarly to Moisy & Jiménez, the domain is divided into boxes of varying size, and the number of boxes containing one or more structures are counted. For simplicity a box size of 2^n was used, where the integer n takes a value between 0 and 5. However, the non-cubic domain size ($51 \times 51 \times 8$ grid points) means that cubic

boxes can only be used up to a size of $r_x = r_y = r_z = 8$. For larger boxes the size in the z direction must remain fixed at $r_z = 8$, while the other dimensions are increased.

These box counts should be compared to a set of randomly distributed structures, which can be calculated using the Poisson expectation of finding empty boxes given the number of structures and total volume (79). This random box count, N_0, has been modified in the current investigation to take account of the largest non-cubic boxes and data distributed across N_{vol} volumes (see Equation 2.15).

$$N_0 = \left(\frac{N_{vol}V_xV_yV_z}{r_xr_yr_z} \right) \left(1 - exp\left(-\frac{r_xr_yr_z}{v_0} \right) \right) \qquad (2.15)$$

V_x, V_y and V_z are the reduced volume size (limited to a maximum of $32 \times 32 \times 8$) in the x, y, and z direction respectively. v_0 is a measure of mean structure separation evaluated as $v_0 = N_{vol}(V_xV_yV_z)/N_{total}$, where N_{total} is the total number of structures across all volumes.

Chapter 3

Numerical Accuracy Simulation

PIV experimental accuracy depends on a large number of factors and is highly dependent on the flow conditions and experimental set-up (100), making quantification through comparison with previous experiments difficult. Furthermore, there are an insufficient number of TPIV accuracy studies, and the only study to perform full parametric analysis (31) uses a measure of accuracy which cannot easily be translated into a useful quantity (see Section 2.3.3). Although representative in terms of measurement accuracy, the flow used for the simulations detailed in Section 2.3.3 is not representative of the experimental flow field, and the displacement evaluation algorithm is not comparable with the more advanced algorithm used to evaluate the experimental results (see Section 2.2.6).

Therefore, in order to obtain experimental uncertainty estimates for DPIV and TPIV experiments a more representative set of simulations have been conducted, using velocity data from a Homogenous Isotropic Turbulence (HIT) DNS simulation (121) to represent the ideal flow field. The following chapter begins by detailing the ideal flow, computational set-up and flow matching strategy. After which the two techniques are assessed using a wide variety of measures.

3.1 Computational Set-up

The DNS simulation solves the Navier-Stokes equations for a 400^3 grid point domain using a spectral approach, achieving a Taylor microscale Reynolds number

88

of $R_\lambda = 141$ (121). The domain size was $L = 2\pi$ and the ratio of this to the integral length scale was $L/\Lambda = 3.8$. This Reynolds number is similar to the lowest Reynolds number experimental results (see Table 4.1), allowing a comparison to be made at this point through a careful choice of set-up parameters.

Flow spatial similarity is achieved by matching the DNS and experimental non-dimensional ratio of domain size to Kolmogorov microscale, L/η. Different sized regions are sampled to represent a range of spatial resolutions, with these sampled fields stretched to fit the DPIV and TPIV volumes of interest (detailed in Table 3.1). It should be noted that particle velocity averaging will take place over the entire interrogation window size and it is therefore appropriate to match this critical dimension to the DNS resolution, which is $\Delta x/\eta = 1.89$. Therefore, the highest resolution case uses an interrogation window size equal to the DNS resolution but with a 50% overlap, giving a vector spacing of $\Delta x/\eta = 0.93$, which is quoted as the equivalent resolution in Table 3.1.

The DNS non-dimensional RMS velocity was scaled to match the experimental RMS displacement of ~ 5 pixels (or voxels). Matching the velocity in this manner will produce similar particle displacements, permitting a reasonable comparison with the experimental results.

In both DPIV and TPIV cases the simulations were set-up by creating sets of artificial images. In the DPIV simulation, the DNS velocity field slice was mapped into image space using a 3rd order polynomial Taylor series expansion calibration function obtained from the D3 case (see Section 2.1.9). Particle centres were randomly distributed in a volume matching the experimental domain size in image space, assuming a light sheet thickness of 1.5mm, with a seeding density of $N_{ppp} = 0.01$ chosen to match the experimental conditions. The velocity at each particle position was obtained from the DNS velocity field, using bi-cubic interpolation to calculate the velocity at positions between grid nodes. These were used to linearly displace the particles, removing those displaced beyond the volume bounds. The original and displaced particle positions in the x and y directions were then used directly as particle image co-ordinates to create two images, with particles modelled as blobs of Gaussian intensity approximately 3 pixels in diameter.

Case Ref.	DNS slice size	Fields	Equiv. Res. $\Delta x/\eta$
$S-D1$	$33 \times 33 \times 1$	100	0.93
$S-D3$	$103 \times 103 \times 1$	100	2.97
$S-D5$	$171 \times 171 \times 1$	100	4.95
$S-D10$	$339 \times 339 \times 1$	100	9.84
$S-T1$	$39 \times 39 \times 7$	50	0.94
$S-T3$	$103 \times 103 \times 21$	50	2.52
$S-T4$	$171 \times 171 \times 35$	44	4.20
$S-T8$	$339 \times 339 \times 69$	5	8.36

Table 3.1: Simulation case set-up details, Case references abbreviations are: S=Simulation; D=DPIV; T=TPIV

In the TPIV simulation, particles centres were randomly distributed in a volume matching the real space experimental domain size of $60 \times 60 \times 12$mm. The velocity at the particle positions was found using tri-cubic interpolation, which as before, was used to linearly displace the particles whilst removing particles displaced beyond the volume bounds. The original and displaced particle positions were then transformed into image space for 4 cameras using a set of polynomial calibration functions obtained from the TPIV set-up (see Section 2.1.9.2). Projecting the particles into image space in this manner will eliminate the calibration error, which has a significant effect on the measurement uncertainty (31; 138). Therefore, disparity maps obtained from the camera self calibration (see Section 2.2.2) were used to displace the particle image co-ordinates, modelling the camera misalignment error based on the known experimental accuracy. Artificial images were created from these particle image co-ordinates using the same modelling assumptions as the DPIV simulation, with a seeding density of $N_{ppp} = 0.02$.

Vector fields were computed from the artificial images using *Davis 7.2* for the DPIV and and an in-house algorithm for the TPIV simulations, with processing options chosen to match the experimental cases (see Section 2.2.6). Where available 50 and 100 velocity fields were taken at regularly spaced intervals throughout the DNS data cube for the DPIV and TPIV simulations respectively, to produce reasonably well converged statistics and permit analysis of quantities such as the energy spectrum. Due to the size of the largest field slices (see Table 3.1) fewer samples could be taken, and most significantly only 5 fields were used for the

TPIV $S - T8$ case. The resolutions quoted in Table 3.1 refer to the spacing between vectors.

To simulate the effect of camera noise on the PIV accuracy, a 10% level of random noise was added to each set of images. The effect of noise is investigated more fully for the $S - D1$ and $S - T1$ cases, where varying levels of noise ranging from 0 to 50% (of the total image intensity) are added to the images.

The $S - D1$ and $S - T1$ cases are the only simulations in which the ideal and PIV velocity data point are similarly spaced. Therefore, a direct comparison can be made between the ideal and PIV velocity, vorticity and dissipation fields using Equation 3.1. This represents the local error, which is the modulus of the difference between ideal and PIV fields at each point in the domain, summed and normalised by the mean ideal values. The local error is representative of the mean error in the instantaneous fields, which will affect visualisation of the instantaneous distribution of quantities in the domain. The necessary modulus use in this error estimate precludes determination of the error direction (under or over prediction).

The global error is the mean ideal value minus the mean PIV value, again normalised by the mean ideal value (Equation 3.2), where N_{ideal} and N_{PIV} are the number of data points for ideal and PIV fields respectively. Positive or negative values for this error indicate PIV under or over-prediction respectively. The global error is representative of of the mean statistical error, relating to statistical quantities such as the mean RMS velocity and mean dissipation rate. Except for the highest resolution cases, the ideal and PIV data points are at different spatial locations. Therefore, as the effect of spatial averaging over the interrogation region cannot be discounted in the lower resolution cases, while it is still possible to assess the global errors, the local errors are not compared.

$$E_{local} = \frac{\sum \left(|u_{ideal} - u_{PIV}| \right)}{\sum |u_{ideal}|} \times 100 \qquad (3.1)$$

$$E_{global} = \frac{\sum u_{ideal}/N_{ideal} - \sum u_{PIV}/N_{PIV}}{\sum u_{ideal}/N_{ideal}} \times 100 \qquad (3.2)$$

Although representative as far as possible, Reynolds number differences between the ideal and real flow and the different forcing methods may affect the

measurement accuracy. Additionally, these simulations neglect the following effects which may be present during experimentation: particle loss and uneven illumination due to laser sheet non-uniformity, sheet misalignment between pulses, and irregular particle shape and size; any inhomogeneity in the particle distribution; persistent non-random noise due to camera hot-pixels and light diffusion due to multiple scattering effects; use of interpolation to displacement particles, particularly using the coarse ideal data in the high resolution case. Furthermore, the velocity matching employed in the current simulation matches the mean value calculated experimentally, which may itself be under-predicted due to spatial averaging effects, affecting accuracy predictions accordingly.

3.2 Simulation Results

DPIV and TPIV accuracy is assessed using a number of measures, visually comparing local fields of velocity, vorticity and squared strain rate, before examining trends in local and global mean quantities. The squared strain rate, S_{ij}^2, is representative of the dissipation rate, and was chosen over the latter quantity to give a more generally applicable representation of the mean local errors. Given that the velocity is expressed as a displacement in pixels, the squared strain rate is therefore non-dimensional. The one dimensional energy spectrum is also calculated and compared, as are topological invariants and divergence for the 3D TPIV fields.

3.2.1 Visual Comparison of Fields

Figures. 3.1 and 3.2 compare the u velocity component and squared strain rate across a range of spatial resolutions for the DPIV simulation. The same colour map scaling is used for each resolution case allowing easy comparison of their relative magnitudes. Flow details are well predicted in the high resolution $S - D1$ case. However, with reducing spatial resolution the fields show the effect of spatial averaging, resulting in smoothing and a loss of fine scale information. This can be further illustrated by directly averaging the ideal field (taking the mean value within an interrogation window sized region at each interrogation window location

on the ideal field) and comparing it to the PIV results (Figure 3.1(i)-(l)). The effect of this direct spatial averaging is extremely close to the DPIV simulation, demonstrating the smoothing source.

The effect of smoothing on the velocity magnitude is slight, with a small reduction in the peak velocities. However, the effect on the vorticity and dissipation fields is more pronounced, and despite similar spatial distributions, the peak magnitude is increasingly under-predicted as the spatial resolution is reduced. Therefore, despite the ability of high resolution PIV to identify the spatial distribution of important quantities such as dissipation, at low resolution spatial averaging significantly reduces the magnitude of these predictions, and alters structure topology.

(a) $S - D1$ ideal	(b) $S - D3$ ideal	(c) $S - D5$ ideal	(d) $S - D10$ ideal
(e) $S - D1$ DPIV	(f) $S - D3$ DPIV	(g) $S - D5$ DPIV	(h) $S - D10$ DPIV
(i) $S - D1$ SF ideal	(j) $S - D3$ SF ideal	(k) $S - D5$ SF ideal	(l) $S - D10$ SF ideal
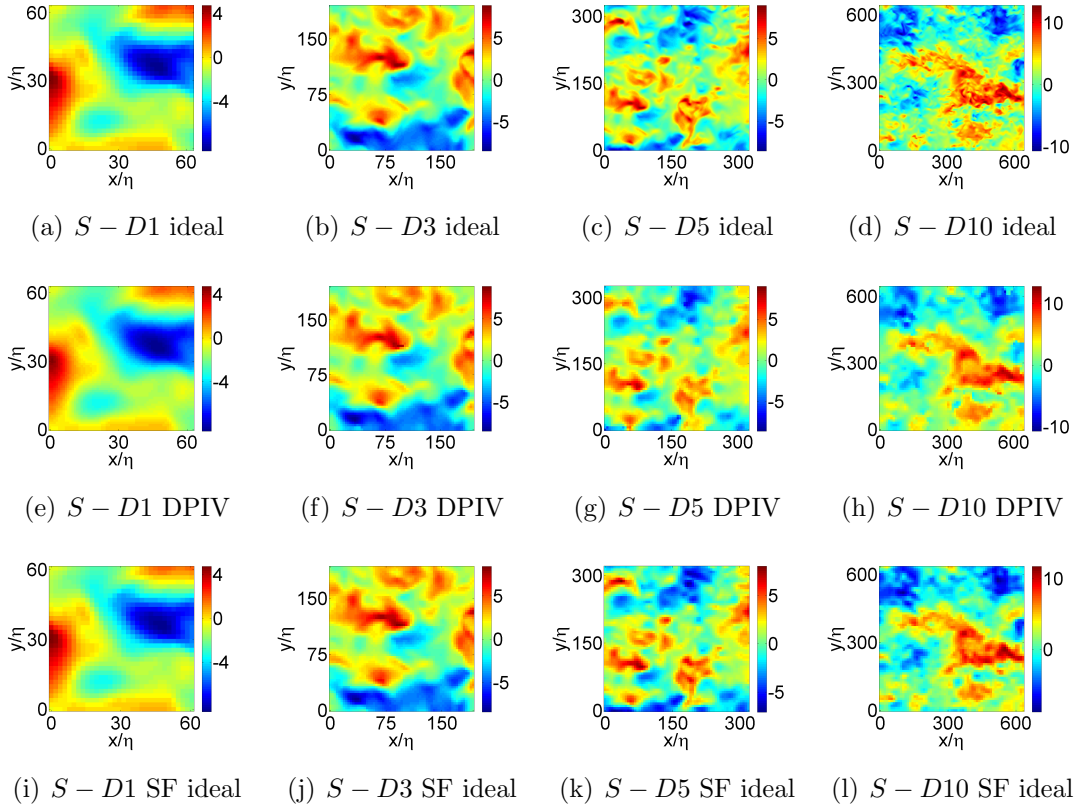
Figure 3.1: DPIV u velocity comparison; displacement in pixels; SF=Spatially filtered

(a) $S - D1$ ideal   (b) $S - D3$ ideal   (c) $S - D5$ ideal   (d) $S - D10$ ideal

(e) $S - D1$ DPIV   (f) $S - D3$ DPIV   (g) $S - D5$ DPIV   (h) $S - D10$ DPIV
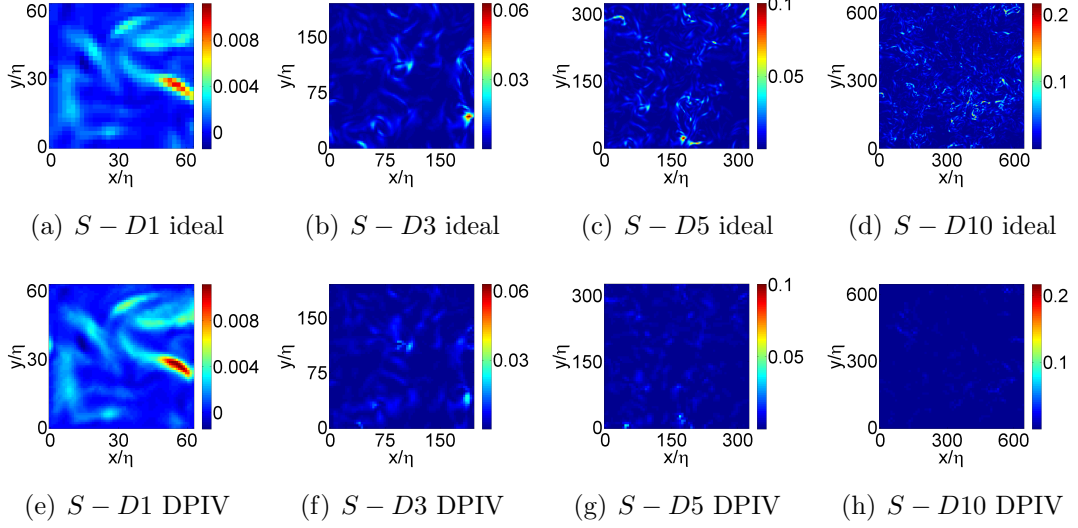
Figure 3.2: DPIV squared strain rate comparison; squared strain rate $(\text{pixels/pixel})^2$

Figure 3.3 shows the velocity, vorticity and squared strain rate error distribution for the $S - D1$ case. The largest errors appear to be concentrated in regions of highest gradient, as a result of correlation peak broadening in these areas (100). However, as the gradient magnitude is higher in these regions the relative gradient error magnitude reduces with increasing gradient magnitude (36). This trend is demonstrated later in Section 3.2.4. The inclusion of fewer points in the finite difference calculation at the edges also results in a loss of accuracy (as discussed in Section 2.4.1), which is remedied by removing these points from the mean local and global estimates.

(a) $S - D1$ u velocity   (b) $S - D1$ v velocity   (c) $S - D1$ vorticity   (d) $S - D1$ dissipation
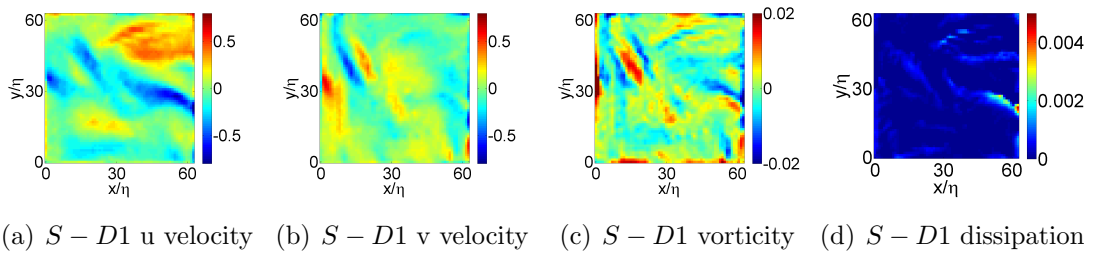
Figure 3.3: Spatial location of errors, absolute error magnitude

Figures 3.4 and 3.5 compare u velocity and squared strain rate slices across a

range of spatial resolutions for the TPIV simulation, with and without Gaussian smoothing. The ideal and TPIV velocity vector locations are not coincident, and therefore slices are taken at the central depth plane of the ideal fields, with cubic interpolation used to obtain values for the TPIV fields at these locations. As the volume edges are trimmed where all camera lines of sight fail to converge, these plots have a smaller domain size in comparison with the DPIV simulation.

Although the main flow features are predicted well the fields are considerably less smooth than the DPIV simulation, as a result of the noise introduced through the TPIV technique (31). This noise will produce significant errors during gradient calculation, and therefore smoothing with a $3 \times 3 \times 3$ Gaussian kernel is recommended for 3D data (36). The kernel is as described in Section 2.4. The smoothed fields show significantly better agreement with the ideal flow, and improved gradient predictions (Figure 3.5). As in the DPIV simulation, the high resolution case predicts flow field magnitude reasonably well, with lower resolution resulting in increasing under-prediction. The under-prediction of gradient quantities is more severe than the DPIV simulation, which may be due to the necessary smoothing.

Figures 3.6(a)-(h) show iso-surfaces of vorticity magnitude above a threshold value of twice the mean DNS data vorticity magnitude (denoted as 'high threshold'). Reducing resolution is again shown to reduce peak vorticity magnitude, resulting in the identification of fewer high vorticity magnitude structures. However, if the known reduction in vorticity magnitude is taken into account and the threshold based instead on the mean TPIV vorticity magnitude (Figure 3.6(i)-(l), with the threshold denoted as 'low threshold'), the same high vorticity structures can be identified at high resolution. As the spatial resolution is reduced spatial averaging effects again remove the finer flow features, with the vorticity distributions in lowest resolution cases bearing little resemblance to the ideal flow distribution.

3.2.2 Local Errors

Figure 3.7 shows mean error variation in the local velocity, velocity gradient, vorticity and dissipation estimates for the $S-D1$ and $S-T1$ cases with increasing

(a) $S - T1$ ideal   (b) $S - T3$ ideal   (c) $S - T4$ ideal   (d) $S - T8$ ideal

(e) $S - T1$ TPIV   (f) $S - T3$ TPIV   (g) $S - T4$ TPIV   (h) $S - T8$ TPIV

(i) $S - T1$ TPIV GS   (j) $S - T3$ TPIV GS   (k) $S - T4$ TPIV GS   (l) $S - T8$ TPIV GS
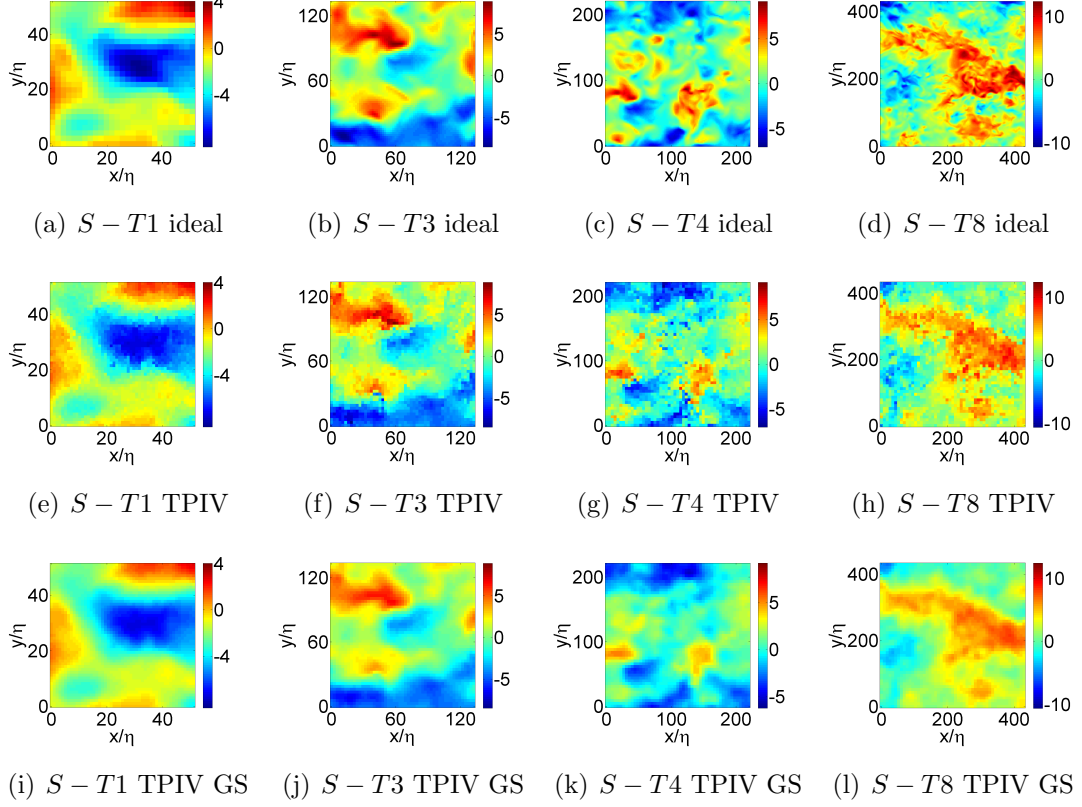
Figure 3.4: TPIV u velocity comparison; displacement in voxels; GS= 2-pass Gaussian smoothing

noise. Figure 3.7(a) shows the absolute velocity error. The accuracy of the DPIV simulations is similar to previous studies (100), taking a value of around 0.1 pixels. Even after smoothing the TPIV accuracy is found to be lower, with a mean error of around 0.2 voxels in the u and v components and 0.3 voxels in the w component, which is again in line with previous experiments (134; 138). Figure 3.7(b) shows the velocity in percentage terms, with a DPIV measurement uncertainty of around $3 \rightarrow 4\%$, and after smoothing around 6% and 11% for the TPIV in-plane and out-of-plane components respectively. The effect of noise is slight below a value of 20%, becoming increasingly significant at higher levels, showing a similar trend to the previous study (see Section 2.3.3). The effect of Gaussian smoothing on the TPIV velocity field is shown to increase accuracy by several percent, and also reduce the measurement uncertainty noise dependence.

(a) $S - T1$ ideal   (b) $S - T3$ ideal   (c) $S - T4$ ideal   (d) $S - T8$ ideal

(e) $S - T1$ TPIV   (f) $S - T3$ TPIV   (g) $S - T4$ TPIV   (h) $S - T8$ TPIV

(i) $S - T1$ TPIV GS   (j) $S - T3$ TPIV GS   (k) $S - T4$ TPIV GS   (l) $S - T8$ TPIV GS
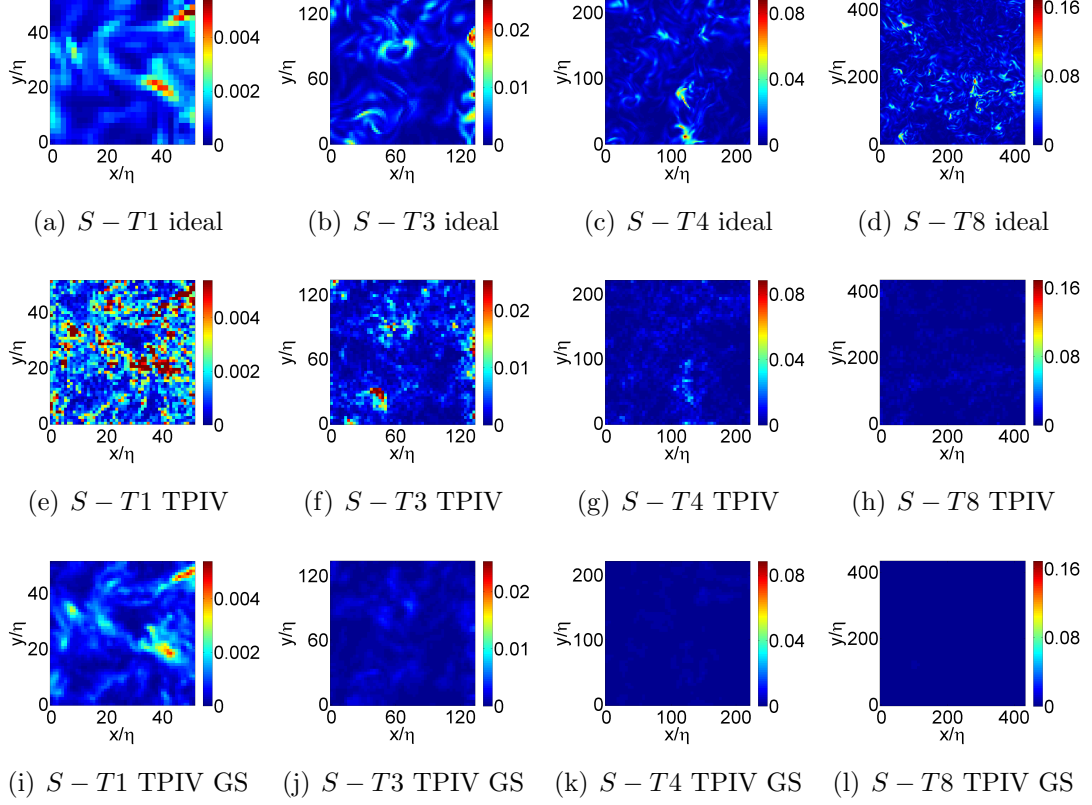
Figure 3.5: TPIV dissipation rate comparison; squared strain rate $(\text{voxels}/\text{voxel})^2$; GS= 2-pass Gaussian smoothing

Figure 3.7(c) shows the measurement uncertainty associated with each gradient component. The TPIV gradient fields are subject to extremely large errors, which can be controlled to a certain degree through smoothing. The DPIV gradient uncertainty is around 20% for reasonable noise levels, which is half that of the smoothed in-plane TPIV gradients. As expected the highest uncertainty is associated with w component velocity gradients, particularly those out-of-plane.

Figures. 3.7(d) and (e) demonstrate gradient error effects on local vorticity and dissipation prediction, producing vorticity errors around 17% and 27%, and dissipation rate errors around 22% and 40% for the DPIV and smoothed TPIV simulations respectively at reasonable noise levels. As with the velocity gradients, smoothing dramatically reduces the measurement uncertainty, and noise only produces a significant accuracy reduction at very high levels. The known ve-

(a) $S - T1$ ideal, high thresh. (b) $S - T3$ ideal, high thresh. (c) $S - T4$ ideal, high thresh. (d) $S - T8$ ideal, high thresh.

(e) $S - T1$ TPIV GS, high thresh. (f) $S - T3$ TPIV GS, high thresh. (g) $S - T4$ TPIV GS, high thresh. (h) $S - T8$ TPIV GS, high thresh.

(i) $S - T1$ TPIV GS, low thresh. (j) $S - T3$ TPIV GS, low thresh. (k) $S - T4$ TPIV GS, low thresh. (l) $S - T8$ TPIV GS, low thresh.
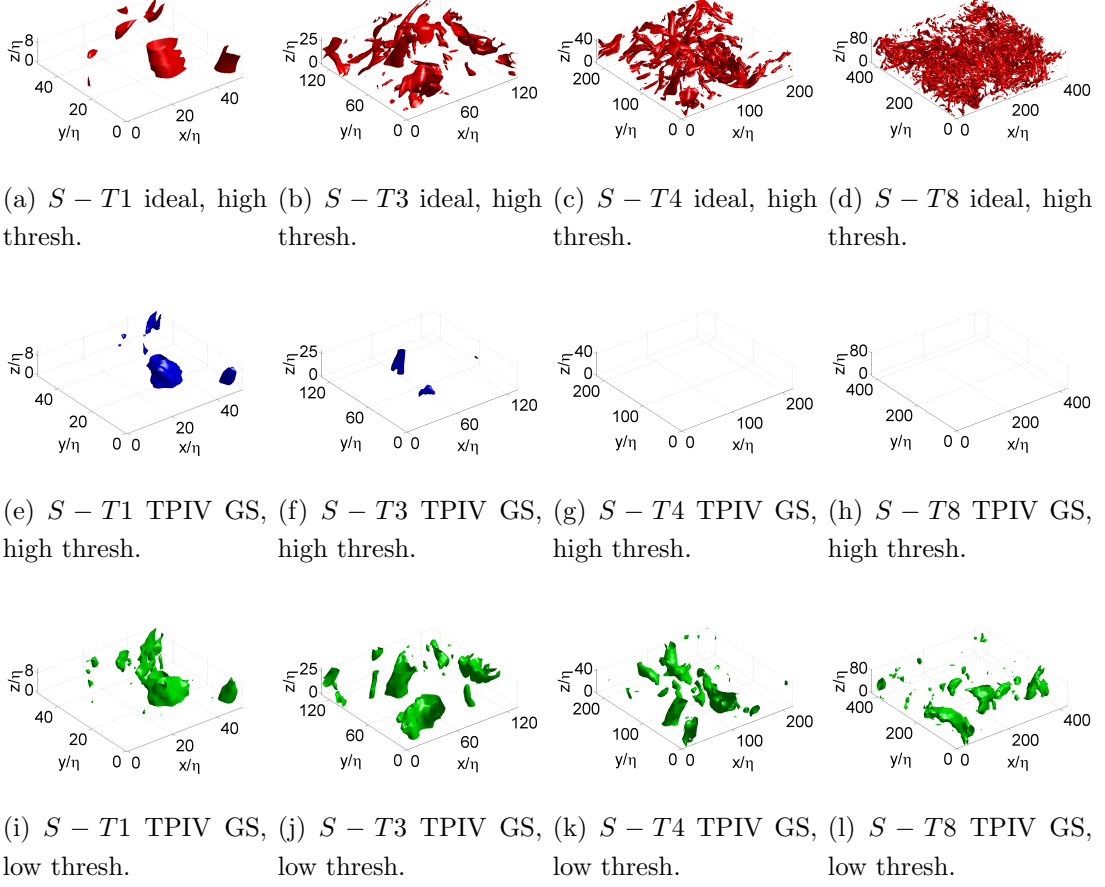
Figure 3.6: TPIV iso-surfaces of vorticity magnitude (voxels/voxel); (a)-(d) $|\omega_{ideal}| > 2\langle|\omega_{ideal}|\rangle$; (e)-(h) $|\omega_{TPIV}| > 2\langle|\omega_{ideal}|\rangle$; (i)-(l) $|\omega_{TPIV}| > 2\langle|\omega_{TPIV}|\rangle$

locity gradient sensitivity to noise and error amplification through use of squared gradient terms produces the high measurement uncertainty, which although substantial, is again in line with previous experiments (36; 134).

3.2.3 Mean Global Errors

Figure 3.8 show the global RMS velocity, velocity gradients and dissipation variation with reducing resolution for the DPIV and TPIV techniques in terms of percentage error magnitude.

Figure 3.8(a) shows that at high resolution, both the DPIV and TPIV results

(a) Velocity (absolute)

(b) Velocity (percentage)

(c) Velocity gradient

(d) Vorticity magntiude

(e) Dissipation

Figure 3.7: Local errors; GS= 2-pass Gaussian smoothing

give good RMS velocity predictions. As resolution is reduced the velocity is increasingly under-estimated, resulting in increasing errors. The DPIV accuracy decrease is slight, showing $\sim 7\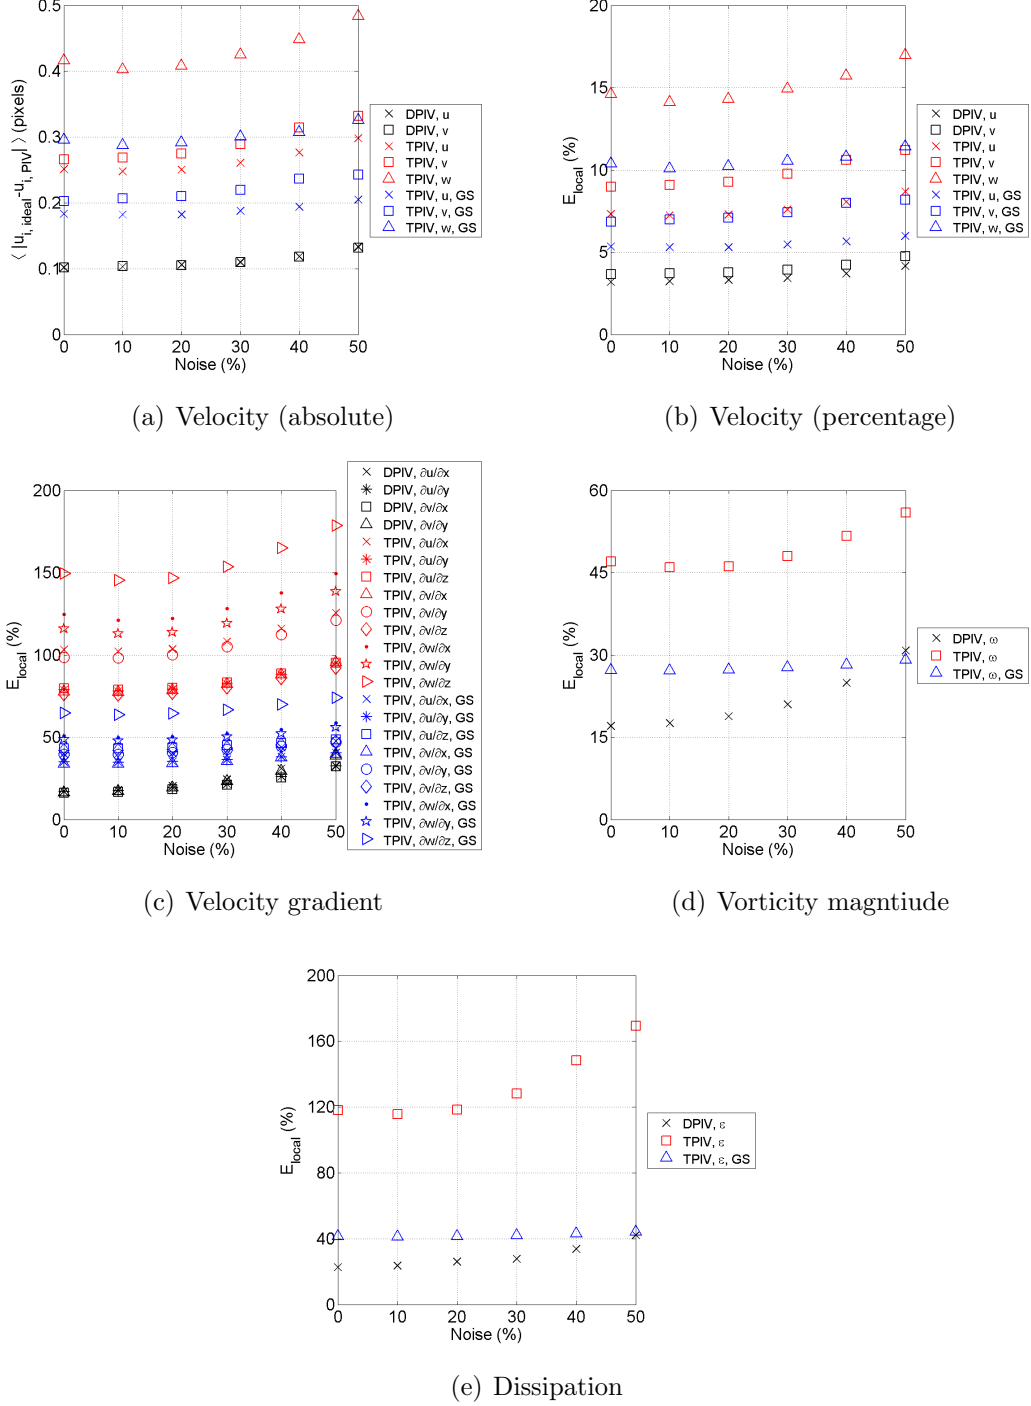%$ error for the $S-D10$ case. The increase in TPIV measurement uncertainty is more significant however, with an error magnitude of $\sim 20\%$ for the smoothed $S-T8$ case. The erroneous over-prediction of velocity in the unsmoothed TPIV data counters the natural under-prediction due to spatial averaging, resulting in lower measurement uncertainty for the unsmoothed results.

This noise related over-prediction error is much more prominent in the unsmoothed field velocity gradients in Figure 3.8(b), resulting in poor accuracy at high measurement resolution. As measurement resolution is reduced the gradient under-prediction due to spatial averaging increases, with the two effects balancing each other in the unsmoothed TPIV results at a resolution of $\Delta x/\eta \sim 3$. Beyond this value the spatial averaging error dominates, resulting in increasing gradient under-prediction. Despite the apparent low error magnitude resulting from this balance, it is not clear that this relationship will hold, and as such the smoothed results are preferable. However, it is unlikely that all noise related gradient over-prediction can be eliminated through smoothing, and therefore care should also be exercised in the interpretation of smoothed dissipation estimates which may also be subject to a balance of error sources.

At the lowest resolution the DPIV and smoothed TPIV results produce gradient errors of around 80% for the lowest resolution cases. The effect of these on the dissipation estimate are shown in Figure 3.8(c), producing errors of up to 90%. This analysis confirms the previous observations (see Sect. 3.2.1) that at low resolution the effects of spatial filtering remove the smallest scale high gradient regions responsible for the majority of the dissipation, resulting in severe under-prediction.

3.2.4 Divergence

Another stringent accuracy test can be conducted by assessing velocity gradient variation from the incompressible flow zero divergence condition ($\nabla \cdot u_i = 0$). Figure 3.9(a) shows the variation in mean normalised divergence ratio, ξ (142) (defined in Equation 3.3), with measurement resolution for the TPIV simulation.

(a) Velocity

(b) Velocity gradient

(c) Dissipation
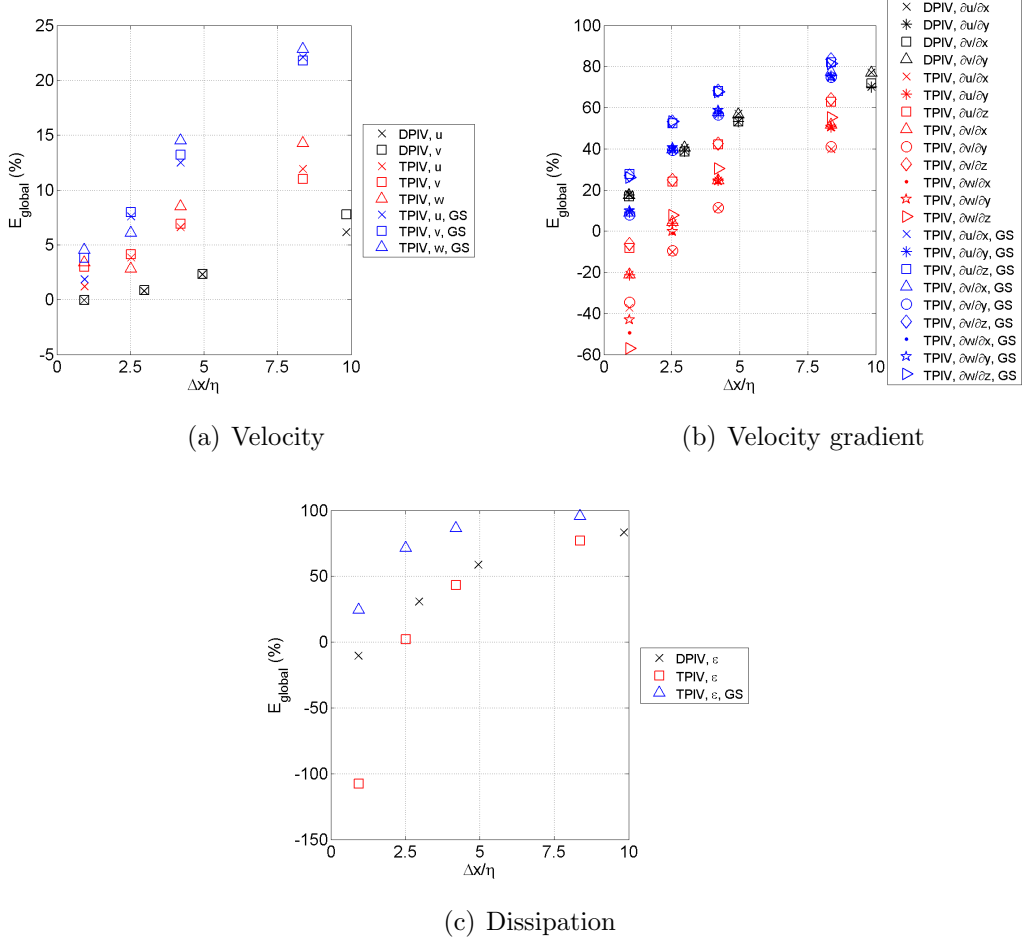
Figure 3.8: Global percentage errors; GS= 2-pass Gaussian smoothing

In comparison with the ideal data (zero divergence) the simulation divergence is high. However, as this accuracy criterion is based on squared velocity gradients it is very stringent, with the unsmoothed divergence values similar to the full resolution value of $\xi = 0.74$ obtained by Zhang et al. (142). Smoothing is shown to significantly reduce the divergence.

$$\xi = \frac{(\partial u/\partial x + \partial v/\partial y + \partial w/\partial z)^2}{(\partial u/\partial x)^2 + (\partial v/\partial y)^2 + (\partial w/\partial z)^2} \tag{3.3}$$

The divergence was also quantified by Mullin et al. (82) for dual plane Stereo-PIV measurements using the local divergence value normalised by the local velocity gradient tensor norm $\nabla \cdot u_i / (\nabla u_i : \nabla u_i)^{1/2}$. The normalised divergence was

101

(a) Mean normalised divergence error

(b) Normalised Divergence PDF

Figure 3.9: Divergence Analysis

found to be normally distributed with a mean of zero, and an RMS variation of 0.35, which is similar to that found in the present investigation after smoothing (see Figure 3.9(b)). The variation in the $S -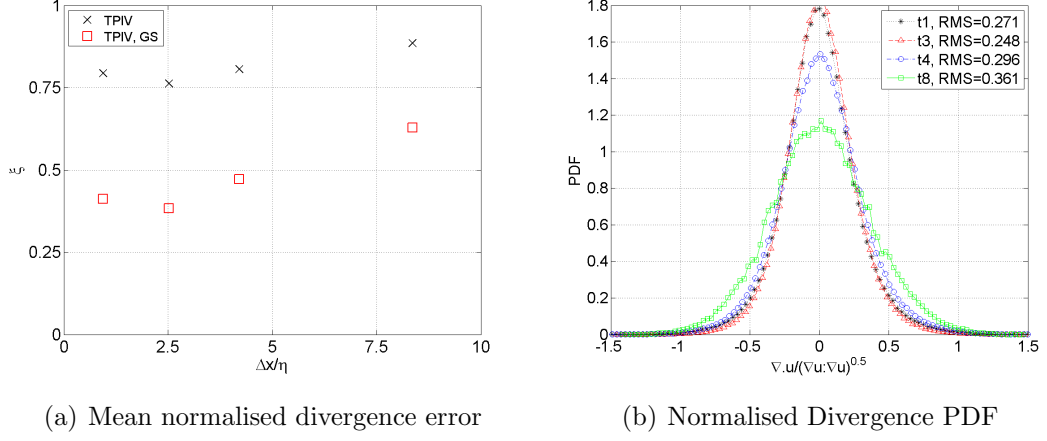 T8$ case is caused by an insufficient quantity of data resulting in poor statistical convergence; making it difficult to draw conclusions from this case.

As PIV error is a function of gradient magnitude, JPDFs were constructed to correlate the normalised divergence values with the velocity gradient magnitude at each point, as shown in Figure 3.10. The ideal data has been filtered using a kernel width equal to the interrogation window size to demonstrate the increase in divergence due to the effects of spatial averaging. The JPDFs show that regions of high gradient magnitude, which may be associated with high vorticity magnitude coherent structures, have lower divergence errors and therefore lower relative uncertainty. The distribution is similar to that presented by Ganapathisubramani et al. (36) for Stereo-PIV measurements of a turbulent jet. As measurement resolution is reduced the characteristic triangular distribution loses definition, becoming increasingly circular as a result of high gradient magnitude region smoothing; a trend which can be independently observed in the ideal data. However, comparison with the ideal data suggests that the divergence error is dominated not by spatial averaging effects but rather through noise associated with TPIV.

