

PAPER • OPEN ACCESS

A versatile scanning method for volumetric measurements of velocity and density fields

To cite this article: J L Partridge *et al* 2019 *Meas. Sci. Technol.* **30** 055203

View the [article online](#) for updates and enhancements.

You may also like

- [Optical Fiber Sensors Vol III: Components and Subsystems](#)
Peter Cockshott
- [Prediction of excess volume in implementing the ABC system for breath-hold treatment: A preliminary study](#)
A T Hoang, W L Jong and N M Ung
- [Accuracy testing of Wolf von Wulffing stand table in estimating teak stand volume](#)
T A Rahmani, F Kristiantoro, B Prihanto et al.

A versatile scanning method for volumetric measurements of velocity and density fields

J L Partridge[✉], Adrien Lefauve[✉] and Stuart B Dalziel[✉]

Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge, CB3 0WA, United Kingdom

E-mail: j.l.partridge@damp.cam.ac.uk

Received 21 December 2018, revised 15 February 2019

Accepted for publication 1 March 2019

Published 10 April 2019



Abstract

Understanding turbulence in a stratified environment requires a detailed picture of both the velocity field and the density field. Experimentally, this represents a significant measurement challenge, especially when full three-dimensional data is needed to accurately characterise the turbulent fields. This paper presents a new approach to obtaining such data through well-resolved, near-instantaneous volume-spanning measurements. This is accomplished using a simultaneous particle image velocimetry and planar-laser-induced fluorescence scanning setup that is highly versatile while allowing precise overlapping light sheets in a volume. Moreover, the approach detailed here allows for complex scan sequences over varying volumes without the need for mechanical adjustment. To demonstrate the technique, preliminary results from the buoyancy-driven exchange flow through an inclined duct connecting two reservoirs of different density fluids are shown.

Keywords: PIV, LIF, 3D-PIV, 3D-LIF

(Some figures may appear in colour only in the online journal)

1. Introduction


Flows of geophysical and environmental interest are frequently turbulent, exhibiting a range of scales and chaotic motions characterised by a large Reynolds number $Re = \frac{UL}{\nu} \gg 1$. Here, U is a typical velocity scale, L is a typical length scale and ν is the kinematic viscosity. Moreover, buoyancy forces (due to density differences) play an important role in these flows and are often characterised by a Richardson number $Ri = \frac{N^2 L^2}{U^2}$, where N is the buoyancy frequency.

Despite advances in experimental measurement techniques, capturing near-instantaneous, simultaneous volumetric density and velocity fields remains challenging. Nevertheless, having such information is desirable in the study of stratified turbulence due to the three-dimensional (3D) nature of the turbulent flow and the dynamical significance of the density

field. Such data would also be helpful if coherent structures are to be extracted from the flow. Coherent structures have offered great insight into unstratified turbulent flow (e.g. wall-bounded turbulence (Jiménez 2018)) and, more recently, stratified flow (e.g. Lucas *et al* (2017)). As well as this, there is still a need for experimental measurements of realisable flows to complement direct numerical simulations (DNS). This is due to the complexity of modelling certain physical boundary conditions and domains, as well as the non-trivial dependence on the Schmidt number $Sc = \frac{\nu}{\kappa}$, where κ is the molecular diffusivity. As demonstrated by Zhou *et al* (2017), the behaviour of stratified flows depend on Sc . However, DNS of high- Sc stratified flows ($Sc = 10^2$ – 10^3 in the salinity-stratified ocean) require sufficient resolution to not only capture the Kolmogorov scale, but also the Batchelor scale, respectively defined as

$$\eta_k = \left(\frac{\nu^3}{\epsilon} \right)^{1/4} \quad \text{and} \quad \eta_b = \frac{\eta_k}{Sc^{1/2}}, \quad (1)$$

where ϵ is the dissipation rate of turbulent kinetic energy. Currently, this resolution requirement makes high- Sc DNS prohibitive, even at the relatively low Re and ϵ that are easily

 Original content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](https://creativecommons.org/licenses/by/3.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

achieved using salt-stratified (high-Sc) laboratory experiments. To this end, we present a new approach to measure simultaneous volumetric density and velocity fields at the high Re and Sc typical of laboratory flows.

The layout of the remainder of this paper is as follows. In section 2 we describe some of the challenges associated with the current methodologies to measure velocity and density fields in stratified flows before describing our new approach in section 3. Our experimental setup, including the specific hardware used to perform the measurements, is detailed in section 4. Illustrative results are presented in section 5. Finally, discussions and conclusions are given in section 6.

2. Measurement challenges

2.1. Velocity measurements

Here we discuss some of the challenges associated with velocity measurement techniques. Two volumetric methods are compared: non-scanned and scanned. In non-scanned approaches the velocity field is calculated by illuminating and imaging (with multiple cameras) the whole volume of interest instantaneously. This differs from scanned techniques, where planar measurements calculated from thin light sheets are combined to build a volume. A full review of the techniques is beyond the scope of this paper and the reader is referred to Discetti and Coletti (2018) for more information.

A challenge when making optical measurements such as particle tracking velocimetry (PTV) or particle image velocimetry (PIV) in stratified flows is dealing with refractive index variations. The refractive index variations can be minimised by choosing two solutions that have different densities but the same refractive index (e.g. McDougall (1979), Dalziel *et al* (1999) and Verso *et al* (2017)), but inevitably small residual refractive index variations remain. Unfortunately, non-scanned volumetric approaches are more susceptible to any residual refractive index mismatch than scanning measurements, making them less suited to the study of stratified flows. This is due to non-scanned approaches relying on accurately triangulating the particles' 3D positions and calculating their 3D displacements at subsequent times. In comparison, Dalziel *et al* (2007) demonstrated that the planar PIV approach, used in the scanning measurements, is more robust to such errors as the pattern matching algorithm is less sensitive to absolute refractive index variations than to their rate of change.

As well as the refractive index issues, spatial resolution is also a limiting factor for non-scanned approaches as all of the particles within the measurement volume have to be distinguishable simultaneously in the camera images. This means that the yield of vectors in the total volume is limited, with either the spatial resolution or spatial extent being compromised. The resolution of a non-scanned system is typically expressed in particles per pixel (ppp). The latest algorithms can achieve 0.1 ppp, yielding one velocity vector per particle (Schanz and Schröder 2016). To compare this with a scanning system, consider a resolution of one velocity vector every 8×8 pixels, a four megapixel camera will give $\sim 6 \times 10^4$ vectors on the plane for PIV compared with $\sim 4 \times 10^5$ vectors

for the 0.1 ppp non-scanned system. However, by building a volume of more than six planes, the scanning approach produces a higher volumetric vector yield. On the other hand, an advantage to the non-scanning approach is that the 3D field is effectively instantaneous: there is no time lag when imaging across the volume and the flow is adequately frozen by the pulse width of the light source.