(a) $S - T1$ SF ideal   (b) $S - T3$ SF ideal   (c) $S - T4$ SF ideal

(d) $S - T1$ TPIV   (e) $S - T3$ TPIV   (f) $S - T4$ TPIV

(g) $S - T1$ TPIV GS   (h) $S - T3$ TPIV GS   (i) $S - T4$ TPIV GS

Figure 3.10: TPIV gradient and divergence magnitude JPDF, contours of base 10 exponentials; SF=Spatially filtered; GS=2-pass Gaussian Smoothing

Another way of expressing the divergence error is through JPDFs of $\partial u / \partial x$ against $-(\partial v / \partial y + \partial w / \partial z)$, in which the degree of divergence is represented by the scatter of data away from the zero divergence line where the components are equal, shown in Figure 3.11. The divergence is quantified through use of a correlation coefficient, Q, between the components. Reducing resolution manifests in the spatially filtered ideal data as an increase in scattering from the zero divergence line, reducing the almost perfect correlation value of $Q = 0.99$ for the $S - T1$ ideal data to $Q = 0.95$ for the $S - T4$ ideal case. However, TPIV noise

is shown to dominate, with a much lower correlation value of $Q = 0.66$ for the smoothed $S - T1$ case. This is lower than the value of 0.82 calculated by Ganapathisubramani et al. ([36]) and the value of 0.7 calculated from the multi-probe hot-wire results of Tsinober et al. ([123]), demonstrating the reduction in accuracy in comparison with hot-wire and stereo-PIV techniques. Noise is controllable using increased smoothing, with the twice smoothed TPIV simulation predicting a similar correlation coefficient to Tsinober et al. ([123]).

(a) $S - T1$ SF ideal   (b) $S - T3$ SF ideal   (c) $S - T4$ SF ideal

(d) $S - T1$ TPIV   (e) $S - T3$ TPIV   (f) $S - T4$ TPIV
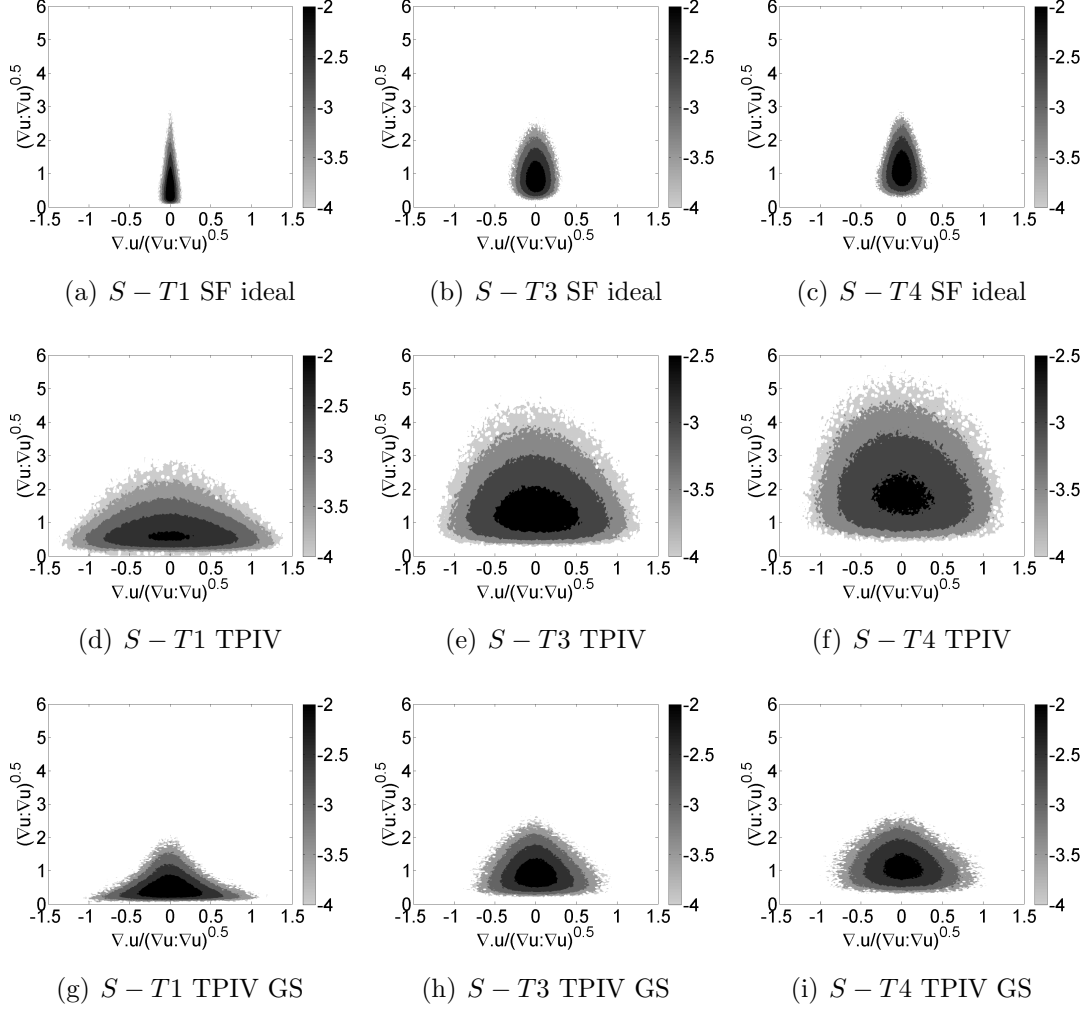
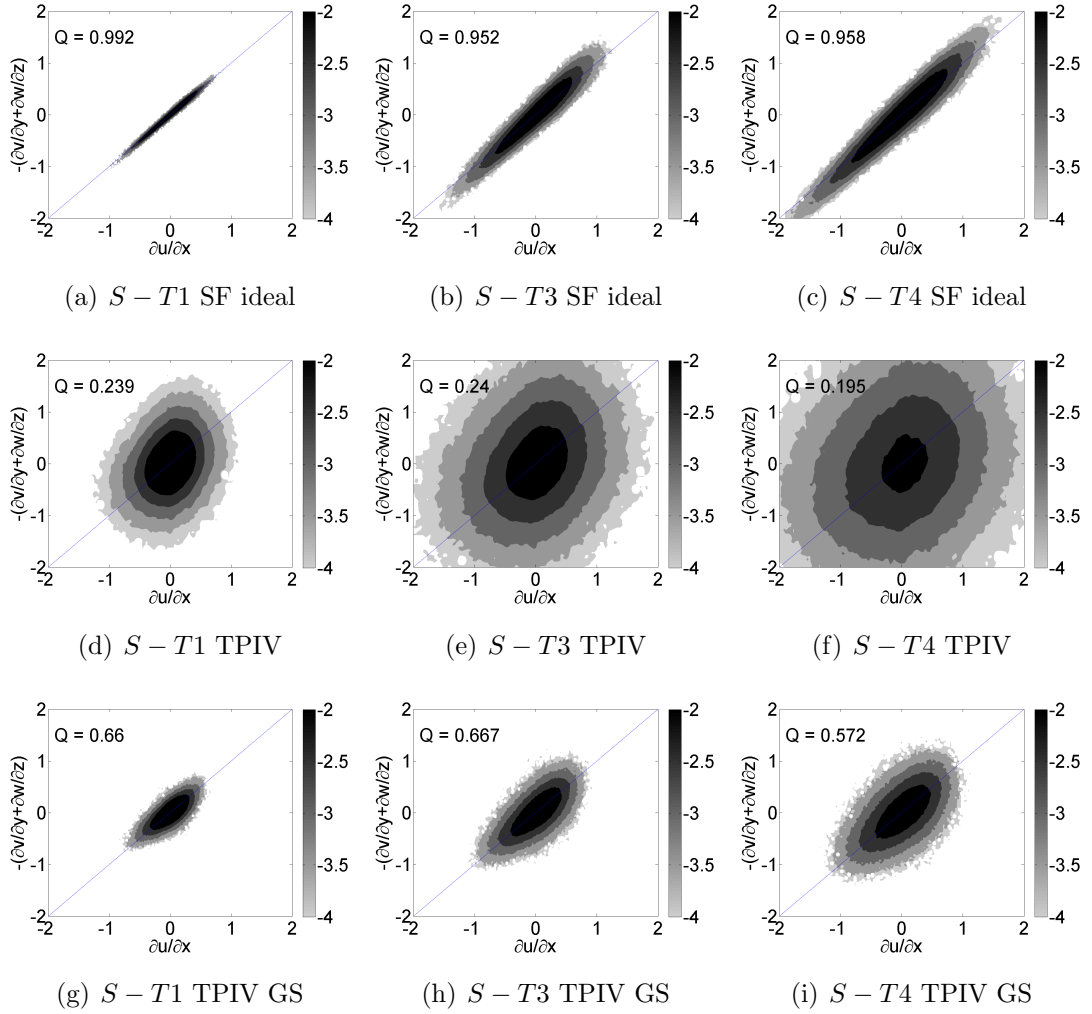(g) $S - T1$ TPIV GS   (h) $S - T3$ TPIV GS   (i) $S - T4$ TPIV GS

Figure 3.11: TPIV divergence correlation JPDF, contours of base 10 exponentials; SF=Spatially filtered; GS=2-pass Gaussian Smoothing

3.2.5 Effect on Energy Spectrum

The 1D energy spectrum can be directly calculated from the ideal and PIV data using the method described in Section 2.4.3.

Figure 3.12(a) shows the effect of aliasing due to a reduction in resolution for the $S - T8$ case with and without windowing function application. Aliasing limits the simulation data to lower wavenumbers and results in a loss of energy, shown by the lower magnitude truncated spectra. The Hann windowing function redistributes energy throughout the spectrum, causing a shift from high to low wavenumbers, and increasing the spectrum gradient. Although this improves the agreement of the ideal data with the Pao spectrum[1] (88) around the dissipation region, it causes the already aliased simulation data to diverge further from this line.

The effect of aliasing on Kolmogorov power law scaling was examined by studying the pre-multiplied spectra shown in Figure 3.12(b). Despite a reasonably high degree of fluctuation in the ideal data, which may be due to insufficient low wavenumber modes (similar fluctuations can be observed in Ishihara et al. (49)), the brief plateau indicates Kolmogorov inertial range scaling. However, energy losses through aliasing in the simulation result in a small persistent negative gradient, the magnitude of which increases rapidly at the lowest wavenumber.

Figures. 3.13(a) and (b) show the effect of increasing noise for the $S - D1$ and $S - T1$ cases with and without the use of the windowing function. Without a windowing function increasing noise appears insignificant, with both ideal and PIV spectra peeling away from the Pao curve at high wavenumbers due to severe spatial truncation effects, as shown by Foucaut et al. (34).

However, contrary to the low resolution case, Figure 3.13(b) shows that windowing function application appears to improve high wavenumber accuracy, increasing the agreement with the Pao curve. Controlling the spatial truncation error through the windowing function allows image noise to dominate the dissipation region accuracy, with higher noise levels causing the spectrum to peel

---

[1]The Pao spectrum is a one-dimensional turbulent energy spectrum for the entire universal equilibrium range of wavenumbers, deduced from the spectral transfer of kinetic energy at large wavenumbers (88). The spectrum has been shown to compare favourably with experimental measurements (88; 104).

(a) Effect of Hann windowing

(b) Pre-multiplied energy spectrum
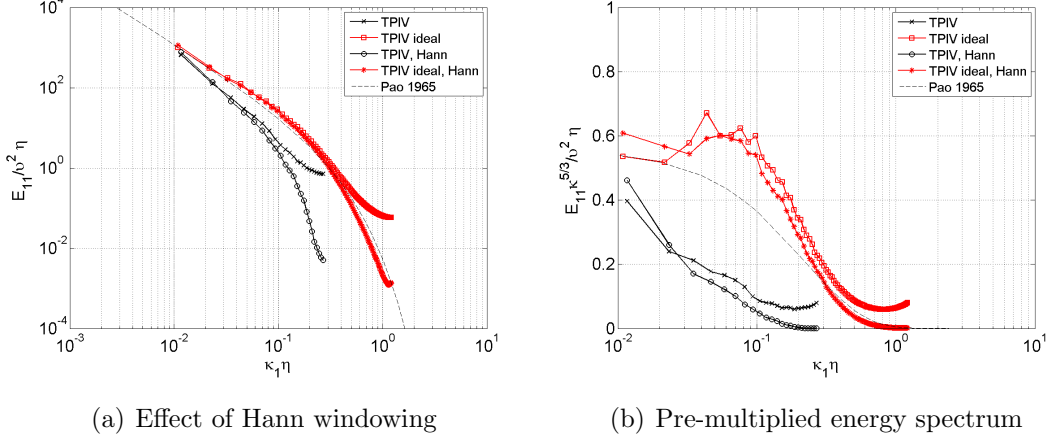
Figure 3.12: Effect of low resolution and windowing functions on the energy spectra, $S-T8$ case

further from the Pao curve; an effect observed by a number of previous studies (34; 48; 130). The higher TPIV noise level is controlled through smoothing, bringing these results in-line with the DPIV results. At the lowest wavenumber the dramatic energy increase through windowing function re-distribution also causes divergence from the Pao result, with similar effects shown by de Jong et al. (24).
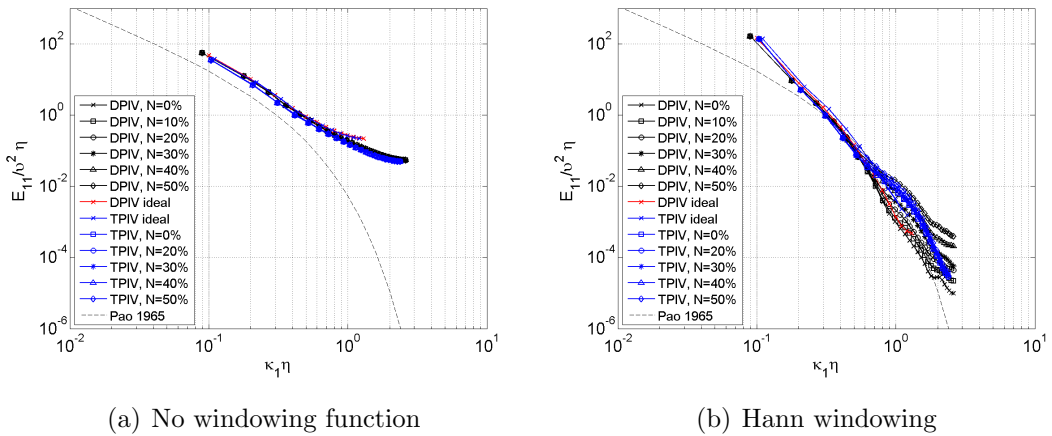

(a) No windowing function

(b) Hann windowing

Figure 3.13: Effect of noise and windowing functions on the energy spectra, $S-T1$ case (N= Noise level)

3.2.6 Topological Analysis

As described by Chong et al. ([22](#)) the flow can be characterised topologically using invariants of the symmetric parts of the velocity gradient tensor. Figure [3.14](#) shows JPDFs of two invariant quantities Q_A^* and R_A^*, normalised by the mean second invariant of vorticity, $\langle \overline{Q_W} \rangle$, as $Q_A^* = Q_A / \langle \overline{Q_W} \rangle$ and $R_A^* = R_A / \langle \overline{Q_W} \rangle^{3/2}$; a plot which has been shown to take a characteristic teardrop shape for a number of DNS studies ([84](#); [85](#)). This distinctive distribution is clearly shown in the ideal data, which becomes fuller as larger volumes are sampled in cases corresponding to lower resolutions. The simulation ability to reproduce this distribution is again highly noise sensitive, requiring smoothing to reduce this sensitivity and improve gradient accuracy (shown in Figure [3.14](#)). Severe noise effects in the unsmoothed results removes much of the distribution definition, leaving it almost circular. The effect spatial averaging due to reducing measurement resolution manifests as gradient magnitude reductions, resulting in smaller lower magnitude distributions.

3.3 Summary of Simulation Results

Despite the close matching of experimental and simulation details, differences between the ideal and experimental Reynolds number and flow fields, and the use of modelling assumptions may affect these accuracy predictions. Due to a number of unconsidered error sources and the possible under-prediction of velocity the error estimates presented in this chapter may be slightly low, with these extra error sources further increasing measurement uncertainty.

At high resolution the ability of both DPIV and TPIV to accurately predict the velocity and gradient fields is demonstrated. Although measurement uncertainty increases in TPIV due to additional noise introduced during tomographic reconstruction, which manifests particularly in sensitive quantities such as the divergence and topological invariants, this can be partially controlled with velocity field smoothing.

The effect of reducing measurement resolution is shown to act in a similar way to spatial averaging over interrogation window sized regions, significantly

Figure 3.14: JPDF of normalised topological invariants Q_A^* and R_A^*, contours of base 10 exponentials

reducing peak gradient magnitude and resulting in increasing under-prediction of important quantities such as dissipation and vorticity. This under-prediction is particularly severe for measurement resolutions lower than $\Delta x / \eta = 3$, suggesting this represents a minimum requirement for estimating these quantities and studying their spatial distribution, with higher resolution desirable.

Direct energy spectrum calculation from PIV data is subject to aliasing and spatial truncation effects, reducing agreement with the expected trend. Furthermore, the use of windowing functions to reduce truncation effects appears to

be non-trivial, with accuracy dependent on the degree of aliasing caused by the spatial resolution.

The quantification of measurement uncertainty in this manner, through the careful matching of experimental conditions and use of a representative turbulent velocity field, has enabled accurate DPIV and TPIV error estimation. These estimates were used in the following chapters, during experimental results analysis.

Chapter 4

Flow Characterisation

Before investigating the nature of the coherent flow structures it is necessary to characterise the flow in terms of both statistical turbulent properties and typical length and time scales, including the variation of these with Reynolds number. It is also important to quantify the level of statistical convergence, relating the variation between the different cases and runs to this and the expected accuracy.

Therefore, the chapter begins with an assessment of the statistical convergence, centring, and resolution, with these used to quantify the measurement uncertainty. The fluctuating velocities and gradients are assessed and used to determine the flow isotropy, after which the flow structure is statistically characterised. Next a variety of dissipation rate calculation methods are compared before using two of these methods, based on scaling argument and structure function methods, to derive other important quantities such as Taylor and Kolmogorov microscales. Finally the chapter concludes with an investigation into the TPIV divergence error, its effect on topological invariants, and the use of a divergence correction scheme. These results provide a basis for the topological investigation in the following chapter.

Case	$\Omega_I R_I$ (mm/s)	Re	R_λ	u_{rms} (mm/s)	v_{rms} (mm/s)	w_{rms} (mm/s)	λ (mm)	η (mm)	ϵ_m (m^2/s^3)	τ (ms)	$\nu \times 10^6$ (m^2/s)
$D1-1$	56	3.77×10^4	240	17.4	10.7	-	16.4	0.536	2.01×10^{-5}	243	1.18
$D1-2$	112	7.55×10^4	334	33.6	21.9	-	11.8	0.327	1.45×10^{-4}	90.5	1.18
$D1-3$	223	1.51×10^5	479	69.1	41.8	-	8.2	0.191	1.26×10^{-3}	30.7	1.18
$D1-4$	335	2.26×10^5	581	102	64.5	-	6.76	0.143	4.01×10^{-3}	17.2	1.18
$D1-5$	447	2.94×10^5	651	131	86.6	-	6.03	0.120	8.67×10^{-3}	11.8	1.22
$D2-1$	56	3.67×10^4	236	17.3	10.7	-	16.6	0.550	1.97×10^{-5}	248	1.22
$D2-2$	112	7.35×10^4	332	34.2	22.2	-	11.8	0.330	1.52×10^{-4}	89.3	1.22
$D2-3$	223	1.47×10^5	432	57.7	43.5	-	9.09	0.222	7.36×10^{-4}	40.7	1.22
$D2-4$	335	2.20×10^5	584	105	62.5	-	6.73	0.142	4.48×10^{-3}	16.5	1.22
$D2-5$	447	2.94×10^5	633	124	84.7	-	6.2	0.125	7.30×10^{-3}	12.9	1.22
$D3-1$	28	1.91×10^4	174	9.04	5.61	-	22.5	0.867	2.82×10^{-6}	643	1.17
$D3-2$	56	3.82×10^4	248	18.3	11.7	-	15.8	0.511	2.34×10^{-5}	223	1.17
$D3-3$	112	7.65×10^4	356	37.8	23.9	-	11	0.296	2.06×10^{-4}	75.2	1.17
$D3-4$	223	1.53×10^5	478	68.1	47.5	-	8.21	0.191	1.21×10^{-3}	31.1	1.17
$D3-5$	335	2.29×10^5	628	117	69.8	-	6.25	0.127	6.18×10^{-3}	13.7	1.17
$D4-1$	112	8.06×10^4	335	31.7	20.8	-	11.7	0.326	1.21×10^{-4}	95.6	1.11
$D4-2$	223	1.61×10^5	501	70.8	43.6	-	7.84	0.178	1.36×10^{-3}	28.6	1.11
$D4-3$	335	2.42×10^5	548	84.7	65.2	-	7.17	0.156	2.32×10^{-3}	21.9	1.11
$D4-4$	447	3.22×10^5	700	138	87.2	-	5.61	0.108	1.01×10^{-2}	10.5	1.11
$T1-1$	28	1.86×10^4	162	7.98	5.4	7.71	24.3	0.972	1.94×10^{-6}	787	1.2
$T1-2$	56	3.72×10^4	224	15.3	10.2	15.6	17.5	0.595	1.38×10^{-5}	295	1.2
$T1-3$	112	7.45×10^4	323	31.8	20.6	30.8	12.2	0.344	1.23×10^{-4}	98.7	1.2
$T1-4$	223	1.49×10^5	458	64.1	41.2	58.5	8.58	0.204	1.00×10^{-3}	34.6	1.2
$T1-5$	335	2.23×10^5	555	94.1	64.7	99.0	7.08	0.153	3.18×10^{-3}	19.4	1.2

Table 4.1: Typical Flow Properties

4.1 Mean Flow Properties, Accuracy and Convergence

This section begins with a mean flow parameter summary, which provides a useful reference for the following analysis. This is followed by an investigation of mean flow and mean spatial distributions of important turbulent quantities. Assessments of flow convergence and measurement resolution are also conducted, with measurement uncertainty quantified through these measures.

4.1.1 Summary of Results

Some typical flow properties are listed in Table 4.1. The dissipation rates, Kolmogorov length and time scales and Taylor microscales listed here were calculated using the scaling argument method (as defined in Equation 1.1, using a constant value of $A = 0.5$). A full comparison of dissipation estimation methods presented later in sections 4.4.2 and 4.4.3.

4.1.2 Mean Flow and Convergence

As detailed in Section 2.1.1 the large scale flow structure at the tank centre is a combination of confined toroidal shear and poloidal pumping; producing a highly unsteady flow with a turnover time proportional to the impeller rotation rate. The mean flow averaged over 1000 fields is shown for a single case in Figure 4.1.
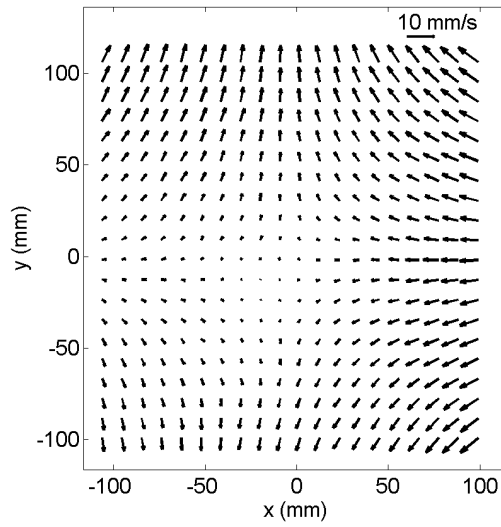

Figure 4.1: Mean flow field, $D1 - 1$, every 3rd vector

Although care was taken to ensure all measurements are performed at the exact geometric centre of the tank, fractional impeller misalignment, camera and light sheet positioning inaccuracies, and a lack convergence at the large scales may still result in discrepancies between experiments at different Reynolds numbers and from different set-ups, which are now quantified.

Statistical convergence is assessed by studying cumulative RMS velocities and mean dissipation rates normalised by their final values and plotted against impeller rotations, as shown in Figure 4.2. In general the level of convergence after 80 impeller rotations is high, with Figure 4.2(a) showing that u_{rms} is converged within $\sim \pm 10\%$ after 20 rotations and $\sim \pm 5\%$ after 60 rotations, which is around twice the v_{rms} variation shown in Figure 4.2(b). Figure 4.2(c) shows the dissipation rate (calculated directly from the strain rate tensor as described in Section 4.4.2) is converged within $\sim \pm 2.5\%$ after 60 impeller rotations. The

velocity and dissipation convergence is indicative of the large and small scale convergence, demonstrating the higher level of the latter, especially in comparison with the u velocity component.

(a) u_{rms} convergence

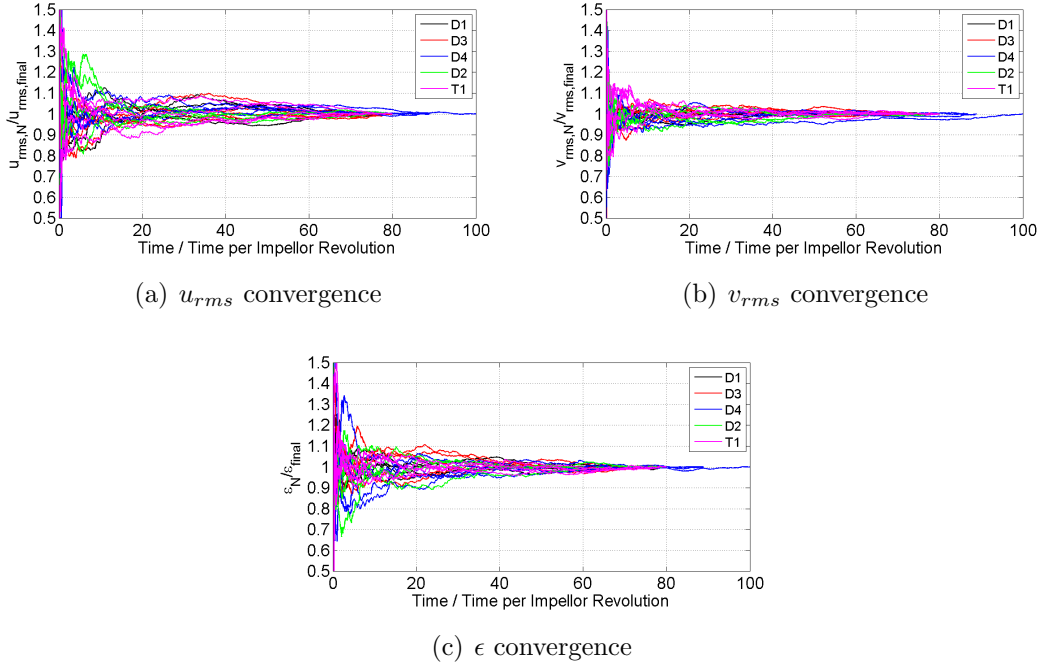(b) v_{rms} convergence

(c) ϵ convergence

Figure 4.2: Normalised convergence plots

Flow convergence within the field of view will not take into account flow centring discrepancies resulting from impeller rotation speed and camera positioning effects, which may cause additional discrepancies between runs and cases respectively. The difference between the flow centres is quantified by plotting stagnation point location, shown in Figure 4.3, with domain size representations for each experimental case.

The significant scatter demonstrates a lack of convergence, with points up to 100mm and 50mm from the flow centre in horizontal and vertical directions respectively. Results from similar cases fail to show any strong Reynolds number trends, indicating that the lack of convergence may dominate over stagnation point movement due to imperfect impeller alignment or velocity mismatch. The inability to separate out convergence effects make evaluation of the other error components difficult.
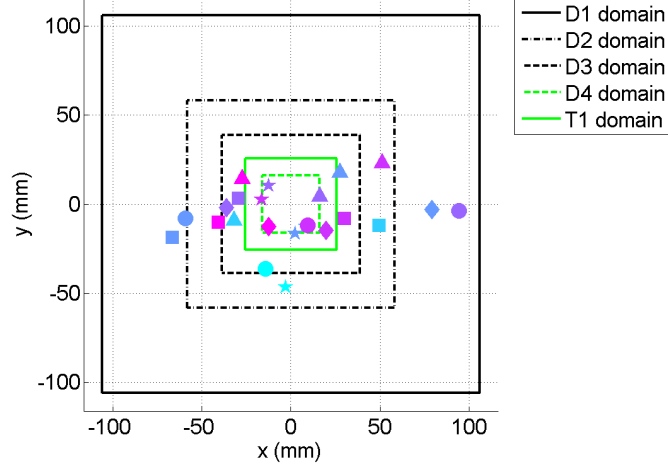
Figure 4.3: Location of stagnation points: $\triangle = D1$; $\square = D2$; $\bigcirc = D3$; $\diamondsuit = D4$; $\bigstar = T1$; colors in order of increasing Reynolds number from light blue to pink

The effect of this centring is dependent on flow inhomogeneity, which was assessed by studying spatial contours of the ensemble averaged RMS velocity normalised its spatial mean, $\sqrt{\langle u^2 \rangle}/\overline{\sqrt{\langle u^2 \rangle}}$, shown in Figure 4.4(a). The $D1$ case was chosen to demonstrate this effect over a large domain. The normalised RMS velocity variation of $\sim 0.9 - 1.1$ demonstrates reasonably high flow homogeneity, similar to that shown by Hwang & Eaton (48) in an acoustically driven near isotropic homogeneous turbulent flow.

Examination of Figures 4.1 and 4.4(a) indicates that the u component of fluctuating velocity magnitude is a function of the mean flow, increasing away from the central stagnation point. Therefore, relating this variation to the flow homogeneity suggests the u and v components of RMS velocity may vary by up to 10% and 5% respectively, taking both flow convergence and centring effects into account.

4.1.3   Experimental Resolution

The spatial resolution was calculated as the interrogation window spacing normalised by the Kolmogorov length scale (calculated from the scaling argument
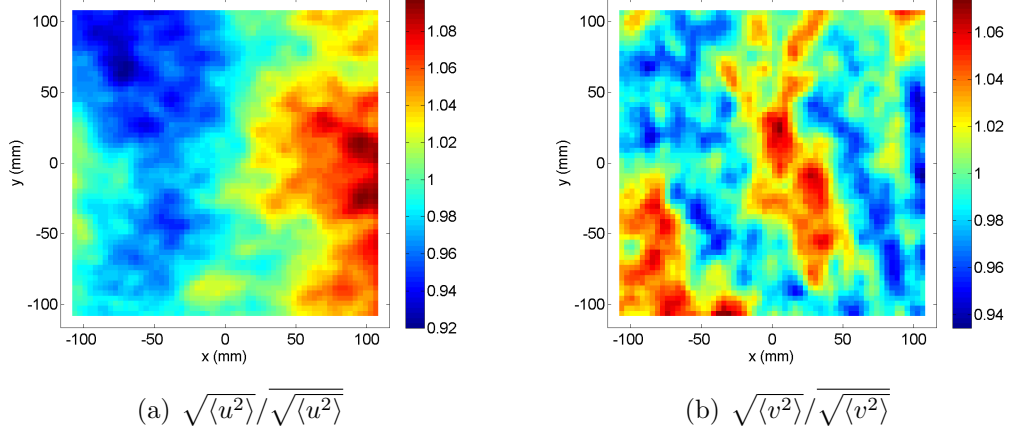
(a) $\sqrt{\langle u^2 \rangle}/\overline{\sqrt{\langle u^2 \rangle}}$ (b) $\sqrt{\langle v^2 \rangle}/\overline{\sqrt{\langle v^2 \rangle}}$

Figure 4.4: Spatial variation in ensemble averaged normalised fluctuating velocity, $D1-1$

method), which was plotted against Reynolds number in Figure 4.5. The Kolmogorov length scale decreases with increasing Reynolds number resulting in a reduction in resolution. However, for the lower Reynolds number runs in the $D3$, $D4$ and $T1$ cases, a resolution $\Delta x/\eta \lesssim 3$ is achieved, which as discussed in Section 3.3 should be suitable for analysis of the flow structure and gradient quantities.

These resolution values will be used in conjunction with the convergence estimate shown in Section 4.1.2 to quantify the measurement uncertainty for each case. A value of measurement uncertainty is interpolated from the global error trends shown in Section 3.2.3 at each resolution for each quantity of interest. This is then added to the estimated convergence error to produce total error estimates, which are plotted as error bars in the following analysis.

4.2 Statistical Flow Properties

Having estimated the measurement uncertainty, the statistical flow properties can now be investigated. This section includes quantification of local and global fluctuating velocities and gradients, including ratios of global properties which are used to assess the level of isotropy at integral and dissipation scales.
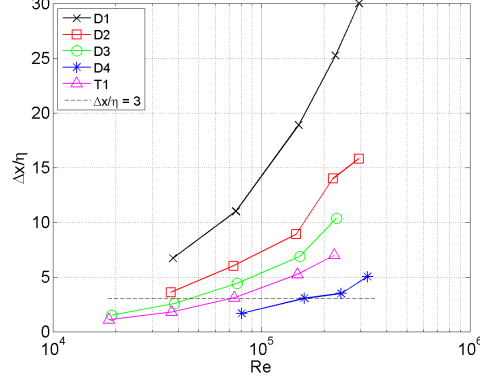
Figure 4.5: Spatial resolution

4.2.1 Fluctuating Velocities

Figure 4.6 shows the probability density functions for the 3 components of fluctuating velocity normalised by impeller rotation velocity, which are fitted to zero mean Gaussian distributions using a least squares approach.

Rotational symmetry similarities can be identified between Figures 4.6(a) and 4.6(c). The mean strain field generated by the poloidal mode has the effect of amplifying the axial vorticity (126) and increasing radial and azimuthal turbulent fluctuations. This effect can be seen by comparing the standard deviations in the radial directions, $\sigma_x = 0.314$ and $\sigma_z = 0.298$, to that of the axial direction, $\sigma_y = 0.189$, producing a ratio ~ 0.6. This result is in broad agreement with Voth et al. (126) who found a ratio of 0.5 between radial and axial standard deviations.

In general the data collapse is reasonable, with the highest deviation seen in Figure 4.6(a) for the central region of the u velocity PDF, which may be due to both poor large scale convergence (Section 4.1.2) and peak-locking effects. The improved collapse shown in Figure 4.6(b) may be due to the reduced level of turbulent fluctuation, and increased convergence in this direction. The w velocity PDF exhibits slightly non-Gaussian behaviour at low velocity, which may be attributed to the lower accuracy of this velocity component measurement and the more significant effect of peak-locking on the TPIV data. The accuracy of the out-of-plane component is lower than the in-plane components due to the camera configuration (see Section 2.1.9.2). Set-up restrictions resulted in lower camera

angles than the optimum $40 - 50°$, which causes the tomographic reconstruction of particles to become stretched in the out-of-plane direction, causing a loss of accuracy during cross-correlation. The TPIV data also contains fewer higher magnitude velocity events, which may be due to the necessary smoothing.
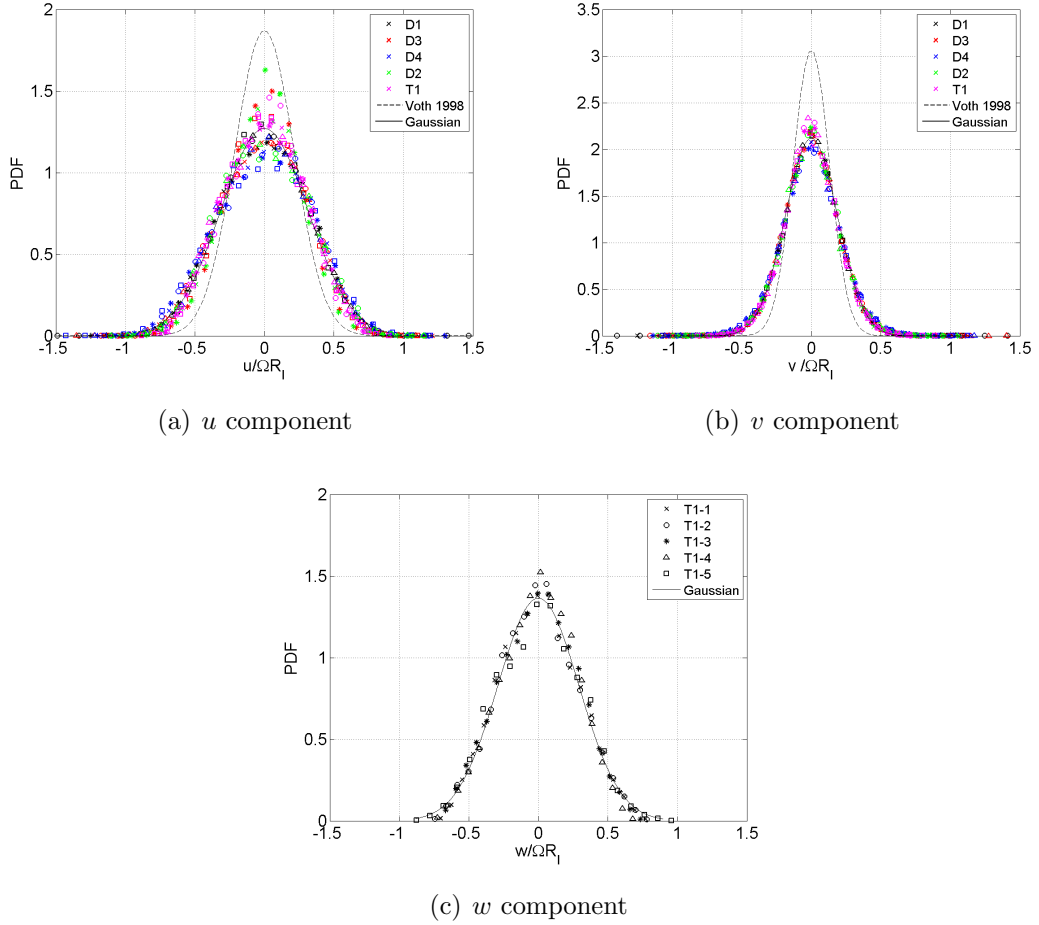

(a) u component

(b) v component

(c) w component

Figure 4.6: Fluctuating velocity PDFs; for (a) and (b): $\times$ = run 1; $\bigcirc$ = run 2; $*$ = run 3; $\triangle$ = run 4; ; $\square$ = run 5

In Figure 4.7 normalised RMS velocity is plotted against Reynolds number. The relatively high degree of scatter is likely a result of the large scale convergence issues (see Section 4.1.2). However, the majority of points are within the measurement uncertainty, and in agreement with other experiments (126; 144) the RMS velocity scales linearly with the impeller rotation frequency. Voth et

al. (127) comment on the importance of this linear scaling, which implies that there are no major changes to the large scale flow structure. The fluctuation magnitude of $\sim 30\%$ in the present investigation is similar to the $\sim 35\%$ measured by Zocchi et al. (144), with slight variations attributable to set-up and geometry differences.
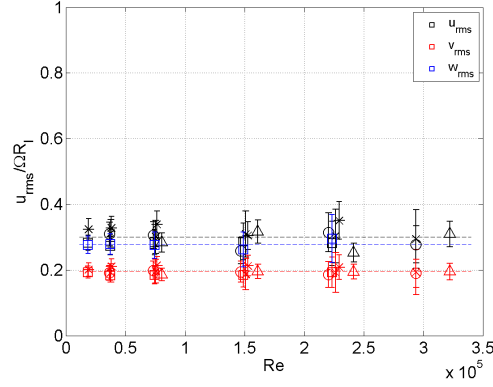

Figure 4.7: RMS velocity, component by colour as indicated, case by symbol: $\times$ = $D1$; $\bigcirc$ = $D2$; $\ast$ = $D3$; $\triangle$ = $D4$; ; $\square$ = $T1$, $[--]$ = mean component value averaged over all cases and runs

4.2.2 Ratio of RMS Velocities and Reynolds Stresses

Figure 4.8 shows the variation of RMS velocity ratios and Reynolds stresses with Reynolds number. The ratio of radial and axial velocity components is ~ 1.5 (See Figure 4.8(a)), which agrees well with the value of 1.48 calculated by Voth et al. (126); referred to as a typical value for homogenous shear flows (38). The TPIV radial component ratio of unity confirms the axial symmetry. The mean Reynolds stress should be zero due to symmetry (127), which is confirmed in Figure 4.8(b).

4.2.3 Velocity Gradients

Velocity gradients were evaluated using a central differencing scheme (detailed in Section 2.4.1) yielding four and nine components for the DPIV results and TPIV results respectively, which are studied both locally and globally.

(a) RMS velocity ratios
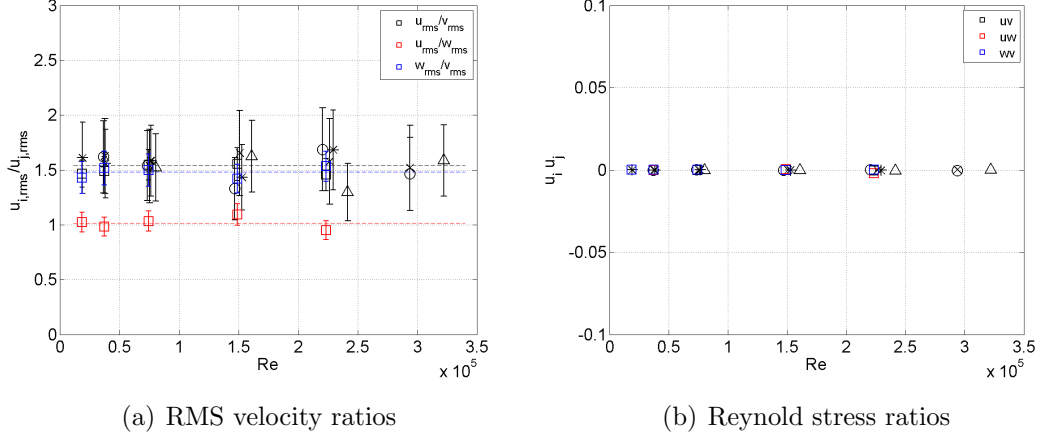
(b) Reynold stress ratios

Figure 4.8: Ratio of RMS velocities and Reynolds stresses, components by colour as indicated, case by symbol: $\times = D1$; $\bigcirc = D2$; $* = D3$; $\triangle = D4$; ; $\square = T1$, $[--] = $ mean component value averaged over all cases and runs

4.2.3.1   Local Velocity Gradients

Figure 4.9 shows PDFs of normalised gradient components for all DPIV and TPIV cases. The RMS vorticity normalisation parameter was chosen to match the results of Jiménez et al.(53). Despite the reasonable collapse of data at low gradient magnitudes, there is some variation in the tails. Comparing the in-plane gradients from DPIV and TPIV results shows that the latter appear to be slightly under-predicted, due to technique differences (see Section 3.2.2).

The five additional TPIV gradient components show similar behaviour to the mean local gradient field errors shown in Section 3.2.2. The $\partial w/\partial x$ and $\partial w/\partial y$ components are similar in magnitude to the other in-plane off-diagonal components ($\partial u/\partial y$ and $\partial v/\partial x$). The out-of-plane gradients, $\partial u/\partial z$ and $\partial v/\partial z$, are smaller in magnitude than the other off-diagonal terms, due to the increasing gradient under-prediction demonstrated in Section 3.2.2. Despite the large expected error in the $\partial w/\partial z$ term, the magnitude of this component is similar to the other diagonal components, due to the balance of noise and spatial filtering effects.

Jiménez et al.(53) compile a variety of numerical and experimental results, demonstrating a slight increase in normalised gradient magnitude with increasing

(a) $\partial u/\partial x$ (b) $\partial u/\partial y$ (c) $\partial u/\partial z$

(d) $\partial v/\partial x$ (e) $\partial v/\partial y$ (f) $\partial v/\partial z$

(g) $\partial w/\partial x$ (h) $\partial w/\partial y$ (i) $\partial w/\partial z$
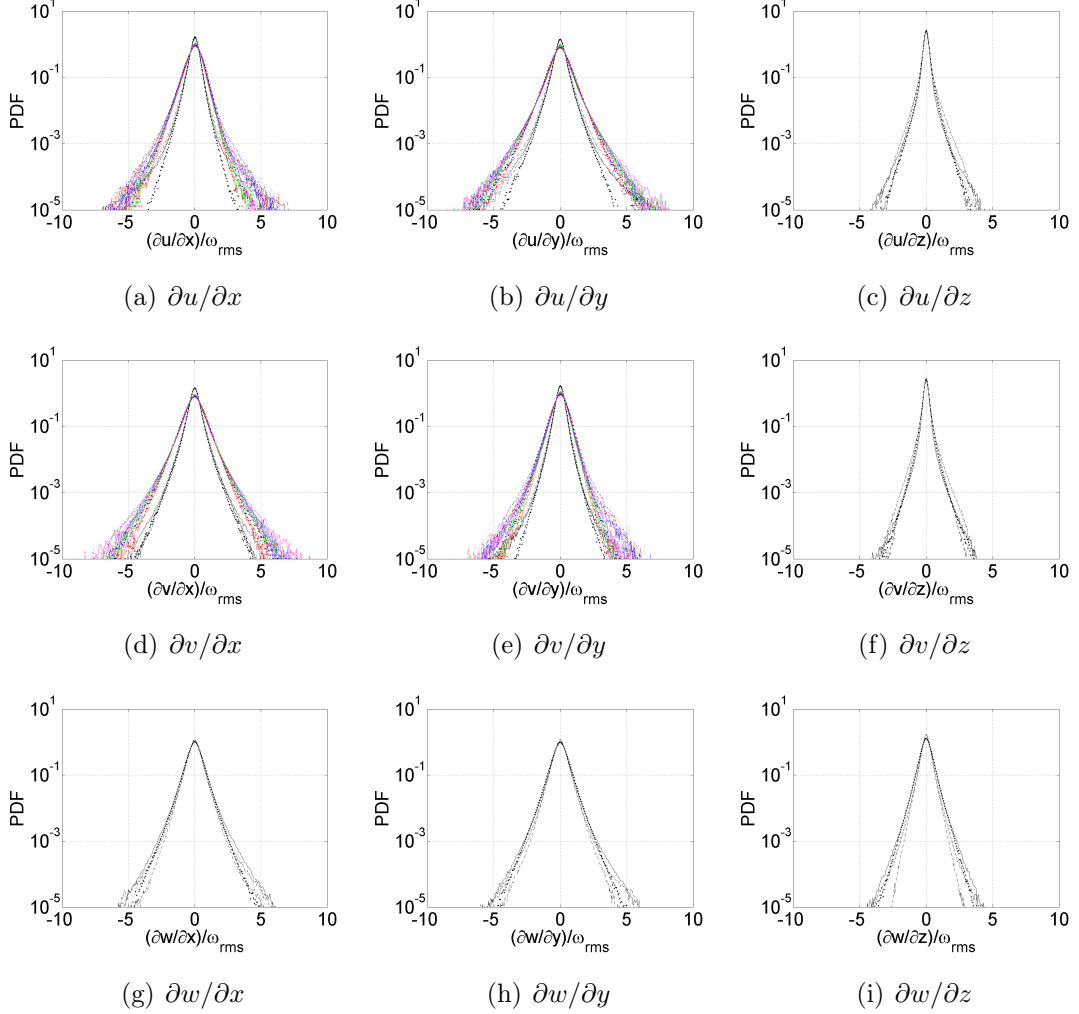
Figure 4.9: Components of velocity gradient; colours representing different cases: $T1$ = black, $D1$ = red, $D2$ = blue, $D3$ = green, $D4$ = magenta: and line types representing run numbers: $[-]$ = run 1, $[\cdots]$ = run 2, $[---]$ = run 3, $[-.-.]$ = run 4, $[***]$ = run 5

Reynolds number. Similar behaviour is also shown experimentally by Belin et al. (7). Notably however, despite the large range of Reynolds numbers covered in the current investigation, no strong Reynolds number trends are shown and the gradients scale directly with the RMS vorticity. The simulations in Section 3.2.3 demonstrated that a reduction in resolution reduces gradient magnitude through spatial filtering, which may counter the expected gradient magnitude increase at

high Reynolds number.