Scanning methods have their own unique challenges. The fundamental issue when using the scanning approach is ensuring sufficient light sheet overlap between a pair of images to permit accurate velocity measurements, i.e. keeping sufficient particle overlap between image pairs. Previous work by Brücker (1997) overcame this issue with the novel use of mirrors on a rotating drum to obtain a series of discrete pairs of light sheets that were coincident in space. However, the use of such a mechanism relies on mechanical reconfiguration to alter the scanning volume and does not lend itself to more complicated scan sequences. More recent studies have opted to scan a light sheet using oscillating mirrors positioned using galvanometers (for accurate and fast positioning). For example, Krug *et al* (2014) produced a light sheet by passing a laser beam through a cylindrical lens before scanning the sheet using an oscillating mirror. The light sheet, while scanning, passed through a second cylindrical lens with the same orientation as the first to further expand the sheet. Ultimately, this approach produced slowly diverging rather than parallel light sheets (as a small angular adjustment was needed to sweep the whole volume of interest) across the volume. A similar approach was demonstrated by Lawson and Dawson (2014) but with the second cylindrical lens set at 90° to make the scanned sheets parallel. Unfortunately, approaches like this are typically limited in spatial extent as they require large, high-quality optics so volumes are typically only a few centimeters deep. The new scanning method detailed in this article allows for easy adjustment of the scanning volume, requiring no mechanical adjustment, as well as capturing data over large ($0(10^{-1})$ m) deep scan volumes.

2.2. Density measurements

Density fields are often measured using planar-laser-induced fluorescence (PLIF). In a standard PLIF setup, a fluorescent dye is added to passively tag one of the solutions and, through careful calibration, the concentration field of the fluoresced dye can be related to the underlying density field (Crimaldi 2008).

Simultaneous velocity and density measurements can be achieved through a careful choice of fluorescent dye, selecting the absorption wavelength of the dye to the wavelength of the laser, and using a dye with a large Stokes shift, such that the fluoresced wavelength is significantly larger than the absorption wavelength (e.g. Webster *et al* (2001)).

As PLIF only requires a single image for each plane, it is relatively straightforward to extend it to 3D using the scanning approach. Numerous studies have made such measurements (e.g. Dahm *et al* (1991) and Tian and Roberts (2003)) but, to date, simultaneous velocity and density measurements in a volume have typically been limited to small spatial extents




Figure 1. Schematic of the scanning PIV/PLIF system showing a plan view of the scanning apparatus: (1) laser source; (2) beam tube to enclose the laser beam; (3) housing containing the oscillating mirror assembly that allows the beam to be displaced vertically; (4) linear guide rail that allows the light sheet-producing-optics to travel in the scanning direction; (5) carriage containing a 45° mirror and the light-sheet-producing optics; (6) light sheet; and (7) motor used to move the traverse.

(e.g. Krug *et al* (2014)). As it is desirable to have simultaneous density and velocity measurements over large volumes, we have extended the scanning approach to allow fast scanning of large volumes that readily allow the simultaneous capture of 3D density fields.

3. Methodology


3.1. Hardware

In the present study, a dual-cavity, pulsed laser is used as the light source. The laser has a maximum repetition rate of f_{laser} for each cavity. To accomplish the scanning, a mirror and the light-sheet-producing optics, consisting of a set of diverging (plano-concave) cylindrical lenses, are positioned on a linear traverse. The layout of all of the components that produce and scan the light sheet is shown in figure 1.


To allow overlapping light sheets, a pair of oscillating mirrors is used to superimpose a small additional translation of the light sheet in the scanning direction. To accomplish this, the beam from the pulsed laser is first directed by a fixed 45° mirror onto the first oscillating mirror. The beam is reflected by this mirror onto the second oscillating mirror before being redirected to the traverse carriage by a second fixed 45° mirror (see figure 2(a)). To facilitate initial beam alignment, the two fixed 45° mirrors are housed on adjustable mounts. Once on the traverse carriage, the beam is reflected by a final 45° mirror that passes the unexpanded beam through a set of diverging cylindrical lenses to form a light sheet perpendicular to the traversing direction (see figures 1 and 2(b)).

Figure 3 shows the traverse position (black line) and light sheet position (filled circles) after taking into account the displacement induced by the oscillating mirrors. The grey line indicates the approximate (due to the finite response time of the mirrors) position of the light sheet if it were a continuous light source. The figure demonstrates how the oscillating mirrors effectively allow the continuous motion of the traverse to be broken down into a series of discrete pairs, thus allowing overlapping light sheets for the PIV analysis.

By positioning the oscillating mirrors as shown in figure 2(a), and rotating them in tandem, the beam is displaced



(a)



(b)

Figure 2. (a) Details of the oscillating mirror assembly: (1) first fixed 45° mirror that directs the incoming beam from the laser into the plane of the sketch; (2) and (3) oscillating mirrors that allow the beam to be displaced vertically a distance D_m ; (4) second fixed 45° mirror that redirects the beam down to the scanning optics. (b) Beam path after exiting the oscillating mirror assembly, looking from the side of (a): (5) beam after exiting the mirror assembly in (a); (6) fixed 45° mirror on the traverse carriage that makes the laser beam vertical (and the vertical beam displacement horizontal); (7) cylindrical lenses to produce a light sheet; (8) 3D view of (7) to illustrate the curvature of the cylindrical lens. The position of the laser beam and the carriage are shown for two sequential laser pulses, with the first and second pulses of the pair shown by the solid and dotted lines, respectively.

vertically before entering the light-sheet-producing optics. The parallel displacement in the vertical, after passing through the 45° mirror at the base of the optics (see figure 2(b)), results in a parallel displacement in the traversing direction. To preserve the parallel nature of the beam displacement, the thickness of the light sheet δz should be controlled using a focusing module (e.g. a set of spherical lenses) between the laser and the oscillating mirror assembly.

For a given oscillating mirror assembly (i.e. a given L_c and H_c the horizontal and vertical distance between the mirror centres, respectively) the resting state angle is prescribed by $\theta_0 = 2\theta_m = \tan^{-1}(H_c/L_c)$. The vertical beam displacement $D_m/2$ can then be determined from

$$\frac{D_m}{2} = L_c \sin(\theta_0 + \theta') - H_c \cos(\theta_0 + \theta'). \quad (2)$$

Here, θ' is the angular perturbation away from the resting state θ_0 (see figure 4). The apertures of the two oscillating mirrors do not need to be equal and, as clear from figures 2(a) and 4, the aperture of the first mirror only needs to be large enough to

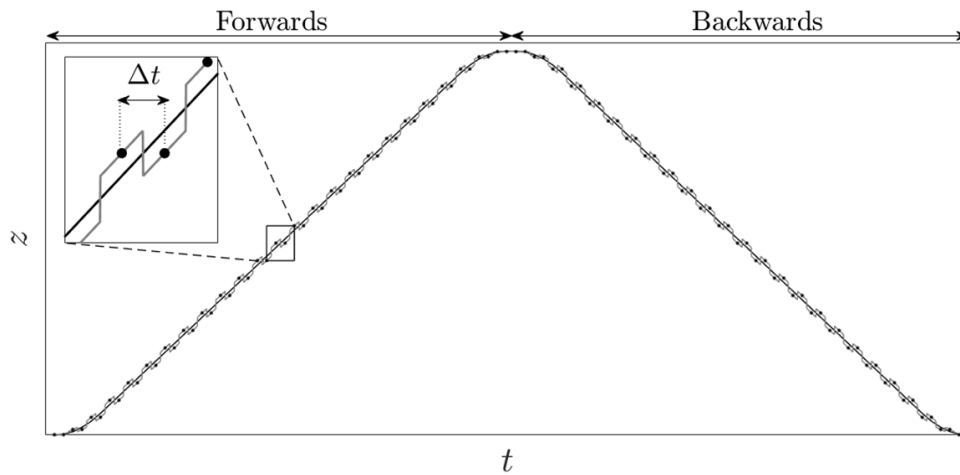


Figure 3. Profile of the traverse position with time (black line) and the position of the light sheet due to the displacement of the mirrors (filled circles) that enable pairs of images at discrete spatial locations. The position of the light sheet if it was being fired continuously, rather than pulsed is also shown (grey line).




Figure 4. Schematic showing the layout of the oscillating mirrors illustrating the positioning and subsequent rotation to achieve a vertically displaced beam.

accommodate the incoming beam. The aperture of the second mirror A_m needs to be larger (in one direction) than the beam diameter to accommodate the beam displacement, and its size (along with the resting angle θ_0) will determine the maximum beam displacement possible. Larger mirrors require larger galvanometers and generally have a slower response. Increasing the distance between the two mirrors allows for a smaller θ' (and faster response) for a given D_m , but with a larger positioning error due to the angular resolution of the galvanometer.

To simultaneously obtain the full velocity field and density field, a set of three cameras are required for the measurements (a typical layout is shown in figure 5). For the present discussion, we keep the cameras stationary as this removes a potential source of noise from the velocity fields. As is the case




Figure 5. Schematic of the scanning PIV/PLIF system showing a top view of the system and the camera layout: (1) the region of interest; (2) the light sheet that is scanned a distance W ; (3) the array of cameras: Cam A and Cam B for particle images and Cam C for the concentration of a scalar field for PLIF.

for non-scanned volumetric methods, having the cameras stationary requires that the lenses have a sufficiently small aperture (large f-number) so that the particles remain adequately focused across the volume of interest. For the stereo PIV measurements, two cameras are typically positioned with some angular separation and fitted with Scheimpflug adapters to provide better focusing on the scanned planes. These cameras need to be fitted with shortpass filters (or, preferably, bandpass filters centred at the wavelength of the laser) to eliminate the signal fluoresced by the dye. For the density measurements, the third camera is ideally positioned with its optical axis normal to the light sheet (to minimise distortion) and needs to be fitted with a longpass filter to remove the light directly scattered from the particles, leaving only the signal from the fluorescent dye.

The pulsed laser, oscillating mirrors, traverse and cameras are all triggered using a hardware-based timing system with all of the components sharing a common clock to keep everything synchronised (see the appendix A for more information).

To perform the PIV analysis, an accurate coordinate calibration procedure is required. This is especially true given




Figure 6. Sketch of experimental setup. Top view: (1) approximate region of the duct where the volumetric measurements are made by the three cameras (A and B for stereo PIV and C for PLIF). Side view: (2) inclined duct that confines the exchange flow between the two reservoirs; (3) reservoir containing the relatively dense fluid; (4) reservoir containing the relatively light fluid (tagged with Rhodamine 6G for PLIF). The laser beam was emitted from the scanning system (5) and illuminated the flow through the base of the duct. Measurement volume: a 3D view of the measurement volume (1) is shown clarifying the orientation of the coordinate system used.

the scanned nature of the measurements, as the optical path between the light sheet and the camera varies in time. The calibration procedure used for the measurements presented is detailed in appendix B.

4. Experimental setup

The methodology outlined in this paper was used to perform volumetric measurements of the buoyancy-driven exchange flow in an inclined duct similar to that studied by Meyer and Linden (2014), Lefauve *et al* (2018) and Lefauve (2018). We opted to use this particular experiment with our new measurement method as it is stratified (hence knowledge of the density field was desired) and has 3D spatial structure (due to the confining duct boundaries). Moreover, turbulent flow states are possible that result in a complex 3D structure of both the velocity field and density field. To this end, we chose our parameters such that we were in the ‘Turbulent Regime’ identified by Meyer and Linden (2014).

4.1. Experimental apparatus

The experimental setup is sketched in figure 6 and consisted of two reservoirs connected only through a duct. Each reservoir had dimensions $1000 \times 200 \times 500$ mm and the duct had a square cross-section of height and width $H = 45$ mm with a total length $L = 1350$ mm. The duct and the reservoirs were made of Perspex (acrylic) and were of good optical quality. The duct passed through a flexible gasket that was located in the central wall separating the two reservoirs. This allowed

the duct to be tilted whilst still maintaining a seal between the reservoirs. The reservoirs were filled with fluids of different densities $\rho_0 \pm \Delta\rho/2$ where $\Delta\rho = 11.7 \text{ kg m}^{-3}$. The duct connecting the two reservoirs was tilted at an angle $\theta = 5^\circ$ from the horizontal (figure 6).

We chose to orientate the x axis with the streamwise direction of the flow, along the duct, with z in the spanwise direction (the scanning direction in this setup). The y axis was then inclined at an angle of 5° from the vertical upwards direction. Note that, rather than using the convention of having the z axis aligned with gravity as is common in stratified flow literature, our chosen orientation matches that used in the methodology of this paper (i.e. z in the scanning direction). This coordinate system had its origin in the middle of the duct, such that $-L/2 \leq x \leq L/2$ and $-H/2 \leq y, z \leq H/2$. Our measurement volume was approximately $6H$ (300 mm) in the streamwise direction and spanned the full vertical and spanwise extents of the duct (both H or 45 mm here). The measurement volume was offset from the centre of the duct (in the relatively dense reservoir, as shown in figure 6) to avoid the flexible gasket that separated the two reservoirs.

4.2. Optical components

For the experiments presented in this paper, a frequency-doubled dual-cavity Litron Nano L100 Nd:YAG laser was used as the light source, with a wavelength of 532 nm. The laser had a maximum repetition rate $f_{\text{laser}} = 100$ Hz for each of the two laser cavities, each outputting 50 mJ per pulse. For the galvanometer mirrors we used two Thorlabs GVS311/M units

attached to a Thorlabs GPS011 driver/power supply. The aperture of the mirrors was $A_m = 10$ mm and the orientations were controlled by an analogue signal that determined the subsequent rotation of the galvanometer. Given that only small rotations of the mirror ($<1^\circ$) were needed in the current system, the highest resolution setting for the controller was used (a ± 10 V input signal corresponded to a $\pm 10^\circ$ rotation). In the arrangement here, the mirrors were mounted so that their resting state was $\theta_0 = 19.9^\circ$ with a horizontal and vertical distance between the centres of $L_c = 188$ mm and $H_c = 68$ mm, respectively. Given this geometric arrangement, the aperture of the mirrors, and the diameter of the beam from the laser (~ 4 mm), the beam could be displaced by ± 2.9 mm (with a relative rotation of $\pm 0.4^\circ$ to the mirrors resting angle θ_m), which gives a maximum displacement between image pairs of $D_m = 5.8$ mm. The angular resolution of the galvanometer introduces negligible error for our geometric arrangement, and is 0.2% of the maximal rotation $\pm 0.4^\circ$ required to have the maximal beam displacement. Note that this maximum allowable beam displacement D_m (see figure 2) set an upper limit on the speed of traversing V for a fixed inter frame time Δt as, to get overlapping light sheets, we require $V < \frac{D_m}{\Delta t}$. For the current system, in single-pulse mode with $\Delta t = 10$ ms, this meant that the oscillating mirrors limited the maximum speed of the traverse to $V_{\max} = 580$ mm s $^{-1}$. Note that, in practice, the limit was less than this due to the limited aperture of the sheet-producing optics and the non-negligible divergence of the laser source.