The non-Gaussian PDF asymmetry is noted by a number of investigators (37; 53; 82). This is quantified by the last two authors by modelling the tails as a straight line exponential decay in semi-logarithmic axes, as in Equation 4.1, where a is the rate of decay, q is any velocity gradient, and $P(q)$ is the probability density function of that velocity gradient (37). Left and right hand side decay constants are calculated for all gradient terms using a least squares approach, with the difference between them expressed as a percentage in Figure 4.10. Diagonal gradient components are coloured separately, with a single colour for all off-diagonal components. Similar to the findings of Ganapathisubramani et al. (37) and Mullin & Dahm (82) substantial differences are found in the scaling of the diagonal components, with much smaller differences in the off-diagonal terms. The cause of this asymmetry is thought to be the departure from local isotropy (82).
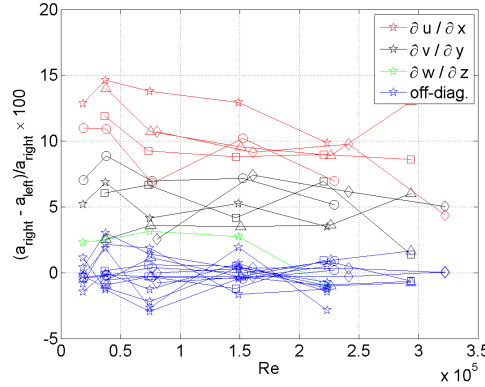
$$a = -\frac{d}{d|q|} \ln P(q) \tag{4.1}$$

Figure 4.10: Percentage change in decay constant for left and right hand PDF tails; symbols denote cases: $\triangle = D1$; $\square = D2$; $\bigcirc = D3$; $\diamondsuit = D4$; $\bigstar = T1$

4.2.3.2 Global Mean Velocity Gradients

According to the Richardson cascade concept (for example see Davidson (23)), energy injected at the large scale should be transferred to increasingly small scales until it is dissipated into heat. Due to the high Reynolds number of the
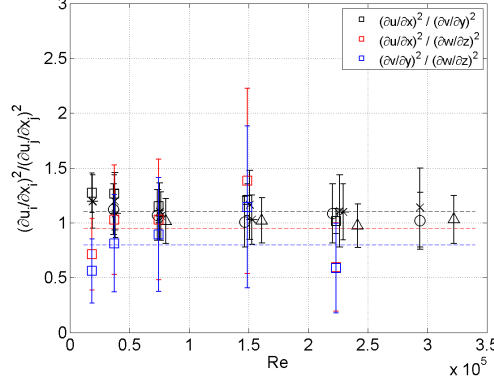
Figure 4.11: Velocity gradient ratio, components by colour as indicated, case by symbol: $\times = D1$; $\bigcirc = D2$; $* = D3$; $\triangle = D4$; ; $\square = T1$, $[--] = $ mean component value averaged over all cases and runs

current measurements there should exist a sufficiently large separation between the injection and dissipation scales, that any anisotropy present at the largest scale should be lost during the transfer process. As the velocity gradients are evaluated over the space of several cells, which in the present investigation are of order η (see Section 4.1.3), global isotropy at the smallest scales can be tested by examining the ratio of mean squared velocity gradients, shown in Figure 4.11.

As expected for isotropic turbulence the ratio of in-plane gradients takes a value of around 1 for all data sets. Despite the reasonably high degree of scatter, at low Reynolds number the ratio appears slightly higher than unity, which may indicate an insufficient separation of scales due to under developed turbulence, similar to the results of Voth et al. (126). However, this small level of anisotropy is within the measurement uncertainty, making more definite conclusions difficult. Additionally, the lower resolution $D1$ and $D2$ cases may not be capturing the smallest flow scales, particularly at the highest Reynolds numbers. However, as the range of scales increases with Reynolds number, despite the failure of these cases to reach right down to the dissipation scale, the flow still appears to have lost much of the anisotropy present at larger scales.

The ratios, $(\partial u/\partial x)^2/(\partial w/\partial z)^2$ and $(\partial v/dy)^2/(\partial w/\partial z)^2$, were also shown for the TPIV data. These both demonstrate a high degree of scatter and take values significantly below 1. This under-prediction is due to the poor out-of-plane mea-

surement accuracy (see Section 3.2.3), and within the measurement uncertainty the gradients are isotropic.

4.3 Statistical Flow Structure Characterisation

In this section the flow structure is characterised statistically using the velocity correlation function, longitudinal structure function, and energy spectrum. In addition to further flow isotropy investigation and integral scale determination, this structural analysis is also useful in assessing the validity of the current measurements through their reproduction of the widely recognised Kolmogorov scaling laws. Conversely, it is also possible to estimate the dissipation rate from these quantities on the assumption of this scaling.

4.3.1 Velocity Correlation Function

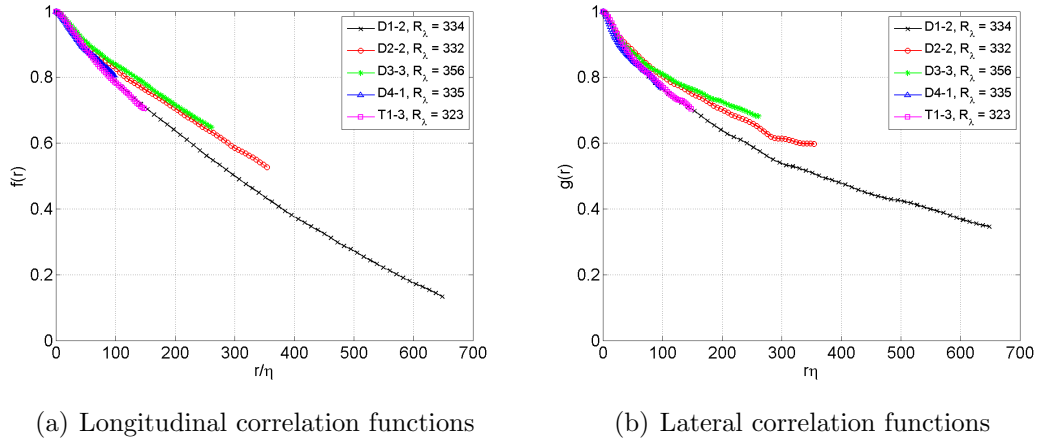(a) Longitudinal correlation functions   (b) Lateral correlation functions

Figure 4.12: Velocity correlation functions

In order to further examine the flow anisotropy and make an estimate of the integral length scale, that is the scale over which velocities are appreciably correlated, it is useful to plot the velocity correlation functions (see Section 1.1.1).

Figure 4.12 shows longitudinal and lateral u velocity correlation functions. The agreement between different cases and runs is reasonable at small distances,

with more significant variation seen as the distance is increased. As velocity is correlated over a range of distances, de Jong et al. (24) note that spatial filtering should affect the correlation less severely at larger separation distances. However, the lack of convergence, particularly of the larger scales appears to dominate the correlation, resulting in the observed variation between cases. Domain size differences are seen through function truncation at $f(r) > 0$, with even the largest domain failing to capture the correlation function at $f(r) = 0$.
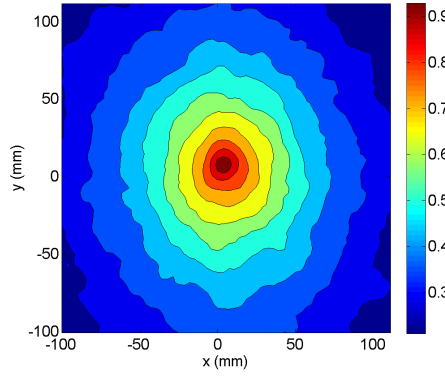

Figure 4.13: Combined 2D velocity correlation function, $D1 - 5$

The combined 2D correlation function of velocity magnitude in Figure 4.13 shows the effect of the flow straining due to toroidal re-circulation, increasing correlation in the y direction over large distances, resulting in an oval distribution.

The combined 3D velocity magnitude correlation function for the TPIV data is plotted on Figure 4.14 using semi-transparent iso-surfaces to display a range of correlation values. Despite the axial symmetry of the current flow, the increased uncertainty of the depth component causes strong additional correlation function deformation; resulting in a prolate spheroid shape.

An integral length scale estimate is made through velocity correlation function integration. The premature function truncation ($f(r) > 0$) due to insufficient domain size requires the modelling of the function tail. Therefore, a decaying exponential was fitted to the available data and extrapolated to $f(r) = 0$ (see Figure 4.15(a)). To decrease the length scale estimate dependancy on extrapolated data, only functions from the largest domain $D1$ case were used to determine the

(a) Entire function

(b) Enlarged centre
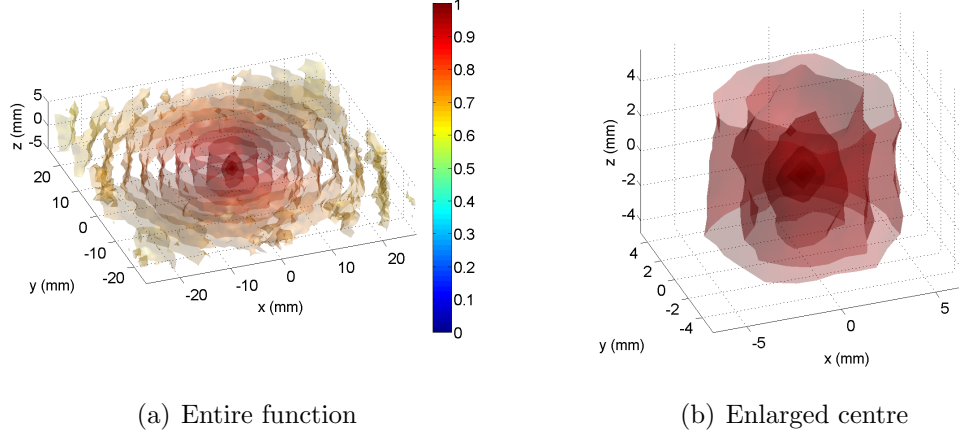
Figure 4.14: Combined 3D velocity correlation function, $T1 - 4$

(a) Example of extrapolated function, $D1$

(b) Length scale estimates
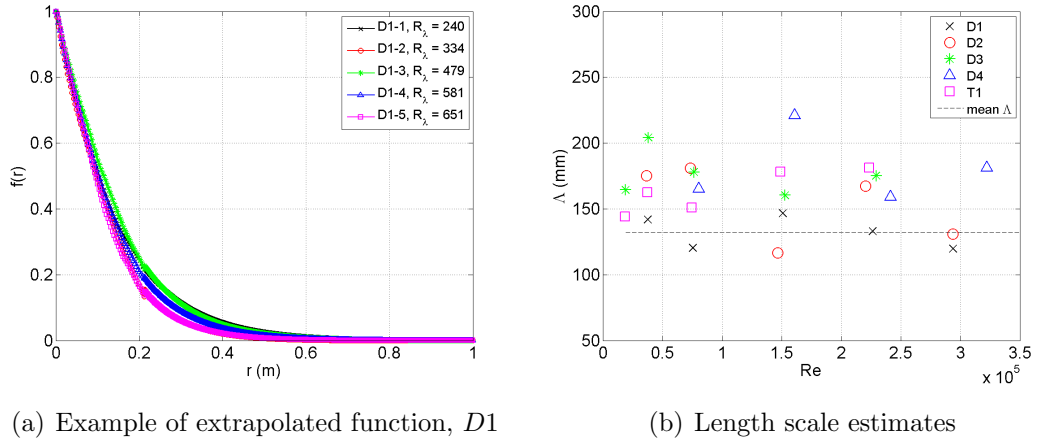
Figure 4.15: Length scale estimation

mean value. Length scale estimates are shown in Figure 4.15(b), with the level of scatter consistent with the statistical convergence. A mean value from the $D1$ case of $\Lambda = 130$mm is calculated, which will be used in all scaling argument estimates.

4.3.2 Structure Function

Figure 4.16 shows the second order longitudinal structure function normalised by the RMS velocity for each case at a single similar Reynolds number and across a

range of Reynolds number for the $D1$ case.

(a) Case comparison

(b) Reynolds number comparison

Figure 4.16: Longitudinal second order structure function

Figure 4.16(a) shows reasonable agreement between cases over small separation distances with variation increasing with distance. The interrelation of this quantity wi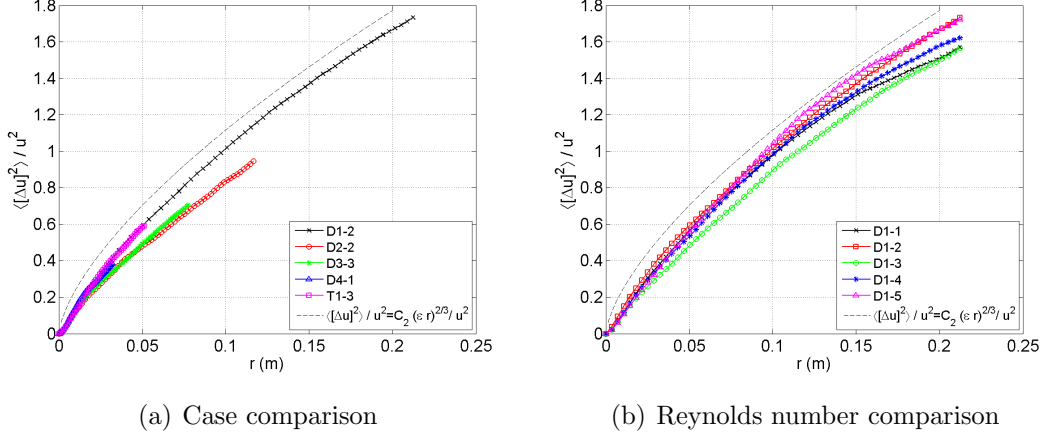th the velocity correlation function ($[\Delta u]^2 = u^2(1 - f)$) means this variation is attributable to similar sources (see Section 4.3.1).

Due to the large domain size the $D1$ case allows the structure function to be calculated over a greater range of separation values, shown in Figure 4.16(b). The variation between different Reynolds numbers is within the limits of the measurement uncertainty, which is set in this case by the RMS velocity normalisation parameter.

To assess agreement with Kolmogorov's two-thirds law the $D1$ case has also been plotted logarithmically and pre-multiplied by $1/C_2(\epsilon r)^{2/3}$ in Figure 4.17, where $C_2 = 2.12$ as $R_\lambda \to \infty$ (114). Gotoh et al. (40) show the exponent C_2 is constant over a value of $R_\lambda = 250$, and therefore a finite Reynolds number correction has not been made. Although modification to the original K41 law usually involve an additional exponent to take account of intermittency, the correction for the 2nd order structure function is usually small, and within the bounds of experimental error (23).

Figure 4.17(a) shows a collapse of data at reasonable separation distances onto the $r^{2/3}$ line (plotted for run $D1 - 2$), which is also shown in Figure 4.17(b),
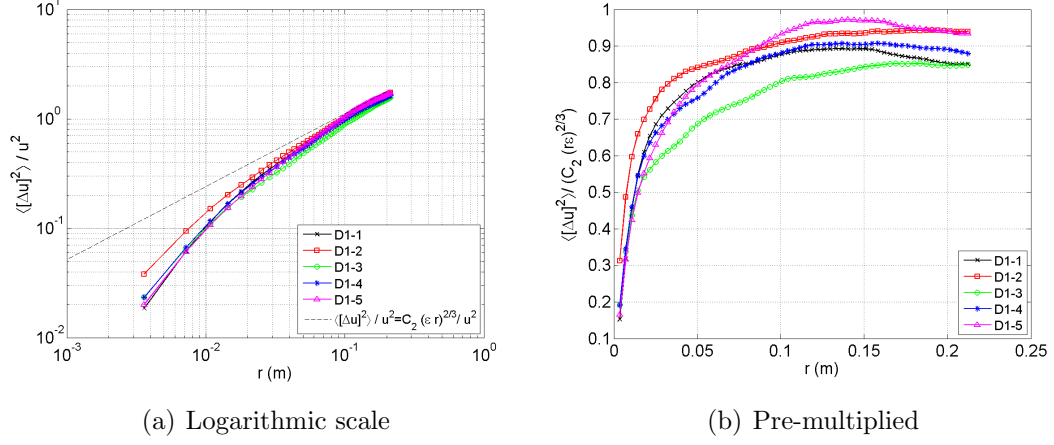
(a) Logarithmic scale

(b) Pre-multiplied

Figure 4.17: Longitudinal second order structure function, Reynolds number comparison

where $r^{2/3}$ scaling is represented by the plateau at a value of unity. The horizontal section of the curves appear slightly later than may be expected, showing a slight positive gradient in the inertial range. Similar behaviour is noted by Gotoh et al. (40), who attribute this to an insufficient number of Fourier modes in the energy spectrum inertial range, applying an exponent correction to their results. Similarly, the large scale convergence issues in the present investigation may produce the weak positive gradients shown in Figure 4.17(b).

Conversely, assuming the flow scales as predicted by Kolmogorov's two-thirds power law, the structure function can be used to estimate the dissipation rate (24). Figure 4.18 shows normalised structure function variation, $(1/r)([\Delta u]^2/C_2)^{3/2}$, with separation distance, r, with the plateau now representing the mean dissipation rate, ϵ. These estimates have been compared with other methods in Section 4.4.2.

4.3.3 Energy Spectrum

The 1D Energy spectrum was calculated directly from the PIV data using the method described in Section 2.4.3. Due to its non-trivial effect, the use of the Hann windowing function is also demonstrated.
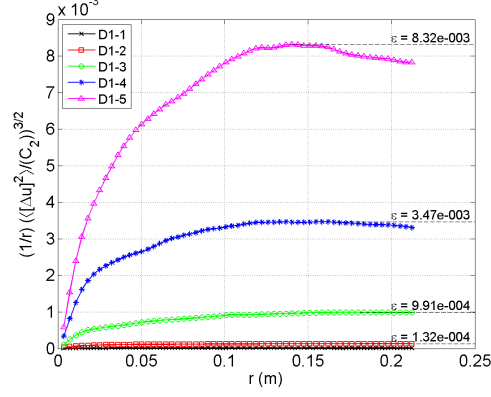
Figure 4.18: Dissipation estimate using the second order structure function

Figure 4.19(a) shows a good collapse of data between all experimental runs and reasonable agreement with the Pao result (88) at low wavenumbers. Measurements positions range from the tail end of the inertial range into the dissipation region. Without a windowing function, similarities are observed between the present results, previous experimental investigations (48; 68), and the accuracy simulation (see Section 3.2.5), with the spectrum peeling away from the Pao curve at high wavenumbers due to spatial truncation effects and experimental uncertainty.

Inertial range agreement with Kolmogorov scaling can be assessed more rigorously through the pre-multiplied spectra shown in Figure 4.19(b), which for clarity shows only the highest Reynolds number results containing the lowest wavenumber components. The failure of these pre-multiplied spectra to plateau demonstrates similarity with the simulation results (see Section 3.2.5), and the absence of Kolmogorov scaling in the inertial range is attributed to an insufficient domain size.

Figures 4.20(a) and 4.20(b) show the effect of Hann windowing on the $D1$ and $D4$ cases. Use of this windowing function produces similar effects to those seen in the accuracy simulation (see Section 3.2.5), improving the dissipation region in the $D4$ results, but causing increased inertial range under-estimation in the $D1$ results. Despite significant deviations from the Pao spectra, the observed errors are consistent with those defined in the numerical simulations, and therefore within the bounds of technique accuracy, these appear to be predicted reasonably.

(a) All cases and runs (b) Pre-multiplied selected cases
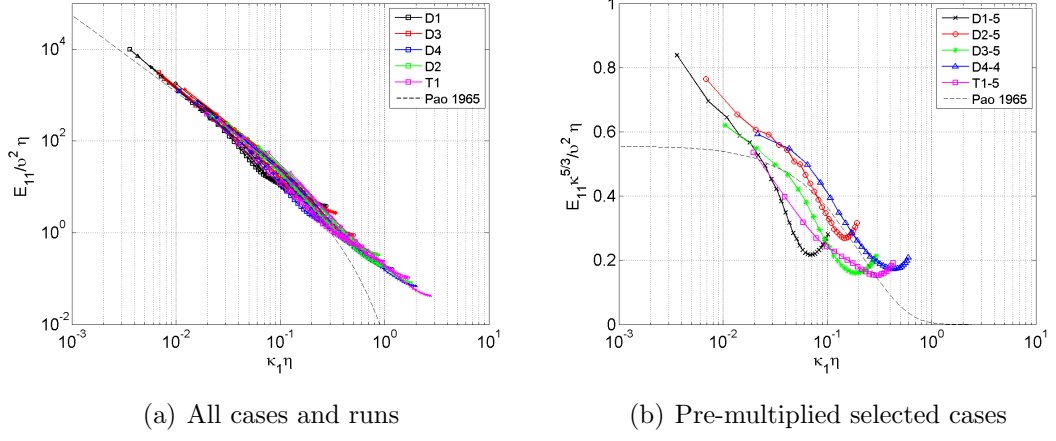
Figure 4.19: Energy spectra without windowing function, for (a): cases by colour as indicated, run number by symbol: $\times$ = run 1; $\bigcirc$ = run 2; $*$ = run 3; $\triangle$ = run 4; ; $\square$ = run 5

(a) case $D4$ (b) case $D1$
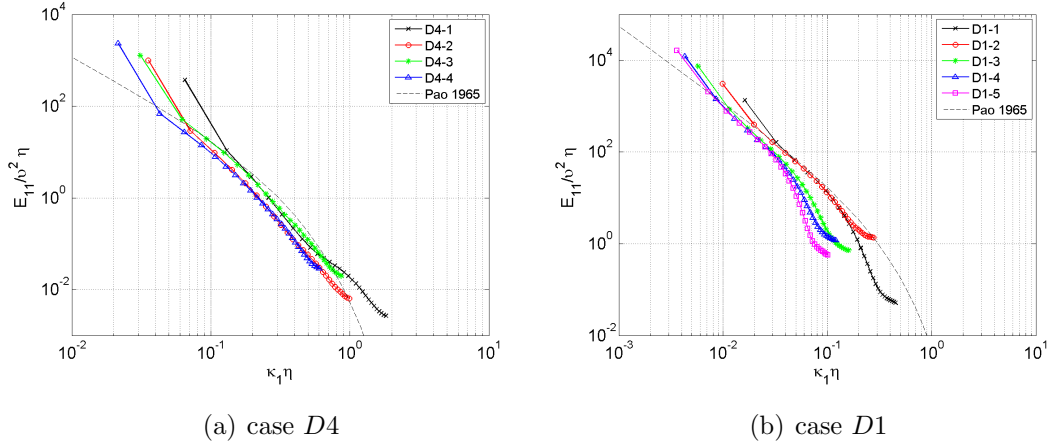
Figure 4.20: Effect of Hann windowing on the energy spectra

The dissipation rate can be estimated through the Energy spectra using Equation 4.2 (144).

$$\epsilon = 15\nu \int\limits_0^\infty \kappa^2 E_{11}(\kappa)d\kappa \tag{4.2}$$

An example of the velocity gradient spectrum, $\kappa^2 E_{11}$, is shown in Figure 4.21.

129

The effect of convergence and noise at both low and high wavenumbers respectively result in dramatic deviations from the ideal distribution, with similar behaviour identified by Ganapathisubramani et al. (36). These points, identified in Figure 4.21, were trimmed before completing the original normalised E_{11} spectrum through a least squares fit of a 3 constant equation with fixed $-5/3$ power law and exponential decay term (Equation 4.3), which gave the best fit over widest range of values. The fitted curve is converted back into dimensional form, before being multiplied by κ^2 (shown non-dimensionally in Figure 4.21). The completed function integral is used to calculate dissipation, which has been compared with other methods in Section 4.4.2.

$$\frac{E_{11}(\kappa)}{\upsilon^2 \eta} = a(\kappa\eta)^{-5/3} e^{-b(\kappa\eta)^c} \tag{4.3}$$
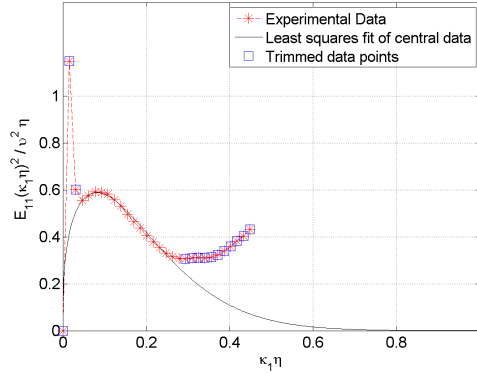
Figure 4.21: Dissipation estimate from energy spectrum

4.4   Estimation of Turbulent Properties and Scales

This section begins with an investigation of local and global dissipation rate estimates. As discussed by de Jong et al. (24), accurate measurement of this quantity is extremely difficult, with no current consensus to preference the multitude of available methods. Therefore, the current investigation compares a number of these methods, before selecting two estimates to derive other quantities of interest such as the Taylor and Kolmogorov microscales.

4.4.1   Local Dissipation Rate Distribution

Figure 4.22 shows PDFs of the normalised dissipation rate, ϵ^* (where $\epsilon^* = \epsilon/\epsilon_m$).
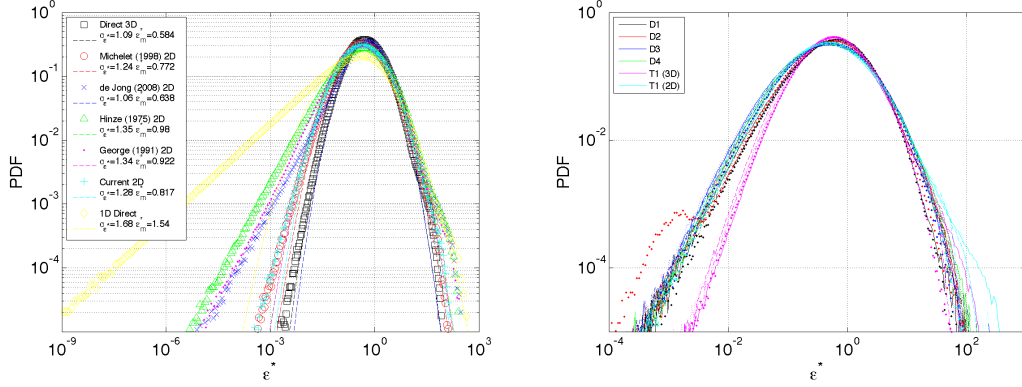
(a) Effect of estimation method on $T1-3$ case (b) All cases and runs with current 2D method

Figure 4.22: Local normalised dissipation rate distribution; for (b): cases by colour as indicated, run numbers by line type: $[—] = $ run 1, $[\cdots] = $ run 2, $[---] = $ run 3, $[-.-.] = $ run 4, $[***] = $ run 5

As discussed in Section 2.4.2, absence of the full VGT requires assumptions of isotropy to complete the dissipation estimate. The replacement of gradient terms according to user preference has lead to a number of slightly different gradient combinations (24; 39; 43; 76), which are compared in Figure 4.22(a) for the $T1-3$ case. In agreement with previous investigations (37; 82) the dissipation rate calculated from the full VGT is distributed log-normally, which is confirmed through a least squares regression (Equation 4.4), where σ_{ϵ^*} is the standard deviation and ϵ_m^* the mean.

$$P\left(\epsilon^*\right) = \frac{1}{\sqrt{2\pi}\sigma_{\epsilon^*}\epsilon^*}exp\left(-\frac{(ln(\epsilon^*)-\epsilon_m^*)^2}{2\sigma_{\epsilon^*}^2}\right) \qquad (4.4)$$

Other methods exhibit greater asymmetry and predict a higher probability of lower dissipation values, caused by the projection of 3D dissipative structures onto lower dimensional space (37). This conclusion is supported by the severe under-prediction and slight over-prediction at the low and high dissipation magnitude ends respectively of the 1D estimate. The use of more gradient terms in the 2D

		Run 1		Run 2		Run 3		Run 4		Run 5	
Case	Method	σ_{ϵ^*}	ϵ_m^*	σ_{ϵ^*}	ϵ_m^*	σ_{ϵ^*}	ϵ_m^*	σ_{ϵ^*}	ϵ_m^*	σ_{ϵ^*}	ϵ_m^*
$D1$	2D	1.06	0.604	1.1	0.606	1.07	0.589	1.07	0.588	1.03	0.565
$D2$	2D	1.23	0.774	1.21	0.753	1.19	0.719	1.16	0.689	1.1	0.634
$D3$	2D	1.27	0.833	1.28	0.842	1.24	0.796	1.22	0.77	1.18	0.721
$D4$	2D	1.24	0.773	1.28	0.828	1.24	0.781	1.2	0.737	-	-
$T1$	3D	0.999	0.471	1.05	0.535	1.09	0.584	1.03	0.522	0.992	0.492
$T1$	2D	1.23	0.737	1.28	0.802	1.28	0.817	1.24	0.767	1.18	0.704

Table 4.2: Log-normal dissipation rate distribution constants

estimates reduce the asymmetry by varying degrees. The estimates proposed by de Jong et al. (24) and Hinze (43) both use only 3 of the available 4 measured gradient terms, which may cause the greater asymmetry in these distributions. The methods of Michelet et al. (76) and the authors own choice of components produce similar results, and give the closest approximation to the log-normal distribution, with the latter used to produce all subsequent 2D estimates.

Figure 4.22(b) shows normalised dissipation rate PDFs for all cases. The data collapses well into two distinct distributions for the 2D and 3D calculation methods. The differences between calculation methods is again illustrated through 2D TPIV estimates, which produce similar results to the DPIV estimates. After normalisation the effect of reducing spatial resolution is slight, although a weak reduction in the high magnitude tail is observed, which may relate to the increasing dissipation rate under-prediction. The effect of spatial resolution will be examined in the following section. Finally, log-normal distribution parameters are included in Table 4.2.

4.4.2 Comparison of Global Dissipation Rate Estimates

As demonstrated in previous work (24; 65) and the present investigation (see Section 3.2.3) accurate dissipation estimation is highly dependent on measurement spatial resolution. Therefore, the following comparison considers resolution effects using two approaches: firstly by dropping an increasing numbers of vectors from the highest resolution $D4$ case to artificially reduce the resolution (shown in Figure 4.23(a)); and secondly by comparing all experimental results, which

include a reasonable range of measurement resolutions (shown in Figure 4.23(b)). The dissipation rates are normalised by the scaling argument estimate (calculated using Equation 1.1), forcing this estimate to a constant value of $A = 0.5$.

Lavoie et al. (65) demonstrate the effect of spatial filtering and finite difference gradient calculation on PIV dissipation rate prediction, by filtering a known spectrum using a function based on interrogation window dimensions. Therefore, the measurement resolution in Figure 4.23 is plotted against the normalised interrogation window width, $\Delta h/\eta$, which is double the interrogation window spacing, $\Delta x/\eta$.
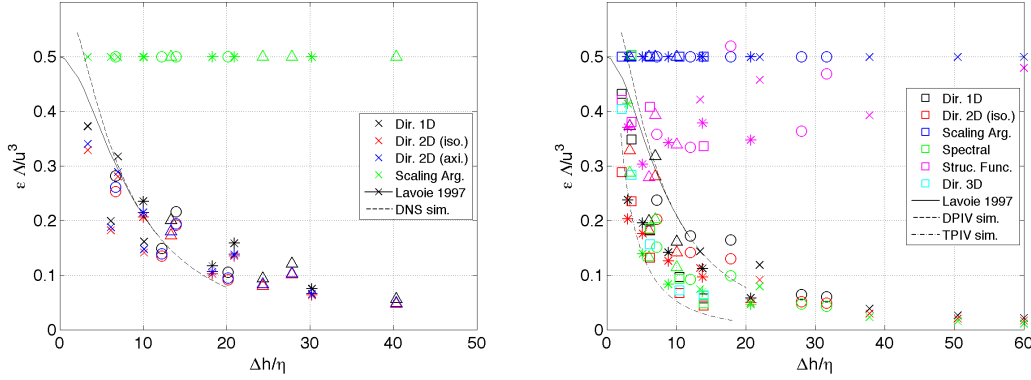
(a) Resolution variation through dropped vectors, $D4$

(b) All results

Figure 4.23: Effect of measurement resolution on dissipation rate estimation; for (a): calculation method by colour as indicated, resolution by symbol: $\times$ = full resolution; $\bigcirc$ = every 2nd point; $*$ = every 3rd point; $\triangle$ = every 4th point; for (b): calculation method by colour as indicated, case by symbol: $\times = D1$; $\bigcirc = D2$; $* = D3$; $\triangle = D4$; $\square = T1$

Figure 4.23(a) shows that direct estimation of the dissipation rate using the 1D or either of the 2D methods produce very similar results. The direct 1D method produces the highest values closely followed by the 2D axisymmetric and isotropic estimates, which are within $\sim 10\%$. Despite considerable scatter the estimates show increasing dissipation under-prediction as the measurement resolution is reduced; in broad agreement with the spectral correction predicted

by Lavoie et al. (65), and the correction ratio derived from the numerical accuracy investigation (see Section 3.2.3).

Again despite considerable scatter, a similar trend in the direct estimates is shown in Figure 4.23(b), with under-prediction increasing as measurement resolution is reduced. The more severe under-prediction of the $D3$ DPIV case is expected to result from RMS velocity over-prediction and corresponding under-prediction of the scaling factor. The 1D and 2D TPIV estimates show reasonable agreement with those from the full 3D method, indicating that the poor prediction of the local dissipation structure (shown in Section 4.4.1) only has a slight effect during mean dissipation rate calculation. This observation is in contrast to Ganapathisubramani et al. (37) who show severe under prediction of the dissipation rate for the 1D method.

Increased TPIV measurement uncertainty was shown in Section 3.2.3 to cause more severe dissipation rate under-prediction in comparison with DPIV. This is reflected in the current results, which show similar under-prediction to the TPIV correction curve. Slightly higher than expected dissipation rate estimates may be a result of increased experimental noise resulting in gradient over-prediction, which compensate slightly (albeit erroneously) for the effects of spatial averaging.

Despite efforts to correct for lost dissipation scale energy through spectra modelling in Section 4.3.3, the energy spectrum estimates are close to the direct method. This under-prediction is expected to stem from spatial averaging effects, which will result in a loss of energy at the smallest scales. The model effectively extends the measured spectrum to account for higher frequency components lost due to insufficient resolution, but fails to take account of the lost energy within the measured spectrum itself. Although Lavoie's spectral correction could be applied to the spectrum directly, this level of manipulation is akin to remodelling the entire spectrum, which although useful in terms of fully demonstrating the source of the energy losses, may be excessive for dissipation rate estimation.

As shown by de Jong et al. (24) the second order structure function is a fairly robust method for dissipation rate estimation, due to its use of relatively insensitive large scale motions in comparison with the highly sensitive local velocity gradient terms used during direct estimation. Therefore, despite considerable scatter ($\approx 0.3 \rightarrow 0.5$), this method results in the most consistent dissipation rate

estimate. Accuracy is inversely proportional to correlation function truncation, and despite exponential modelling to complete the function, better estimates are obtained for cases with larger domains, irrespective of spatial resolution.

Normalising the structure function dissipation estimate by R_I^4/ν^3 permits comparison with the results of Zocchi et al. (144), shown in Figure 4.24(a). The measurement uncertainty for this method scales with truncation error and is estimated based on the convergence of the pre-multiplied functions (similar to Figure 4.17(b)), and within the bounds of uncertainty the dissipation rates follow the expected Re^3 scaling. The current results appear to predict slightly higher dissipation rates than Zocchi et al., which may be attributed to geometry differences between set-ups.

It is possible to correct for the known spatial averaging effects through the correction ratio's calculated from the results in Section 3.2.3, which are shown in Figure 4.24(b). Despite increased scatter in comparison with Figure 4.24(a) the corrected estimates predict Re^3 scaling. As measurement resolution is reduced the correction ratio tends to zero (see Figure 4.23(b)) dramatically increasing the correction magnitude and its associated error, resulting in erroneously high dissipation rates for low resolution high Reynolds number results. However, as the direct method measurement uncertainty was calculated from the dissipation rate under-prediction, it is difficult to estimate the uncertainty in these corrected results, and therefore no error bars are included in Figure 4.24(b).

4.4.3 Taylor Microscale Estimate

In the current investigation two approaches have been used to calculate this length scale: the first involves fitting a parabola to the velocity correlation function close to $r = 0$ (see Section 1.1.2); the second relates it directly to local velocity gradients through the dissipation rate defined by Equation 1.4 (23).

A least squares approach was used to fit parabolas to the velocity correlation function close to $r = 0$ using between $N_{LS} = 5 - 10$ points in the regression. The Taylor microscale is a function of the number of points, N_{LS}, used during the regression, and therefore to estimate this value as $r \to 0$ a linear regression of

(a) Structure Function Method

(b) Corrected Direct Method
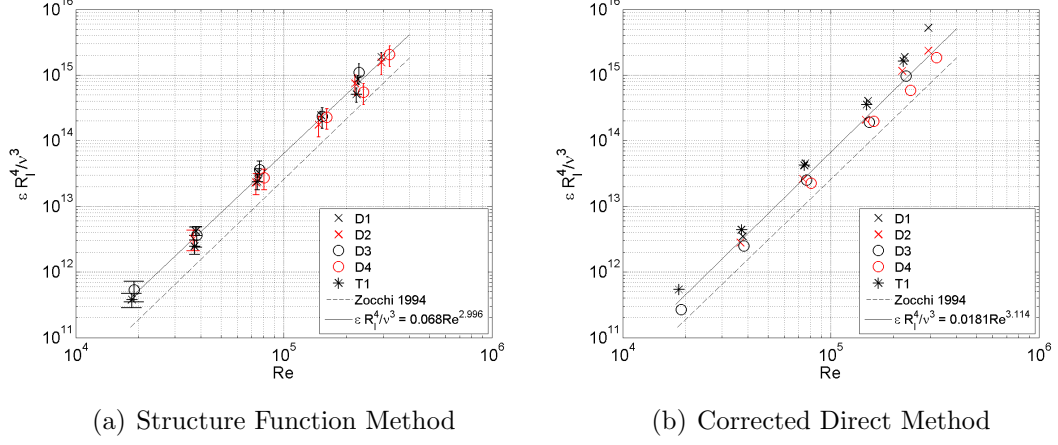
Figure 4.24: Different dissipation estimates in comparison to the scaling argument method

the microscale variation with N_{LS} has been calculated (shown in Figure 4.25(a)), with the microscale estimated as $\lambda_{N_{LS}=0}$.

The Taylor microscale is also calculated using the scaling argument dissipation estimate, and Figure 4.25(b) shows a parabola comparison between this and the best fit method for the $D1-2$ case. Both demonstrate reasonable agreement with the velocity correlation function around $r = 0$, with the best fit method predicting a slightly higher value for this case.

(a) Effect of N_{LS} on least squares fit

(b) Calculated parabola
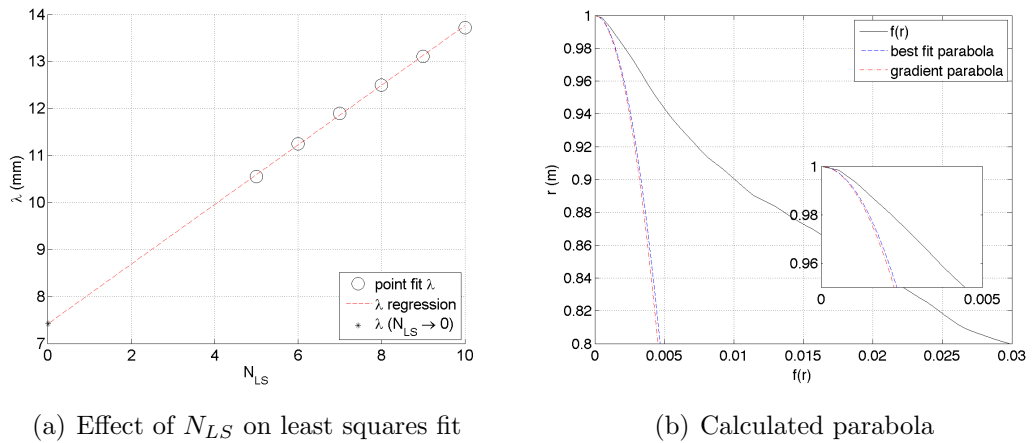
Figure 4.25: Parabolic fit method for estimating Taylor microscale, $D1-2$

136

The two methods are compared in Figure 4.26. The u_{rms}/U Reynolds number scaling shown in Figure 4.7 accounts for the expected gradient based Taylor microscale $Re^{-1/2}$ scaling. The parabola best fit method produces a much higher degree of scatter, with only the higher resolution DPIV case predicting the correct scaling. Therefore, the lower resolution results may contain an insufficient number of points close to the high curvature $r = 0$ region where the parabola fit is conducted, resulting in the observed over-prediction. Despite the high TPIV resolution, these results severely over-predict the microscale, which may be due to increased noise in the velocity correlation function.

Figure 4.26(b) compares values of the Taylor microscale calculated using different dissipation rate estimates from the scaling argument and structure function methods. There is good agreement for the majority of runs, and only the increased dissipation under-prediction due to structure function truncation in the higher resolution cases results in slight microscale over-prediction.

(a) Parabolic fit method   (b) Structure function method

Figure 4.26: Different Taylor microscale estimation methods; for (a) and (b): calculation method by colour as indicated, case by symbol: $\times = D1$; $\bigcirc = D2$; $* = D3$; $\triangle = D4$; $\square = T1$

Figure 4.27 shows Taylor microscale Reynolds number variation with Reynolds number for the structure function and scaling argument methods. The predictions again follow the expected $Re^{1/2}$ scaling, and demonstrate that the original

Reynolds number range of $Re \approx 2 \times 10^4 \rightarrow 2 \times 10^5$ is equivalent to $R_\lambda \approx 160 \rightarrow$ 700.
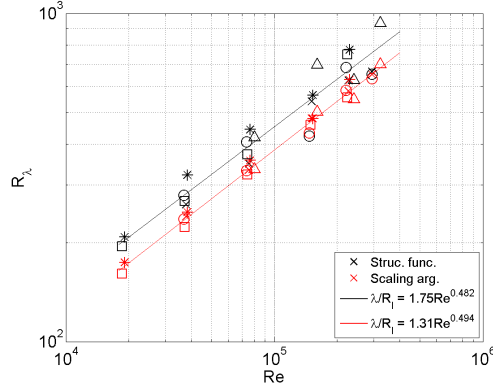
Figure 4.27: Different estimates of the Taylor microscale Renyolds number, R_λ; calculation method by colour as indicated, case by symbol: $\times = D1$; $\bigcirc = D2$; $*$ $= D3$; $\triangle = D4$; $\square = T1$

4.4.4   Kolmogorov Microscales

It is also useful to compute the Kolmogorov length and time scales, which give an idea of the smallest length and time scales within the flow. These have been calculated from Equation 1.2 using the dissipation rate estimates from the structure function and scaling argument methods. These length and time scales estimates are functions of the fourth and square root of dissipation rate respectively, appreciably reducing measurement uncertainty and resulting in good agreement between cases and the expected Reynolds number scaling. Again the difference between scaling argument and structure function methods manifests in the latter predicting slightly higher microscale values, although it should be noted that these differences are within the measurement uncertainty.