To produce a scanning light sheet from the nominally circular beam emitted from the laser, a system of cylindrical lenses was mounted on the traversing carriage¹. The traverse carriage was mounted on an Igus DryLin SAW rail system. The system was chosen due to its sliding plastic bearings, avoiding a metal on metal contact that is prone to introducing vibrations into the system, and as it was impervious to salt and water. This rail was also chosen for its relatively wide separation between the rails that minimised any roll of the carriage during motion. To move the carriage, a stepper motor was chosen for the simplicity of positional control without relying on a separate position resolver (as, for example, would be the case with a conventional servo motor). The traverse carriage was attached to a stainless steel lead screw of pitch 2 mm, diameter 10 mm, and with a maximum travel of 500 mm. Due to the inertia of the carriage and friction from the rails, the carriage could not simply be put instantaneously into motion at a constant velocity and so a constant acceleration/deceleration phase was implemented. For the experimental results shown here, the total time spent accelerating/decelerating was $\sim 15\%$ of the scanning period. As a consequence, the spacing of

measurements in z was non-uniform at the beginning and end of the scan.

Inevitably, vibrations were introduced into the traverse system. However, the vibrations introduced into the carriage because of the stepper motor were found to be negligible in practice (micro stepping was used to help minimise this). Moreover, to check if the traverse system deteriorated over time, a three-axis accelerometer was attached to the traverse carriage to monitor the vibration levels during the scanning. Even after multiple uses in the lab environment, the accelerometer data showed that the vibration levels of the carriage had not increased. To minimise any residual vibrational feedback from the complete traversing system onto the experiment itself, the traversing system and laser were mounted on a large (0.6 m \times 1.2 m) honeycomb optical breadboard (Thorlabs PBG51507).

Three Teledyne Dalsa Falcon2 8M cameras were used to image the flow. These ten-tap CameraLink CMOS cameras had a maximum resolution of 3320×2502 pixels. However, given the aspect ratio of the duct, a reduced resolution for each of the cameras (3320×824 pixels) was used. This reduced resolution also allowed for the 200 frames per second needed to match the maximum speed of the lasers in double-pulse mode. The three cameras (two for PIV and one for PLIF) were fixed in position to one side of the duct (as shown in figure 6), and the angular offset between the two PIV cameras was chosen to be $\sim 80^\circ$. To improve focusing across the image, both PIV cameras were fitted with Scheimpflug adapters. It was not deemed necessary to install liquid filled prisms in between the stereo PIV cameras and the volume of interest but we note that in some situations this could be advantageous, as discussed by Prasad and Jensen (1995). The PLIF camera had its optical axis normal to the light sheet and so, unlike the PIV cameras, did not require a Scheimpflug adapter. The PIV cameras were fitted with Micro Nikkor 60 mm f/2.8D lenses at aperture f/8 and the PLIF camera had a Nikkor 50 mm f/1.2D lens at aperture f/1.2 (so that only a low concentration of fluorescent dye was required in the dyed layer, consistent with the linear relation assumed in the calibration calculation, see appendix B.4). Locating the PIV cameras 0.6 m away from the measurement volume provided an adequate field of view (the width of the section imaged was approximately 250 mm) and allowed the images to be adequately focused across the scan. Due to the increased aperture, the PLIF camera was positioned 1 m away from the measurement volume (the width of the section imaged was approximately 320 mm for the PLIF camera).

As it was not feasible to only seed particles in the region of interest, a transparent Perspex (acrylic) box filled with water (but with no particles) was positioned between the inner wall of the reservoir and the outer boundary of the duct itself in the optical path between the cameras and the light sheet. This ‘optical box’ removed a large portion of noise present in the raw images for all cameras that would otherwise occur due to the optical path between the light sheet and cameras encountering a large number of particles when high seeding densities were used.

¹ Note that for historical reasons a Dantec focusing module (9080X0911) was attached to the cylindrical lenses in the current arrangement. Due to the spherical nature of the focusing lenses, a simple parallel displacement of the beam would not provide parallel light sheets. Instead, the rotation of the individual mirrors was fine tuned so that the laser entered the optics at a slight angle thus allowing nearly parallel light sheets with good overlap between image pairs. In an ideal system, the focusing module would be placed before the oscillating mirrors.

4.3. Experimental protocol

Prior to the experiment, the system was calibrated as discussed in appendix B.1. The calibration refinement step was typically performed with particle images captured during the experiment with a second set of refinement images captured between Cam C (the PLIF camera with the filter removed) and Cam A either before or after the experiment. For the experiment, both of the reservoirs were filled with fresh water and allowed to settle overnight to help degas the fluid, to minimise bubble formation during the experiment, and reach the ambient temperature of the laboratory. The duct was then inclined to the desired angle (as determined by a Digi-Pas DWL-280Pro digital inclinometer) and the end open to the reservoir that would contain the denser fluid was temporarily sealed (to avoid unwanted exchange of fluid between the reservoirs before the start of the experiment).

The density difference between the reservoirs was achieved using two salt solutions, NaNO_3 and NaCl , with their densities measured at a temperature of 20°C (the same temperature as the laboratory environment) using an Anton Paar DMA 5000 density meter. This particular combination of salts was selected to allow matching of refractive index and as they, to a good approximation, mix linearly and have similar diffusivities at the low concentrations required here (Olsthoorn and Dalziel 2017). Although the heat of dilution of NaNO_3 is significantly greater than that of NaCl , for the concentrations used here the temperature change was not dynamically significant to the mixing process, and the associated refractive index change was small enough to not cause an issue. We matched the refractive indices of the solutions at 532 nm (the wavelength of the laser source) with a relative error of $\Delta n/n \approx 10^{-4}$ and verified this by a handheld refractometer (Reichert Technologies Goldberg).

For PIV, polyamide particles with a diameter of $50\ \mu\text{m}$ and density between $1020\text{--}1030\text{ kg m}^{-3}$ were used to seed the flow. The particles were chosen due to their small ratio of settling to mean flow velocities $V_p/(\sqrt{g'H}) = 4.17 \times 10^{-4}$ while still being large enough to provide a clear particle image. The polyamide particles were added to the flow with a small amount ($\sim 5\text{ ml}$) of MagnumTM Rinse Aid to prevent aggregation of particles. Enough seeding particles were added to the flow to ensure that there were always $\gtrsim 5$ particles in any interrogation window.

At this stage, the first sequence of calibration images for the PLIF measurement could be captured, as discussed in appendix B.4. Note that before recording any calibration images for the PLIF measurements, or beginning capture during the experiment, the laser and associated optical components were allowed to warm up for at least 30 s . This avoided a non-trivial time-varying change in spatial structure and magnitude of the light sheet. Such changes would otherwise hamper the PLIF calculation. After this first sequence of calibration images had been captured, rhodamine 6G was added to the reservoir containing the relatively light NaCl solution and mixed giving a final concentration of rhodamine 6G in the reservoir of $C_1 = 15\ \mu\text{g l}^{-1}$. Adding the dye to this reservoir avoided having the light sheet pass through a layer of dyed

fluid before entering the measurement volume. After ensuring the dyed fluid was mixed into the reservoir and the duct, the second sequence of PLIF calibration images were captured.

The experiment could then be started by simply removing the temporary seal from the end of the duct. Initially, there was a transient stage of the experiment as a gravity current of denser un-dyed NaNO_3 fluid propagated through the duct into the opposite reservoir. After this period of transient flow had finished we began taking the 3D measurements reported here.