4.5   Flow Divergence

Divergence errors are likely to arise from a combination of measurement uncertainty, finite differencing error and spatial filtering of the velocity field. Therefore,

(a) Kolmogorov length scale   (b) Kolmogorov time scale

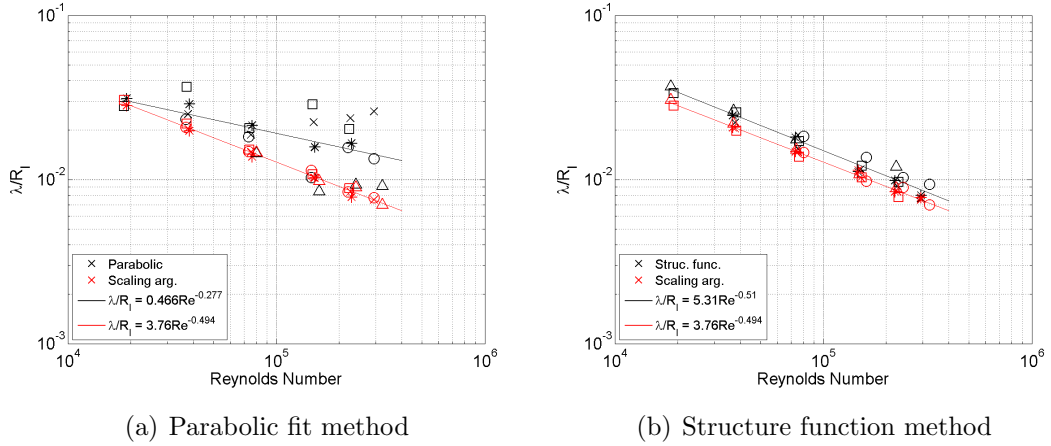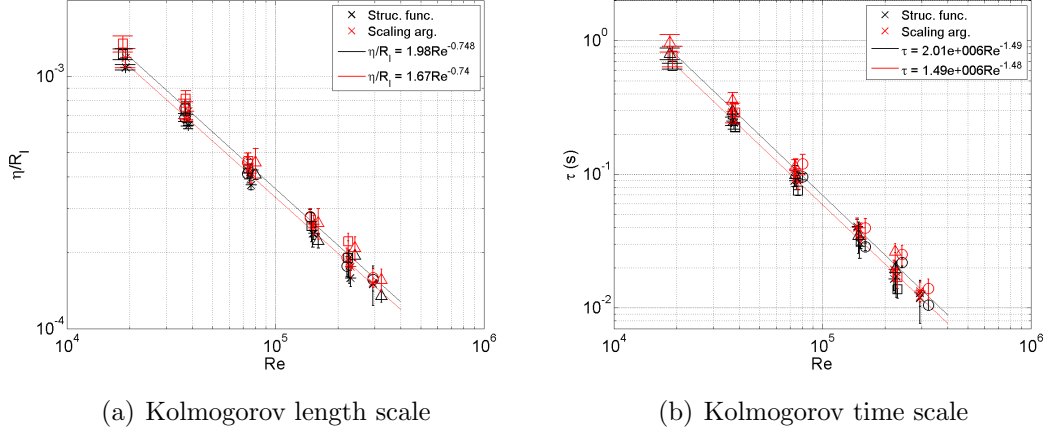Figure 4.28: Comparison of Kolmogorov microscale estimates; for (a) and (b): calculation method by colour as indicated, case by symbol: $\times = D1$; $\bigcirc = D2$; $*$ $= D3$; $\triangle = D4$; $\square = T1$

this section begins with a further TPIV accuracy quantification through the flow divergence, using a number of measures introduced in Section 3.2.4. Several divergence correction schemes are also assessed, as is the effect of divergence and noise on the statistical topological invariants.

4.5.1 Divergence Correlation

Figure 4.29 shows JPDFs of $\partial u/\partial x$ against $-(\partial v/\partial y + \partial w/\partial z)$, in which the degree of divergence is represented by the scatter of data away from the zero divergence line where these components balance. Data from the second cross-correlation pass is also plotted to demonstrate the discretisation error effect on the divergence. The second and third passes are denoted by IW_2 and IW_3 respectively, where $IW_2 = 1.56\ IW_3$. Due to the small volume thickness, it is not possible to study a wider range of discretisation levels.

Gradient components for the smaller window size shown in Figures 4.29(a) to (c) produce a poor correlation, resulting in a rounded distribution and a low correlation value, Q. As shown in Figure 4.29(d) to (f), increasing interrogation window size reduces divergence and improves this correlation value. A similar effect is shown by Zhang et al. (142) who conclude that this discretisation effect

139

(a) $T1-1$, IW_3 (b) $T1-3$, IW_3 (c) $T1-5$, IW_3

(d) $T1-1$, IW_2 (e) $T1-3$, IW_2 (f) $T1-5$, IW_2
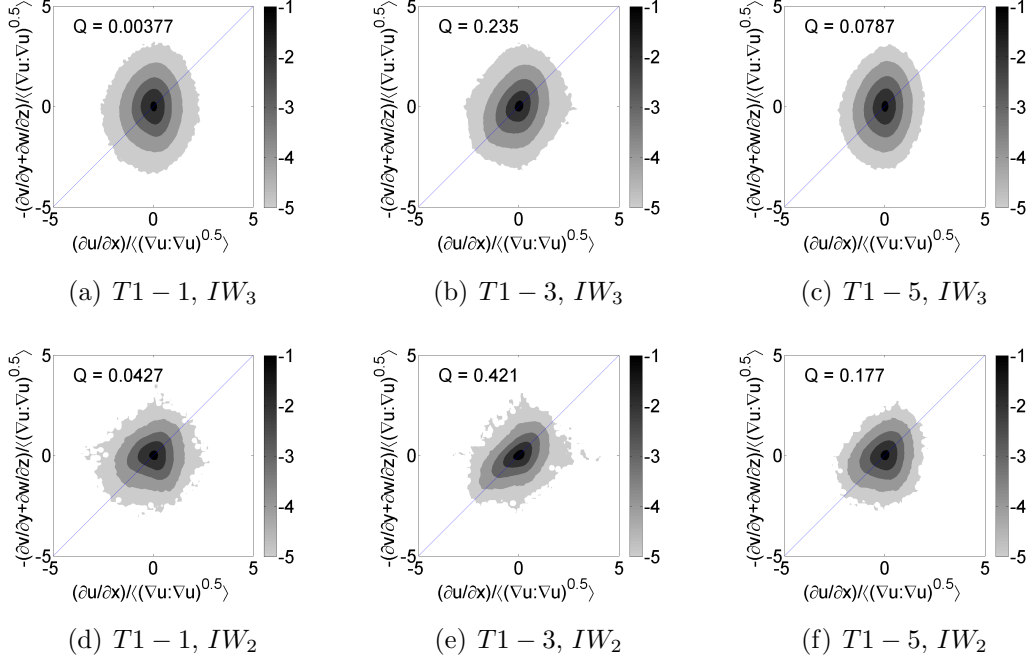
Figure 4.29: Divergence correlation JPDF, contours of base 10 exponentials

results from the non-uniform particle distribution in the volume; a problem which is exacerbated at low seeding densities, where the mean particle position (interrogation region centroid) and centrally positioned velocity vector may be different. Due to the random particle distribution method used in the numerical simulation this problem may not have been accurately modelled, which might explain the lower correlation values found in the experimental data. The correlation values are significantly lower than previous investigations (36; 123), especially the highest and lowest Reynolds number cases. However, as discussed in Section 3.2.4 this accuracy reduction is related to the noise introduced in this full 3D method, as opposed to higher accuracy single point hot-wire and planar stereo PIV measurements.

4.5.2 Normalised Local Divergence

Although regions of high gradient magnitude may produce relatively large absolute divergence contributions, these may be small in relation to the gradients,

140

and therefore it is also useful to study the normalised local divergence.

(a) Gradient magnitiude normalised divergence

(b) Normalised divergence ratio, ξ PDF
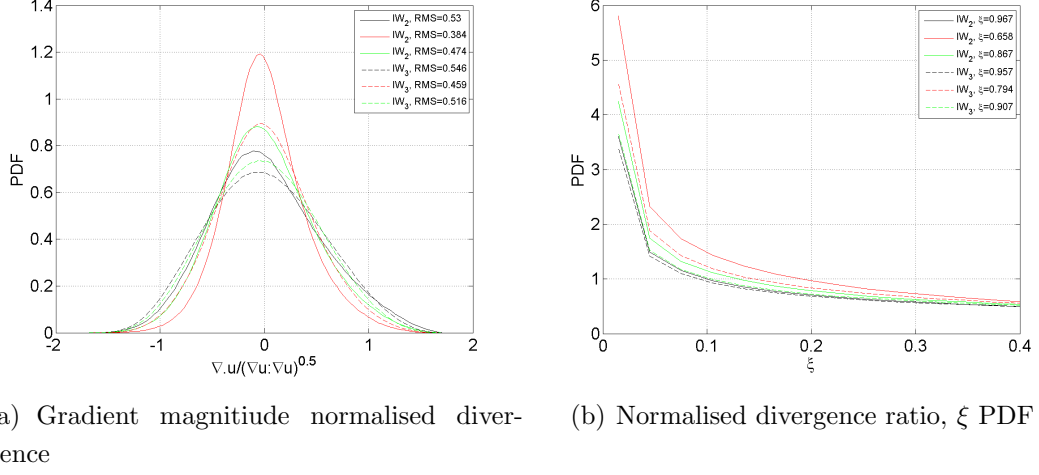
Figure 4.30: Gradient and divergence magnitude JPDF

Figure 4.30(a) shows the effect of discretisation on a PDF of local divergence normalised by the local norm of the velocity gradient tensor, $\chi = \nabla \cdot u/(\nabla u : \nabla u)^{1/2}$. The larger window reduces the RMS variation slightly, although except for the $T1-3$ case, the divergence is significantly higher than the value of 0.35 found in dual plane stereo-PIV measurements by Mullin & Dahm (82). A similar reduction in divergence error is observed from the PDF of the normalised divergence ratio, ξ (see equation 3.3), shown in Figure 4.30(b). Again however, only the $T1-3$ case gives a value at the smallest interrogation volume size, IW_3, similar in magnitude to the value of 0.74 calculated by Zhang et al. (142).

4.5.3 Divergence Location

Lower magnitude gradient regions will be more susceptible to inherent PIV errors, and therefore although the flow contains a reasonable amount of divergence the location of this should be quantified.

Figure 4.31 illustrates this relationship through the JPDF of normalised divergence against velocity gradient magnitude. The characteristic pyramid shape indicates that regions of low and high gradient magnitude are associated with the highest and lowest divergence respectively. Again comparing Figures 4.31(a) to
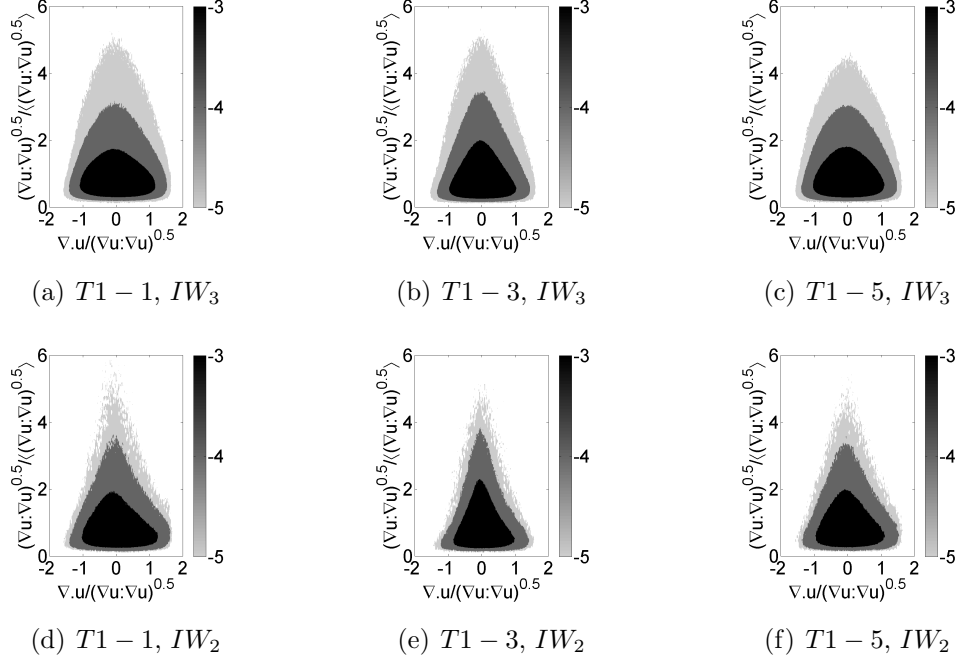
Figure 4.31: Gradient and divergence magnitude JPDF, contours of base 10 exponentials

(c) and 4.31(d) to (f) shows the effect of discretisation, with larger window sizes reducing the divergence error and resulting in a sharper distribution peak.

The current study is focused on high gradient magnitude regions associated with coherent structures, which this analysis indicates will have reasonably low divergence errors. Therefore this association is now examined through two different divergence reduction methods.

4.5.4 Divergence Thresholding

The effect of thresholding on the divergence correlation, and gradient magnitude JPDFs is shown in Figures 4.32 and 4.33 respectively. The thresholding is based on the normalised divergence ratio, $\xi < \xi_{rms}$, removing around 35% of the data with the lowest gradient SNR.

This dramatically increases the correlation values in Figure 4.32, which are shown to be particularly sensitive to these normalised divergence values. As ex-

(a) $T1-1$, IW_2 (b) $T1-3$, IW_2 (c) $T1-5$, IW_2
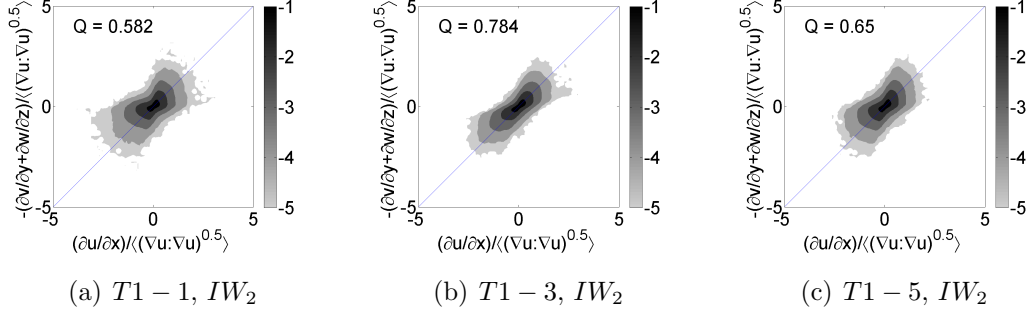
Figure 4.32: Divergence correlation JPDF conditioned on low normalised divergence, $\xi < \xi_{rms}$, contours of base 10 exponentials

(a) $T1-1$, IW_2 (b) $T1-3$, IW_2 (c) $T1-5$, IW_2
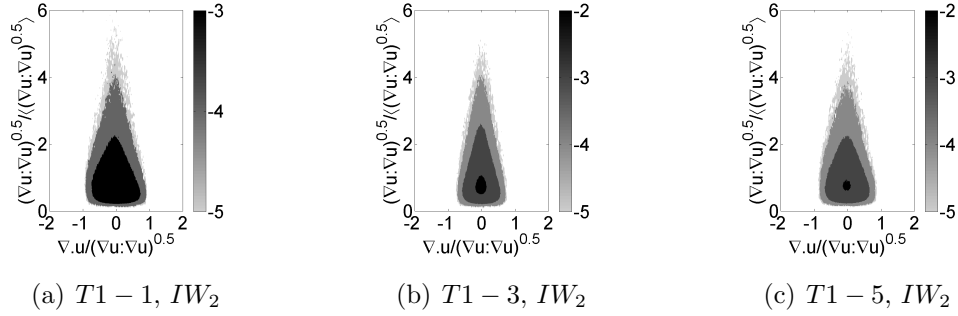
Figure 4.33: Gradient and divergence magnitude JPDF conditioned on low normalised divergence, $\xi < \xi_{rms}$, contours of base 10 exponentials

pected Figure 4.33 shows that the removal of these points does not significantly affect the high gradient magnitude regions, modifying the distribution shape towards that shown by Ganapathisubramani et al. (36).

4.5.5 Divergence Correction

Another illustration of this association was performed by removing divergence using the method of convex projections (111). In this method a vector field correction containing the same divergence as the original vector field is calculated and subtracted from the original field. The problem is re-cast in terms of a three-dimensional Poisson equation with the solution to this providing the closest divergence free solution to the original field in a least squares sense. The domain is zero-padded at the edges, and the dirichlet condition is applied at these bound-

aries, with the solution equal to zero at these points. Although at low resolution the velocity fields should not be entirely divergence free due to the effects of spatial filtering, this inherent divergence error was shown to be small for resolutions $\Delta x/\eta \leq 8$ (see Section 3.2.4).

Application of this process for a single $T1 - 3$ case velocity field reduced the normalised divergence ratio from $\xi = 0.76$ to $\xi = 0.02$, through changes of 5.8, 7.6, and 6.7% in the u, v and w components respectively. The magnitude of the velocity correction is similar to the estimated error for this case.

(a) Original, $\xi = 0.76$　　　　　　(b) Divergence corr., $\xi = 0.02$
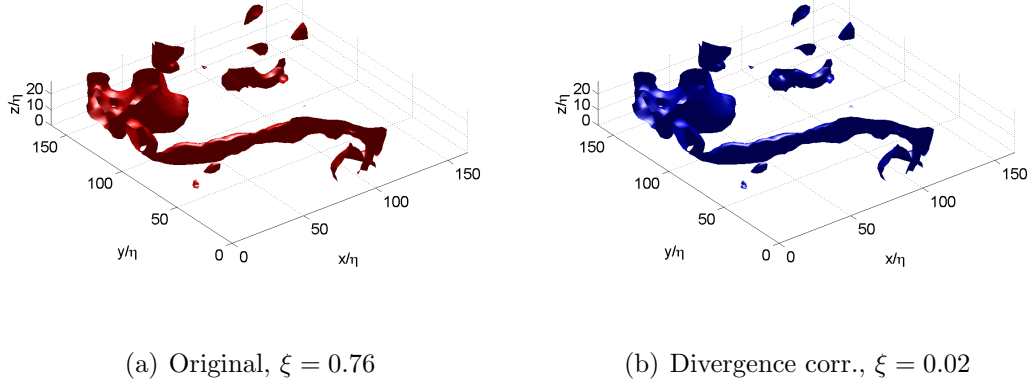
Figure 4.34: Iso-surfaces of vorticity magnitude $|\omega| > \omega_{rms}$

(a) Original, $\xi = 0.76$　　　　　　(b) Divergence corr., $\xi = 0.02$
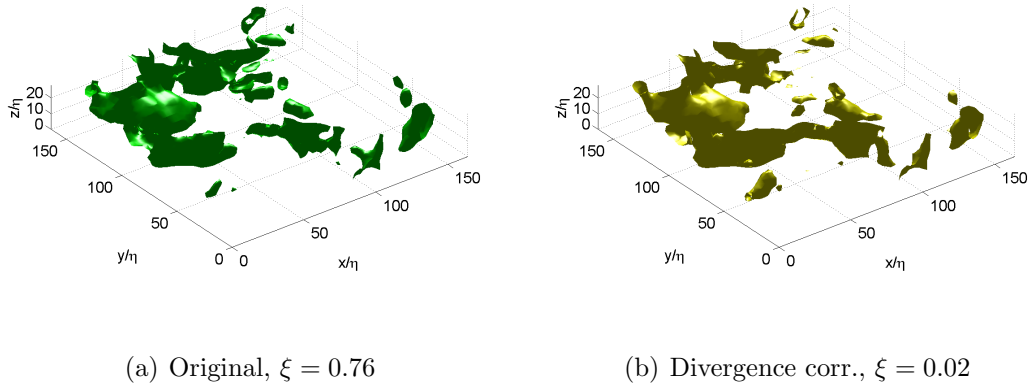
Figure 4.35: Iso-surfaces of dissipation rate magnitude $\epsilon > 2\epsilon_m$

		Original			Thresholding			Corrected		
	Window	ξ	χ_{rms}	Q	ξ	χ_{rms}	Q	ξ	χ_{rms}	Q
$T1-1$	IW_2	0.967	0.530	0.0427	0.323	0.313	0.582	0.0492	0.0683	0.980
	IW_3	0.957	0.546	0.00377	0.331	0.325	0.538	0.139	0.132	0.928
$T1-2$	IW_2	0.932	0.506	0.117	0.302	0.298	0.625	0.0493	0.0657	0.980
	IW_3	0.928	0.527	0.058	0.311	0.311	0.576	0.129	0.123	0.936
$T1-3$	IW_2	0.658	0.384	0.421	0.193	0.228	0.538	0.042	0.0567	0.984
	IW_3	0.794	0.459	0.235	0.249	0.27	0.682	0.116	0.111	0.945
$T1-4$	IW_2	0.787	0.440	0.255	0.243	0.262	0.711	0.0577	0.0699	0.976
	IW_3	0.853	0.486	0.156	0.273	0.287	0.649	0.129	0.119	0.940
$T1-5$	IW_2	0.867	0.474	0.177	0.275	0.28	0.650	0.0692	0.0781	0.971
	IW_3	0.907	0.516	0.0787	0.308	0.309	0.580	0.154	0.136	0.910

Table 4.3: Divergence Values

Figures 4.34 and 4.35 show divergence correction does not significantly affect the high magnitude vorticity and dissipation distributions, reaffirming that the divergence is not strongly associated with these high gradient regions. Furthermore, the correction changes the mean dissipation rate estimate by less than 2%. This implies that the current investigation of these high gradient regions will be largely unaffected by the divergence error present. Mean divergence values at different resolutions and after both thresholding and correction schemes are listed in Table 4.3.

4.5.6 Effect of Divergence and Resolution on Topological Invariants

A number of JPDFs have been constructed using the invariants introduced in Section 1.1.6.2 in order to study the flow topology. The first second and third invariants, P_A, Q_A, and R_A are calculated from the summation of multiple first, second and third order gradient term combinations respectively, and therefore are likely to be highly sensitive to noise. Furthermore, although regions containing high vorticity and dissipation were shown to be relatively insensitive to the divergence error, this may still have a significant effect on the higher order gradient products and lower magnitude gradient regions. Although the numerical

(a) IW_2, orig.   (b) IW_2, div. thresh.   (c) IW_2, div. corr.

(d) IW_3, orig.   (e) IW_3, div. thresh.   (f) IW_3, div. corr.
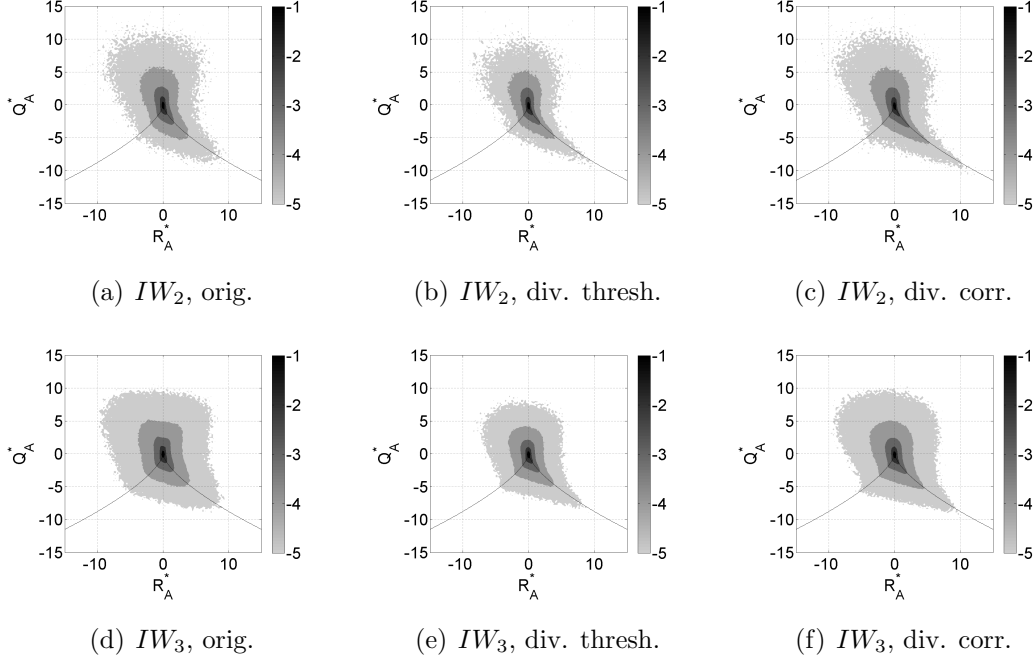
Figure 4.36: The effect of resolution and divergence on the Q-R plane for the $T1 - 3$ case, contours of base 10 exponentials

accuracy investigation demonstrated the manifestation of these effects as a loss of distribution definition (see Section 3.2.6), the higher divergence associated with the experimental results requires further analysis.

Figure 4.36 shows JPDFs of the non-dimensional invariants Q_A^* and R_A^* at two different resolutions for the $T1 - 3$ case. Divergence effects are assessed using two methods: firstly thresholding by $P_A^* = P_A/\langle \overline{Q_W} \rangle^{1/2}$ with the limits $|P_A^*| < (2/3)P_{A,rms}^*$, which excludes around 35% of the highest divergence data; and secondly through the divergence correction scheme introduced in Section 4.5.5.

Comparing distributions at different resolutions demonstrates the effect of noise, which causes blurring of the higher resolution contours with an increased tendency towards SN/S/S topology. The significant reduction of this SN/S/S preference in the lower resolution results indicates this is a noise related arte-fact. To confirm this, a 1×10^7 point Monte Carlo simulation was conducted to demonstrate the effect of random gradient field errors ($-1 < \partial u_i/\partial x_j < 1$) on the topological invariants, shown in Figure 4.37. The substantial effect of noise
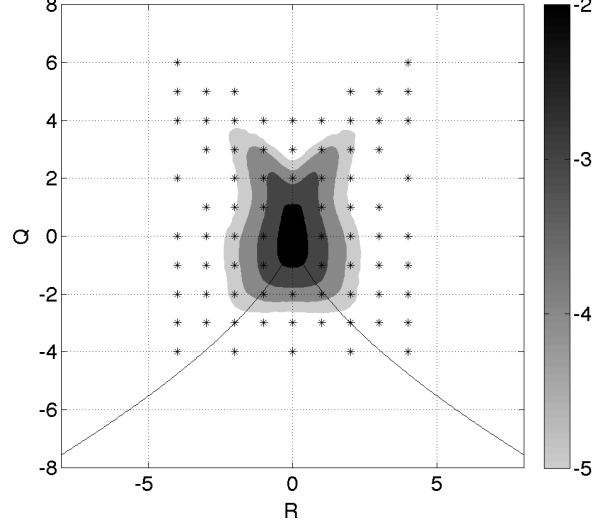
Figure 4.37: Error in Q-R plane due to random gradient field noise, $-1 < \partial u_i / \partial x_j < 1$, contours of base 10 exponentials, black stars represent combinations of the maximum possible errors

in the SN/S/S quartile strengthens the argument that this is the cause of the invariant plot deformation.

Both thresholding and divergence correction reduce distribution blur, particularly around the teardrop tip, improving the agreement with previous investigations (84; 85).

Divergence effects are investigated further by studying trivariate JPDFs of the normalised invariants P_A^*, Q_A^*, and R_A^*, shown in Figure 4.38. The 9 subplots correspond to different P_A^* bin locations, with the upper and lower bounds and the percentage of points contained in each bin indicated at the top of each subplot. These 9 bins account for all points within $\pm 2P_{A,rms}^*$, which is approximately 96% of the total number of points.

As the magnitude of P_A^* increases the expected teardrop shape appears to rotate about the origin. This behaviour is explainable through the definitions of these two invariants (see Equations 1.14 to 1.16), and from knowledge of the relative invariant magnitudes (in general $P_A^* << Q_A^*$). If P_A^* and Q_A^* are both positive the first terms in the invariant R_A^* will become $R_A^* = -P_A^{*3} + 3P_A^* Q_A^* \ldots$,
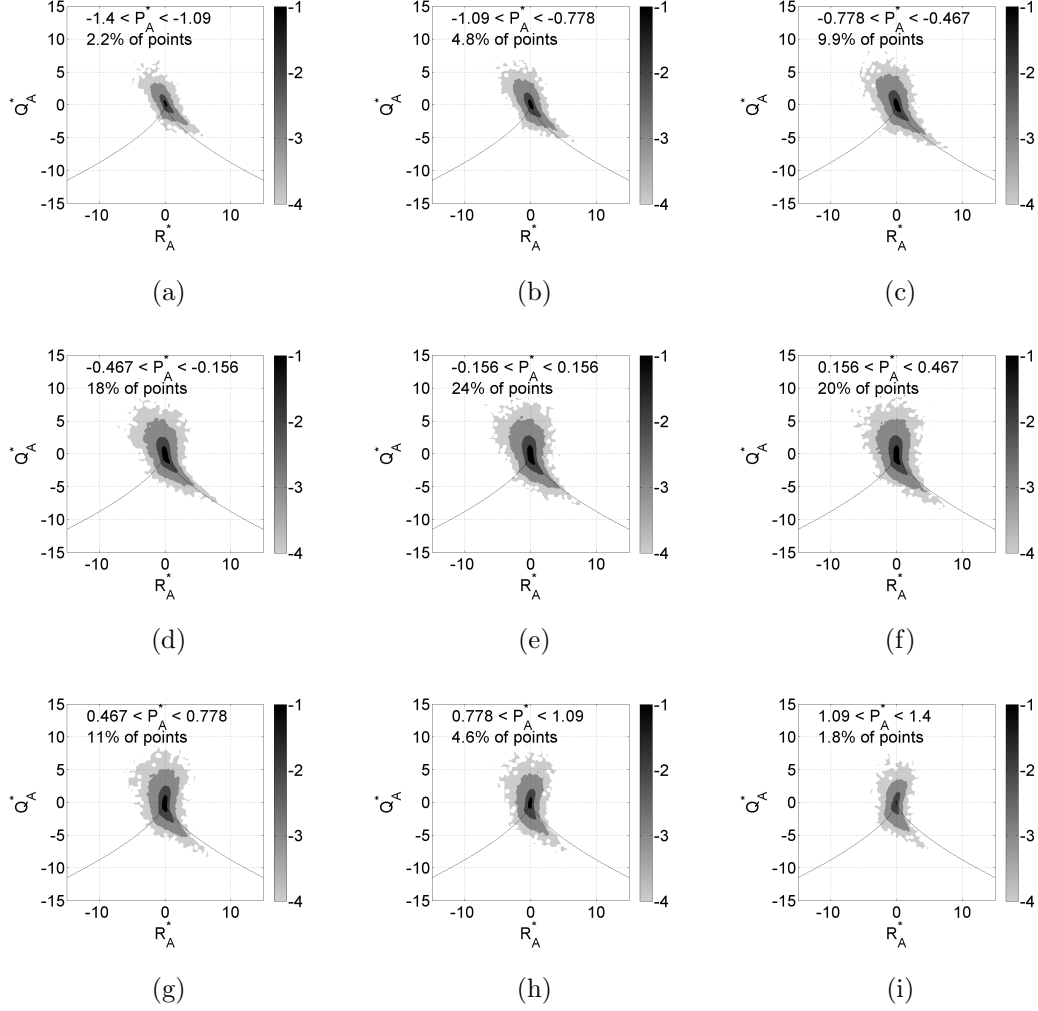
Figure 4.38: Trivariate JPDF slices of normalised topological invariants P_A^*, Q_A^* and R_A^* for high resolution case, contours of base 10 exponentials

and as $3Q_A^* P_A^* >> P_A^{*3}$, R_A^* will tend towards positive values. On the other hand if P_A^* is positive but Q_A^* is negative the leading terms become $R_A^* = -P_A^{*3} - 3P_A^* Q_A^* \ldots$, meaning R_A^* will tend towards negative values. As the magnitude of P_A^* is increased the effect of these tendencies is also increased, resulting in an increasing net clockwise rotation on the $Q_A^* - R_A^*$ plane. Conversely, it can be shown that negative values of P_A^* will produce an anti-clockwise rotation. The addition of these high divergence regions are responsible for the $Q_A^* - R_A^*$ plane blurring shown in Figure 4.36. The divergence threshold applied previously

effectively removes points from the first and last 3 subplots (Figure 4.38(a) to (c) and (g) to (i)), where the tip is furthest from the $\Delta_A = 0$ line, resulting in the sharper distribution shown in Figure 4.36(b).

Although thresholding results on a normalised divergence criterion (results with low gradient SNR) was also tested, no significant improvements were observed. It is likely that regions containing low SNR correspond to small gradients and also low magnitude invariants, and therefore any distortion of the JPDF is likely to occur at its centre where effects are masked by noise.

Therefore, due to this extreme sensitivity to noise and divergence the slightly lower resolution results (from the second cross-correlation pass, IW_2) will be used in Sections 5.2.1 and 5.3.4, with a comparison made in each case against the divergence free results.

Chapter 5

Coherent Structure Characterisation

Having quantified measurement accuracy and resolution, characterised the Reynolds number scaling and nature of the statistically steady flow, and demonstrated the consistency of the DPIV and TPIV data, the latter will now be used to investigate the dissipation scale coherent flow structures. As described in Section 1.1.6, the use of full 3D volumetric data permits a more in depth study of flow topology, which should allow a more intuitive understanding of the flow. Additionally, the availability of the full velocity gradient tensor gives access to statistical topological analysis techniques, which are free from the assumptions of isotropy.

The chapter begins with a visual characterisation of these structures, before statistically quantifying these observations using a variety of methods. The role of the strain field in the formation of these structures is then assessed, and a number of formation mechanism examples reviewed.

5.1 Coherent Structure Visualisation

Before the visual characterisation a comparison of different vortex identification schemes and thresholding effects is conducted, in order to select an appropriate scheme. Next, some observations are made with regard to the size, shape and distribution of coherent flow structures across a range of Reynolds numbers.

5.1.1 Vortex Identification

As described in Section 1.1.7, different authors employ a range of vortex identification schemes, and at present there exists no commonly agreed standard. Therefore, in order to choose a method for the present investigation different schemes were compared and their relative performance assessed.

(a) Enstrophy, $\Omega > n\Omega_{rms}$

(b) Second Invariant, $Q_A > nQ_{Arms}$

(c) Swirling Strength, $\lambda_{ci} > n\lambda_{ci,rms}$

(d) Discriminant, $\Delta_A > n\Delta_{A,rms}$

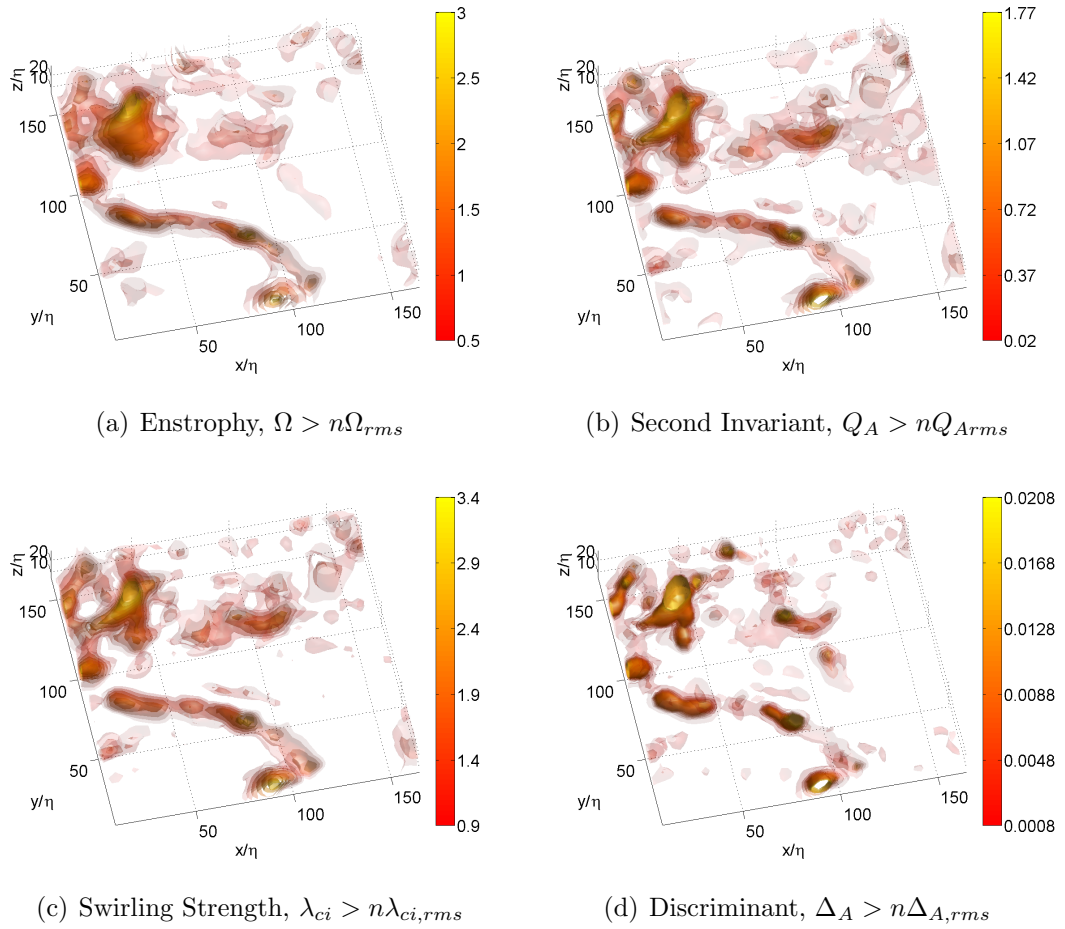Figure 5.1: Comparison of vortex identification methods; colour bars of isosurface level, n

A comparison of four different identification methods is presented in Figure 5.1 for a single example field from the $T1 - 3$ case. Iso-surfaces are plotted across a range of threshold levels with the colour map indicating the level, n, which is a multiple of global RMS property variation, with each scheme requiring a different

set of thresholds. To aid viewing, the lowest threshold levels are made transparent with reducing transparency as the threshold level in increased, allowing a range of iso-surface levels to be compared on the same plot, which is extremely useful.

In agreement with the study by Dubief & Delcayre (27), there is reasonable agreement between the majority of schemes, with similar structures identified by all. In particular, enstrophy, swirling strength and positive Q_A criterions give very similar flow structure distributions. Although similarities are also seen in Figure 5.1(d) for the discriminant criterion, this method shows the greatest differences from the other techniques, with a less fluid transition between strong and weak regions.

The techniques can be grouped into two distinct categories by threshold level sensitivity. The enstrophy and swirling strength methods are the most sensitive, with moderate threshold changes producing significant structural changes, whereas the positive Q_A and discriminant criteria are relatively insensitive, with large threshold changes required to produce structural effects. For example, to produce similar structural variations using the enstrophy and Q_A criterions, threshold ranges of around 6 and 90 times the lowest values respectively are required.

If a well defined minimum threshold value existed insensitivity may be viewed as an advantage, however both the positive Q_A and discriminant schemes require a 'slightly positive' threshold, with this value significantly altering the distribution. From another perspective a threshold sensitive method may be preferable, as this would allow the viewer to visually sort vortices by strength; a possible advantage in characterising coherent structures which possess a correlation between topology and strength.

Another consideration is the accuracy of the different criteria with respect to random gradient field errors ($-1 < \partial u_i / \partial x_j < 1$), which are assessed through a 1×10^6 Monte Carlo simulation. The mean error magnitude is first normalised by the mean criterion value, and second by the enstrophy criterion error, with this relative error shown in Figure 5.2. Although the enstrophy criterion is the most accurate for modest noise levels, the swirling strength and positive Q_A both shown similar accuracy. However, due to the use of cubed gradient terms during its calculation, the accuracy of the discriminant criterion is significantly lower,
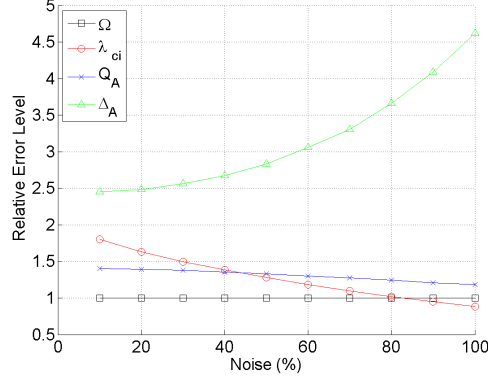
Figure 5.2: Relative accuracy of vortex identification schemes

which may explain the slightly different distribution of structures predicted by this scheme.

Therefore due to its accuracy, threshold sensitivity, and simplicity the enstrophy criterion has been adopted in the present investigation.

5.1.2 Coherent Structure Definition

To define the term 'coherent structure' in the present investigation a further thresholding study was conducted. Figure 5.3 shows enstrophy iso-surfaces for four different threshold ranges, the last three of which are below the level $\Omega < \Omega_{rms}$, which has been previously used as a division between the incoherent background and coherent flow structure (53).

Figure 5.3(a) shows high enstrophy iso-surfaces above the coherent threshold. Figure 5.3(b) shows the greatest dependence on this coherent flow structure, occurring predominantly on its periphery and emphasising the absence of a sharp coherent to background cut-off. As the threshold level is further reduced the distribution becomes an increasingly complex tangle of inter-connected regions as shown in Figure 5.3(c). Although these lower enstrophy regions are given some structure by default, existing only where higher enstrophy does not, the flow structure is visibly less well defined. Finally at the lowest background level shown in Figure 5.3(d) the distribution becomes spotty, with noise expected to dominate the distributions.

(a) $\Omega_{rms} > \Omega > 6\Omega_{rms}$

(b) $0.2\Omega_{rms} > \Omega > \Omega_{rms}$

(c) $0.03\Omega_{rms} > \Omega > 0.2\Omega_{rms}$

(d) $0 > \Omega > 0.03\Omega_{rms}$
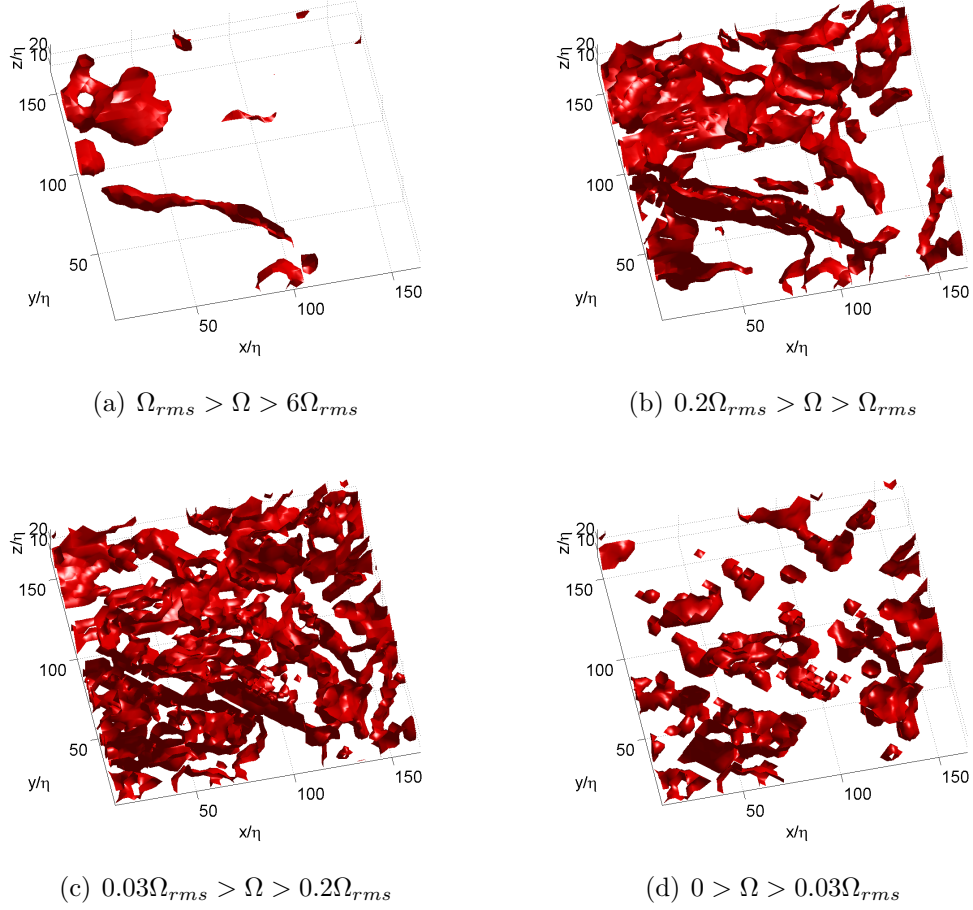
Figure 5.3: Isosurfaces of enstrophy for a range of threshold levels

Despite structure arising from its enforced containment, the complexity of the background enstrophy distribution would make studying its instantaneous distribution extremely difficult. Therefore, the current study will focus on coherent structures in regions where $\Omega > \Omega_{rms}$. Although this may exclude lower enstrophy structures and truncate others, a threshold at some value is necessary, to simplify the distribution enough for investigation. Jiménez et al. (53) suggest that high enstrophy structures are simply intense realisations of the background vorticity, and therefore through their study it may also be possible to understand the weaker background structures.

5.1.3   Visual Characterisation of Coherent Structures

In order to qualitatively analyse coherent structures two fields were selected from each different Reynolds number run, shown in Figure 5.4. Iso-surfaces of enstrophy and dissipation rate across a range of coherent threshold values are plotted, again using transparency to increase viewability. The enstrophy and dissipation rate use different colour maps, representing threshold levels of $\Omega > n\Omega_{rms}$ and $\epsilon > n\epsilon_m$ across the range, $1 < n < 4$.

A number of shapes can be identified from Figure 5.4 with a predominance of short tube-like and ribbon-like structures for both enstrophy and dissipation. The diameter and thickness of these varies from approximately $2 \rightarrow 20$ grid points, with the length of the ribbon-like structures varying from around $4 \rightarrow 51$ grid points (up to the maximum domain size). Accounting for spatial resolution reduction at high Reynolds numbers produces structure diameter and thickness variations of $2 \rightarrow 140\eta$, and ribbon-like structures length variation of $4 \rightarrow 360\eta$. The wrinkled surface of the majority of structures is probably due to noise rather than structure deformation by smaller scales, with similar artefacts present in the accuracy simulation (see Section 3.2.1).

The majority of structures appear to be orientated with the out-of-plane z-component, with a similar bias also seen during the accuracy simulations, arising from a combination of sources: firstly the under prediction of gradient terms shown in Section 4.2.3.1 gives preference to the out-of-plane component of vorticity; and secondly the smaller volume size in the out-of-plane direction severely truncates structures exceeding this size. For example, if tube-like structures are orientated out-of-plane their length will shortened to the volume thickness, but if these are aligned in-plane and their diameter exceeds the volume thickness then they may be truncated into ribbon-like structures. Similarly, unless sheet-like structures lie completely in-plane these will also be truncated into ribbon-like structures. The shape of structures is quantified in Section 5.2.3.2.

Furthermore, structure size may also be grid limited, with even the highest resolution case ($\Delta x/\eta \approx 1$) identifying structures of order Δx. Therefore, even at this resolution the measurements may not be capturing the very smallest scales. This observation is investigated further in Section 5.2.3.1.