4.4. Measurement resolution and processing

The in-plane resolution of the measurements for the setup discussed here were primarily set by the laser system and the choice of camera. The resolution in the scanning direction was predominantly set by the laser as the maximum number of velocity fields that could be obtained in 1 s was 100 (achieved by firing the two cavities of the laser out of phase at their maximum repetition rate 100 Hz). Moreover, the precise choice of laser and optical components used to produce the light sheet controlled the thickness of each plane δz , and the subvolume over which each of the initial 2D pixel displacements were determined. For the measurements herein, given the small depth required to image the duct (45 mm) and the typical thickness of the light sheet in the current setup ($1\text{--}2\text{ mm}$), we chose to discretise the domain into ~ 40 planes. This avoided unnecessary overlapping of light sheets given that the uncertainty in the position of the velocity fields in the z direction will be the same order as the light sheet thickness. To adequately resolve the particle displacements, the system was in double-pulse mode so there was a spacing between light sheets of $\Delta t = 7\text{ ms}$ between the pairs of analysed frames.

The resolution of the in-plane measurements in (x, y) was set by the resolution of the camera, the seeding density of the PIV particles, and the PIV algorithm used. For the results shown here, all of the raw images were processed using DigiFlow (Dalziel *et al* 2007). The processing used the DigiFlow 2017a PIV algorithm, selecting an initial interrogation window of height and width 31 pixels and a spacing of 12 pixels (both horizontally and vertically) equivalent to a 60% overlap of interrogation windows. The algorithm underlying these PIV calculations has some important differences from most PIV implementations. First, as introduced for synthetic schlieren by Dalziel *et al* (2000), the pattern matching kernel is based on an L_1 norm measure of the differences between images in the interrogation windows and is used in place of the more common L_2 norm of a cross-correlation function. This kernel is applied iteratively, utilising displacement information from previous passes to advect the image pair captured at $t = t_i \pm \Delta t/2$ to their anticipated state at $t = t_i$. The size, shape, and weighting profile of each interrogation window are adjusted (based on their information content and the spatial gradients in the displacement field) during this process to faithfully capture high gradients while achieving a low level of noise. Additionally, a strategy of weighting or removing anomalous pixels increases further the robustness of the results to particles entering or leaving the light sheet.

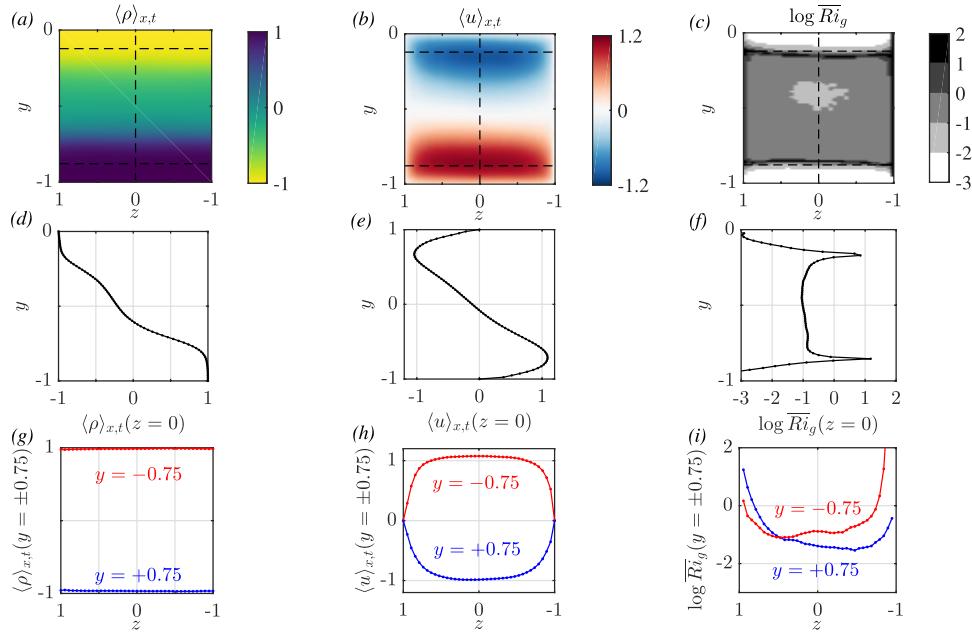


Figure 7. Means from the volumetric data. Each column corresponds to a different quantity: ρ , u (streamwise velocity), and \overline{Ri}_g . The first row (a)–(c) shows the $\langle \cdot \rangle_{x,t}$ averaging (or combination of averages in (c) to form \overline{Ri}_g) of each quantity on a (y, z) slice. The second row (d)–(f) shows a single vertical profile of the same data at the centre of the duct ($z = 0$) that corresponds to the dashed line shown in (a)–(c). The final row (g)–(i) shows the variation in the spanwise direction z of the same data for two different vertical locations, again shown by the dashed lines in (a)–(c).

The effective resolution is not constant across the field of view due to the distortion across the images, caused by the angles of the cameras relative to the light sheet. The resolution also varies as the light sheet is scanned, a consequence of the cameras being stationary, so there is higher resolution when the light sheet is closest to the cameras. However, with the setup detailed here, the loss of resolution because of the scanning was negligible (a change in pixels of $<5\%$ for a given physical length over the depth of the scan) compared to the distortion of the images (a change of $\sim 20\%$ across the image). Therefore, the maximum resolution of the velocity measurements in the (x, y) plane reported here was one velocity vector every 0.51 mm and the minimum resolution was one velocity vector every 0.63 mm. The total yield of vectors in a scan, acquired every 0.77 s, was approximately $400 \times 100 \times 40$ (x, y, z) . Indicative of the quality of the vector fields, the reconstruction error field $E(\mathbf{x})$ (see equation (B.12)) had an average in the (x, y) plane of $\langle E \rangle_{x,y} < 6.7 \times 10^{-2}$ (spanning all z in the scan) with a maximum standard deviation $\sigma_{\max} = 4.6 \times 10^{-2}$, both in units of pixel/frame. It was observed that the error systematically increased in z (increasing distance from the cameras). However, it is worth noting that, due to size constraints of the current dual-plane calibration target, the calibration could only be performed over half of the spanwise extent of the duct (the side situated closest to the cameras). This is a possible cause of the increasing trend observed as a similar trend is also found if we simply go between the forward and inverse mappings as defined in appendix B.2.

The density fields in world coordinates were calculated as outlined in appendix B.4. Note that, for the data shown here, only the first of the two PLIF images was used to reconstruct the density field. As no interrogation window is required for

the PLIF calculation, the resolution of the PLIF measurements is higher than those of the velocities. In the current setup, there was a density measurement in the (x, y) plane every 0.1 mm. In total, there were approximately $3000 \times 500 \times 40$ (x, y, z) density measurements every 0.77 s, given by the resolution of the camera and the number of frames in a scan. For both the velocity and density fields, the resolution in the scanning direction z is 1.26 mm with an accuracy set by the light sheet thickness δz (1–2 mm).

Before presenting the data, the density fields were processed to remove line artifacts from (x, y) planes (due to light rays passing through residual air bubbles, clusters of particles, and imperfections in the tank surfaces) using a three-step approach. As all of the rays emanated from some virtual origin below the tank and spread with the light sheet, the first step straightened the lines by mapping the (x, y) planes into a ‘ray coordinate’ system producing planes with vertical line artifacts. The second step removed these vertical line artifacts from the planes using a wavelet method described by Münch *et al* (2009). The planes, now without line artifacts, were mapped back to world coordinates. Finally, after removing the line artifacts, the density data was interpolated onto the lower resolution grid of the velocity data after being median filtered over a suitably sized window.