(a) $T1-1$, example1

(b) $T1-1$, example2

(c) $T1-2$, example1

(d) $T1-2$, example2

(e) $T1-3$, example1

(f) $T1-3$, example2

Figure 5.4: Isosurfaces of enstrophy, $\Omega > n\Omega_{rms}$ (red-yellow colour map), and dissipation rate, $\epsilon > n\epsilon_m$ (blue-green colour map)

(g) $T1 - 4$, example1

(h) $T1 - 4$, example2

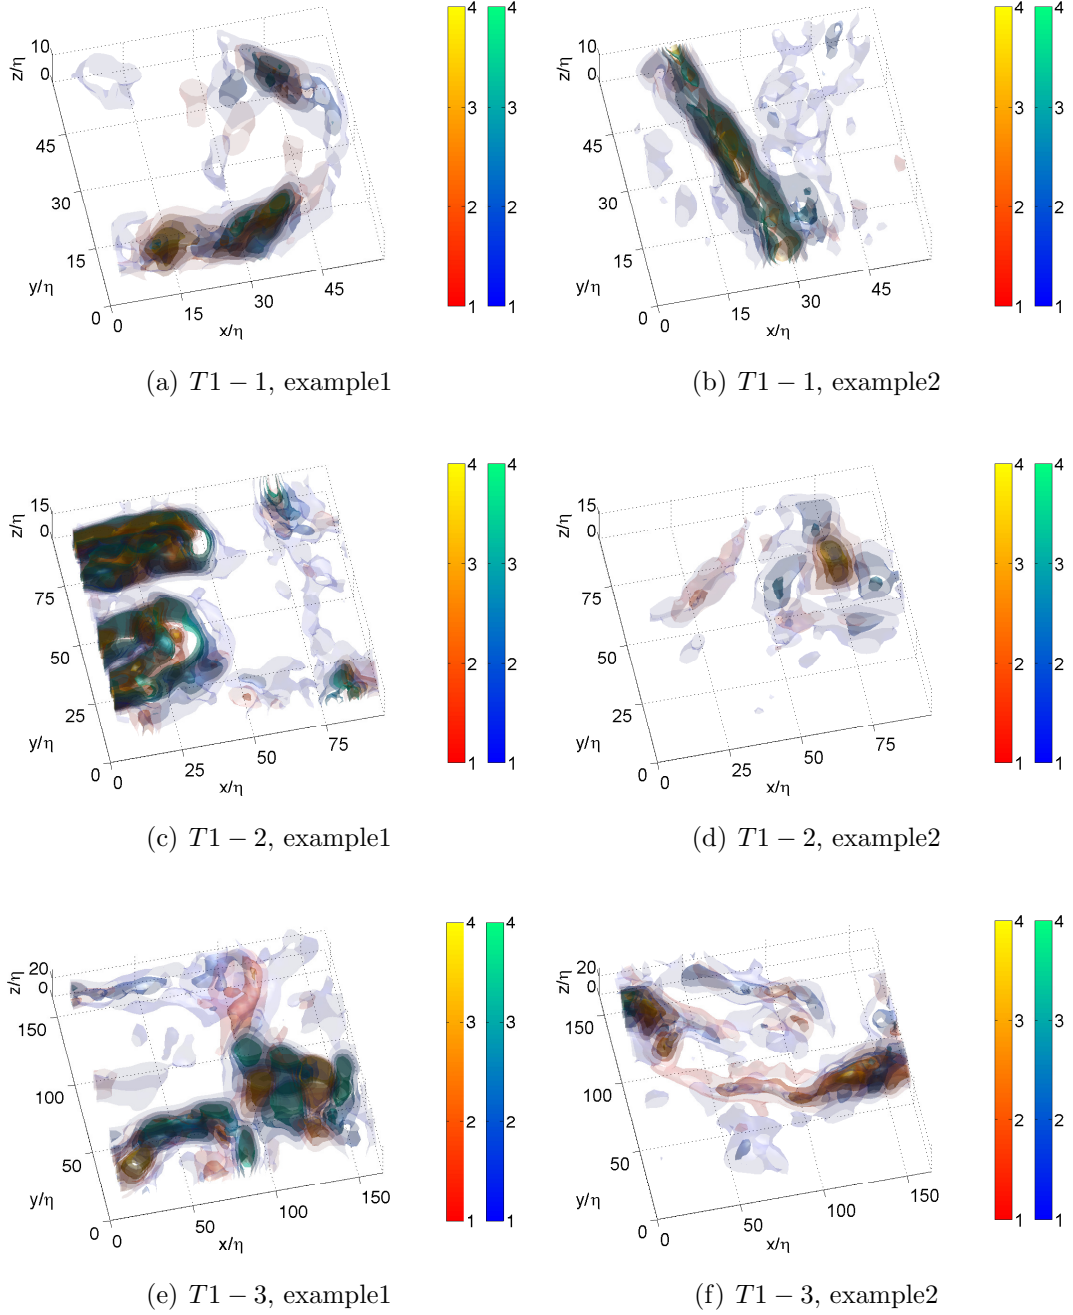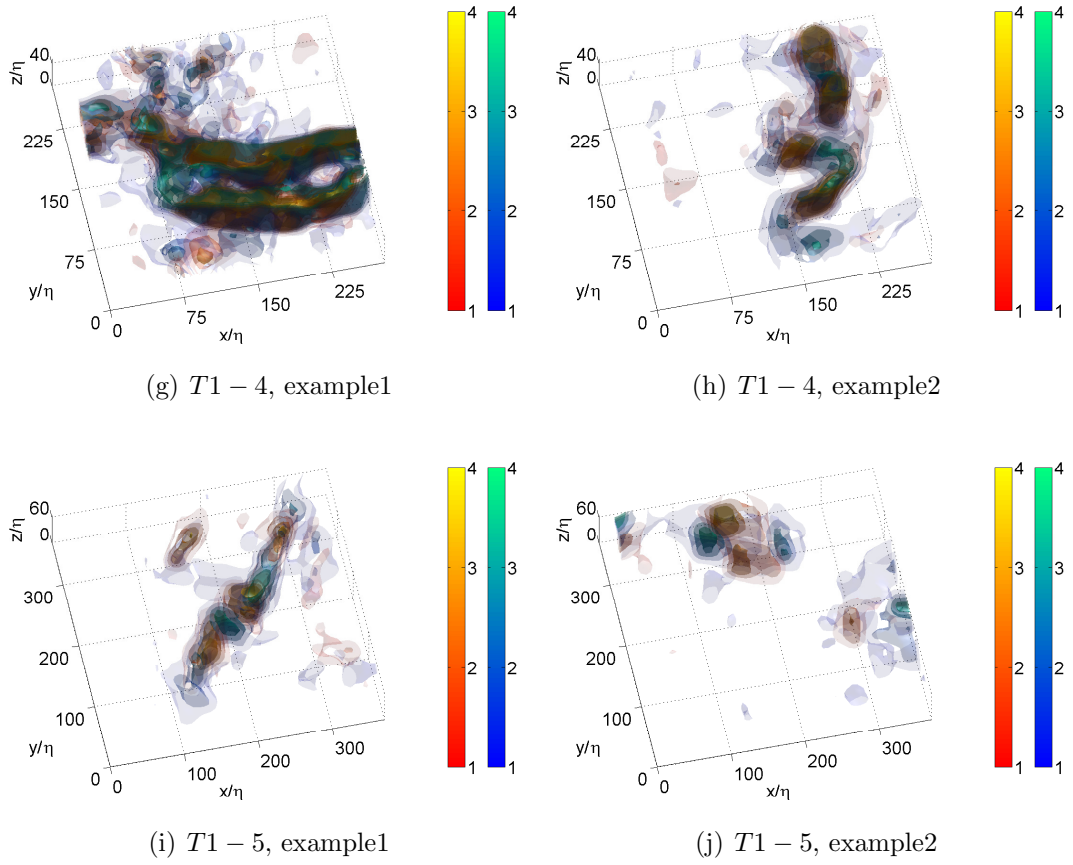(i) $T1 - 5$, example1

(j) $T1 - 5$, example2

Figure 5.4: Isosurfaces of enstrophy (continued), $\Omega > n\Omega_{rms}$ (red-yellow colour map), and dissipation rate, $\epsilon > n\epsilon_m$ (blue-green colour map)

Enstrophy and dissipation disassociation has been shown by a number of authors to occur around tube-like structures (37; 58; 125), and is often likened to a Burgers vortex where a core of enstrophy is surrounded by a sheet of dissipation (20). A number of examples can be seen in Figure 5.4, where tube-like enstrophy structures are surrounded by ribbons and tubes of dissipation. Further to this Ganapathisubramani et al. (37) note that the most intense dissipation occurs when several structures are clustered close together, with crumpled sheets of dissipation occurring between them. Again numerous examples of this can be seen, where strong dissipation structures are sandwiched between strong enstrophy structures.

However, in other cases, within moderate-to-intense ribbon-like structures regions of high enstrophy and dissipation are often coincident. These regions are consistent with the intense shear layers described by some authors (53; 99), which are thought of as either a precursor to the tube-like structures or a consequence of them, occurring on their periphery. The contribution of these structures to the global enstrophy and dissipation is quantified in Section 5.2.2.

The strength of the structures also varies considerably, with some structures just visible above the lowest threshold while others contain regions over 4 times this strength. In some cases stronger regions simply represent a structure's core, with these surrounded by a series of concentric tubes or ribbons. Alternatively, more than one high enstrophy core can be nested within larger lower magnitude ribbon or tube-like structures, forming an embedded cluster of structures.

Figure 5.4 shows that the structures have a strong tendency to cluster together in this manner, forming groups of tube and ribbon-like structures, with isolated structures occurring more rarely. The clusters often appear to be located at the periphery of larger tube and ribbon-like structures which advect through the volume. This behaviour is similar to that observed by Jiménez et al. (53) who noted that the majority of structures tend to occur at the interface between generally empty larger scale eddies, and quantified by Moisy & Jiménez (79) who found that structures form clusters at a range of scales equal to the inertial range. Similar sized clusters (in terms of grid points) can be observed throughout the range of Reynolds numbers, with this qualitatively similar behaviour perhaps indicative of scale-invariance, as observed in other investigations (56).

The clusters tend to be sparsely distributed, surrounded by lower background vorticity and dissipation, demonstrating the intermittency phenomenon. It is also interesting to note that there are several examples of long chains of tube-like structures, which bear resemblance to the breakdown of a vortex sheet through the Kelvin-Helmholtz instability. Also of interest are the closely grouped bundles of tube-like structures, which bear some similarity to the intertwined 'burst' vortex filaments observed by Cadot et al. (20). The clustering is quantified in Section 5.2.4, with the latter observations investigated further in Section 5.3.5.

Finally, it should be noted that the fields in Figure 5.4 represent a tiny fraction of the data reviewed, and a better understanding of the typical flow structures, noise related structural wrinkling, clustering, and intermittency can be gained by studying the time dependant nature of these structures, and the reader is referred to the video files in Appendix D.

5.2 Statistical Characterisation of the Coherent Structures

Observations from the visual characterisation are quantified statistically in this section using a variety of methods. First the spatial distribution of vorticity and dissipation are examined through JPDFs of their topological invariants, and through conditional averaging the relative impact of coherent structures on global flow properties is assessed. Next, the structure geometry is characterised using the vortex line tracking and box counting algorithms introduced in Section 2.4.4 with the latter also used to quantify clustering.

5.2.1 Dissipation and Enstrophy Distribution

Figure 5.5 shows JPDFs of normalised strain rate and rotation rate invariants $-Q_S^*$ against Q_W^*; quantities which are analogous to the dissipation and enstrophy density respectively (92). The self-similar flattened triangular distribution indicates that although there is a weak correlation between these quantities there is also a certain degree of separation at the highest values. Despite reasonable general agreement, the ratio of $(Q_W)_{MAX}/(-Q_S)_{MAX}$ for the $T1-1$ and $T1-5$

(a) $T1-1$ (b) $T1-3$ (c) $T1-5$

(d) $T1-1$, div. corr.   (e) $T1-3$, div. corr.   (f) $T1-5$, div. corr.
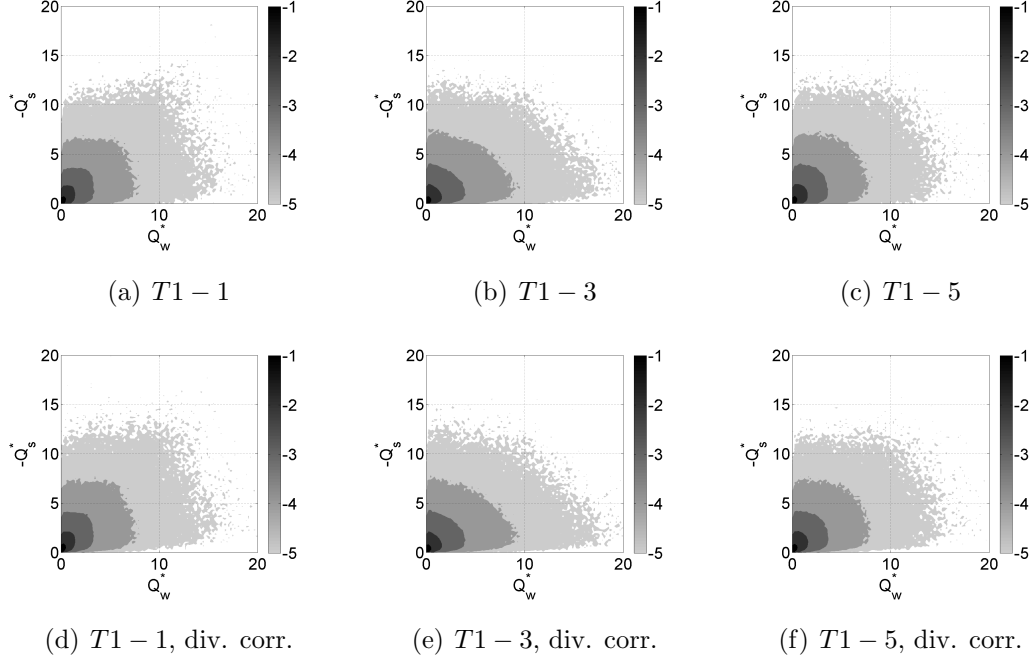
Figure 5.5: JPDF of normalised invariants $-Q_S^*$ and Q_W^* at three Reynolds numbers, with and without divergence correction, contours of base 10 exponentials

cases is ~ 1.5, which is slightly less than the ratio ~ 2 predicted by other studies (53; 85; 141). The $T1-3$ case predicts a slightly higher ratio (~ 1.7) which suggests that the under-prediction in other cases is a result of noise, which may also explain the squarer shaped distributions at the lowest and highest Reynolds number.

The separation of enstrophy and dissipation shown in Figure 5.5 is similar to the DNS results of Ooi et al. (85). The high Q_W^* and low $-Q_S^*$ region is interpreted as solid body rotation with disassociated dissipation, which the authors note may result in extended life times; an observation confirmed in the present study by the time-resolved videos in Appendix D, which show long structure life times in comparison with the Kolmogorov time scale. These regions may therefore correspond to the tube-like structures observed in Section 5.1.3, which also demonstrated a separation of enstrophy and dissipation. Similarly the regions of moderate enstrophy and moderate to high dissipation may be representative of the ribbon-like structures observed in Section 5.1.3, where significant enstro-

phy and dissipation was found to be coincident. Characterisation of the coherent structures in this manner would be extremely useful, and to this end these relations are investigated further during the classification of structure geometry in Section 5.2.3.2.

It is also interesting to note, in contrast to the $Q_A - R_A$ distribution (see Figure 4.36) the original and divergence free fields produce very similar distributions, due to the smaller divergence dependence of the 2nd invariant Q compared with the 3rd invariant R (see Equations 1.14 to 1.16). This reduced dependence helps to explain the earlier observation in Section 4.5.5, where removing the divergence had minimal effect on the vorticity and dissipation fields.

5.2.2 Prevalence and Significance of Coherent Structures

| (a) Volume Fraction | (b) Percentage Contribution |
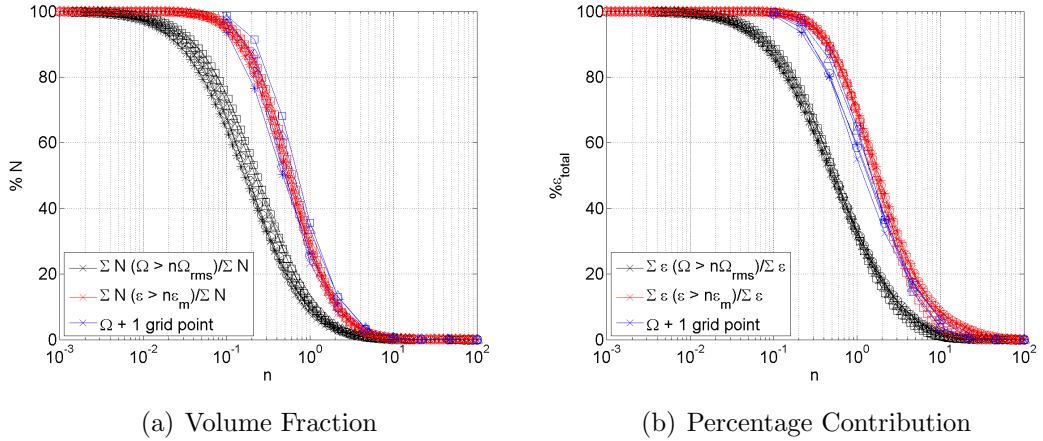
Figure 5.6: Intermittency and significance of high intensity structures; threshold by colour as indicated, case by symbol: $\times = T1 - 1$; $\bigcirc = T1 - 2$; $* = T1 - 3$; $\triangle = T1 - 4$; $\square = T1 - 5$

In order to quantify the intense structure intermittency and significance, their volume fraction and percentage contribution to the total dissipation rate are shown in Figure 5.6. The volume fractions were obtained by calculating the number of points greater than a given enstrophy or dissipation threshold compared to the total number of grid points. Similarly the dissipation rate percentage contri-

bution was found by integrating the dissipation of points above given enstrophy and dissipation thresholds in comparison to the total integrated dissipation.

Figure 5.6(a) shows that the volume fraction occupied by intense enstrophy structures is very low, with only around 10% of the flow exceeding a threshold of $n = 1$, and 2% exceeding a threshold of $n = 3$. The fraction occupied by intense dissipation structures is higher with around 30% of the flow exceeding a threshold of $n = 1$ and around 6% exceeding $n = 3$. These low volume fractions indicate an extremely intermittent distribution of these intense structures, particularly the enstrophy field. The Reynolds number dependence of the volume fraction found by previous authors (52; 59) cannot be seen in the present results, where instead there is a good collapse of data. The absence of this trend is expected to result from the loss of spatial resolution with increasing Reynolds number, with the collapse of data attributable to the previously observed scale invariance.

Figure 5.6(b) shows the most intense 10% of enstrophy structures are coincident with around 30% of the total dissipation, which is a significant fraction. However, from the previous volume visualisation (see Section 5.1.3) it was clear that regions of high enstrophy and dissipation are often highly correlated, with a tendency to disassociate in tube-like structures. Therefore, when determining the significance of these high enstrophy structures, the high dissipation regions surrounding these cannot be ignored. For example, Figure 5.6(b) shows that regions of dissipation greater than the mean value are responsible for the bulk of the dissipation ($\sim 70\%$), with even a higher threshold of $n = 2$ only reducing the contribution to around 40%. These values are extremely similar to those found by Ruetsch & Maxey (103). Given the high spatial correlation between these fields, it is possible that the majority of the dissipation is caused by these intense structures.

To quantify this another threshold was introduced, which includes regions surrounding high enstrophy structures. Points over a given enstrophy threshold were identified, and the volume fraction and dissipation contribution from these points and all neighbouring points calculated, shown in blue in Figure 5.6. By including the surrounding regions, it can be shown that approximately 25% of the flow gives rise to 60% of the total dissipation; a significant proportion. The reasonable agreement between this method and the standard dissipation threshold, in

addition to the previous flow structure observations in Section 5.1.3, indicates the high dissipation regions which are responsible for the majority of the dissipation are spatially connected to the high enstrophy structures. Further quantification of enstrophy location with respect to specific structures will be conducted in Section 5.2.3.2.

5.2.3   Geometry Characterisation

In order to characterise the size, strength and shape of the coherent flow structures two different algorithms are applied, following the methodology of Jiménez et al. (53), and Moisy & Jiménez (79).

5.2.3.1   Vortex Line Tracking Method

As described in Section 2.4.4.1, the vortex tracking algorithm of Jiménez et al. (53) was used to identify coherent structures in the statistically converged data sets, then calculate their radius and circulation.

Figures 5.7(a) and (b) show PDFs of structure radius in terms of grid and Kolmogorov units. The radius estimate in terms of grid units shown in Figure 5.7(a) are extremely similar. While it would be reasonable to expect structure radii in the lower resolution cases to demonstrate some grid dependence due to spatial averaging, the collapse of data indicates that spatial resolution may limit the radius estimates in all cases (although the grid dependence of the highest resolution case cannot be determined). It should also be noted that this grid limited behaviour may also arise from the smoothing operation. Therefore, when the radius estimates are expressed in terms of Kolmogorov units, the combination of constant grid unit size and reducing Kolmogorov length scale result in larger radius estimates for the higher Reynolds number cases, as shown in Figure 5.7(b). Therefore, for all cases except the highest resolution $T1-1$ case the radius estimates may not be representative of the dissipation scale coherent structures, and instead may represent some other slightly larger scale. The mean radius estimates are listed in Table 5.1.

Previous investigations indicate the radius of intense vortex tubes is around $3-5\eta$ (53). However, the known grid dependence of the $T1-2$ case ($\Delta x/\eta \approx 1.75$)

(a) Radius in grid units, $R_{1/e}/\Delta x$

(b) Radius in Kol. units, $R_{1/e}/\eta$
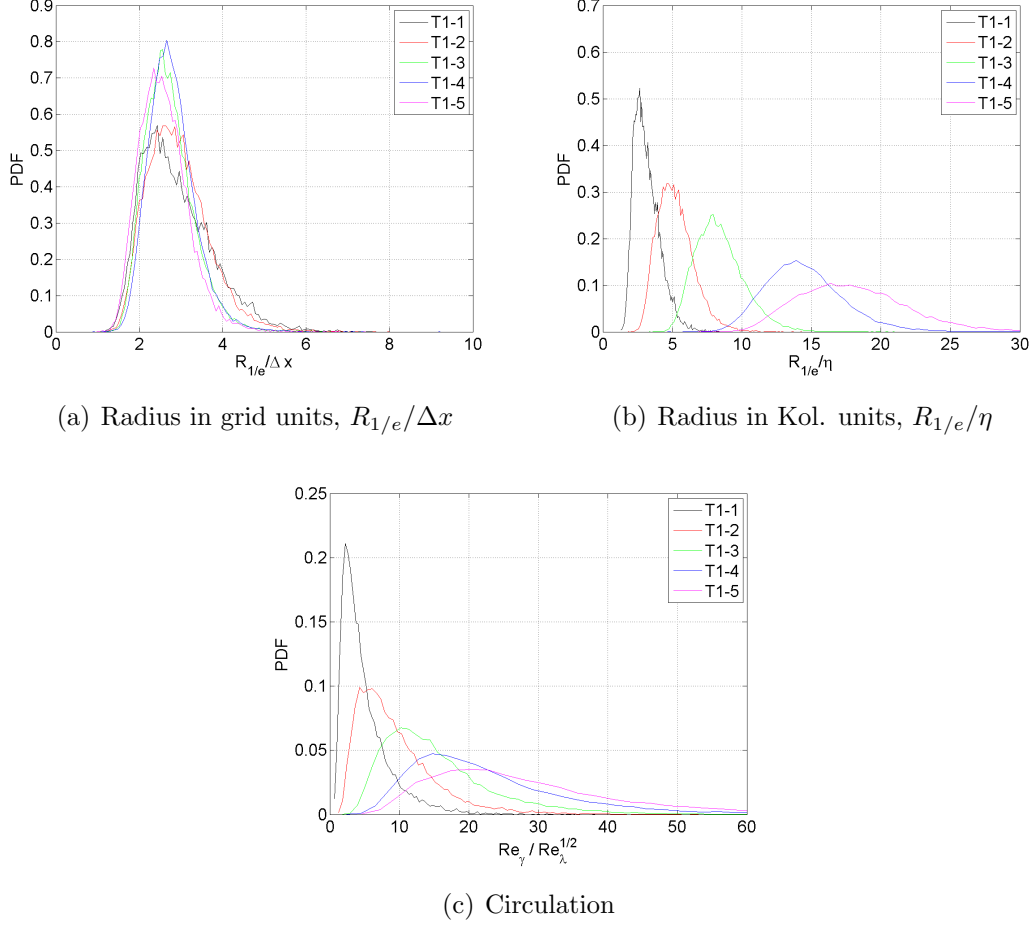
(c) Circulation

Figure 5.7: Radius and Circulation PDFs from vortex tracking analysis

indicates radius estimates from other similar resolution investigations may also be grid dependent. In particular, high Reynolds number DNS studies typically use a resolution based on the maximum wave number $\kappa_{max}\eta \approx 1.5$, which equates to a physical grid spacing of $\Delta x/\eta \approx 2.1$ (93). The experimental study by Ganapathisubramani et al. (37) also found similar radius structures using a PIV vector spacing of $\Delta x/\eta \approx 1.5$. The similarity between these grid spacings and that of the $T1-2$ case, may explain the similar radius estimate from this case. Therefore, in characterising the size of the most intense enstrophy regions higher resolution may be required.

Figure 5.7(c) shows a PDF of circulation normalised by $R_{\lambda}^{1/2}$. The circulation

Case	R_λ	res. $(\Delta x/\eta)$	N_{struc}	$R_{1/e}/\Delta x$	$R_{1/e}/\eta$	$Re_\gamma/R_\lambda^{1/2}$	Re_γ
$T1-1$	161	1.09	2420	2.9	3.17	5.02	64
$T1-2$	224	1.79	3673	2.92	5.22	9.43	141
$T1-3$	323	3.10	4968	2.72	8.42	15.5	278
$T1-4$	458	5.24	6328	2.78	14.6	23.1	494
$T1-5$	555	6.99	5099	2.59	18.1	29.5	695

Table 5.1: Table of filament properties from vortex tracking approach

increase, at a rate greater than the Taylor scale Reynolds number, is expected to again stem from the over-prediction of radius, which may dominate over the known under-prediction of vorticity in Equation 2.14. The mean normalised and actual circulation values are included in Table 5.1. Jiménez et al. (53) present a compilation of normalised circulation values which vary from $Re_\gamma/R_\lambda^{1/2} \approx 16 \rightarrow 21$ for Reynolds numbers of $R_\lambda = 36 \rightarrow 168$. Again, the $T1-2$ and $T1-3$ cases from the current investigation appear closest to these values, which may be in part due to the similar resolution of these cases, although the under-estimation of vorticity results in smaller normalised circulation values.

It should be noted that the algorithm also contains a significant bias towards radius over-estimation. For each coherent structure identified the algorithm computes the radius based on the variation of vorticity within a specified perimeter, under the assumption that the highest enstrophy will be concentrated in tube-like structures and with radial symmetry. Therefore, the inclusion of ribbon-like structures will produce lower vorticity variation in one direction, increasing the radial estimate. Although this weakness is noted in Jiménez et al. (53), the resulting bias towards larger radius structures is not explicitly stated. In order to consider the geometry of individual flow structures a different approach is required.

5.2.3.2   Box Counting Method

Structural classification was performed on the statistically converged data sets using the box counting method detailed in Section 2.4.4.2, with the number of classifiable structures shown in Table 5.2. A PDF of the shape parameter, H, is shown in Figure 5.8. Recalling that $0 > H > 1$, with values approaching 0 and
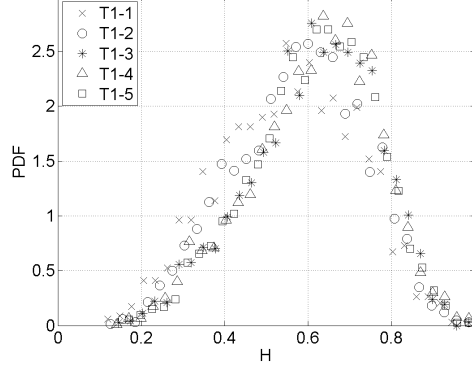
Figure 5.8: Shape Parameter PDF

Case	$T1-1$	$T1-2$	$T1-3$	$T1-4$	$T1-5$
Classifiable structure #	1205	2222	2470	3677	2545

Table 5.2: Number of classifiable structures

1 representing sheet-like and tube-like structures respectively, the mean value of $\sim 0.6 \rightarrow 0.7$ indicates the majority of flow structures tend to be tube-like. The good agreement between cases again suggests scale-invariance, with only the $T1-1$ case showing exception. However, this case contains half the number classifiable structures, and therefore these differences may be due to insufficient convergence.

Figures 5.9(a) to (c) and Figures 5.9(d) to (f) show JPDFs of shape parameter variation with normalised enstrophy and dissipation respectively at three Reynolds numbers. These enstrophy and dissipation values represent the mean of each classifiable structure. Despite considerable scatter due to the low number of data points available, the skewed triangular distribution peak shown in Figures 5.9(a) to (c) indicates that the highest enstrophy structures tend to be tube-like. However, the majority of the tube-like structures are relatively low enstrophy, with the JPDF peak just over the enstrophy threshold. These findings are in agreement with Moisy & Jiménez (79) who showed that increasing the vorticity threshold resulted in increasingly tube-like structures.

Figures 5.9(d) to (f) show that the mean dissipation within ribbon-like structures is larger than that associated with tube-like structures, implying that vortex

(a) $\Omega - H$ JPDF $T1 - 1$ (b) $\Omega - H$ JPDF $T1 - 3$ (c) $\Omega - H$ JPDF $T1 - 5$

(d) $\epsilon - H$ JPDF $T1 - 1$ (e) $\epsilon - H$ JPDF $T1 - 3$ (f) $\epsilon - H$ JPDF $T1 - 5$
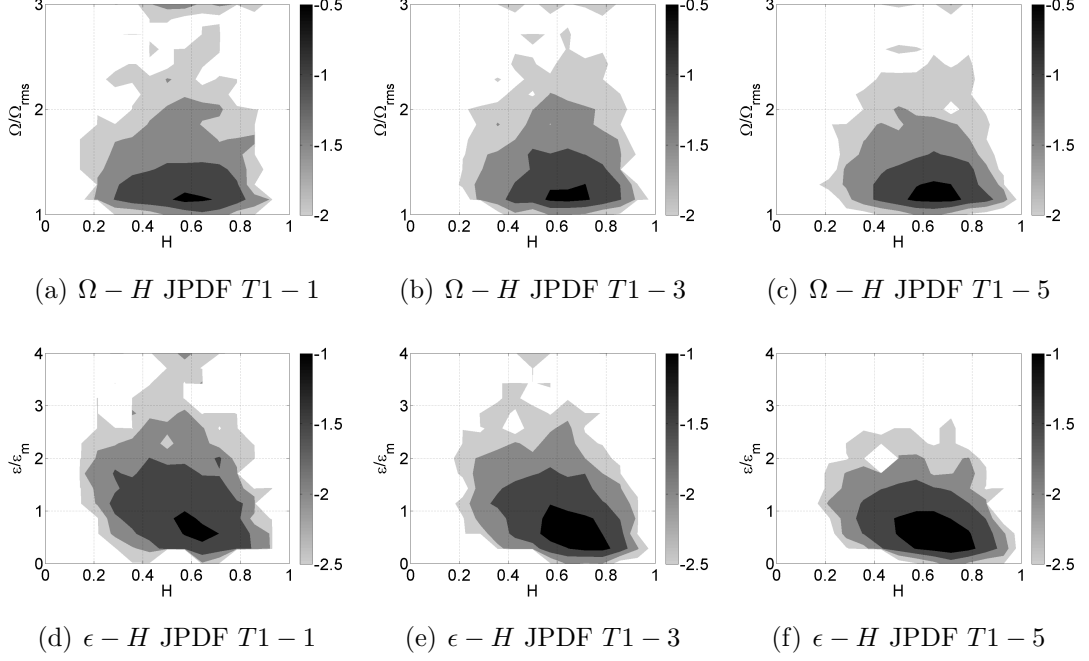
Figure 5.9: JPDF of shape parameter against normalised enstrophy and dissipation for classifiable structures, contours of base 10 exponentials

(a) $T1 - 1$ (b) $T1 - 3$ (c) T1-5
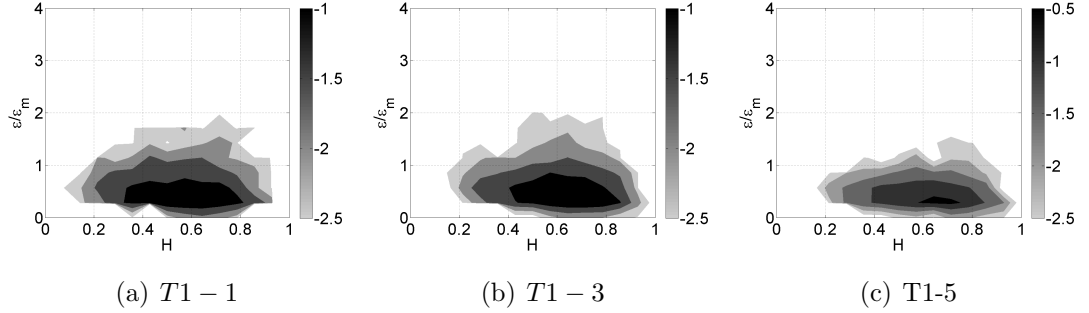
Figure 5.10: JPDF of shape parameter against normalised dissipation for regions surrounding the classifiable structures, contours of base 10 exponentials

sheets may be responsible for the majority of the dissipation. However, this definition does not account for dissipation surrounding the tube-like structures which may also contribute considerably, and therefore as in Section 5.2.2 it is useful include this.

The box count is extended to include points immediately surrounding each

structure, although for clarity the surrounding regions are considered separately here. Where adjacent structures share the same surrounding point, the volume fraction and dissipation contribution at that point are divided equally between the structures. Figures 5.10(a) to (c) show JPDFs of the shape parameter variation with the surrounding region normalised dissipation, which in comparison with Figures 5.9(d) to (f) appears much flatter and takes lower mean dissipation values. The small positive gradient indicates that slightly stronger regions of dissipation surround the tube-like structures, although this trend is very weak.

(a) Vol. Frac. Structures

(b) Perc. Cont. Structures

(c) Vol. Frac. Surrounding

(d) Perc. Cont. Surrounding

Figure 5.11: Volume fraction and percentage dissipation contribution of the classifiable structures and surrounding regions

The dissipation values shown in these JPDFs represent the mean contribution

of each structure, and therefore it is also useful to consider the volume fraction and total dissipation percentage contribution of different structures, 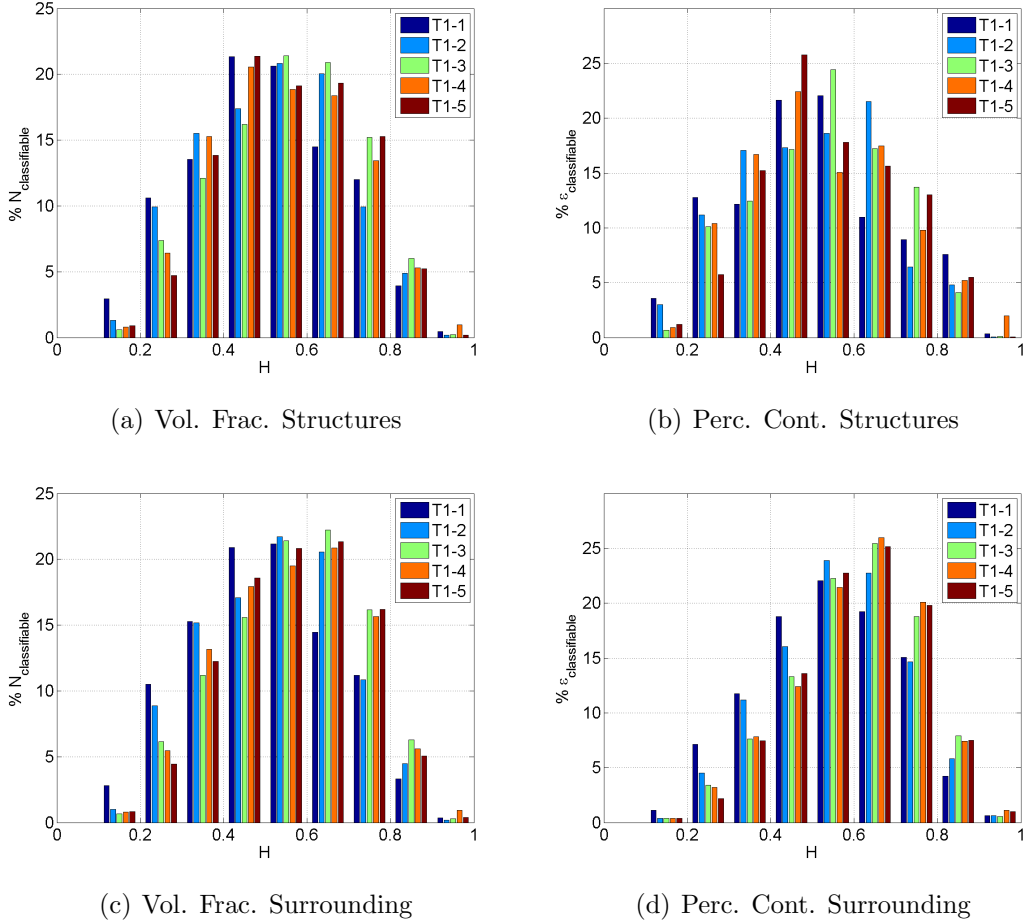and their surrounding regions. To this end, points within the structures and surrounding regions were divided into bins according to their parent structure's shape parameter, shown in Figure 5.11.

The volume fraction and percentage contribution are expressed as a percentage of the total number of points within the classifiable structures and the surrounding regions respectively, and therefore the sum of the bins for each case is equal to 100%. Although the classifiable structures represent a subset of the total number of coherent structures, assuming that the alignment is random the remaining unclassified structures should produce similar statistics. Therefore, these contributions are representative of the volume and dissipation fractions calculated in Section 5.2.2 for the same enstrophy threshold ($\Omega > \Omega_{rms}$).

Figure 5.11(a) shows a more symmetric distribution in comparison with Figure 5.8, which implies that despite being less numerous the ribbon-like structures generally occupy a larger volume in comparison with the tube-like structures. The global dissipation contribution in Figure 5.11(b) shows a similar relatively symmetric distribution. The stronger dissipation associated with more ribbon-like structures only has a small influence, slightly skewing the distribution towards lower shape parameter values.

Figure 5.11(c) shows the volume fraction distribution of regions surrounding different shape high enstrophy structures, which show a similar distribution to the structure distribution in Figure 5.11(a). Figure 5.11(d) shows the dissipation contribution from these surrounding regions, with a clear skew towards tube-like structures. Therefore, despite equal volume fractions, the higher dissipation surrounding the tube-like structures gives rise to a greater total contribution.

This analysis indicates that the significant global dissipation contribution from regions of high enstrophy shown in Section 5.2.2 must be attributed to the combination of ribbon and tube-like structures. Ribbon-like structures tend to be coincident regions of high enstrophy and dissipation, are relatively large in volume, and are surrounded by regions of weaker dissipation. In contrast, tube-like structures tend to be smaller in volume but more numerous, are coincident with moderate dissipation, and are surrounded by regions of higher dissipation.

These observations may be considered contrary to a number of previous studies (19; 52; 53) which suggest that the high enstrophy tube-like structures contribute little to the global dissipation rate.

These observations seem to suggest that tube-like structures contribute more significantly to the global dissipation rate in comparison with a number of previous studies (19; 52; 53). However, these differences may arise from threshold chosen, the inclusion of the region surrounding the high enstrophy structures, and also the consideration of both tube and ribbon-like structures, with the present study supporting the spatial dissipation distribution study by Ganapathisubramani et al. (37).

Visualisation of this method is also useful for observing the interaction of different structures and the evolution of these over time, and the reader is referred to video files in Appendix D.

5.2.4 Clustering Behaviour

In order to quantify the clustering of structures observed in Section 5.1.3 the box-counting method detailed in Section 2.4.4.3 has been applied. Figure 5.12(a) shows normalised box counts plotted against the cube root of box volume, $r = (r_x r_y r_z)^{1/3}$, which is representative a box length scale. At small box sizes the box counts follow the random distribution. However, as box length scale increases the box counts drop below the random distribution, indicating that the structures occupy a smaller volume of space in comparison with the random set, and are therefore clustered.

The Reynolds number of the highest resolution case ($R_\lambda = 162$) is very close to the investigation by Moisy & Jiménez (79) ($R_\lambda = 168$), and therefore a comparison can be made at this point using the clustering fraction shown in Figure 5.12(b). Due to the limited domain size in the present study the clustering behaviour over the entire inertial range cannot be examined, and therefore a comparison is made at the largest box size for this case. In the present investigation a length scale of 20η corresponds to a clustering fraction of 0.24, which is very close to the value of 0.25 found by Moisy & Jiménez (79). Slight differences may arise from the use of non-cubic boxes in the current investigation, and necessary

(a) Box counts (b) Clustering fraction
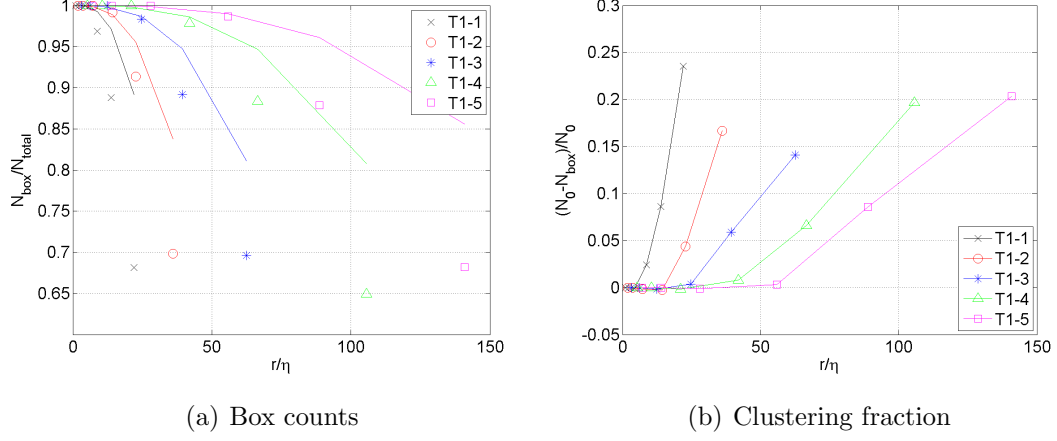
Figure 5.12: Box counting method clustering; (a) Symbols represent box counts, N_{box}, solid lines represent random box count, N_0

use of a lower threshold due to a lower number of high enstrophy structures in the smaller data set (the data set used in the study by Moisy & Jiménez is over 6 times larger).

The loss of fine scale information through spatial filtering demonstrated in Section 3.2 produces qualitatively similar high enstrophy and dissipation distributions across all cases (see Section 5.1.3). Therefore, due to the loss of this fine-scale information, the clustering behaviour in the lower resolution cases may not be representative of complete flow field, instead relating only to structures remaining after spatial filtering. However, as these structures are representative of slightly larger scales, the similar clustering behaviour between cases again supports the idea of scale-invariance across a range of scales.

While it is extremely interesting to observe the agreement with the previous investigation and quantification of scale-invariance, the limited domain size prevents determination of a maximum clustering fraction.

5.3 Statistical Flow Topology

Next the role of the strain field in the formation of coherent structures is investigated by decomposing this into its principal components, and studying the

location and alignment of these with respect to other quantities of interest such as the vorticity and dissipation rate. These relationships are also investigated through topological invariants, providing a comparison across a range of methods. Finally, with consideration of the underlying strain field, some examples of structure formation are presented from the time-resolved measurements.

5.3.1 Predominant Strains and their Spatial Distribution

Betchov (9) suggested that the local flow topology is dependent on the sign of the intermediate rate of strain, β; with negative values indicating bi-axial compression and stretching and vortex-line topology, and positive values indicating sheet-forming bi-axial extension coupled with compression (see Section 1.2.3.2). This not only provides a convenient approach for topological analysis, but also theoretical predictions of expected mean topologies. A number of statistical studies have been made to assess these predictions by comparing ordered strain rate tensor eigenvectors, $\alpha > \beta > \gamma$, with vorticity and dissipation (4; 37; 82).

Figures 5.13(a) and 5.13(b) show the three principal rate of strain eigenvalues, normalised by the mean strain rate $e_m = \langle \overline{e} \rangle$, where $e = \sqrt{\alpha^2 + \beta^2 + \gamma^2}$, for the original and divergence free results. There is an excellent collapse of data across different Reynolds numbers, and good agreement is shown with previous studies (37; 82). The longer tail and lower peak of the γ PDFs compared to the α PDFs indicates the flow preference for large compressions over extensions.

Incompressibility requires that $\alpha + \beta + \gamma = 0$, and therefore $\alpha \geq 0$, $\gamma \leq 0$. Continuity violations are clearly seen in the original data as negative and positive points in the α and γ strain rate PDFs respectively (see Figure 5.13(a)). These are largely corrected in the divergence free results (Figure 5.13(b)), which also show a better collapse of data. Averaging over all cases and conditioning for positive values of β produces strains in the ratio $3.2 : 1 : -3$ for the original fields, and $3.3 : 1 : -4.3$ for the divergence corrected fields, the latter of which are consistent with the ratio of $3 : 1 : -4$ found in other numerical and experimental studies (4; 37; 123). Persistent incompressible flow violations in the divergence corrected data together with the use of an imperfect divergence correction scheme are thought to account for the slight over-prediction of the strain rate ratios.

(a) Strain rate tensor eigenvalue PDF

(b) Strain rate tensor eigenvalue PDF

(c) Normalised intermediate rate of strain eigenvalue PDF
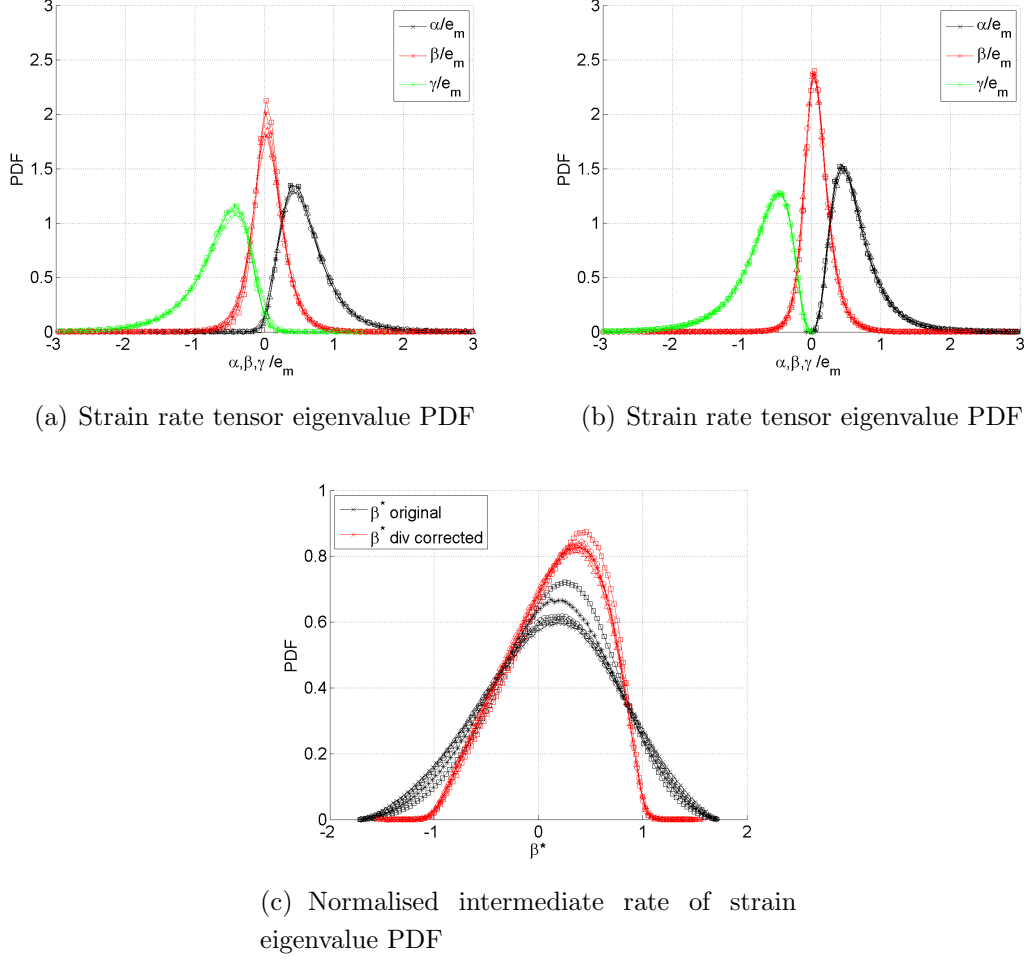
Figure 5.13: Strain rate eigenvalue PDFs, case by symbol: $\times = T1 - 1$; $\bigcirc = T1 - 2$; $* = T1 - 3$; $\triangle = T1 - 4$; $\square$ = T1-5

Figure 5.13(c) shows the intermediate rate of strain for both original and divergence free data, normalised as in Ashurst et al. (4) as $\beta^* = \sqrt{6}\beta/e$, where incompressibility now demands that $-1 > \beta^* > 1$. The skewed distributions indicate a predominance $\beta^* > 0$, with increasing asymmetry and sharper peaks after divergence correction, with a peak shift from $\overline{\beta^*} \approx 0.25 \rightarrow 0.4$.

The gradient measurement uncertainty dependence of β^* was quantified by Lund & Rogers (69) who showed that a gradient error corresponding to a measurement uncertainty divergence correlation coefficient value of $Q = 0.7$ shifted the peak value from $\overline{\beta^*} \approx 0.55$ to $\overline{\beta^*} \approx 0.25$. The use of an isotropic diver-

gence correction was shown to aid in the partial recovery of the true distribution, although the peak value was still under-estimate ($\overline{\beta^*} \approx 0.4$). Larger gradient uncertainty errors in the present study give rise to more significant divergence errors (see Section 4.5), resulting in the substantial under-prediction of the mean strain rate in the original data. The slight under-prediction of the divergence free data is consistent Lund & Rogers (69) resulting from imperfect divergence correction.

(a) $T1-1$ (b) $T1-3$ (c) $T1-5$

(d) $T1-1$, div. corr.   (e) $T1-3$, div. corr.   (f) $T1-5$, div. corr.

Figure 5.14: JPDF of normalised intermediate rate of strain eigenvalue and strain rate magnitude, contours of base 10 exponentials

Figure 5.14 shows a JPDF of the normalised intermediate rate of strain eigenvalue, β^*, and normalised strain rate magnitude, $e/\langle \overline{e} \rangle$ for the original and divergence free data. To aid comparison similar contour levels are used to the experimental study by Ganapathasubramani et al.(37). Although the distribution appears relatively unchanged with increasing Reynolds number the divergence correction has a significant effect, again shifting the distribution peak towards increasingly positive β^* values.

This correlation between high magnitude strain and positive β^* indicates the sheet forming preference of the large strains. Again, the slightly lower positive peak of β^* in comparison with previous studies (4; 37) does not necessarily indicate a preference for lower magnitude strain rates, but is more likely a result of imperfect divergence correction. This is evidenced by the lowest divergence case $(T1 - 3)$ most closely matching the previous value of ≈ 0.5.

(a) $T1 - 1$ (b) $T1 - 3$ (c) $T1 - 5$

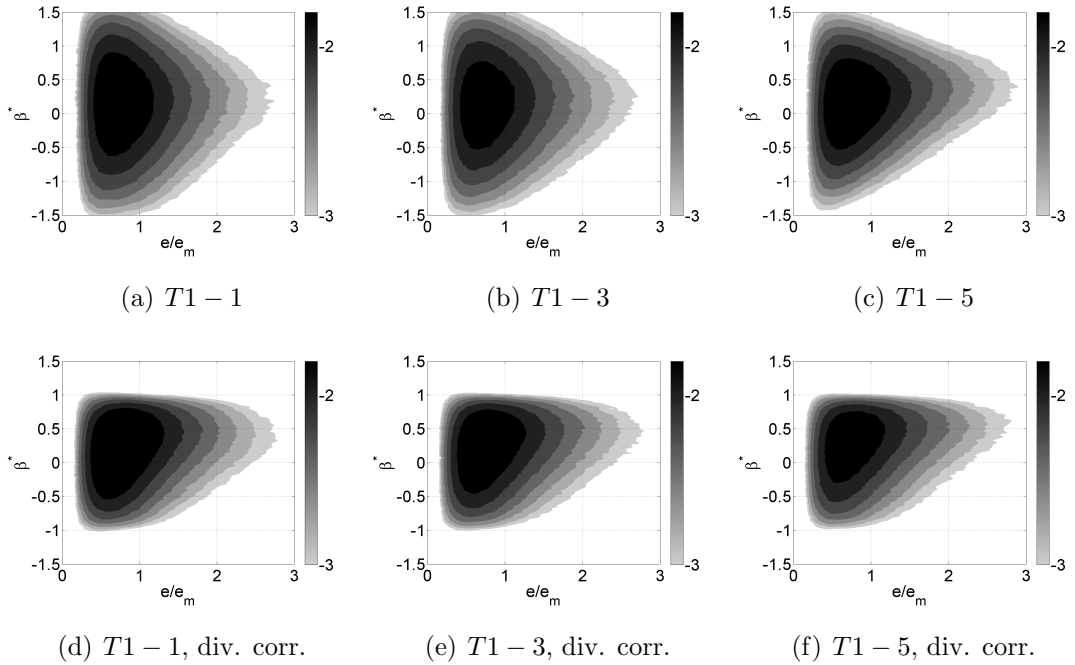(d) $T1 - 1$, div. corr.   (e) $T1 - 3$, div. corr.   (f) $T1 - 5$, div. corr.

Figure 5.15: JPDF of normalised intermediate rate of strain eigenvalue and strain rate magnitude, contours of base 10 exponentials

Figures 5.15 and 5.16 show JPDFs of vorticity and dissipation rate magnitude against the non-dimensional intermediate rate of strain β^*. Similar differences are observed between the original and divergence free data, with a shift towards increasingly positive values of β^* in the latter. Regions of intense vorticity and dissipation show a preference for positive values of β^*, which are associated with sheet forming strains. The more defined, less rounded peak shown for the dissipation distributions in Figure 5.16 demonstrates the stronger preference of this quantity in comparison with the vorticity.

(a) $T1-1$ (b) $T1-3$ (c) $T1-5$

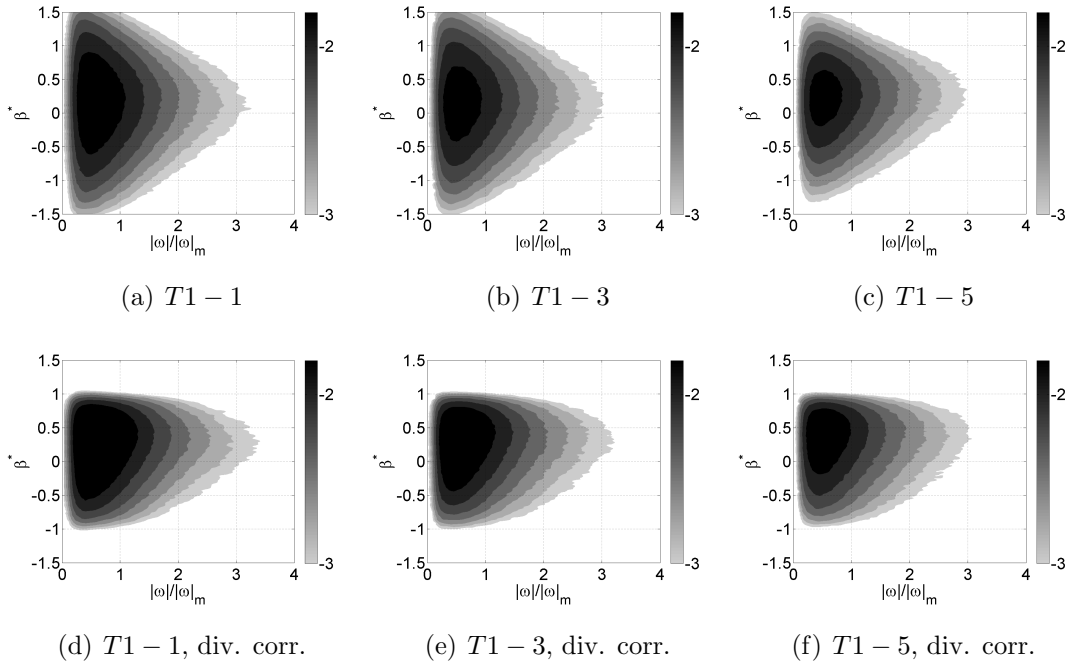(d) $T1-1$, div. corr.   (e) $T1-3$, div. corr.   (f) $T1-5$, div. corr.
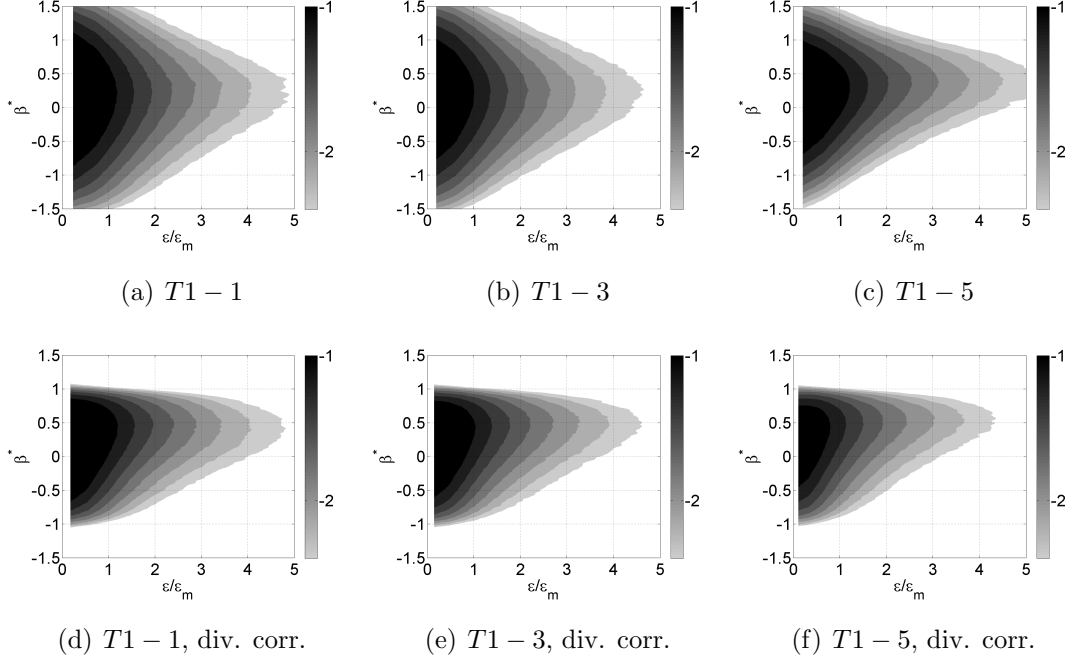
Figure 5.16: JPDF of normalised intermediate rate of strain eigenvalue and strain rate magnitude, contours of base 10 exponentials

These preferences can be further clarified by examining β^* PDFs conditioned on increasing vorticity and dissipation magnitude, shown in Figure 5.17 for the lowest divergence $T1-3$ case after divergence correction. In Figure 5.17(a) higher vorticity regions show no change in their preference towards positive values of β^*, indicating the high vorticity structure may be independent of the local strain rate. In contrast, conditioning on increasing dissipation rate changes the flow preference towards regions of high positive β^*, shown in Figure 5.17(b). These findings are in agreement with previous experimental and numerical studies (4; 37; 69), which re-enforce the observation that the dissipation is concentrated in regions of sheet-forming strain.

It is also interesting to further study the ratio of the principal strains defined earlier. The continuity equation, $\alpha + \beta + \gamma = 0$, can be arranged into the form $\alpha/\beta = -(\gamma/\beta + 1)$, and therefore it is possible to characterise the principal strains using either α/β or γ/β, with the former chosen and plotted as a PDF in Figure 5.18. The ratio limits $(-2 > \alpha/\beta > 1)$ can be derived from earlier

(a) Conditioned on $|\omega|$ (b) Conditioned on ϵ
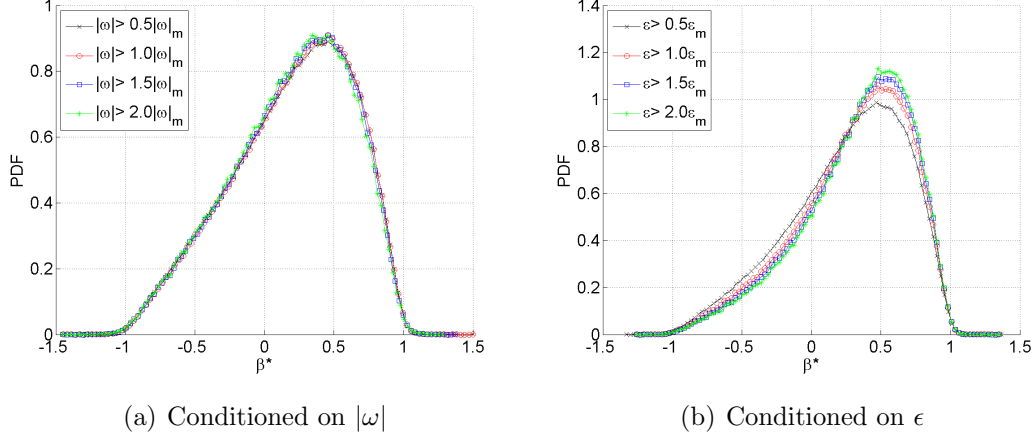
Figure 5.17: Conditional PDFs of normalised intermediate rate of strain eigenvalue on divergence corrected data for the $T1 - 3$ case

definitions of these quantities, with divergence accounting for values outside this range. Recalling that $\alpha > 0$, implies that regions where $\alpha/\beta > 0$ correspond to regions where $\beta > 0$. Therefore, the larger peak for $\alpha/\beta > 0$ demonstrates the flow preference for regions of sheet-forming strain. Due to the non-Gaussian PDF distribution the location of this peak does not match the mean ratio calculated previously, instead corresponding to the ratio $2 : 1 : -3$, with this ratio representing instead the most probable strain state. It is also worth noting that the peak for regions of axial strain ($\alpha/\beta < 0$) also shares this ratio. Therefore, while the ratio of the mean values ($3 : 1 : -4$) may provide a useful check against other numerical and experimental results, it may not be linked to the physical flow behaviour. The effect of thresholding is again demonstrated, with enstrophy and dissipation thresholds applied to Figures 5.18(a) and (b) respectively. In similarity with Figure 5.17, high enstrophy regions show no change in their preference for regions of $\beta > 0$, but regions of high dissipation show significant preference.

From earlier visual observations and geometry characterisation, moderate to high enstrophy is concentrated in both ribbon-like and tube-like structures, with the latter containing the highest enstrophy. Therefore, from the analysis in this section, it is possible that these moderate to high enstrophy ribbon-like structures may be formed by sheet-forming strains. However, the formation of the

177

(a) Conditioned on Ω (b) Conditioned on ϵ
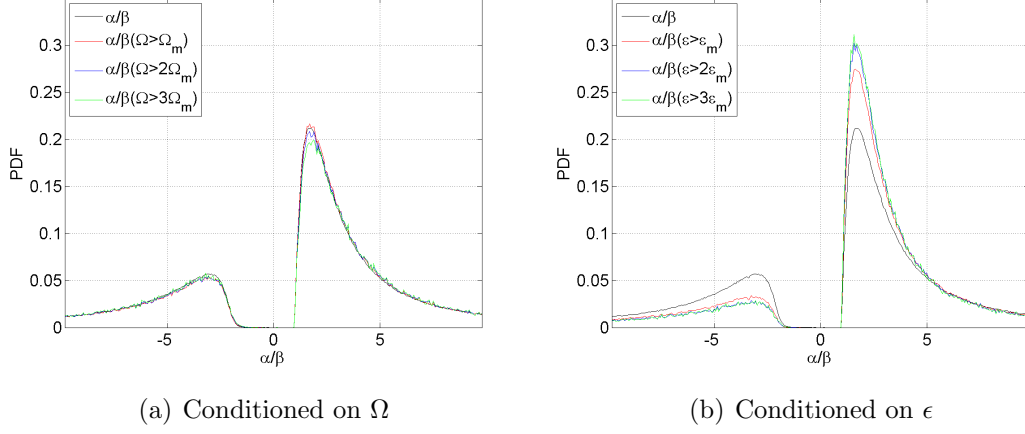
Figure 5.18: Conditional PDFs of principal strain eigenvalue ratio α/β on divergence corrected data for the $T1 - 3$ case

highest vorticity tube-like structures may not be dependent on the presence of a favourable $\beta^* < 0$ tube-forming strain field and may result from other mechanisms, such as the unsteady roll up of moderate intensity vorticity sheets proposed by Ruetsch & Maxey (103).

This conclusion can also be reached through the volume fractions calculated previously. For example, if the formation of tube-like structures were strain field dependent, then the predominant state of bi-axial strain should result in a prevalence of sheet-like structures. However, as demonstrated in Section 5.2.3.2 the high volume fraction of tube-like structures implies that this is not the case. Therefore, alternative formation mechanisms have been investigated further in Section 5.3.5.

5.3.2 Orientation of the Principal Strains

As detailed in Section 1.2.3.2 previous studies have found alignment between regions of high vorticity and the intermediate rate of strain. This can be calculated through the dot product of the unit vorticity vector and principal strain rate eigenvectors, returning an angle, $\cos\phi$, in the range $0 - 1$ corresponding at these extremes to perpendicular and parallel alignment respectively.

Figure 5.19 shows alignment angle PDFs for each principal rate of strain. The collapse of data from different Reynolds number runs is very good, all showing a large peak in $\hat{\beta} \cdot \hat{\omega}$ at $|\cos\phi| = 1$, confirming the strong parallel alignment preference of the intermediate rate of strain with the vorticity vector shown in previous studies (4; 37; 82). The peak in $\hat{\gamma} \cdot \hat{\omega}$ at $|\cos\phi| = 0$ indicates a weaker compressive strain alignment tendency, which is perpendicular to the vorticity vector, whereas the flatness of $\hat{\alpha} \cdot \hat{\omega}$ shows no strong alignment preference. These results show less divergence dependence than previous measures, which may be due to the use of divergence-correction-insensitive vorticity.

(a) Original data
(b) div. corr.
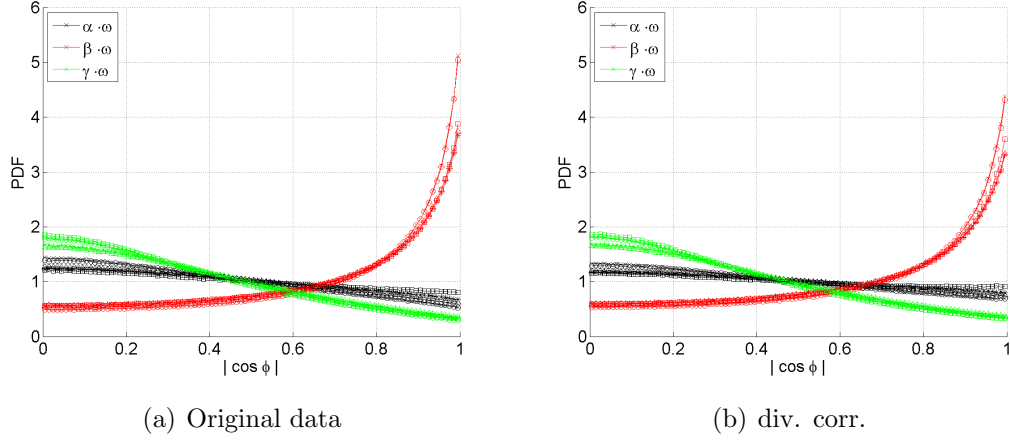
Figure 5.19: Alignment of principal strain rate eigenvectors with vorticity vector, case by symbol: $\times = T1 - 1$; $\bigcirc = T1 - 2$; $* = T1 - 3$; $\triangle = T1 - 4$; $\square = T1 - 5$

JPDFs of intermediate rate of strain alignment and vorticity magnitude are shown in Figure 5.20 demonstrating the distribution of this alignment. As shown in previous studies (4; 37), in regions of high enstrophy the vorticity vector tends to be aligned parallel with the intermediate rate of strain. Again with no significant differences between the original and divergence free data.

Again the alignment of high enstrophy regions with sheet-forming as opposed to tube-forming strains provides further evidence that the formation of these structures may be independent of the strain field. Therefore, to explain the correlation and alignment of these quantities Jiménez (51) considered that rather than the strain field producing these high enstrophy regions, instead the high

(a) $T1 - 1$ (b) $T1 - 3$ (c) $T1 - 5$

(d) $T1 - 1$, div. corr.   (e) $T1 - 3$, div. corr.   (f) $T1 - 5$, div. corr.
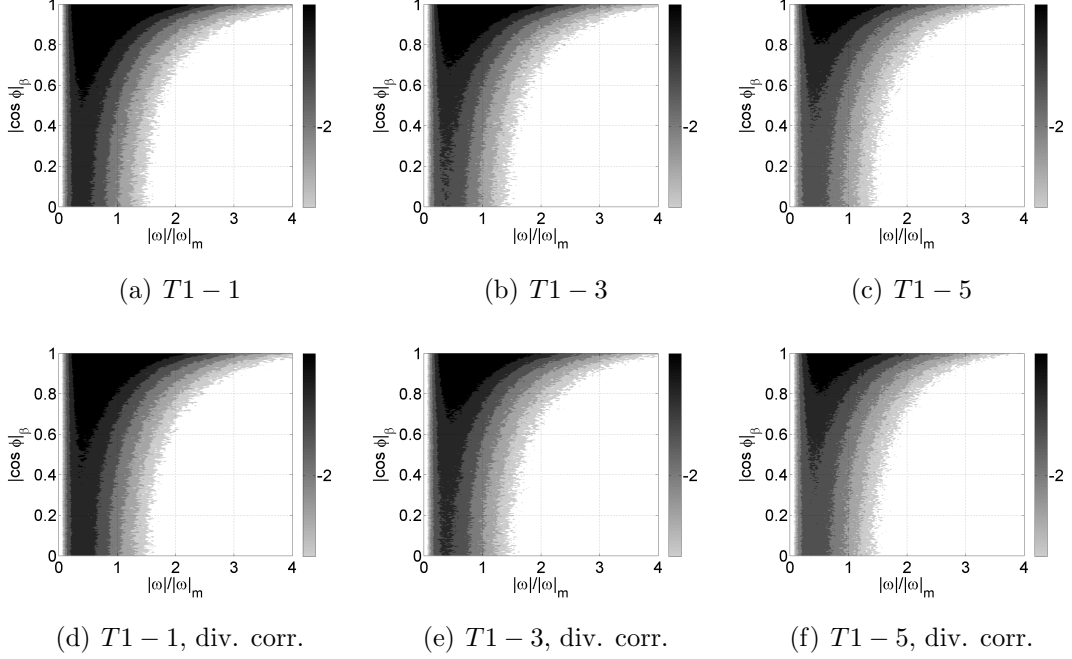
Figure 5.20: JPDF of intermediate rate of strain vorticity alignment and vorticity magnitude, contours of base 10 exponentials

enstrophy regions may produce the strain field. By only considering the dominant axial component of vorticity in these tube-like structures, Jiménez (51) idealised the problem in 2D. Here the largest and smallest principal strains induced by the vorticity are normal to the vortex axis and therefore the intermediate strain is automatically aligned with the vorticity vector. As the strength of these tube-like structures is much greater than the background vorticity, this strain field may dominate over other strains present, resulting in the alignment shown. In the real flow the situation is unlikely to be this simple, however this idea is consistent with the findings of the current investigation.

5.3.3 Visualisation of Strain Rate Distribution

In order to visually interpret the statistical correlation and alignment preferences of the strain rate, vorticity and dissipation fields these are plotted for a single example field from the $T1 - 3$ case in Figure 5.21.

(a) Enstrophy and Dissipation

(b) Vortex stretching strain field

(c) Sheet forming strain field

(d) Alignment of vorticity and β

(e) Alignment enstrophy correlation

(f) Orientation of vorticity

Figure 5.21: Comparison of strain rate distribution and alignment where: (a) enstrophy, $\Omega > n\Omega_{rms}$ (red-yellow colour map)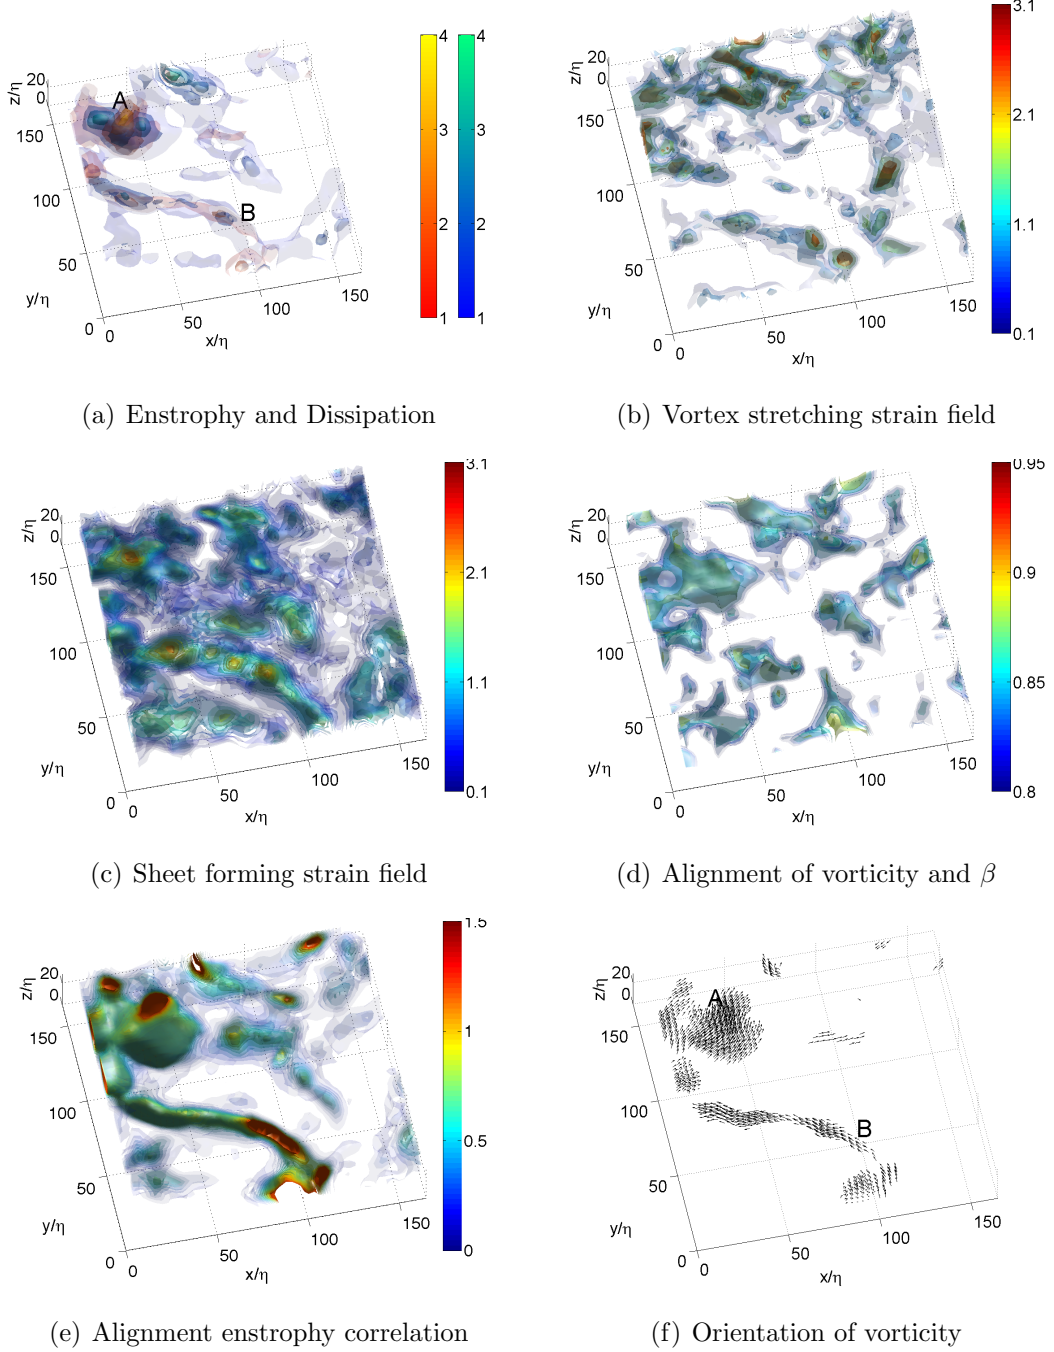, and dissipation rate, $\epsilon > n\epsilon_m$ (blue-green colour map); (b) $\beta < -n\beta_{rms}$; (c) $\beta > n\beta_{rms}$; (d) $cos(\phi)_\beta > n$; (e) $cos(\phi)_\beta \Omega > n$

181

Figure 5.21(c) shows a strong correlation between regions of positive strain rate, $\beta > 0$, and intense vorticity and dissipation, compared to Figure 5.21(b) where the regions of $\beta < 0$ are less well correlated with these fields. This is in agreement with the previous statistical analysis and other investigations (37; 53). It is also interesting to note that the sign of the intermediate rate of strain fluctuates somewhat along the long structure, B.

Figure 5.21(d) shows iso-surfaces of strong alignment between the vorticity vector and β. For reference the orientation of the vorticity vectors themselves are shown for all points over the threshold, $\Omega > \Omega_{rms}$, in Figure 5.21(f). Although there is a strong alignment of these quantities for the strong structure A, the alignment with structure B is quite patchy. The strong alignment in regions of high enstrophy is emphasised by studying a plot of the correlation $cos(\phi)_\beta \Omega$, shown in Figure 5.21(e).

This visual demonstration again shows that only the most intense enstrophy structures are well aligned with the intermediate rate of strain. This can be explained through the previous reasoning, where the strain field is a result of the tube-like structures, resulting automatically in the favourable alignment. Although the weaker structure B also gives rise to an aligned strain field, when combined with the surrounding strain the combined field is no longer aligned with the vorticity throughout this structure. Indicating that this alignment may only be an artefact of strong enstrophy.

5.3.4   Study of Topological Invariants

The strain field role can also be quantified through topological invariants. Figure 5.22 shows JPDFs of normalised velocity gradient tensor invariants Q_A^* against R_A^*, which can be used to categorise the flow in terms of four non-degenerate flow topologies. The near origin maximum shows the majority of the gradients are low magnitude. The characteristic distribution deformation associated with experimental divergence and noise (see Section 4.5.6) is visible on the uncorrected distributions and especially Figures 5.22(a) and 5.22(c). As shown in Section 4.5.6, modifying the fields towards a divergence free condition significantly alters the

(a) $T1-1$ (b) $T1-3$ (c) $T1-5$

(d) $T1-1$, div. corr.   (e) $T1-3$, div. corr.   (f) $T1-5$, div. corr.

Figure 5.22: JPDF of normalised invariants Q_A^* and R_A^* at three Reynolds numbers, with and without divergence correction, contours of base 10 exponentials

distribution, shifting it towards the distinctive self-similar teardrop shape associated with a predominance of local stable focus/stretching (vortex stretching) and unstable node/saddle/saddle (bi-axial strain) flow topologies, closely following the $\Delta_A = 0$ line in the latter mode.

Despite noise related artefacts there is a high level of agreement between the distributions at different Reynolds numbers, with each predicting similar magnitude contour levels. Ooi et al. (85) conclude that the largely similar collapse of data from a range of studies may be indicative of a form of invariant universality in $Q_A - R_A$ space, which within the bounds of the experimental accuracy, the present study appears to support.

To assess strain rate correlation in invariant space Figure 5.23 shows JPDFs of normalised strain rate tensor invariants Q_S^* against R_S^*. These distributions again show reasonable agreement with previous work (85), and a high level of similarity between Reynolds numbers. The inclusion of the 3rd invariant R_S^* again introduces significant divergence errors, which manifest in the distribution

(a) $T1 - 1$ (b) $T1 - 3$ (c) $T1 - 5$

(d) $T1 - 1$, div. corr.   (e) $T1 - 3$, div. corr.   (f) $T1 - 5$, div. corr.
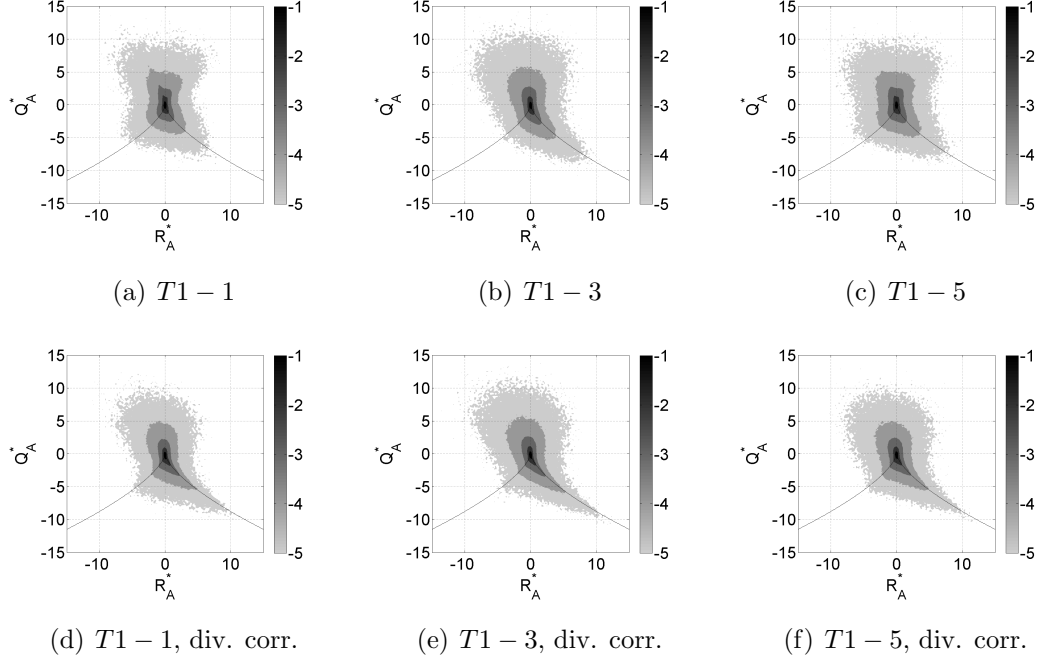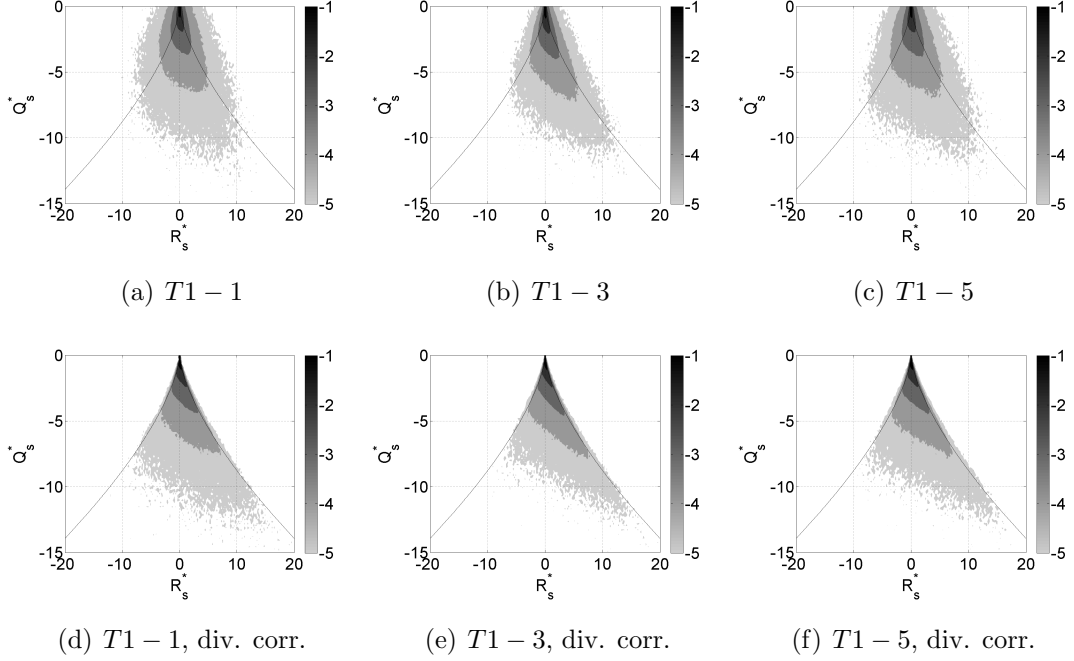
Figure 5.23: JPDF of normalised invariants Q_S^* and R_S^* at three Reynolds numbers, with and without divergence correction, contours of base 10 exponentials

crossing the bounding $\Delta_A = 0$ line. After divergence correction these incidences are almost completely eliminated, dramatically sharpening the peak at high values of R_S^* and $-Q_S^*$.

The second invariant R_S^* can be related to intermediate rate of strain through the signum function, $sgn(\beta) = sgn(R_S^*)$ (85), and as such takes the same sign. Therefore, the $R_S^* > 0$ preference is associated with the preference $\beta > 0$, indicating the majority of the flow ($\sim 65\%$) is subject to bi-axial straining. Again, in agreement with the findings in Section 5.3.1 there is also a strong correlation between high values of R_S^* and $-Q_S^*$ indicating the dissipation is concentrated in the most intense regions of bi-axial, sheet forming strains.

The strain field effect on regions of intense vorticity and dissipation can also be examined through the stretching of vorticity, using the quantity Σ (see Equation 1.17). Figures 5.24 and 5.25 show JPDFs of this normalised ratio Σ^* against the normalised strain rate and rotation invariants $-Q_S^*$ and Q_W^* respectively. As Σ^* is a function of the third invariant, R, once again divergence has a significant

(a) $T1 - 1$ (b) $T1 - 3$ (c) $T1 - 5$

(d) $T1 - 1$, div. corr.   (e) $T1 - 3$, div. corr.   (f) $T1 - 5$, div. corr.

Figure 5.24: JPDF of normalised invariants Σ^* and $-Q_S^*$ at three Reynolds numbers, with and without divergence correction, contours of base 10 exponentials

impact on the distribution at low values of $-Q_S^*$ shown in Figure 5.24, with the corrected fields agreeing more closely with previous investigations. Due to the shape of the distribution in Figure 5.25 the effect of divergence is less pronounced, although the corrected fields give rise to a slightly tighter set of contours.

The preference towards positive Σ^* values indicates a predominance of vorticity stretching, with the majority of this stretching taking place in regions of low level background enstrophy (Figure 5.25) and moderate to high strain (Figure 5.24). Jiménez et al. (53) explain this lack of correlation between high enstrophy and vorticity strain rate as evidence that the high enstrophy structures do not stretch themselves, and after reaching a sufficiently high degree of rotation de-couple from their parent strain fields and experience little more stretching. Again therefore, these findings are in agreement with sections 5.3.1 and 5.3.2, suggesting that the highest enstrophy structures are commonly formed through mechanisms other than vortex stretching.

The alignment of the vorticity vector and the intermediate rate of strain can

(a) $T1 - 1$ (b) $T1 - 3$ (c) $T1 - 5$

(d) $T1 - 1$, div. corr.   (e) $T1 - 3$, div. corr.   (f) $T1 - 5$, div. corr.
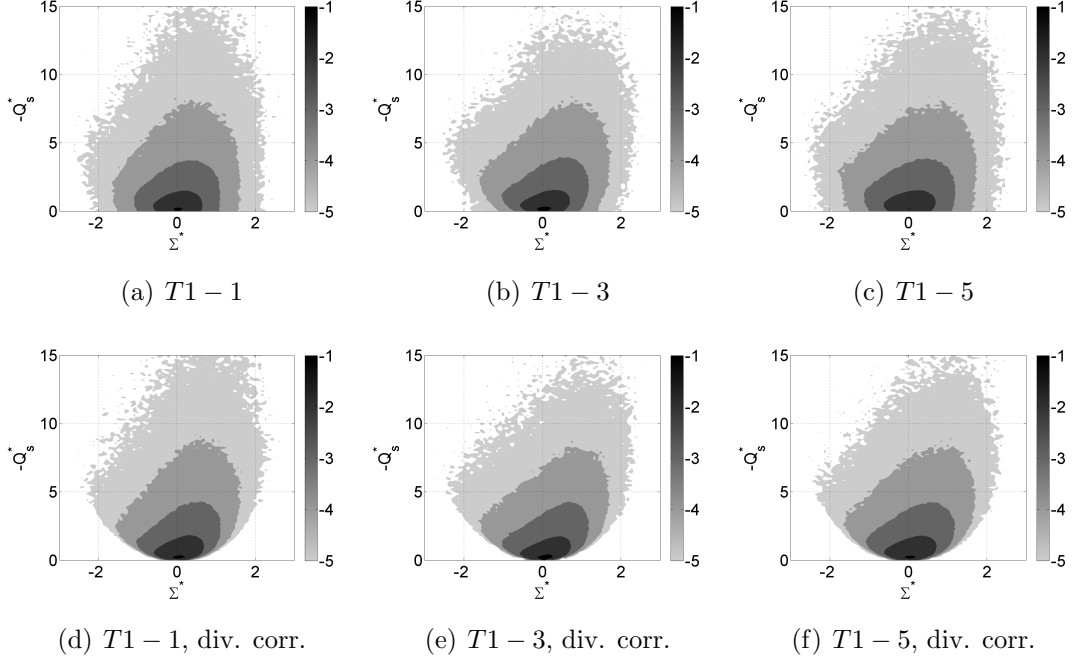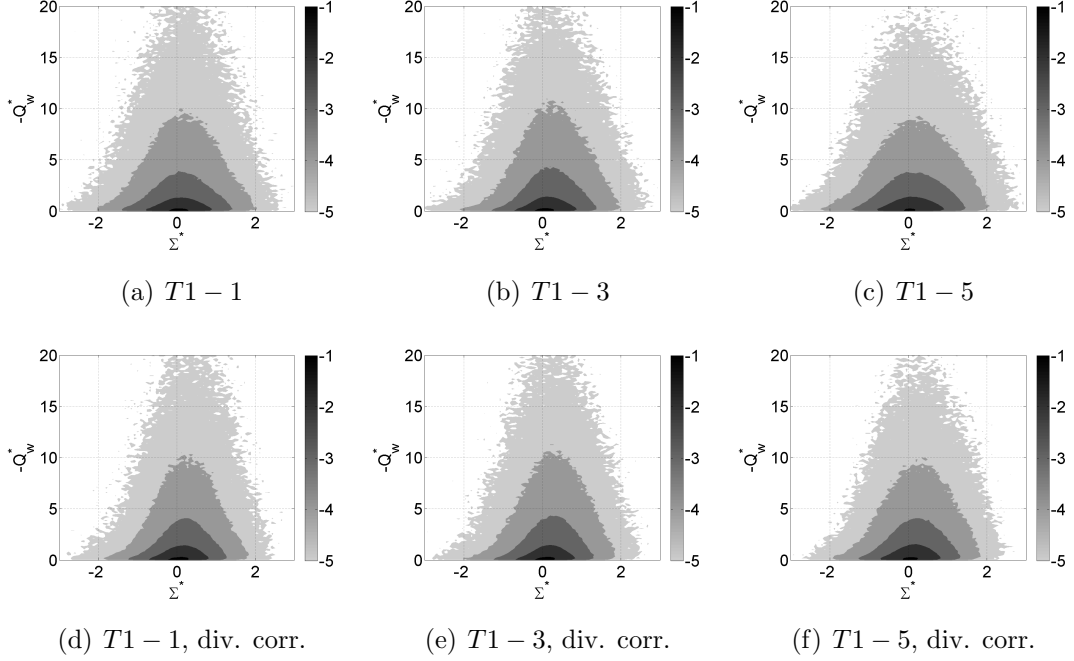
Figure 5.25: JPDF of normalised invariants Σ^* and $-Q_S^*$ at three Reynolds numbers, with and without divergence correction, contours of base 10 exponentials

also be illustrated in the $Q_A - R_A$ plane by plotting the normalised invariants Q_A/β^2 against $R_A/|\beta^3|$ (suggested by M.S. Chong, personal communication), as shown in Figure 5.26.