5. Experimental results

For the results presented here, we choose to non-dimensionalise velocities \mathbf{u} by $\sqrt{g'H}$, where $g' = g\Delta\rho/\rho_0$ is the reduced gravity, and normalise all lengths by $H/2$. As a result, with the duct angle fixed at $\theta = 5^\circ$, we can construct a Reynolds

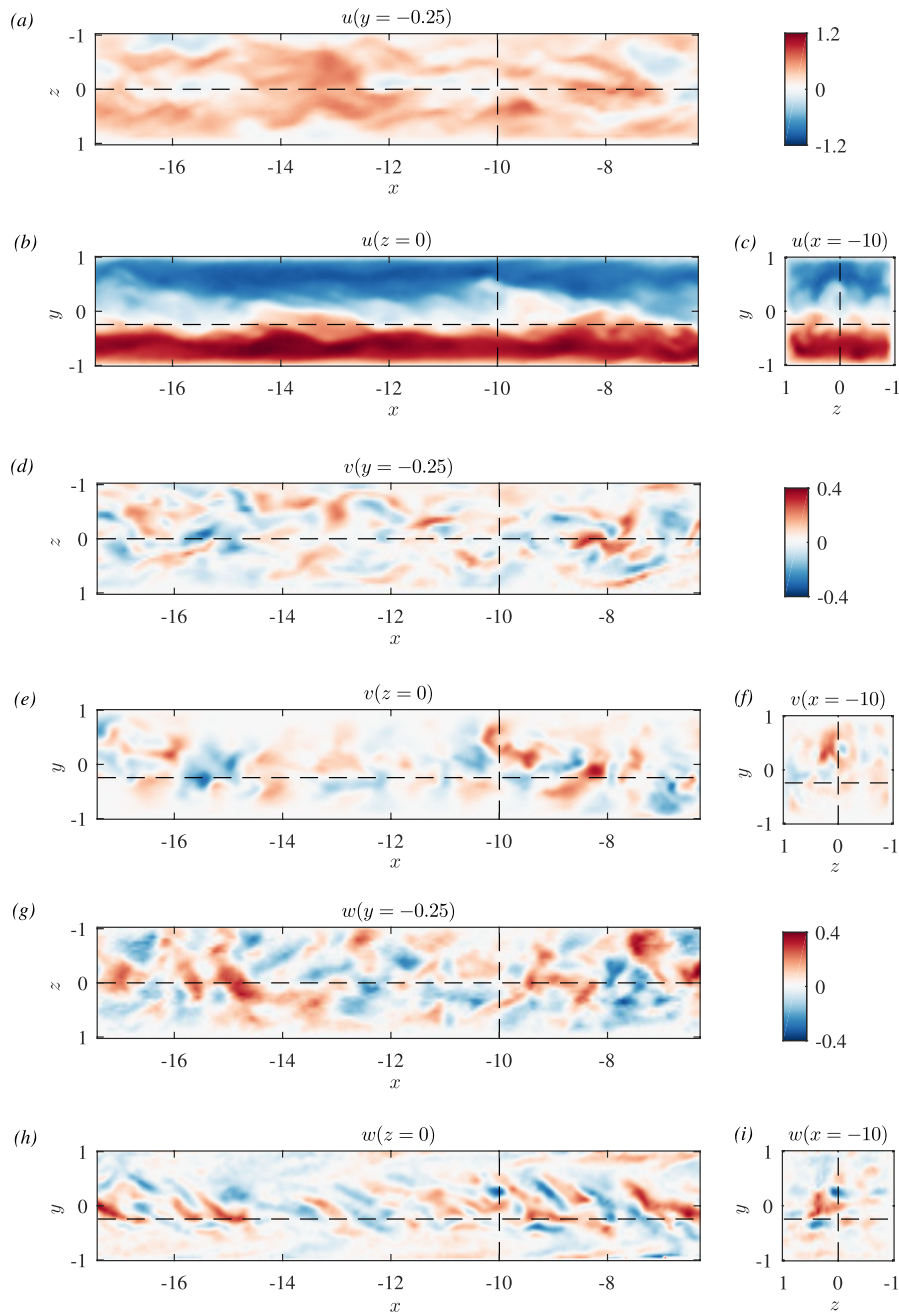


Figure 8. Instantaneous snapshots of all of the velocity components u (a)–(c), v (d)–(f), and w (g)–(i) with no filtering of the data. On (a), (d) and (g) the $y = -0.25$ horizontal plane; (b), (e) and (h) streamwise vertical $z = 0$ plane; and ((c), (f) and (i) the vertical $x = -10$ plane. The locations of the various slices are indicated by the dashed lines in (a)–(i). Note the different scale as typically $|u| > |v|, |w|$.

number $\text{Re} = \sqrt{g'H}H/(2\nu) = 1516$ to characterise the flow. The natural time scale to non-dimensionalise time is the advective time $\frac{1}{2}\sqrt{H/g'}$. Finally, the non-dimensional density field is $\tilde{\rho} = 2(\rho - \rho_0)/\Delta\rho$ such that $-1 \leq \tilde{\rho} \leq 1$. Therefore, our final measurement volume in non-dimensional units was $-17.4 \leq \tilde{x} \leq -6.3$, $-1 \leq \tilde{y} \leq 1$, $-1 \leq \tilde{z} \leq 1$ and a total of 561 advective time units were captured with a non-dimensional time between volumes of 2.40. The tilde indicating non-dimensional variables will be dropped henceforth.

To get an overall impression of the flow, time-averaged quantities of \mathbf{u} and ρ are shown in figure 7. In comparison to a laminar exchange flow, where there are two layers separated

by a sharp density interface, mixing is evident in this relatively high Re number flow. Here, thin and approximately well-mixed layers are confined to the horizontal boundaries at $y = \pm 1$. These layers are separated by a weakly-stratified interior (figures 7(a) and (d)). Mean profiles of the streamwise velocity u are also shown in figures 7(b) and (h) and indicate the spanwise variation present within the flow due to the confining lateral boundaries at $z = \pm 1$. In the same geometry, the importance of this confinement to the flow at lower Re has been investigated by Lefaue *et al* (2018), who found a non-trivial modification to the classical Holmboe instability with significant 3D structure. These results demonstrate that,




Figure 9. Instantaneous snapshots of density ρ (a)–(c) and enstrophy $|\omega|^2$ (d)–(f). On (a) and (d) the $y = -0.25$ horizontal plane; (b) and (e) streamwise vertical $z = 0$ plane; and ((c) and (f) the vertical $x = -10$ plane. The locations of the various slices are indicated by the dashed lines in (a)–(f).

even at relatively low Re , 3D spatial measurements provide meaningful physical insight that would be lost in single plane measurements.

Figure 7(g) also shows that there is no significant spanwise variation of the density, and therefore no mean pressure gradients in the spanwise direction due to horizontal density gradients. Finally, indicative of the stability of the flow, the mean gradient Richardson number $\overline{Ri}_g(y, z)$ is also shown in figures 7(c), (f), and (i). Note that here

$$\overline{Ri}_g(y, z) = \frac{1}{4} \frac{\langle N^2 \rangle_{x,t}}{\langle S^2 \rangle_{x,t}}, \quad (3)$$

where

$$\langle N^2 \rangle_{x,t} = \left(\frac{g'}{H} \right) \frac{\partial \langle \rho \rangle_{x,t}}{\partial y} \quad (4)$$

is the buoyancy frequency associated with the averaged density profile $\langle \rho \rangle_{x,t}$ and

$$\langle S^2 \rangle_{x,t} = 4 \left(\frac{g'}{H} \right) \left(\frac{\partial \langle u \rangle_{x,t}}{\partial y} \right)^2 \quad (5)$$

is the squared shear associated with the mean profiles of streamwise velocity $\langle u \rangle_{x,t}$, with $\langle \cdot \rangle_{x,t}$ indicating the average in the streamwise direction x and time. From figures 7(c) and (f) it is evident that there is an approximately constant

region where $\overline{Ri}_g(y, z) \simeq 0.12$ that is confined vertically to the centre of the duct, associated with the weakly-stratified interior and approximately constant shear in this region (see figures 7(d) and (e)). A low value of $\overline{Ri}_g(y, z)$ is expected given the turbulent nature of the flow, where values are typically $\overline{Ri}_g(y, z) < 0.25$ (Holt *et al* 1992).

Instantaneous data are shown in figures 8 and 9. Significant 3D variation is evident in all components of velocity \mathbf{u} throughout the entire domain, as shown in figures 8(a)–(h). At the same instant in time and on the same planes, the density field ρ and the enstrophy $|\omega|^2$ are shown in figures 9(a)–(e). These figures highlight the necessity of time-resolved 3D fields in analysing such flows, especially if coherent structures are to be extracted from the complex flowfield. Furthermore, the enstrophy could not be calculated without knowing gradient information of all components of \mathbf{u} .

To demonstrate the near-instantaneous nature of the volumetric data, isosurfaces of enstrophy and density are shown in figure 10. To be able to capture structures, we only need to resolve the Eulerian timescales of the flow. This is distinct from needing to resolve the Lagrangian timescale between pairs of images to accurately determine the particles displacements (which are Lagrangian tracers) that would require a much faster frequency of acquisition over each volume (for example, as required in conventional non-scanned volumetric




Figure 10. Instantaneous isosurfaces of (a) $\rho = 0.5$ and (b) $|\omega|^2 = 12$ in (b). For the data shown here, an isotropic 3D Gaussian filter with $\sigma_{\text{filt}} = 1$ vector spacing was used to smooth the data with a central difference scheme used to compute the derivatives of the velocity field. For both plots, the vertical y axis has been stretched by a factor of two to aid visualisation and the vertical and spanwise extents of the data have been limited to $[-0.75, 0.75]$ to avoid the signal at the boundaries obscuring the view.

methods). Crucially, quantities in the scan direction (including derivatives) are reliably resolved as the time between velocity/density subvolumes are only separated by $2\Delta t$ (therefore only double the Δt needed to resolve the particle displacements) and the data are effectively skewed across the scan direction. The distorted nature of the data could be advected using the velocity information but this step has not been carried out on the data presented here.

A further measure of the quality of the data can be inferred from $\zeta = \|\mathbf{u} - \mathbf{u}_{\text{div}}\|_1$ the L_1 norm of the difference between \mathbf{u} and a divergence free field \mathbf{u}_{div} calculated using the method described by Wang *et al* (2017). For the data shown here, the mean of this positive quantity over the whole volume and all time is $\bar{\zeta} = 2.5 \times 10^{-2}$ with standard deviation $\sigma = 2.2 \times 10^{-2}$ (all in non-dimensional units) illustrating that the correction is small (note that the value of the mean corresponds to ~ 0.17 pixel/frame, the same order as our error estimate in section 4.4). Given that all experimental data will have some non-zero divergence, and the fact that the measurements are near-instantaneous rather than truly instantaneous, we anticipate that $\zeta > 0$. However, our measurements are validated as only a small numerical correction is required and we note that the qualitative difference between the \mathbf{u} fields shown in this paper and the corresponding \mathbf{u}_{div} fields is almost imperceptible.

6. Discussion and conclusions

Depending on the flow of interest, there are two enhancements that extend the methodology presented in this paper. These enhancements make use of the same fundamental approach but facilitate its use with ‘slower flows’ (flows that are slow

compared to the camera frame rate) and flows with a strong out-of-plane motion.

For flows that are slow relative to the camera frame rate, the methodology outlined so far would be slaved to the time between light sheets. Slow flows require a larger Δt so that the particle displacements are sufficient to obtain accurate velocity measurements. Therefore, for a given number of light sheets in a scan N the time taken to scan the volume will be $N\Delta t$. However, so far we have only discussed the simple ‘mode 0’ operation of the system, where the mode number is given by m in

$$z_k = Z_{2k} = Z_{2(k+m)+1}. \quad (6)$$

Here, z_k defines the z position where two light sheets overlap, with Z_k the z position of the k th laser pulse. Therefore, for mode 0, overlapping pairs of images each comprise of an even-numbered frame and its immediately following odd number, e.g. $z_0 = Z_0 = Z_1$, $z_1 = Z_2 = Z_3$, etc. However, for slower flows it is beneficial to operate in higher modes where the images that are spatially coincident are separated in frame number. Operating in higher modes allows for greater temporal resolution for the complete scan that would otherwise be slave to the Δt required to obtain a sufficient particle displacement between images. Essentially, the effective Δt used in the PIV calculation can be increased without having to decrease the scan rate or the number of subvolumes the volume is discretised into. However, operating in higher modes does make it necessary to displace the laser beam further and therefore requires bigger oscillating mirrors and larger aperture cylindrical lenses to accommodate the larger amplitude beam displacement.

In traditional PIV methods (planar or stereo), strong out-of-plane motion can cause errors in the PIV algorithm (due to

substantial loss of particles between frames) or demand a high acquisition rate (to minimise the loss of particles between frames) that increases the noise of the measurements due to the inevitable lower in-plane pixel displacement. A potential use of the system, be it scanning or not, is the use of the mirrors on galvanometers to allow thin light sheets in configurations that would otherwise require a thick light sheet to accommodate the out-of-plane motion w . If it known beforehand that the flow to be measured has a unidirectional mean flow in the out-of-plane direction \bar{w} , the mirrors can be used to displace the light sheet in the z direction a suitable amount to accommodate the motion, i.e. $D_m \sim \bar{w}\Delta t$. The anticipated out-of-plane motion could be determined by an initial experiment, with a thick light sheet, and then repeated with a thinner light sheet to improve the accuracy of the measurements.

To conclude, we have presented a highly versatile scanning PIV/PLIF technique. Our approach allows for fast, repeated scans of a measurement volume. The volume is assembled from a series of subvolumes that allow for the simultaneous measurement of all three-components of the velocity field and the density field. The size of the plane (x, y) can be readily adjusted by simply changing the cylindrical lenses, and the distance in the scanning direction can be adjusted ‘on-the-fly’ by suitable commands to the traverse. Moreover, the system lends itself to complex scan sequences, e.g. varying volume sizes in the scanning direction, or several scans followed by fixed plane measurements with higher temporal resolution.

The novel addition to the setup is the use of two mirrors on galvanometers to position the light sheet during the scanning. These mirrors allow thin, overlapping light sheets between pairs of images despite the light-sheet-producing optics being continuously translated by a traverse system to scan a large volume.

Ultimately, the setup detailed in this paper can scan a volume of approximately 0.1 m^3 in less than a second, where the time taken is limited by the repetition rate of the laser and the desired number of frames in the scanning direction. Within each scanned volume, the total yield of velocity vectors is $O(10^6)$ with $O(10^7)$ simultaneous density measurements.