The original fields show significant correlation blurring due to divergence, which predominantly acts through the R_A term (Figure 5.26(a) to (c)). However, after divergence correction the correlation is significantly sharper (Figure 5.26(d) to (f)). The strong alignment with the $Q_A/\beta^2 + R_A/|\beta^3| = -1$ line for all cases again demonstrates the strong alignment between the vorticity vector and intermediate rate of strain. The correlation is widest at low gradient magnitude and pinches at higher magnitude, demonstrating the increased preference of the higher enstrophy structures to take this alignment. These findings are similar to those shown in DNS data by Horiuti(45), representing a neat way of demonstrating the flow alignment preference.

(a) $T1 - 1$ (b) $T1 - 3$ (c) $T1 - 5$

(d) $T1 - 1$, div. corr.   (e) $T1 - 3$, div. corr.   (f) $T1 - 5$, div. corr.
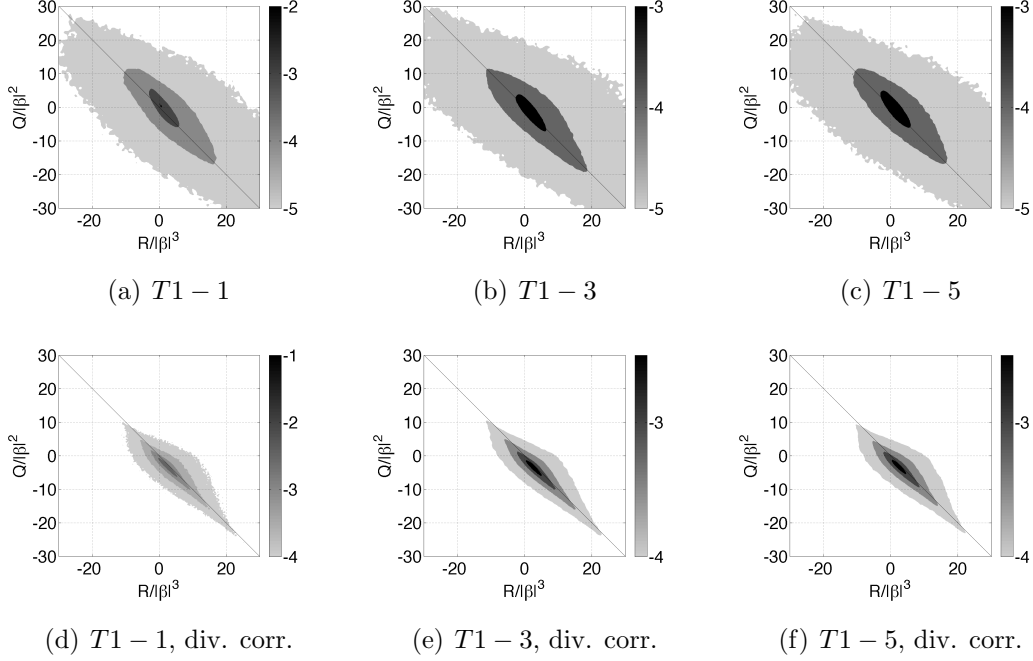
Figure 5.26: JPDF of normalised invariants Q_A/β^2 and $R_A/|\beta^3|$ at three Reynolds numbers, with and without divergence correction, contours of base 10 exponentials

5.3.5 Coherent Structure Evolution

In the previous sections the relationships between fields of strain, vorticity and dissipation have been quantified. There is a predominance of bi-axial strain, with a tendency for both the high enstrophy and dissipation fields to be coincident with these sheet forming strains. These strains may give rise to the moderate to high enstrophy ribbon-like structures identified in Section 5.1.3, which are also coincident with regions of dissipation. However, the fact that the tube-like high enstrophy structures are also coincident with these sheet forming strains has led some authors to suggest that these regions do not simply form under a favourable strain field, and are a result of some other unsteady process.

The use of single instantaneous data sets forces a number of previous investigations to stop at this point (4; 37; 53), leaving these important formation mechanisms to conjecture. However, the quasi time-resolved data available from the current investigation permits the study of coherent structure formation.

Therefore, in this subsection the evolution of coherent structures is demonstrated through the enstrophy field evolution and a number of visual examples.

5.3.5.1 Evolution of Enstrophy

The calculation of each term in the enstrophy transport equation (defined in Section 1.1.4) permits an assessment of the enstrophy field evolution, and identification of the dominant source of these changes. Figure 5.27 shows mid z-plane slices of normalised enstrophy, dissipation, enstrophy transport equation components, and two different component combinations, which illustrate the equivalent rate of change and the stretching/diffusion balance. In order to calculate the enstrophy rate of change component the quasi-time resolved results were required. Therefore, the most temporally well resolved $T1-1$ case was used and additional noise was controlled through Gaussian temporal smoothing with a kernel width of 5 frames. The ratio of measurement frequency to Kolmogorov time for this case is around 12, which is sufficient for this analysis. It should be noted that these first and second order enstrophy gradients are likely to contain significant noise, which may preclude definitive conclusions.

The enstrophy and dissipation distributions in Figure 5.27(a) and (b) show a high enstrophy core sandwiched between two high dissipation regions. The five components of the enstrophy transport equation are shown in Figure 5.27(c) to (g). Comparing Figure 5.27(a) and (c) it is clear that the greatest change in enstrophy occurs in the regions of high enstrophy, which after comparison between Figure 5.27(c) and (d) can be mainly attributed to advection, caused by large scale motions. However, if a frame of reference moving with the vortex core is considered then it would be appropriate to discard these advection effects, and concentrate on the other enstrophy evolution contributions.

Figure 5.27(e) and (f) show that vortex stretching dominates over contraction in the vortex core region, and that these regions of stretching are coincident with regions of significant viscous diffusion. The balance of these components is fundamental to Burger's vortex model (16), with these observations adding support for the similarity of these typical structures to Burger's vortices. Figure 5.27(h) demonstrates in this example that stretching dominates over diffusion, which may

(a) Enstrophy   (b) Fluid dissipation   (c) Rate of change

(d) Advection   (e) Stretching/Compression   (f) Viscous Diffusion

(g) Visc. enstrophy diss.   (h) Stretch/diff. balance   (i) Equivalent rate of change
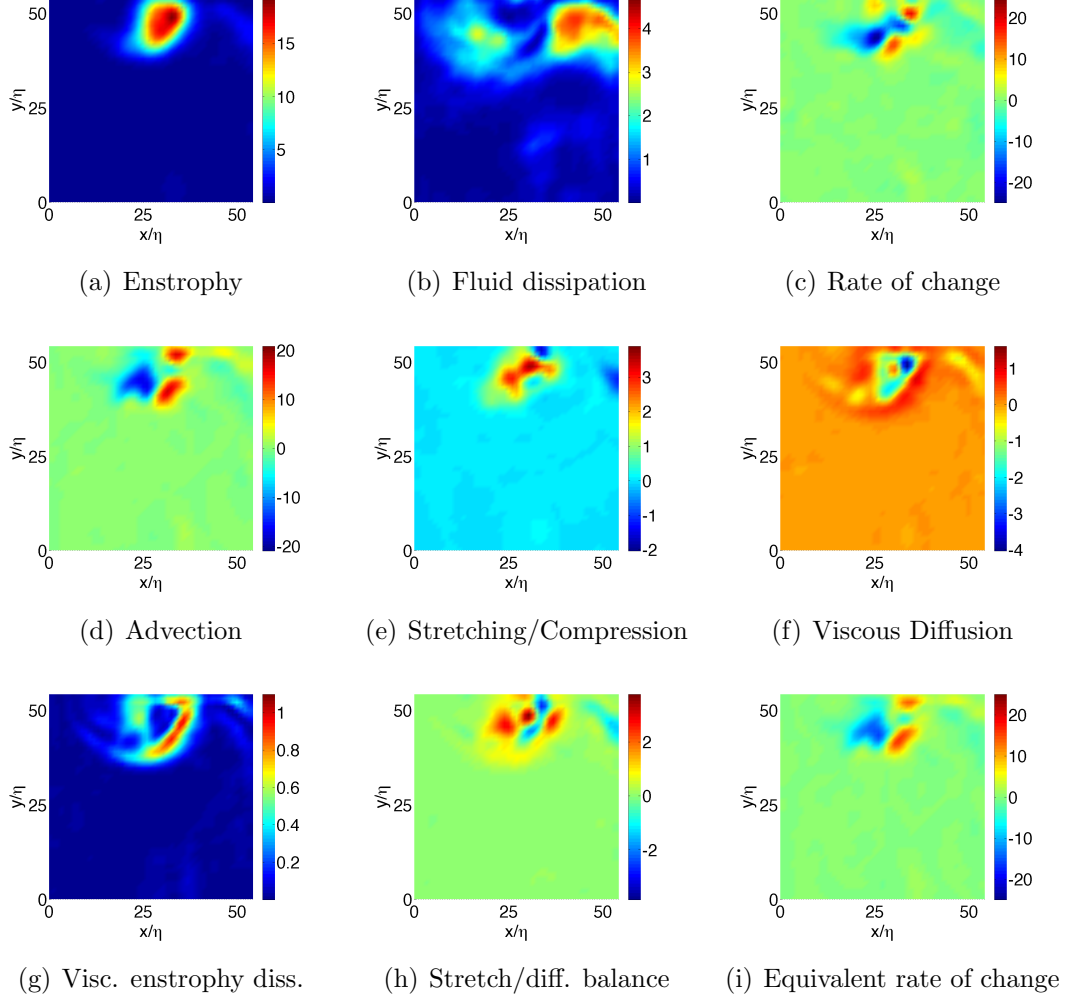
Figure 5.27: Colour maps of enstrophy, dissipation, and terms relating to the evolution of enstrophy; (a) Ω/Ω_m ; (b) ϵ/ϵ_m; (c) Φ/Ω_m (s^{-1}); (d) $-\Upsilon/\Omega_m$ (s^{-1}); (e) Π/Ω_m (s^{-1}); (f) Ψ/Ω_m (s^{-1}); (g) Ξ/Ω_m (s^{-1}); (h) $(\Pi + \Psi)/\Omega_m$ (s^{-1}); (c) Φ_{eq}/Ω_m (s^{-1})

be expected given the roll up and thinning of this tube-like structure (shown later in Section 5.3.5.3).

Figure 5.27(g) shows that enstrophy dissipation through viscosity occurs around the vortex core, with the longevity of the vortices explained through the relatively low magnitude of this. Finally the accuracy of temporal calculation is assessed by comparing the equivalent enstrophy rate of change, shown in Figure 5.27(i),

to that obtained through direct calculation, shown in Figure 5.27(c). Despite slight differences in magnitude the distributions identify similar regions containing significant rates of change, indicating that the temporal measurements may be sufficient for this analysis. Furthermore, this also indicates in the absence of sufficient temporal resolution for the higher Reynolds number cases, the equivalent rate of change may be a reasonable substitute, although the significant noise makes a more definitive conclusion difficult.
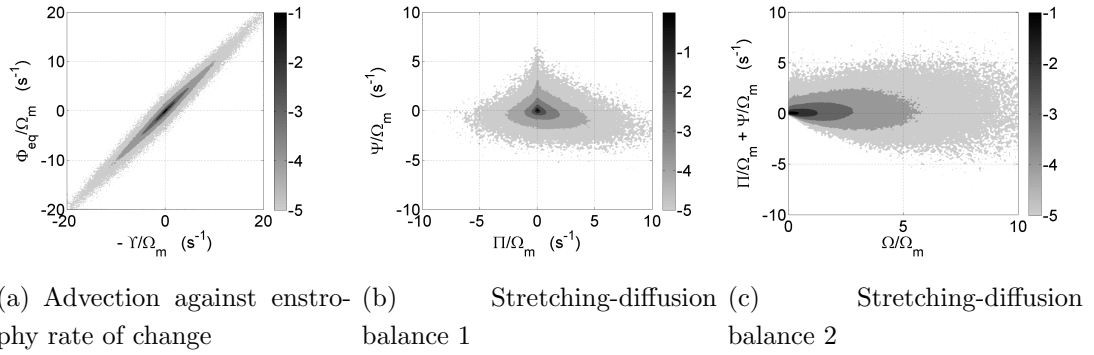

(a) Advection against enstrophy rate of change

(b) Stretching-diffusion balance 1

(c) Stretching-diffusion balance 2

Figure 5.28: JPDFs of enstrophy transport equation components from $T1 - 1$ case with divergence correction (pass IW_2); contours of base 10 exponentials

These trends can be generalised by studying JPDFs from the statistically resolved $T1 - 1$ case data, shown in Figure 5.28. Figure 5.28(a) shows a JPDF of the normalised negative advection term $-\Upsilon/\Omega_m$ against the normalised equivalent rate of change of enstrophy, Φ_{eq}/Ω_m. The high degree of correlation is clear from the close collapse of the data onto the line $-\Upsilon/\Omega_m + \Phi_{eq}/\Omega_m = 1$, demonstrating that as seen previously the advection accounts for the majority of the enstrophy evolution.

The stretching/diffusion balance was investigated further in Figure 5.28(b), which shows a JPDF of normalised stretching/compression term Π/Ω_m against the normalised viscous diffusion Ψ/Ω_m. The skewed JPDF shows a flow preference for enstrophy stretching over compression, with significant correlation between enstrophy increase through stretching and enstrophy reduction through viscous diffusion. Considering these components in the absence of advection demonstrates the dominance of the stretching over the diffusion, especially at

high magnitudes. Therefore, regions exist where enstrophy stretching is greater than diffusion, leading to the intensification of enstrophy. It is also worth noting that the magnitude of both the stretching and diffusion terms are significantly lower than the advection component and therefore the total enstrophy change.

Figure 5.28(c) shows a JPDF of the normalised stretching-diffusion balance $(\Pi + \Psi)/\Omega_m$ against normalised enstrophy Ω/Ω_m. The JPDF distribution indicates that there exist regions where stretching dominates over diffusion and vice versa, however at high enstrophy magnitude, there is a reasonable balance between these components. Therefore, the balance of stretching and diffusion in the highest enstrophy magnitude structures provides further evidence that these tube-like regions may be similar in structure to Burger's vortex model.

5.3.5.2 Vortex Merger

(a) $t/\tau = 0$ (b) $t/\tau = 0.84$ (c) $t/\tau = 1.68$

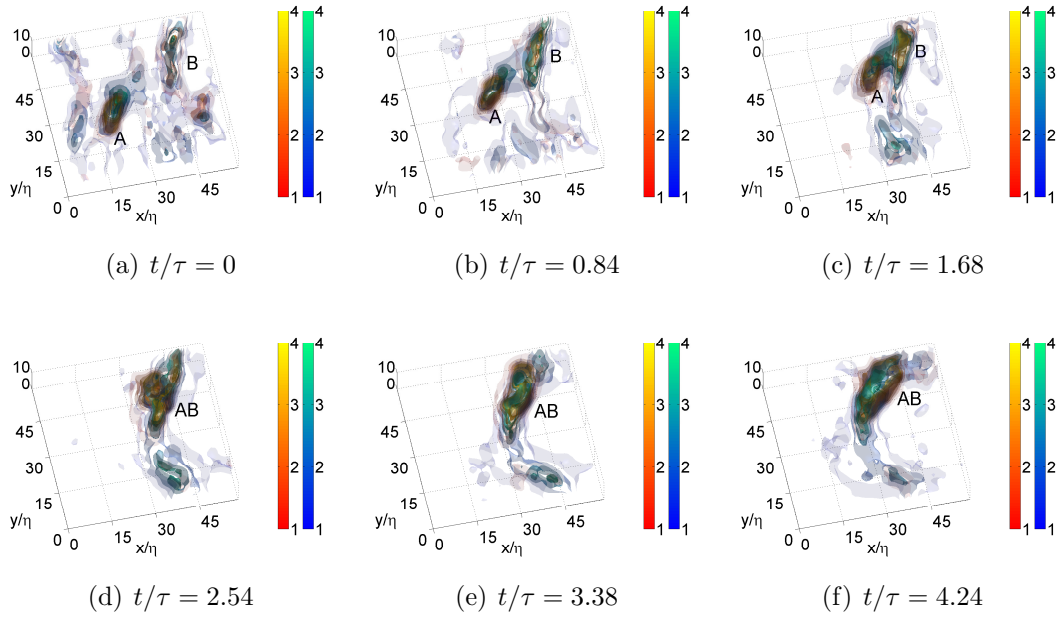(d) $t/\tau = 2.54$ (e) $t/\tau = 3.38$ (f) $t/\tau = 4.24$

Figure 5.29: Merging of coherent structures; case $T1 - 1$; Time-series at 1.5Hz interval ($t/\tau = 0.84$) ; Isosurfaces of enstrophy, $\Omega > n\Omega_{rms}$ (red-yellow colour map), and dissipation rate, $\epsilon > n\epsilon_m$ (blue-green colour map)

Figure 5.29 shows an example of vortex merger for the $T1 - 1$ case, with

two coherent structures combining to form a single structure (separate to the traditional definition of vortex reconnection (60)). Figure 5.29(a) shows tube and ribbon-like structures, A and B, which begin at a separation distance of half the volume width. Figures 5.29(b) to (c) show structure A being advected towards B, which remains relatively stationary. As the structures approach, a high dissipation region is formed between the two.

As the structures' outer layers begin to interact in Figures 5.29(d) to (e) the structure is redefined as AB, although the high magnitude core regions are still separately identifiable at these points. However, by Figure 5.29(f) the structures have completely merged into a single large high enstrophy ribbon-like structure, coincident with a region of high dissipation. This structure appears to be part of a larger scale circular structure, which covers a large proportion of the volume.

5.3.5.3 Roll up of Vortex Sheet

(a) $t/\tau = 0$ (b) $t/\tau = 0.42$ (c) $t/\tau = 0.84$

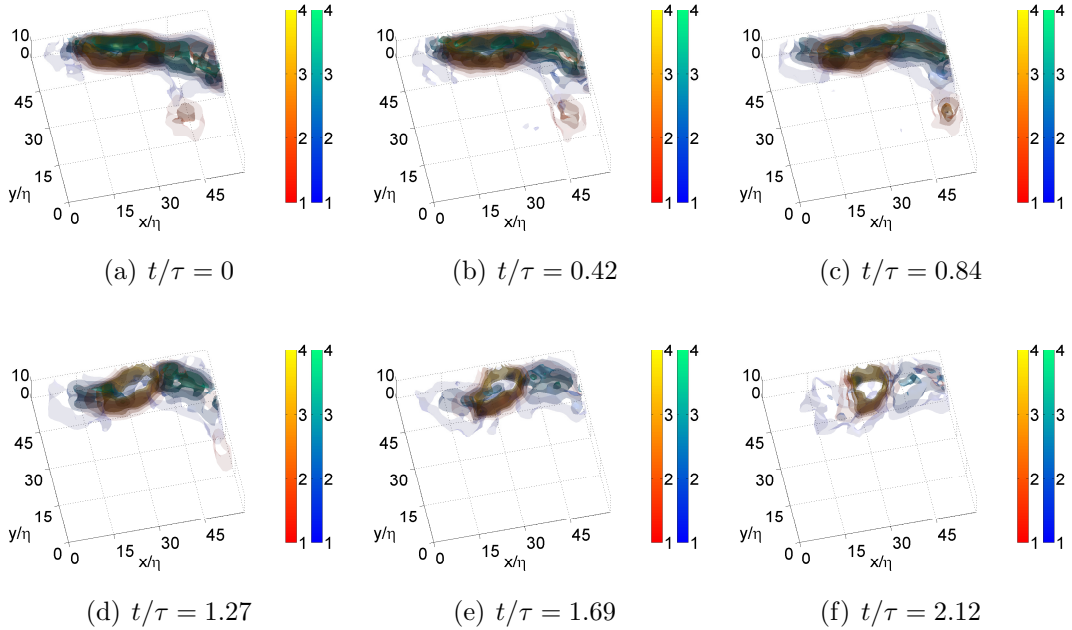(d) $t/\tau = 1.27$ (e) $t/\tau = 1.69$ (f) $t/\tau = 2.12$

Figure 5.30: Roll up of ribbon-like structure into a tube-like structure; case $T1-1$; Time-series at 3Hz interval ($t/\tau = 0.42$); Iso-surfaces of enstrophy, $\Omega > n\Omega_{rms}$ (red-yellow colour map), and dissipation rate, $\epsilon > n\epsilon_m$ (blue-green colour map)

Figure 5.30 demonstrates the roll up of a ribbon-like vorticity structure into a tube-like structure for the $T1 - 1$ case. Figure 5.30(a) shows a ribbon-like high enstrophy structure, which is coincident with a region of high dissipation. The structure spans the domain experiencing some truncation at the edges, with a high magnitude enstrophy core of fluid at around $\Delta x/\eta = 30$.

The roll up of the high enstrophy magnitude core is shown in Figures 5.30(b) to (f). The core breaks away from the other high enstrophy structures, rotating and changing shape during roll up from ribbon to tube-like. During roll up the high magnitude dissipation is expelled from the core of the structure, becoming disassociated in Figure 5.30(f).

5.3.5.4 Clustered Structure Interactions

The evolution of smaller scale structures which are clustered in the periphery of larger scale vortex structures is shown in Figure 5.31 for the $T1 - 2$ case.

(a) $t/\tau = 0$ (b) $t/\tau = 1.58$ (c) $t/\tau = 2.03$

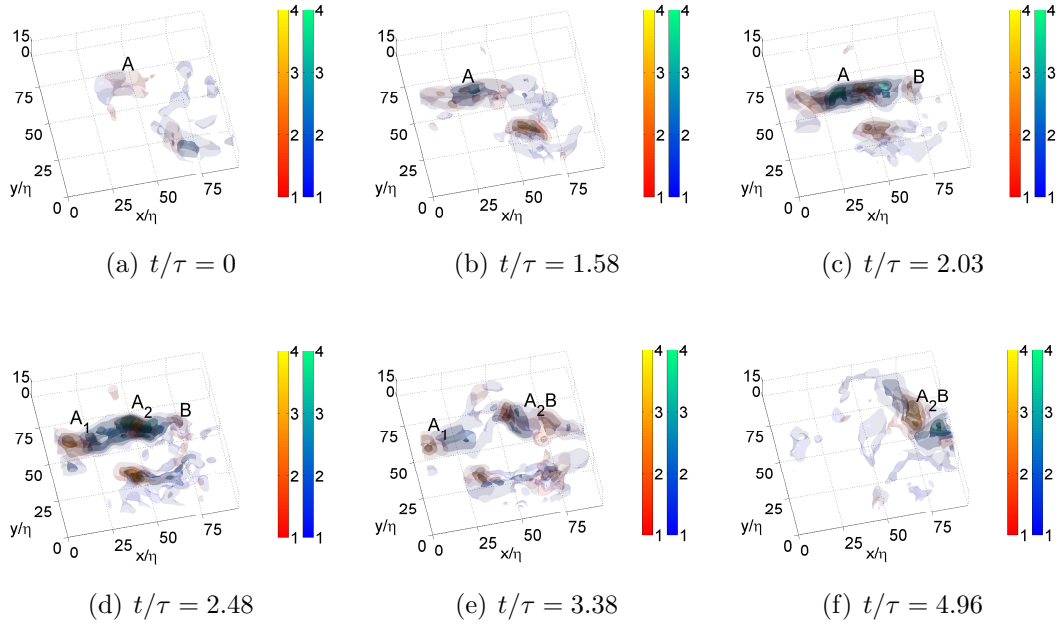(d) $t/\tau = 2.48$ (e) $t/\tau = 3.38$ (f) $t/\tau = 4.96$

Figure 5.31: Evolution of clustered coherent structures; case $T1 - 2$; Time-series at mixed interval; Isosurfaces of enstrophy, $\Omega > n\Omega_{rms}$ (red-yellow colour map), and dissipation rate, $\epsilon > n\epsilon_m$ (blue-green colour map)

Figure 5.31(a) shows the cluster after it enters the volume, with a slightly stretched tube-like structure identified as structure A. As the cluster is advected across the volume, structure A is drawn out into a more ribbon-like structure as shown in Figure 5.31(b). During this process the enstrophy within the structure intensifies, and becomes concentrated in two distinct cores in Figure 5.31(c), while another tube-like structure B is formed. The dissipation is also increased, with local maxima occurring between vortex cores.

The ribbon-like structure A then breaks into two separate tube-like structures, A_1 and A_2, as shown in Figure 5.31(d). The first of these, structure A_1, is advected out of the volume as the larger scale vortex appears to disintegrate (Figure 5.31(e) to (f)). Structure A_2 interacts with structure B, combining to form the new structure A_2B, which at the instance of joining is ribbon-like structure containing the two original vortex cores, shown in Figure 5.31(e). Figure 5.31(f) shows that after continued interaction the cores become completely merged, with the vorticity rolled up into a single strong tube-like structure.

The examples in this subsection demonstrate the stretching of vorticity into a ribbon-like structures, the subsequent breakdown of this into tube-like structures, and also the interaction and roll up of vorticity into strong tube-like structures. Formation mechanisms such as these have been theorised by a number of authors (53; 77; 89; 103), and therefore their demonstration here is an important step in understanding the flow physics at these scales.

The observed formation of small scale ribbon-like structures on the periphery of larger scale structures, and the breakdown of these into yet smaller scale vortex-tubes provide evidence of the Richardson cascade; in which energy from the larger scale is transferred to increasingly small scales. Tsinober (122) notes that the cascade process does not necessarily require intermediate stages, with the interaction of vortices at the integral scale capable of directly creating dissipative scale motions. Evidence of this is provided by Shen & Warhaft (109) who demonstrate a direct coupling of large and small scales. In the current investigation the substantial scale difference between relatively large motions and smaller structures in the boundaries between these appears to support this idea of direct energy transfer. However, similar clustering behaviour was identified in all cases

(see Section 5.2.4) suggesting a range of scales exist, as would be expected from a more conventional cascade interpretation.

The example of vortex merger shows that small structures can combine to form larger structures, which may appear to contradict the standard cascade theory. However, as noted by Tsinober (122), energy can be exchanged in both directions (from large to small scale and vice versa) but dissipation will only occur at the small scales. Therefore, energy injected at the integral scale may be subject to numerous complex interactions, but once it inevitably reaches the smallest scale it will be dissipated through viscosity.

Although it would be extremely useful to quantify these formation mechanisms statistically, the present data set contains too few examples of these processes, and statistical convergence may require a data set orders of magnitude larger.

Chapter 6

Conclusions

This report has introduced the concept of coherent dissipation scale turbulent structures, discussed the current understanding of these and their significance, and described measurement techniques capable of fully capturing these motions. The unique experimental facility brings the small scale motions within the reach of Tomographic PIV, a three-dimensional extension to standard planar PIV, with a review of this technique highlighting its enormous potential. TPIV software was developed and tested both numerically and experimentally, with parametric studies allowing comparison with previous investigations and quantification of measurement uncertainty. Next, after detailing the experimental set-up, a series of high Reynolds number DPIV and TPIV measurements were conducted in the large mixing tank facility, and analysed using a range of statistical measures.

6.1   TPIV Development and Implementation

During the development of TPIV software a greater technique understanding was realised, with the code adopting a number of programming approaches to reduce the significant run times. Numerical testing supported previous parametric studies, and extending this investigation into 3D coupled with a more applicable PIV accuracy criterion addresses the shortage of literature in this area. In comparison with previous findings the current results indicate the technique is less sensitive to image noise, slightly wider camera angles give optimum accuracy, and the benefit of a fifth camera is marginal. Novel findings regarding initialised

intensity fields were also presented, resulting in significant computational cost reductions, accelerated solution convergence, and increased robustness with respect to camera misalignment. Interestingly the initialised fields obtained from this new MFG method were close enough to the iterated solution to allow direct cross-correlation; allowing extremely rapid processing with only a small accuracy penalty. The reconstruction and calibration software was also tested experimentally through a range of studies, the most applicable of which was the study of a laminar vortex ring. This test captured the main vortex ring flow details well, also demonstrating the new MFG method on experimental data.

DPIV and TPIV measurement uncertainty was quantified through a series of simulations, which mimicked experimental conditions through a judicious choice of ideal flow field, spatial and temporal scaling, and tomographic reconstruction parameters. Additional TPIV noise from the reconstruction process manifests particularly in sensitive gradient quantities, but was shown to be partially controllable through velocity field smoothing. Reducing measurement resolution was shown to act as a spatial filter, resulting in peak gradient magnitude, and corresponding dissipation and vorticity under-prediction; with a minimum resolution of $x/\eta = 3$ recommended. Furthermore, typical PIV aliasing and spatial truncation effects were demonstrated to reduce agreement with the expected spectral energy distribution. This study was instrumental in defining the current measurement accuracy, and contributes to the dearth of TPIV accuracy literature.

6.2 Basic Flow Characterisation

DPIV and TPIV measurements made in the large mixing tank over a range of Reynolds numbers were analysed using a number of different methods. Initially, flow convergence and centring analysis determined that despite adequate small-scale convergence the largest scale motions may not be fully resolved. These convergence estimates were combined with resolution estimates to fully quantify the measurement uncertainty, and although undesirable this lack of large scale convergence is unlikely to affect the well converged smallest scales, which are the focus of the current study. TPIV accuracy is also assessed through the highly noise sensitive flow divergence, which is in line with other previous full

3D methods. Crucially, high gradient magnitude regions associated with coherent structures contain relatively low divergence values, meaning these are less sensitive to the divergence errors.

During flow characterisation similarities were observed between the current flow and previous matching geometries, with RMS velocity scaling as expected with impeller rotation frequency, and comparable large scale anisotropy. Global mean gradient ratios confirm that within the measurement accuracy the small scales are isotropic. Although non-Gaussian asymmetry in the gradient PDFs illustrate the departure from local isotropy, the trend of increasing intermittency with increasing Reynolds number shown in previous studies is not reproduced, possibly due to spatial filtering effects at lower resolutions.

Velocity correlation function analysis confirmed flow anisotropy and demonstrated the erroneously large out-of-plane TPIV correlation, due to insufficient angular resolution. Longitudinal structure function and 1D energy spectrum analysis demonstrated reasonable Kolmogorov scaling, with discrepancies attributable to large scale convergence and known PIV issues as observed in previous experimental results and the present investigation's accuracy study.

In agreement with previous results, the local dissipation calculated from the full VGT was distributed log-normally. However, the distribution of individual gradient components differ from this, and therefore estimates from one or two dimensional data completed with assumptions of isotropy tend to be poor; showing significant over-prediction of both small and large dissipation regions. Comparison of different isotropic gradient combinations demonstrate the estimate accuracy increases with the number of measured gradient components. The method of Michelet et al. (76) and the author's own choice of components provide the best estimates, with the arbitrary selection of the latter perhaps indicating the futility of more considered component selection.

The second order structure function method gives the most robust dissipation rate estimate, due to its reliance on larger, less noisy scales of motion. The direct calculation and energy spectrum methods under-predict this quantity, due to spatial averaging of the smallest scales. However, the quantification of this under-prediction from the simulation allows reasonable estimate correction. Importantly, for all of these characterisation results, the collapse of data is consis-

tently within the measurement uncertainty, demonstrating the expected Reynolds number scaling. Furthermore, good agreement between DPIV and TPIV results demonstrate the latter's validity.

6.3 Dissipation Scale Topology and Dynamics

After analysing a range of vortex identification schemes an enstrophy criterion was selected for its accuracy, threshold sensitivity, and simplicity. Visual observation suggests coherent structures take the form of both tubes and ribbons, with the latter likely to represent sheets truncated by the domain. High enstrophy and dissipation regions are well correlated, often coincident in ribbon-like structures or disassociated across a small spatial separation in tube-like structures; a trend also observed in JPDFs of these quantities. The limited domain size made classification of clustering behaviour difficult, however, within the available volume, a similar pattern to previous work was demonstrated (79).

Thresholding analysis showed that despite their small volume fraction (10% above $\Omega > \Omega_{rms}$), these coherent structures account for a significant proportion of the global dissipation (30%). If the surrounding regions are included to account for the known disassociation of dissipation and enstrophy, then 25% of the flow accounts for 60% of the dissipation. It has been argued that high enstrophy structures occupy too little of the volume to contribute significantly towards the total dissipation. However, these conclusion are usually based on the threshold level (for example Ruetsch & Maxey (103) reach this conclusions considering structures above $\Omega > 9.5\langle\Omega\rangle$). In the present study structures of strength $\Omega > \Omega_{rms}$, which are shown to be coherent over many turnover times, give rise to a significant proportion of the global dissipation in relation to their small volume fraction, and therefore on this basis the contribution of these structures cannot be considered insignificant. Higher thresholds only identify the structure cores, neglecting a significant proportion of surrounding fluid, which given the known disassociation of enstrophy and dissipation is likely to bias conclusions relating to their significance.

Coherent structure size characterisation proved inconclusive, with the vortex line tracking algorithm identifying similar diameter structures in terms of grid

units for all cases. This lack of grid independence has implications for other investigations with similar spatial resolution, including numerous DNS studies. Analysing the shape and significance of different coherent structures showed that both tube-like and ribbon-like structures contribute significantly towards the total dissipation. Tube-like structures are more prevalent, but due to their smaller volume occupy a similar total volume fraction to the larger ribbon-like structures. Ribbon-like structures tend to be associated with higher dissipation, although if the surrounding fluid is considered, that around tube-like structures tends to be marginally higher.

Principal rate of strain analysis demonstrates a flow preference towards biaxial strain, which is consistent with theoretical predictions (9). After divergence correction the ordered rates of strain are measured in the ratio $3.3 : 1 : -4.3$, in agreement with previous numerical and experimental studies. Although both high enstrophy and dissipation regions tend to be coincident with sheet forming strains, only the latter shows a direct connection, with the former somewhat independent of the strain field. Furthermore, despite the flow preference for sheet-forming strains, the larger number of tube-like structures implies the formation of these is not dependent on a favourable strain field, relying instead on some other unsteady mechanism such as the breakdown of vortex sheets via the Kelvin-Helmholtz instability.

In agreement with previous findings, in regions of high enstrophy the intermediate rate of strain and vorticity vector tend towards parallel alignment. Vincent & Meneguzzi (125) argue that this alignment occurs during vortex sheet formation, with subsequent sheet breakdown into vortex tubes resulting in the coincidental alignment of these structures too. However, as noted by Jiménez (51), strong vortex tubes give rise to an aligned bi-axial strain field. Therefore, while the mechanism proposed by Vincent & Meneguzzi may exist, it may not be the dominant process, and given the stability and longevity of tube-like structures the coherence of the strain field exerted by these may dominate over other more unsteady strains.

Topological invariant analysis was also used to assess fields of vorticity and strain, and despite higher uncertainty due to noise and divergence effects, reasonable agreement is seen between the present and previous studies, supporting

the concept of invariant universality perceived by Ooi et al. (85). Invariants were also used to demonstrate similar bi-axial strain and dissipation correlation, and the lack of vortex stretching in high enstrophy structures. These reinforce the previous conclusions that while high magnitude dissipation occurs predominantly in regions of sheet forming strain, despite their coincidence, tube-like structures remain independent (or perhaps the source) of these strains.

The self-similarity of both topological and strain and vorticity based JPDF distributions, and close agreement between different Reynolds number cases indicate similar behaviour across a wide range of scales; a feature observed during the visual comparison. As noted by Jiménez et al. (53) this scale-invarinace implies the coherent structures may simply be intense realisations of the background vorticity.

Examples of energy transfer are illustrated through the evolution of coherent structures, including vortex stretching, sheet breakdown, and vorticity roll up and merger. The formation and subsequent breakdown of high enstrophy structures in the periphery of larger scale motions illustrates several large to small scale energy transfer mechanisms; consistent with Richardson's original cascade theory. The substantial separation between parent vortex sizes and that of their children offers some support to the direct coupling concept of Tsinober (122), in which integral and dissipation scale motions may be directly linked. However, the domain size is insufficient to simultaneously render both of these scales, with the size of the largest instantaneous vortices difficult to estimate. Furthermore, these examples only provide qualitative evidence, and even if a method had been defined to characterise these mechanisms, the quantity of data is grossly insufficient for adequate statistical resolution of the dominant energy transfer process.

6.4 Final Remarks

To date, knowledge of the coherent dissipation scale structures has been largely gained computationally through DNS, with experimental support somewhat limited and almost exclusively reliant on assumptions of isotropy or Taylor's hypothesis. Therefore, the similarity between the previous numerical studies and the

current experimental results, which are performed at reasonably large Reynolds numbers and are independent of these assumptions, is extremely useful.

At the end of the Chapter 1 a number of questions were posed regarding the small scale structures. The present study offers some response to these, such as: defining the distributions of enstrophy and dissipation and quantifying the significance of the observed correlation; explaining the prevalence of tubes over sheets through an unsteady formation mechanism that permits consistency with the global predominance of bi-axial strain; and providing evidence of the formation processes and other energy transfer mechanisms. However, the inherent difficulty of measuring small turbulent scales at high Reynolds numbers preclude more definitive answers to other questions, such as: the Reynolds number scaling of the local enstrophy and dissipation distribution; the presence of vortex instability region at $R_\lambda > 700$; and the level of coupling between the integral and dissipation scale.

Determination of these issues await future investigations, which may build on the current understanding and conclusively identify the source and transfer mechanisms responsible for dissipative intermittency; paving the way for the creation of physically accurate models, and perhaps resolving the still open question of small scale universality.

6.5 Future Experimental Work

The fundamental TPIV technique is still the subject of active research, with increased accuracy and reduced computational costs desirable. Reductions in processing time are being sought through alternative reconstruction algorithms, for example after substantially reducing the size of the reconstruction problem using the MFG method (or equivalent) it is possible to apply either a simultaneous reconstruction algorithm (5), or calculate the volume intensity directly using least squares regression. The benefit these methods has yet to be definitively proven, with further investigation required.

The TPIV limitations experienced during the present investigation stem primarily from hardware constraints (as described in section 2.1.3), with improvement requiring an increase in the SNR through increases in either the light source

power or camera sensitivity. This would permit the extension of the results in one of three directions: accuracy, spatial resolution, or domain size.

As PIV accuracy is highly dependent on the number of particles within the interrogation region, reducing the seeding particle diameter while maintaining a similar SNR would allow a larger number of particles per volume, thereby reducing measurement uncertainty. Similarly with more particles, maintaining the same number of particles per interrogation region as the current investigation would allow smaller regions to be used, permitting an increase in the spatial resolution at the same accuracy level. Alternatively using similar seeding particles to the current investigation, a more powerful laser pulse could be spread over a larger volume while maintaining a similar energy density, allowing a larger volume to be investigated at a similar spatial resolution.

Capturing a larger volume at a similar spatial resolution would be useful in better understanding the connection between large and small scale, with important implications for the concept of universality, and higher resolution would permit a more conclusive investigation into the grid dependence of the smallest structures.

Use of a higher repetition rate laser would improve time resolution, permitting a more detailed investigation into the formation of various coherent structures, and through calculation of the material derivative, possible statistical characterisation of their evolution (32). Capturing a significant quantity of data, and applying this approach or other characterisation methods would allow statistical prediction of the dominant energy transfer process. While any conclusions based on a user defined characterisation will be dependent on the method itself, this type of approach offers the greatest potential to understand the physical energy transfer mechanisms; and therefore to model them successfully.

Appendix A

Number of Calculations and Memory Requirements

Number of Calculations

The computational cost of the MFG method can be compared to an iterative calculation using the MART algorithm (Equation 1.20) by considering the number of floating point operations required. In 3D, assuming a pixel to voxel ratio of 1, each voxel should receive an intensity contribution from approximately 4 pixels in each image. The MFG method updates every voxel with each correlated pixel, requiring 2 numerical operations: first to multiply the current pixel-voxel weighting value, W_{ij}, by the current pixel intensity, p_i; second to multiply this product with the previous voxel intensity. A further operation is required to normalise the voxels, raising them to the power $1/N_{cam}$. The number of required calculations can be expressed as, $F_{MFG} = N_{vox}(8N_{cams} + 1)$, where N_{vox} is the total number of voxels in the volume.

The iterative calculation is more involved, and crucially requires calculation of the current object projection, $W_i f$, for each voxel update. This calculation alone involves 2 numerical operations for each voxel in the current pixel's line of sight. The number of voxels can be estimated by assuming approximately 4 voxels will contribute to the pixel in each slice of the volume, resulting in $4L$ voxels, where L is the average length of the line of sight through the volume in voxels. Elimination of low or zero intensity voxels is employed in the current program to accelerate

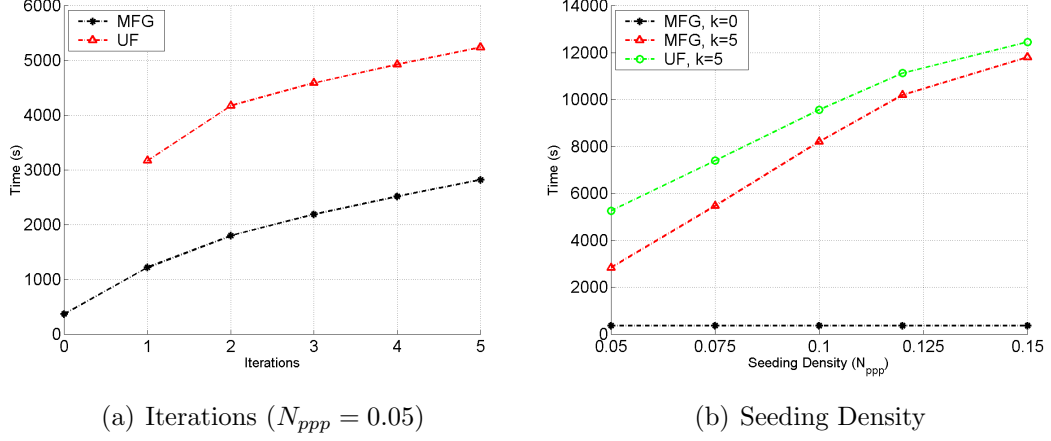(a) Iterations ($N_{ppp} = 0.05$) (b) Seeding Density

Figure A.1: Computational costs of the large volume reconstruction

the algorithm, and due to its significant impact on the total cost, this must be considered. As such, the value of N_{vox} will change after each iteration and L will change during each iteration. For cost estimation the former should be averaged over the entire iterative calculation, as should the latter in addition to being averaged over each iteration, with these averaged values denoted by $\overline{N_{vox}}$ and $\overline{\overline{L}}$. Therefore, combined with an additional 4 numerical operations in the MART algorithm, the iterative cost can be expressed as, $F_{MART} = 4k_{it}N_{cams}\overline{N_{vox}}(8\overline{\overline{L}}+4)$, where k_{it} is the total number of iterations.

The following example uses a $500 \times 500 \times 150$ voxel volume, imaged by three 500×500 cameras at $0°$ and $\pm30°$. Estimates of the $\overline{N_{vox}}$ variable are made by examining the number of zero and low intensity voxels after successive iterations. Dividing these by the total number of camera pixels, permits estimation of L. A uniform field at $k_{it} = 0$ results in initial values of $N_{vox} = 37.5 \times 10^6$ and $L \approx 150$. N_{vox} reduces accordingly for successive iterations in comparison to the initial value: 18%, 11%, 9%, 8%, 8%. After averaging appropriately, values of $\overline{N_{vox}} \approx 11 \times 10^6$ and $\overline{\overline{L}} \approx 30$ are obtained. Therefore, in this example, the MFG method alone performs approximately 175 times fewer calculations than the uniform guess solution after 5 iterations.

If used before the MART algorithm, the MFG method dramatically reduces the initial value of N_{vox} to around 40% of the uniform field method, with the

following reductions after successive iterations: 15%, 10%, 9%, 8%, 8%. After calculating average values of $\overline{N_{vox}} \approx 6 \times 10^6$ and $\overline{\overline{L}} \approx 20$, the MFG method prior to 5 MART algorithm iterations is shown to reduce the number of calculations by a factor of approximately 2.5. It is possible that further cost reductions could be made by using higher threshold values after the MFG to exclude more voxels from the MART iterations (for example, $k_{it} = 0$ solutions in section 2.3.3 are obtained by cross-correlating only around 10% of the total voxels).

Actual program run-times are shown in Figure A.1. These run times relate to a single large volume reconstruction, using a 64 bit, 3.4GHz *Intel Xeon* workstation. For a seeding density of $N_{ppp} = 0.05$, the MFG method more than halves the computational cost, for a given number of iterations (see Figure A.1(a)). As seeding density is increased, the benefit of the MFG method becomes less pronounced, as fewer voxels are excluded from the calculation as a result of higher image densities (see Figure A.1(b)). Using the MFG method without any MART algorithm iterations (the $k_{it} = 0$ solution) is shown to be significantly faster than the iterated solutions. The large magnitude of the speed increase possible with this new method alone, albeit offset by a slight decrease in accuracy (see section 2.3.3), may be particularly useful for time-resolved measurements, requiring hundreds of velocity fields.

It should be noted that the cost of this method is constant in Figure A.1(b) due to a previous programming practice of performing the first guess using all of the information in each image (including the zero intensity lines of sight). The current implimentation makes use of image sparsity, allowing the cost of this method to become dependent on the seeding density, which decreases it considerably, in particular for the lower seeding densities.

Memory Requirements

In addition to the huge number of calculations that must be performed, the reconstruction procedure also requires a significant amount of memory. This requirement stems from the need to relate every voxel to every pixel in the W matrix. Despite the extreme sparsity of this matrix, a huge number of points require storage. The number of points requiring storage can be approximated

by considering that each voxel will be related to 4 pixels per camera in 3D. The number of correlation values can be calculated as $N_{Wvals} = 4N_{cams}N_{vox}$.

However, storing these values in sparse matrix form requires two additional integer values indicating the m, and n position of each correlation value. By using an ordered list system this requirement can be be reduced to a single integer, however this still doubles the number of Bytes required for storage.

So the current example will require storage of 4.5×10^8 correlation values, and as many index values. Conversion to a memory requirement can be made through use of the REAL*4 and INTEGER*4 FORTRAN variables which both require 4 Bytes for every stored value, resulting in a RAM requirement of 3.6GBytes. This requirement is for the W matrix alone, and although smaller by comparison, there are still a number of other very large matrices which require storage.

Currently standard workstations have approximately $2 \rightarrow 4$GBytes of RAM. Therefore, in the absence of sufficient computational resources these weighting values will have to be recalculated every time they are needed, which will significantly increase the cost of the MFG method, but have a smaller relative impact on the MART algorithm. Thus, in order to obtain the maximum speed benefit from the MFG method, a workstation with a large amount of RAM memory is recommended, allowing the W matrix to be stored.

Appendix B

Reference Tables

Table B.1: Comparison of TPIV setups

	Cam. #; Res. (pix^2); Ang. (°)	Vol. (vox^3); Vol. (mm^3)	AR	IW Vol. (vox^3);(mm^3)	Overlap (%); Vectors	Particles	Seeding Density (N_{ppp})	Acc. (pix)
Elsinga *et al.*(30)	4 ; 700^2 ; 30	$700^2 \times 140$; $35^2 \times 7$	5	40^3 ; 2^3	50 ; 36 × 36 × 7	N/A	0.05	-
Elsinga *et al.*(29)	4 ; 1024^2 ; 37	$695 \times 580 \times 227$; $31 \times 26 \times 10$	3	50^3 ; 2.2^3	50 ; 28 × 24 × 5	1μm water droplets	-	0.2
Michaelis *et al.*(75)	4 ; 2048^2 ; 20	$1920 \times 1200 \times 384$; $80 \times 50 \times 16$	5	48^3 ; 2^3	75 ; 141 × 89×29	$30\mu m$ hollow glass spheres	0.085	-
Schröder *et al.*(107)	4 ; 896x860 ; 22.5	$816 \times 465 \times 720$; $34 \times 19 \times 30$	1.8	48^3 ; 2^3	75 ; 46 × 24×41	$1 - 2\mu m$ olive oil droplets	0.059	0.2
Elsinga *et al.*(28)	4 ; 1376×1040 ; 20	$820 \times 297 \times 640$; $30 \times 11 \times 24$	2.4	50^3 ; 1.85^3	75 ; 80 × 60×25	1μm water drops	0.05	0.1

# Appendix C

Additional PIV Set-up Details

The following procedure and settings were used to control the camera and laser timing.

Generator 1 creates a TTL top hat 'control pulse' at the desired measurement frequency, which was used as the master camera trigger input. On triggering, all cameras record 2 frames at a specified frame rate and exposure time, after which they return to standby to await the next trigger input. The frame rate is set to the reciprocal of the time increment, $1/\delta t$, and the exposure time the maximum value of the frame rate reciprocal. In the TPIV experiment three additional cameras are synchronised with the master camera.

To ensure the laser pulses are captured a delayed TTL pulse is triggered from the control signal, and used to trigger the first laser head. The delay time is set to half the camera frame rate reciprocal, placing the pulse in the centre of the exposure. The positive edge of this 5V TTL signal is used to externally trigger the laser flashlamp. The Q-switch is set to the internal trigger mode, to give the maximum pulse energy.

To fire the second laser head a second 5V TTL delayed pulse is triggered from the first laser head signal, with the delay set to the time increment, δt. In conjunction with the frame rate settings, this places the second pulse in the centre of the second frame exposure. Therefore, this delay setting alone controls the time increment and its accuracy, which could be measured to $\pm\delta t/1000$.

Figure C.1(a) shows the cable configuration required for this timing set-up, while a pulse diagram is shown in Figure C.1(b). The camera software set-up
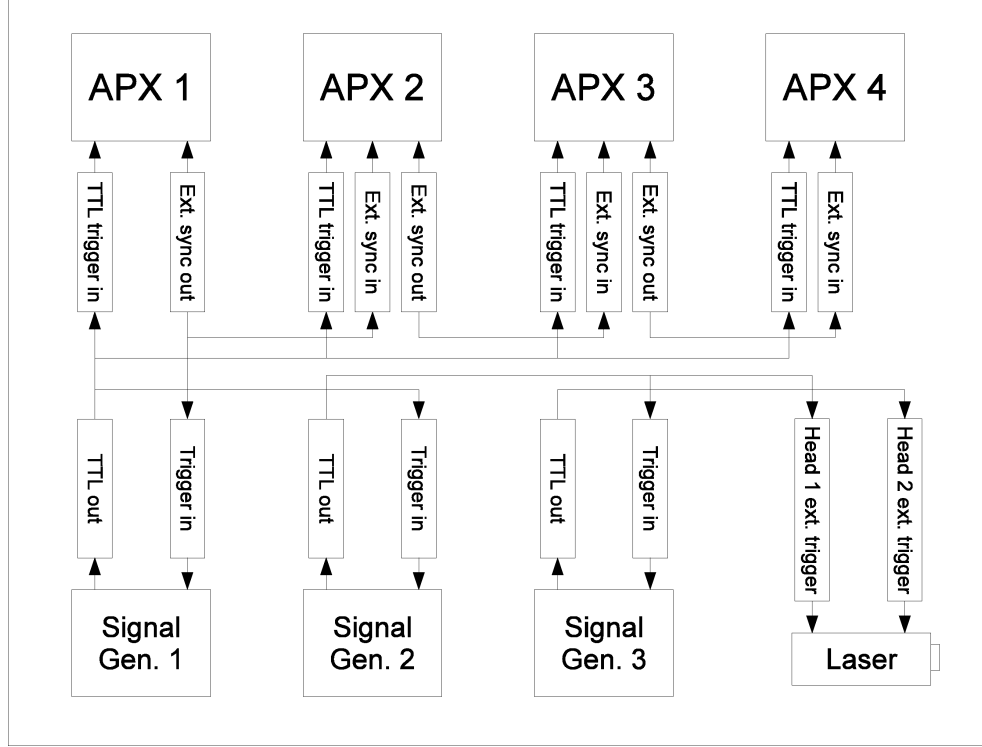
(a) Camera and laser wiring diagram
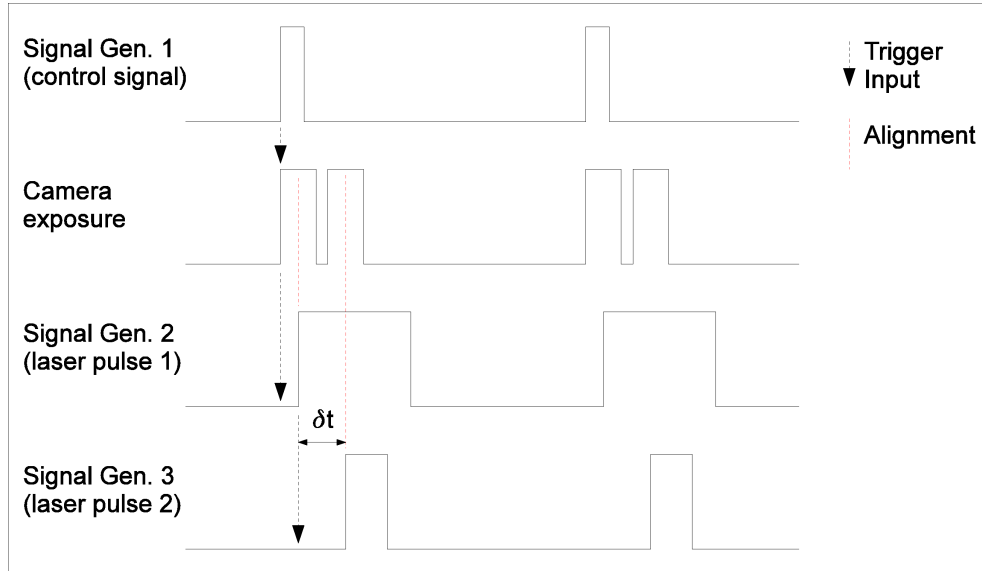
(b) Pulse Diagram

Figure C.1: Wiring and pulse diagrams

	Cam. 1	Cam. 2	Cam. 3	Cam. 4
Status	master	slave	slave	slave
Frame rate	δt	δt	δt	δt
Trigger Mode	Random 2	Random 2	Random 2	Random 2
General in	Trig. pos.	Trig. pos.	Trig. pos.	Trig. pos.
ext_sync_in	disable	cam_sync_pos	cam_sync_pos	cam_sync_pos
ext_sync_out	v_sync_pos	v_sync_pos	v_sync_pos	v_sync_pos

Table C.1: Photron FASTCAM Viewer V 2.4.5.2 Software set-up parameters

details can be found in Table C.1, and frequency settings are listed in Table C.2.

Case	Ω_I (rpm)	Ω_I (rad/s)	Re	δt (ms)	Measurement Freq. (Hz)
D1-1	0.67	0.0698	3.77×10^4	30	0.139
D1-2	1.33	0.1396	7.55×10^4	20	0.279
D1-3	2.67	0.2793	1.51×10^5	10	0.555
D1-4	4.00	0.4189	2.26×10^5	6	0.833
D1-5	5.33	0.5585	2.94×10^5	3	1.11
D2-1	0.67	0.0698	3.67×10^4	30	0.139
D2-2	1.33	0.1396	7.35×10^4	16	0.279
D2-3	2.67	0.2793	1.47×10^5	8	0.555
D2-4	4.00	0.4189	2.20×10^5	5	0.833
D2-5	5.33	0.5585	2.94×10^5	2	1.11
D3-1	0.33	0.0349	1.91×10^4	30	0.167
D3-2	0.67	0.0698	3.82×10^4	22	0.25
D3-3	1.33	0.1396	7.65×10^4	8	0.50
D3-4	2.67	0.2793	1.53×10^5	5	1.00
D3-5	4.00	0.4189	2.29×10^5	3	1.33
D4-1	1.33	0.1396	8.06×10^4	4	0.278
D4-2	2.67	0.2793	1.61×10^5	2	0.555
D4-3	4.00	0.4189	2.42×10^5	1	1.11
D4-4	5.33	0.5585	3.22×10^5	0.5	2.00
T1-1	0.33	0.0349	1.86×10^4	30	0.0714
T1-2	0.67	0.0698	3.72×10^4	15	0.143
T1-3	1.33	0.1396	7.45×10^4	8	0.278
T1-4	2.67	0.2793	1.49×10^5	4	0.555
T1-5	4.00	0.4189	2.23×10^5	2	1.00

Table C.2: PIV timing parameters

Appendix D

Movie Files

A DVD containing movie files can be found at the back of this report. The movies are organised into 3 folders for: vortex ring movies (see Section 2.3.6); mixing tank coherent structure visualisation movies (see Section 5.1.3); and mixing tank shape parameter movies (see Section 5.2.3.2).

The compressed movies (Cinepak codec) are suitable for viewing directly from the DVD using *Windows Media Player* or *QuickTime.* The uncompressed movies are of higher quality, however these files have significantly larger file sizes ($\sim 1GB$) and have been zipped for storage on the DVD. Therefore, these files should be copied to a computer hard disk drive with sufficient free space before unzipping.

The following section details the included files and their description.

Set 1; Vortex ring movies (compressed)

Three vortex ring movie files are included (see Section 2.3.6), demonstrating the vortex ring as it passes through the volume of interest from different camera positions. Iso-surfaces of vorticity magnitude (65% maximum vorticity magnitude) are plotted for the three cases: *1.* Blue iso-surfaces for MFG method alone (MFG $k_{it} = 0$); *2.* Red iso-surfaces for the uniform field initial guess and 5 MART iterations (UF $k_{it} = 5$); *3.* Green iso-surfaces for the MFG and 5 MART iterations (MFG $k_{it} = 5$). The file names for these movies are:

- Vortex ring movie Isometric.avi (AVI 1.8MB)
- Vortex ring movie X-Y.avi (AVI 1.9MB)

- Vortex ring movie Y-Z.avi (AVI 1.4MB)

Set 2; Coherent Structure Visualisation

Movies have been compiled using the time-resolved TPIV mixing tank measurements for each of the 5 runs to allow coherent structure visualisation (see Section 5.1.3). This folder is sub-divided into 3 folders containing: compressed movies; uncompressed movies; and 3D movies from the $T1-1$ case.

Iso-surfaces of enstrophy, $\Omega > n\Omega_{rms}$ (red-yellow colour map), and dissipation rate, $\epsilon > n\epsilon_m$ (blue-green colour map) are plotted on all movie files, and the Kolmogorov time, τ, is indicated.

Set 2a; Compressed movies

Colour bars were not included on these movies. The 4 isosurface levels for enstrophy and dissipation range from $n = 1 \rightarrow 4$ as in Figure 5.4. The file names for these movies are:

- Enstrophy_and_Dissipation_T1_1_compressed.avi (AVI 17.7MB)
- Enstrophy_and_Dissipation_T1_2_compressed.avi (AVI 20.4MB)
- Enstrophy_and_Dissipation_T1_3_compressed.avi (AVI 22.1MB)
- Enstrophy_and_Dissipation_T1_4_compressed.avi (AVI 22.0MB)
- Enstrophy_and_Dissipation_T1_5_compressed.avi (AVI 18.9MB)

Set 2b; Uncompressed movies

Iso-surfaces of enstrophy and dissipation are shown in the main window, with colour bars indicating the threshold level n. The 4 subplots at the bottom of the window illustrate a mid-z-plane slice of the normalised enstrophy, $\Omega^* = \Omega/\Omega_m$, enstrophy advection, $\Upsilon^* = \Upsilon/\Omega_m$, enstrophy stretching or compression, $\Pi^* = \Pi/\Omega_m$, and enstrophy diffusion, $\Psi^* = \Psi/\Omega_m$. The file names for these movies are:

- Enstrophy_and_Dissipation_T1_1.avi (zipped AVI 39.5MB)
- Enstrophy_and_Dissipation_T1_2.avi (zipped AVI 29.5MB)
- Enstrophy_and_Dissipation_T1_3.avi (zipped AVI 36.0MB)

- Enstrophy_and_Dissipation_T1_4.avi (zipped AVI 42.7MB)
- Enstrophy_and_Dissipation_T1_5.avi (zipped AVI 28.8MB)

Set 2c; 3D movies from T1-1 case (uncompressed)

A set of 3D movies was created by combining two slightly offset views of the volume, using perspective projection. Applying the *ColorCode 3D* anaglyph system, red and green channels from the left-hand image are overlaid with a monochrome blue channel created from the right-hand image using the RGB channel weighting [0.15 : 0.15 : 0.7]. The movies should be viewed using the 3D glasses provided. The file names for these movies are:

- Enstrophy_and_Dissipation_T1_1_3D_1.avi (zipped AVI 19.0MB)
- Enstrophy_and_Dissipation_T1_1_3D_2.avi (zipped AVI 32.7MB)
- Enstrophy_and_Dissipation_T1_1_3D_3.avi (zipped AVI 27.0MB)

Set 3; Shape parameter movies

Movies have been compiled using the time-resolved TPIV mixing tank measurements for each of the 5 runs in addition to the box counting algorithm (see section 2.4.4.2) to allow visualisation of the shape parameter (similar to Figure 2.25) and structure evolution. The movies show iso-surfaces of high enstrophy ($\Omega > \Omega_{rms}$), with structures coloured according to their shape parameter value. The shape parameter varies between $0 > H > 1$ with values approaching 0 and 1 representing sheet-like and tube-like structures respectively. Structures which have a mean alignment below the threshold are also included, but these are plotted using a black wire mesh. The Kolmogorov time, τ, is indicated. The file names for these movies are:

- Shape_parameter_T1_1.avi (zipped AVI 17.2MB)
- Shape_parameter_T1_2.avi (zipped AVI 13.6MB)
- Shape_parameter_T1_3.avi (zipped AVI 16.0MB)
- Shape_parameter_T1_4.avi (zipped AVI 20.3MB)
- Shape_parameter_T1_5.avi (zipped AVI 13.9MB)

References

[1] P. ABRY, S. FAUVE, P. FLANDRIN, AND C. LAROCHE. Analysis of pressure fluctuations in swirling turbulent flows. *Journal de Physique II*, **4**:725–733, May 1994. 26, 45

[2] R. J. ADRIAN. Twenty years of particle image velocimetry. *Experiments in Fluids*, **39**:159–169, August 2005. 30, 33, 52, 70

[3] M.P. ARROYO AND C.A. GREATED. Stereoscopic Particle Image Velocimetry. *Measurement Science and Technology*, **2**:1181–1186, 1991. 31

[4] W. T. ASHURST, A. R. KERSTEIN, R. M. KERR, AND C. H. GIBSON. Alignment of vorticity and scalar gradient with strain rate in simulated Navier-Stokes turbulence. *Physics of Fluids*, **30**:2343–2353, August 1987. 9, 22, 172, 173, 175, 176, 179, 187

[5] C. ATKINSON AND J. SORIA. An efficient simultaneous reconstruction technique for tomographic particle image velocimetry. *Experiments in Fluids*, pages 116–+, August 2009. 202

[6] F BELIN, P. MAURER, P. TABELING, AND H. WILLAIME. Observation of intense filaments in fully developed turbulence. *Physics of Fluids*, **9**:3843–3850, 1996. 18, 26, 45

[7] F BELIN, P. MAURER, P. TABELING, AND H. WILLAIME. Velocity gradient distributions in fully developed turbulence: an experimental study. *Journal de Physique II*, **6**:573–583, 1996. 18, 26, 120

[8] P. S. BERNARD AND J. M. WALLACE. *Turbulent Flow: Analysis, Measurement, and Prediction.* John Wiley & Sons, Hoboken, New Jersey, 2002. 5, 7

[9] R. BETCHOV. An inequality concerning the production of vorticity in isotropic turbulence. *Journal of Fluid Mechanics*, **1**:497–504, 1956. 20, 22, 172, 200

[10] D. BONN, Y. COUDER, P. H. J. VAN DAM, AND S. DOUADY. From small scales to large scales in three-dimensional turbulence: The effect of diluted polymers. *Physical Review E*, **47**:28–+, January 1993. 21, 27, 45

[11] O. N. BORATAV AND R. B. PELZ. Structures and structure functions in the inertial range of turbulence. *Physics of Fluids*, **9**:1400–1415, May 1997. 28

[12] M. E. BRACHET. Direct simulation of three-dimensional turbulence in the Taylor-Green vortex. *Fluid Dynamics Research*, **8**:1–8, October 1991. 27

[13] M. E. BRACHET, M. MENEGUZZI, A. VINCENT, H. POLITANO, AND P. L. SULEM. Numerical evidence of smooth self-similar dynamics and possibility of subsequent collapse for three-dimensional ideal flows. *Physics of Fluids*, **4**:2845–2854, December 1992. 20

[14] BRIGHAM. *The Fast Fourier Transform and it Applications.* Prentice Hall, Inc, New Jersey, 1988. 83

[15] C. BRÜCKER. Digital-Particle-Image-Velocimetry (DPIV) in a scanning light-sheet: 3D starting flow around a short cylinder. *Experiments in Fluids*, **19**:255–263, August 1995. 33

[16] J.M. BURGERS. A Mathematical Model Illustrating the Theory of Turbulence. *Advances in Applied Mechanics*, **1**:171–199, 1948. 9, 188

[17] C. BYRNE. Iterative algorithms in tomography. Technical report, Department of Mathematical Sciences, University of Massachusetts Lowell, 2005. 37, 65

[18] C.L. BYRNE. Accelerating the EMML and Related Iterative Algorithms by Rescaled Block Iterative Methods. *IEEE Transactions on Image Processing*, **7**:100–109, January 1998. 37, 65

[19] O. CADOT, D. BONN, AND S. DOUADY. Turbulent drag reduction in a closed flow system: Boundary layer versus bulk effects. *Physics of Fluids*, **10**:426–436, February 1998. 28, 170

[20] O. CADOT, S. DOUADY, AND Y. COUDER. Characterization of the low-pressure filaments in a three-dimensional turbulent shear flow. *Physics of Fluids*, **7**:630–646, March 1995. 8, 11, 15, 18, 19, 21, 22, 26, 27, 30, 45, 46, 158, 159

[21] P. CHAINAIS, P. ABRY, AND J.F. PINTON. Intermittency and coherent structures in a swirling flow: A wavelet analysis of joint pressure and velocity measurements. *Physics of Fluids*, **11**:3524–3539, 1999. 45

[22] M. S. CHONG, A. E. PERRY, AND B. J. CANTWELL. A general classification of three-dimensional flow fields. *Physics of Fluids*, **2**:765–777, May 1990. 11, 13, 107

[23] P.A. DAVIDSON. *Turbulence, an Introduction for Scientists and Engineers.* Oxford University Press, Oxford, UK, 2004. 7, 8, 9, 10, 12, 22, 23, 121, 126, 135

[24] J. DE JONG, L. CAO, S. H. WOODWARD, J. P. L. C. SALAZAR, L. R. COLLINS, AND H. MENG. Dissipation rate estimation from PIV in zero-mean isotropic turbulence. *Experiments in Fluids*, **46**:499–515, March 2009. 82, 83, 106, 124, 127, 130, 131, 132, 134

[25] L. DE ROSE AND D. PADUA. A matlab to fortran 90 translator and its effectiveness. Technical report, Center for Supercomputing Research and Development Coordinated Science Laboratory, University of Illinois, 1996. 65

[26] S. Douady, Y. Couder, and M. E. Brachet. Direct observation of the intermittency of intense vorticity filaments in turbulence. *Physical Review Letters*, **67**:983–986, August 1991. 19, 21, 25, 27, 45, 46

[27] Y. Dubief and F. Delcayre. On coherent-vortex identification in turbulence. *Journal of Turbulence*, **1**:11–+, December 2000. 11, 15, 16, 152

[28] G.E. Elsinga, D.J. Kuik, B.W. van Oudheusden, and F. Scarano. Investigation of the three-dimensional coherent structures in a turbulent boundary layer with Tomographic-PIV. In *45th AIAA Aerospace Sciences Meeting and Exhibit, Reno, Nevada*, January 2007. 39, 41, 42, 66, 71, 208

[29] G.E. Elsinga, F. Scarano, B. Wieneke, and B.W. van Oudheusden. Assessment of Tomo-PIV for Three-Dimensional Flows. In *6th International Symposium on Particle Image Velocimetry, Pasadena, California, USA*, September 2005. 39, 40, 63, 208

[30] G.E. Elsinga, F. Scarano, B. Wieneke, and B.W. van Oudheusden. Tomographic Particle Image Velocimetry. In *6th International Symposium on Particle Image Velocimetry, Pasadena, California, USA*, September 2005. 35, 38, 39, 43, 64, 72, 208

[31] G.E. Elsinga, F. Scarano, B. Wieneke, and B.W. van Oudheusden. Tomographic Particle Image Velocimetry. *Experiments in Fluids*, **41**:933–947, 2006. 32, 34, 35, 61, 63, 66, 71, 73, 76, 88, 90, 95

[32] G.E. Elsinga, F. Scarano, B. Wieneke, and B.W. van Oudheusden. On the time evolution of the local flow topology in a turbulent boundary layer. In *6th International Symposium on Turbulence and Shear Flow Phenomena, Seoul, Korea*, June 2009. 203

[33] S. Fauve, C. Laroche, and B. Castaing. Pressure fluctuations in swirling turbulent flows. *Journal de Physique II*, **3**:271–278, March 1993. 45

[34] J. M. FOUCAUT, J. CARLIER, AND M. STANISLAS. PIV optimization for the study of turbulent flow using spectral analysis. *Measurement Science and Technology*, **15**:1046–1058, June 2004. 83, 105, 106

[35] U. FRISCH. *Turbulence.* Cambridge University Press, Cambridge, UK, 1995. 2, 8, 27, 43

[36] B. GANAPATHISUBRAMANI, K. LAKSHMINARASIMHAN, AND N. T. CLEMENS. Determination of complete velocity gradient tensor by using cinematographic stereoscopic PIV in a turbulent jet. *Experiments in Fluids*, **42**:923–939, June 2007. 94, 95, 98, 102, 104, 130, 140, 143

[37] B. GANAPATHISUBRAMANI, K. LAKSHMINARASIMHAN, AND N. T. CLEMENS. Investigation of three-dimensional structure of fine scales in a turbulent jet by using cinematographic stereoscopic particle image velocimetry. *Journal of Fluid Mechanics*, **598**:141–175, 2008. 9, 19, 24, 25, 121, 131, 134, 158, 164, 170, 172, 174, 175, 176, 179, 182, 187

[38] S. GARG AND Z. WARHAFT. On the small scale structure of simple shear flow. *Physics of Fluids*, **10**:662–673, March 1998. 118

[39] W. K. GEORGE AND H. J. HUSSEIN. Locally axisymmetric turbulence. *Journal of Fluid Mechanics*, **233**:1–23, December 1991. 82, 131

[40] T. GOTOH, D. FUKAYAMA, AND T. NAKANO. Velocity field statistics in homogeneous steady turbulence obtained using a high-resolution direct numerical simulation. *Physics of Fluids*, **14**:1065–1081, March 2002. 126, 127

[41] G.T. HERMAN. *Image Reconstruction from Projections: The Fundamentals of Computerized Tomography.* Academic Press, London, 1980. 36, 37

[42] K.D. HINSCH. Holographic Particle Image Velocimetry. *Measurement Science and Technology*, **13**:R61–R72, 2002. 32

[43] J. O. HINZE. *Turbulence.* McGraw-Hill, New York, 1975. 7, 81, 82, 131, 132

[44] T. Hori and J. Sakakibara. High-Speed Scanning Stereoscopic PIV for 3D Vorticity Measurement in Liquids. *Measurement Science and Technology*, **15**:1067–1078, 2004. 33

[45] K. Horiuti. A classification method for vortex sheet and tube structures in turbulent flows. *Physics of Fluids*, **13**:3756–3774, December 2001. 186

[46] J. C. R. Hunt, A. A. Wray, and P. Moin. Eddies, streams, and convergence zones in turbulent flows. In *Center for Turbulence Research Report CTR-S88*, pages 193–208, December 1988. 13

[47] A. K. M. F. Hussain. Coherent structures and turbulence. *Journal of Fluid Mechanics*, **173**:303–356, December 1986. 8

[48] W. Hwang and J. K. Eaton. Creating homogeneous and isotropic turbulence without a mean flow. *Experiments in Fluids*, **36**:444–454, 2004. 106, 114, 128

[49] T. Ishihara, T. Gotoh, and Y. Kaneda. Study of High-Reynolds Number Isotropic Turbulence by Direct Numerical Simulation. *Annual Review of Fluid Mechanics*, **41**:165–180, January 2009. 105

[50] J. Jeong and F. Hussain. On the identification of a vortex. *Journal of Fluid Mechanics*, **285**:69–94, 1995. 11, 13, 14, 15

[51] J. Jiménez. Kinematic alignment effects in turbulent flows. *Physics of Fluids*, **4**:652–654, April 1992. 179, 180, 200

[52] J. Jiménez and A. A. Wray. On the characteristics of vortex filaments in isotropic turbulence. *Journal of Fluid Mechanics*, **373**:255–285, October 1998. 25, 26, 30, 162, 170

[53] J. Jiménez, A. A. Wray, P. G. Saffman, and R. S. Rogallo. The structure of intense vorticity in isotropic turbulence. *Journal of Fluid Mechanics*, **255**:65–90, 1993. 25, 26, 27, 28, 29, 83, 119, 121, 153, 154, 158, 160, 163, 165, 170, 182, 185, 187, 194, 201

[54] C.J. Kähler and J. Kompenhans. Fundamentals of Multiple Plane Stereo Particle Image Velocimetry. *Experiments in Fluids*, **29**:S70–S77, 2000. 31

[55] L. Kajitani and D. Dabiri. A full three-dimensional characterization of defocusing digital particle image velocimetry. *Measurement Science and Technology*, **16**:790–804, March 2005. 34

[56] Y. Kaneda and T. Ishihara. High-resolution direct numerical simulation of turbulence. *Journal of Turbulence*, **7**:20–+, 2006. 9, 17, 23, 24, 25, 28, 158

[57] R. D. Keane and R. J. Adrian. Optimization of particle image velocimeters: II. Multiple pulsed systems. *Measurement Science and Technology*, **2**:963–974, October 1991. 70

[58] R. M. Kerr. Higher-order derivative correlations and the alignment of small-scale structures in isotropic numerical turbulence. *Journal of Fluid Mechanics*, **153**:31–58, April 1985. 22, 23, 25, 158

[59] S. Kida. Description and Dynamics of Vortical Structures of Turbulence. In D. Reguera, L. L. Bonilla, and J. M. Rubí, editors, *Coherent Structures in Complex Systems*, **567** of *Lecture Notes in Physics, Berlin Springer Verlag*, pages 3–+, 2001. 3, 17, 22, 23, 26, 27, 162

[60] S. Kida and M. Takaoka. Vortex reconnection. *Annual Review of Fluid Mechanics*, **26**:169–189, 1994. 192

[61] S. K. Kim and S. K. Chung. An investigation on airflow in disordered nasal cavity and its corrected models by tomographic PIV. *Measurement Science and Technology*, **15**:1090–1096, June 2004. 39

[62] A.N. Kolmogorov. Local structure of turbulence in an incompressible fluid for very large reynolds numbers. *Akademiia Nauk SSSR Doklady*, **30**:299–303, 1941. 1, 7

[63] A. La Porta, G.A. Voth, F. Moisy, and E. Bodenschatz. Using cavitation to measure statistics of low-pressure events in large-reynolds-number turbulence. *Physics of Fluids*, **12**:1485–1496, 2000. 18, 19, 45, 46

[64] L.D. Landau and E.M. Lifshift. *Fluid Mechanics.* Pergamon Press, Oxford, 1959. 2

[65] P. Lavoie, G. Avallone, F. de Gregorio, G. P. Romano, and R. A. Antonia. Spatial resolution of PIV for the measurement of turbulence. *Experiments in Fluids*, **43**:39–51, July 2007. 132, 133, 134

[66] T.T. Lim and T.B. Nickels. Vortex rings. In S.I. Green, editor, *Fluid Vorticies*, Kluwer Academic Publishers, Dordrecht, 1995. 79

[67] S. J. Lin and G. M. Corcos. The mixing layer - Deterministic models of a turbulent flow. III - The effect of plane strain on the dynamics of streamwise vortices. *Journal of Fluid Mechanics*, **141**:139–178, April 1984. 20

[68] S. Liu, J. Katz, and C. Meneveau. Evolution and modelling of subgrid scales during rapid straining of turbulence. *Journal of Fluid Mechanics*, **387**:281–320, May 1999. 128

[69] T. S. Lund and M. M. Rogers. An improved measure of strain state probability in turbulent flows. *Physics of Fluids*, **6**:1838–1847, May 1994. 173, 174, 176

[70] T. S. Lundgren. Strained spiral vortex model for turbulent fine structure. *Physics of Fluids*, **25**:2193–2203, December 1982. 23, 28

[71] H.G. Maas, A. Gruen, and D. Papantoniou. Particle Tracking Velocimetry in Three-Dimensional Flows. *Experiments in Fluids*, **15**:133–146, 2004. 34

[72] M. V. MELANDER AND F. HUSSAIN. Cut-and-connect of two antiparallel vortex tubes. In *Stanford Univ., Studying Turbulence Using Numerical Simulation Databases, 2. Proceedings of the 1988 Summer Program p 257-286 (SEE N89-24538 18-34)*, pages 257–286, December 1988. 14

[73] A. MELLING. Tracer Particles and Seeding for Particle Image Velocimetry. *Measurement Science and Technology*, **8**:1406–1416, 1997. 50

[74] H. MENG AND F. HUSSAIN. Holographic particle velocimetry - A 3D measurement technique for vortex interactions, coherent structures and turbulence. *Fluid Dynamics Research*, **8**:33–52, October 1991. 32

[75] D. MICHAELIS, C. POELMA, F. SCARANO, AND B. WESTERWEEL, J.AND WIENEKE. A 3D Time-Resolved Cylinder Wake Survey by Tomographic PIV. In *12th International Symposium on Flow Visualization, Göttingen, Germany*, September 2006. 39, 41, 208

[76] S. MICHELET, Y. ANTOINE, F. LEMOINE, AND M. MAHOUAST. Mesure directe du taux de dissipation de l'énergie cinétique de turbulence par vélocimétrie laser bi-composante : validation dans une turbulence de grille. *Academie des Science Paris Comptes Rendus Serie B Sciences Physiques*, **326**:621–626, October 1998. 82, 131, 132, 198

[77] H. K. MOFFATT. Fixed points of turbulent dynamical systems and suppression of nonlinearity Comment 1. In J. L. LUMLEY, editor, *Whither Turbulence? Turbulence at the Crossroads*, **357** of *Lecture Notes in Physics, Berlin Springer Verlag*, pages 250–257, 1990. 194

[78] H. K. MOFFATT, S. KIDA, AND K. OHKITANI. Stretched vortices - the sinews of turbulence; large-Reynolds-number asymptotics. *Journal of Fluid Mechanics*, **259**:241–264, January 1994. 20

[79] F. MOISY AND J. JIMÉNEZ. Geometry and clustering of intense structures in isotropic turbulence. *Journal of Fluid Mechanics*, **513**:111–133, August 2004. 29, 83, 84, 86, 87, 158, 163, 166, 170, 199

[80] A.S. MONIN AND A.M. YAGLOM. *Statistical Fluid Mechanics II.* MIT Press, Cambridge, 1975. 2

[81] N. MORDANT, P. METZ, O. MICHEL, AND J.-F. PINTON. Measurement of Lagrangian Velocity in Fully Developed Turbulence. *Physical Review Letters*, **87**(21):214501–+, November 2001. 18

[82] J. A. MULLIN AND W. J. A. DAHM. Dual-plane stereo particle image velocimetry measurements of velocity gradient tensor fields in turbulent shear flow. II. Experimental results. *Physics of Fluids*, **18**(3):035102–+, March 2006. 19, 24, 54, 101, 121, 131, 141, 172, 179

[83] F. NATTERER. Numerical Methods in Tomography. *Acta Numerica*, **8**:107–141, 1999. 37

[84] P. O'NEILL AND J. SORIA. The relationship between the topological structures in turbulent flow and the distribution of a passive scalar with an imposed mean gradient. *Fluid Dynamics Research*, **36**:107–120, March 2005. 12, 107, 147

[85] A. OOI, J. MARTIN, J. SORIA, AND M. S. CHONG. A study of the evolution and characteristics of the invariants of the velocity-gradient tensor in isotropic turbulence. *Journal of Fluid Mechanics*, **381**:141–174, February 1999. 11, 12, 107, 147, 160, 183, 184, 201

[86] S. OTT, J. MANN, AND J. ANDERSEN. Experimental Study of Relative Turbulent Diffusion. Technical report, RisøNational Laboratory, Denmark, 1999. 3, 34, 58, 62

[87] G. PAN AND H. MENG. Digital in-line holographic piv for 3d particulate flow diagnostics. In *4th International Symposium on Particle Image Velocimetry, Göottingen, Germany*, September 2001. 32

[88] Y.-H. PAO. Structure of Turbulent Velocity and Scalar Fields at Large Wavenumbers. *Physics of Fluids*, **8**:1063–1075, June 1965. 105, 128

[89] T. Passot, H. Politano, P. L. Sulem, J. R. Angilella, and M. Meneguzzi. Instability of strained vortex layers and vortex tube formation in homogeneous turbulence. *Journal of Fluid Mechanics*, **282**:313–338, 1995. 20, 21, 194

[90] B. R. Pearson and R. A. Antonia. Reynolds-number dependence of turbulent velocity and pressure increments. *Journal of Fluid Mechanics*, **444**:343–382, October 2001. 27

[91] F. Pereira and M. Gharib. Defocusing digital particle image velocimetry and the three-dimensional characterization of two-phase flows. *Measurement Science and Technology*, **13**:683–694, May 2002. 33

[92] A. E. Perry and M. S. Chong. Topology of Flow Patterns in Vortex Motions and Turbulence. *Applied Scientific Research*, **53**:357–+, December 1994. 159

[93] S.B. Pope. *Turbulent Flows*. Cambridge University Press, Cambridge, 2000. 17, 164

[94] D. H. Porter, P. R. Woodward, and A. Pouquet. Inertial range structures in decaying compressible turbulent flows. *Physics of Fluids*, **10**:237–245, January 1998. 28

[95] A. K. Prasad. Stereoscopic particle image velocimetry. *Experiments in Fluids*, **29**:103–116, December 2000. 31, 52

[96] A. K. Prasad and K. Jensen. Scheimpflug stereocamera for particle image velocimetry in liquid flows. *Applied Optics*, **34**:7092–+, October 1995. 56

[97] Y. Pu and H. Meng. An advanced off-axis holographic particle image velocimetry (HPIV) system. *Experiments in Fluids*, **29**:184–197, December 2000. 32

[98] D. I. Pullin and P. G. Saffman. On the Lundgren-Townsend model of turbulent fine scales. *Physics of Fluids*, **5**:126–145, January 1993. 28

[99] A. PUMIR. A numerical study of pressure fluctuations in three-dimensional, incompressible, homogeneous, isotropic turbulence. *Physics of Fluids*, **6**:2071–2083, June 1994. 158

[100] M. RAFFEL, C. WILLERT, AND J. KOMPENHANS. *Particle Image Velocimetry: A Practical Guide.* Springer, New York, 1998. 31, 50, 52, 70, 81, 88, 94, 96

[101] L. F. RICHARDSON. Atmospheric Diffusion Shown on a Distance-Neighbour Graph. *Royal Society of London Proceedings Series A*, **110**:709–737, April 1926. 1

[102] S. K. ROBINSON. Coherent motions in the turbulent boundary layer. *Annual Review of Fluid Mechanics*, **23**:601–639, 1991. 11

[103] G. R. RUETSCH AND M. R. MAXEY. Small-scale features of vorticity and passive scalar fields in homogeneous isotropic turbulence. *Physics of Fluids*, **3**:1587–1597, June 1991. 9, 162, 178, 194, 199

[104] S. G. SADDOUGHI AND S. V. VEERAVALLI. Local isotropy in turbulent boundary layers at high Reynolds number. *Journal of Fluid Mechanics*, **268**:333–372, 1994. 105

[105] F. SCARANO. REVIEW ARTICLE: Iterative image deformation methods in PIV. *Measurement Science and Technology*, **13**:1–+, January 2002. 70

[106] T. SCHMITT, J.N. KOSTER, AND H. HAMACHER. Particle Design for Displacement Tracking Velocimetry. *Measurement Science and Technology.*, **6**:682–689, 1995. 51

[107] A. SCHRÖDER, R. GEISLER, G.E. ELSINGA, F. SCARANO, AND U. DIERKSHEIDE. Investigation of a Turbulent Spot Using Time-Resolved Tomographic PIV. In *13th International Symposium on Applications of Laser Techniques to Fluid Mechanics, Lisbon, Portugal*, June 2006. 35, 39, 41, 208

[108] Z.-S. SHE, E. JACKSON, AND S. A. ORSZAG. Intermittent vortex structures in homogeneous isotropic turbulence. *Nature*, **344**:226–228, March 1990. 20

[109] X. SHEN AND Z. WARHAFT. The anisotropy of the small scale structure in high Reynolds number ($R_\lambda \sim 1000$) turbulent shear flow. *Physics of Fluids*, **12**:2976–2989, November 2000. 17, 19, 194

[110] E. D. SIGGIA. Numerical study of small-scale intermittency in three-dimensional turbulence. *Journal of Fluid Mechanics*, **107**:375–406, June 1981. 16, 25

[111] PATRICE SIMARD AND GUY E. MAILLOUX. Vector field restoration by the method of convex projections. *Computer Vision, Graphics and Image Processing*, **52**:360–385, 1990. 143

[112] S. M. SOLOFF, R. J. ADRIAN, AND Z.-C. LIU. Distortion compensation for generalized stereoscopic particle image velocimetry. *Measurement Science and Technology*, **8**:1441–1454, December 1997. 56, 58

[113] K. R. SREENIVASAN. Fractals and multifractals in fluid turbulence. *Annual Review of Fluid Mechanics*, **23**:539–600, 1991. 17

[114] K. R. SREENIVASAN. On the universality of the Kolmogorov constant. *Physics of Fluids*, **7**:2778–2784, November 1995. 126

[115] K. R. SREENIVASAN. An update on the energy dissipation rate in isotropic turbulence. *Physics of Fluids*, **10**:528–529, February 1998. 7, 43

[116] K. R. SREENIVASAN AND R. A. ANTONIA. The Phenomenology of Small-Scale Turbulence. *Annual Review of Fluid Mechanics*, **29**:435–472, 1997. 2, 20, 22, 25, 26

[117] M. STANISLAS, K. OKAMOTO, AND C. KÄHLER. REVIEW ARTICLE: Main results of the First International PIV Challenge. *Measurement Science and Technology*, **14**:63–+, October 2003. 69

[118] M. Stanislas, K. Okamoto, C. J. Kähler, and J. Westerweel. Main results of the Second International PIV Challenge. *Experiments in Fluids*, **39**:170–191, August 2005. 69

[119] G. Stolovitzky, P. Kailasnath, and K.R. Sreenivasan. Kolmogorov's refined similarity hypotheses. *Physical Review Letters*, **69**:1178–1181, August 1992. 2

[120] P. Tabeling, G. Zocchi, F. Belin, J. Maurer, and H. Willaime. Probability Density Functions, Skewness, and Flatness in Large Reynolds Number Turbulence. *Physical Review E*, **53**:1613–1621, February 1996. 18, 26

[121] M. Tanahashi, T. Miyauchi, and J. Ikeda. Fine scale structure in turbulence, fluid mechanics and its application. In *IUTAM Symposium on Simulation and Identification of Organized Structures in Flows*, **52** of *Fluid Mechanics and Its Applications*, pages 131–140, 1999. 88, 89

[122] A. Tsinober. *An Informal Introduction to Turbulence.* Kluwer Academic Publishers, Dordrecht, 2001. 194, 195, 201

[123] A. Tsinober, E. Kit, and T. Dracos. Experimental investigation of the field of velocity gradients in turbulent flows. *Journal of Fluid Mechanics*, **242**:169–192, September 1992. 24, 104, 140, 172

[124] R. Verzicco and J. Jiménez. On the survival of strong vortex filaments in 'model' turbulence. *Journal of Fluid Mechanics*, **394**:261–279, September 1999. 26

[125] A. Vincent and M. Meneguzzi. Dynamics of vorticity tubes in homogeneous turbulence. *Journal of Fluid Mechanics*, **258**:245–254, January 1994. 21, 25, 28, 158, 200

[126] G. A. Voth, A. La Porta, A. M. Crawford, J. Alexander, and E. Bodenschatz. Measurement of particle accelerations in fully developed turbulence. *Journal of Fluid Mechanics*, **469**:121–160, October 2002. 34, 116, 117, 118, 122

[127] G.A. VOTH, K. SATYANARAYAN, AND E. BODENSCHATZ. Lagrangian acceleration measurements at large reynolds numbers. *Physics of Fluids*, **10**:2268–2280, 1998. 18, 19, 45, 47, 57, 118

[128] J. WESTERWEEL. Efficient detection of spurious vectors in particle image velocimetry data. *Experiments in Fluids*, **16**:236–247, February 1994. 70

[129] J. WESTERWEEL. Fundamentals of digital particle image velocimetry. *Measurement Science and Technology*, **8**:1379–1392, December 1997. 70

[130] J. WESTERWEEL, A. A. DRAAD, J. G. T. VAN DER HOEVEN, AND J. VAN OORD. Measurement of fully-developed turbulent pipe flow with digital particle image velocimetry. *Experiments in Fluids*, **20**:165–177, January 1996. 106

[131] B. WIENEKE. Stereo-PIV using self-calibration on particle images. *Experiments in Fluids*, **39**:267–280, August 2005. 31

[132] B. WIENEKE. Volume Self-Calibration for Stereo-PIV and Tomographic-PIV. In *7th International Symposium on Particle Image Velocimetry, Rome, Italy*, Sept. 2007. 61, 63

[133] B. WIENEKE AND D. MICHAELIS. Simultaneous 4-Plane PIV Measurements Using Tomographic Reconstruction. In *12th International Symposium on Flow Visualization, Göottingen, Germany*, September 2006. 39, 40

[134] B. WIENEKE AND S. TAYLOR. Fat-Sheet PIV with Computation of Full 3D-Strain Tensor using Tomographic Reconstruction. In *13th International Symposium on Applications of Laser Techniques to Fluid Mechanics, Lisbon, Portugal*, June 2006. 35, 39, 40, 96, 98

[135] C. E. WILLERT. Stereoscopic digital particle image velocimetry for application in wind tunnel flows. *Measurement Science and Technology*, **8**:1465–1479, 1997. 31, 58

[136] C. E. Willert and M. Gharib. Three-dimensional particle imaging with a single camera. *Experiments in Fluids*, **12**:353–358, April 1992. 33

[137] J. Willneff and A. Gruen. A new spatio-temporal matching algorithm for 3d-particle tracking velocimetry. In *The 9th International Symposium on Transport Phenomena and Dynamics of Rotating Machinery, Honolulu, Hawaii*, Febuary 2002. 34

[138] N. A. Worth and T. B. Nickels. Acceleration of Tomo-PIV by estimating the initial volume intensity distribution. *Experiments in Fluids*, **45**:847–856, November 2008. 65, 74, 90, 96

[139] N. A. Worth and T. B. Nickels. An Experimental Demonstration of Accelerated Tomo-PIV. In *iTi Conference on Turbulence, Bertinoro, Italy*, October 2008. 79, 80

[140] N.A. Worth and T.B. Nickels. A computational study of tomographic reconstruction accuracy and the effects of particle blocking. In *5th Joint ASME/JSME Fluids Engineering Conference, San Diego, California*, July 2007. 49, 50, 71

[141] P. K. Yeung and S. B. Pope. Lagrangian statistics from direct numerical simulations of isotropic turbulence. *Journal of Fluid Mechanics*, **207**:531–586, 1989. 160

[142] J. Zhang, B. Tao, and J. Katz. Turbulent flow measurement in a square duct with hybrid holographic PIV. *Experiments in Fluids*, **23**:373–381, 1997. 100, 101, 139, 141

[143] J. Zhou, S. Balachandar, and R. J. Adrian. Mechanisms for Generating Coherent Packet of Hairpin Vortices in Near-Wall Turbulence. *APS Meeting Abstracts*, pages 5–+, November 1997. 14

[144] G. Zocchi, P. Tabeling, J. Maurer, and H. Willaime. Measurement of the scaling of the dissipation at high Reynolds numbers. *Physical Review E*, **50**:3693–3700, November 1994. 2, 18, 43, 45, 47, 117, 118, 129, 135