We have demonstrated the capability of our technique by taking measurements of a stratified shear flow. As the purpose of this article was to describe the method, we refer the reader to Lefauve *et al* (2018) and Lefauve (2018) for further details of the dynamics of the flow. Measurements of this nature are useful in the investigations of such flows because of their inherent 3D nature, due to turbulence, and the importance of the density field in such problems.

Acknowledgments

We gratefully acknowledge the skills and expertise provided by the technical staff of the G K Batchelor Laboratory and Dr Mark Hallworth. This work was supported by EPSRC Programme Grant EP/K034529/1. Finally, we wish to thank the authors of Wang *et al* (2017) who graciously gave us their code for the velocity divergence correction.

Appendix A. Timings

Typical pulse sequences for the various components are shown in figures A1 and A2. Figure A1 shows the pulse sequence in single-pulse mode (firing the two laser cavities together) and figure A2 shows the double-pulse mode (laser cavities are fired out of phase).

The signals to the cameras control the exposure time and image readout (triggered by the negative edge) so the cameras’ global shutters are open for a time t_e . Given that the flow field is not illuminated except for the short (typically $O(10)$ ns) pulse from the lasers, the effective timing between images captured is determined purely by the laser pulse timing. In single-pulse mode, the maximum temporal resolution of the measurements is determined by the smaller of f_{laser} and the maximum frame rate of the cameras. In double-pulse mode, the maximum temporal resolution is determined by the smaller of $2f_{\text{laser}}$ and the maximum frame rate of the cameras. In the scanning system (in either pulse mode), the time between each measurement subvolume is $2\Delta t_c$, with Δt_c the time between camera frames.

The signal to the mirrors is a square wave alternating between the two mirror positions to translate the light sheet (forward and back) to compensate for the traverse motion. As, in practice, the inertia of the oscillating mirrors means their orientations cannot change instantaneously, the phase of the mirror signal is set so that the mirrors can settle in their required orientations prior to the laser being pulsed. In other words, the position of the laser pulse has to fall somewhere between the start of the current mirror position and the end of the current frame, i.e. $t_{\text{mirr}} \neq 0$ and $t_{\text{off}} \neq 0$ as shown in figures A1 and A2.

Typical signals of the traverse speed, traverse direction and sync pulse over a complete scan are shown in figure A3. The signals are periodic and allow the traverse carriage holding the optics that produce the light sheet to continuously sweep back and forth through the measurement volume. The ‘direction’ signal controls the direction of the traverse, and the ‘sync pulse’ signal is pulsed at the start of every forward scan to initiate the capture of a series of images on all cameras. Note that the scan sync pulse only initiates saving on all cameras, to ensure all recorded sequences start at the beginning of the forward scan, but is periodic allowing, for example, every n th scan to be saved.

Appendix B. Calibration

B.1. Stereo PIV calculation

To gain quantitative measurements, a calibration is required for the stereo PIV measurements. For the scanning system, accurate calibration mappings are needed at all values of z , the scanning direction, within the volume of interest. Moreover, for stereo PIV measurements, gradient information in the z direction is necessary to recover the third component of velocity w . In non-scanning systems this gradient information is obtained by either traversing a calibration

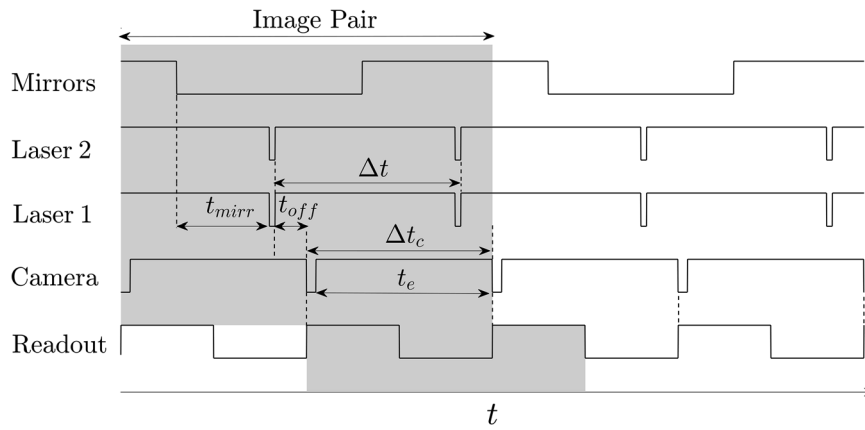


Figure A1. Schematic of the timing sequence sent to the various components of the system when in single-pulse mode. The laser pulse is triggered from the rising edge (+ve) of the pulse train to the lasers. An image pair (one plane of velocity) is indicated by the shaded region.

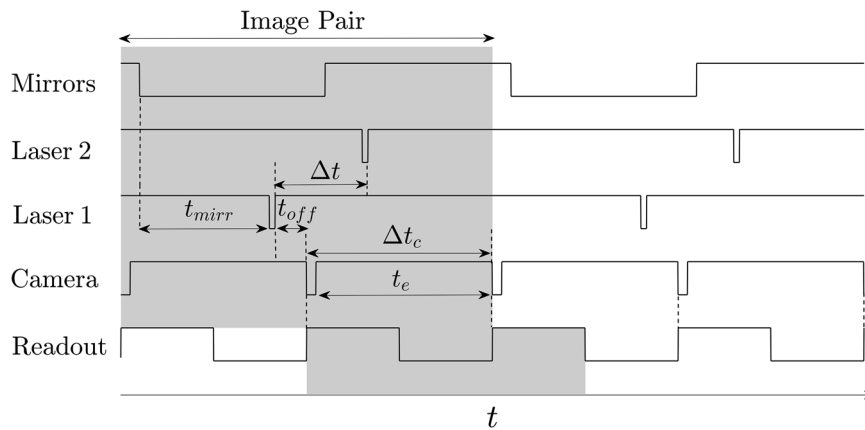


Figure A2. Schematic of the timing sequence sent to the various components of the system as in figure A1 but now in double-pulse mode.

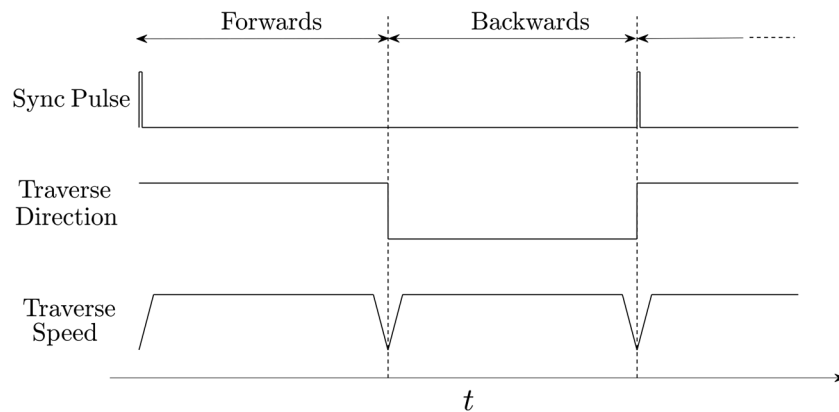


Figure A3. Schematic showing the timing signal over the course of a complete scan. Acceleration/deceleration phases are shown for the traverse, along with the direction signal that controls the direction of travel for the traverse. A sync pulse is shown that is used to trigger the recording of a sequence of images for all cameras at the start a scan.

target (a plane of dots or patterns with known spacing) to several z positions, or by using a multi-plane target where the target has two (or more) planes of dots or patterns at known z locations (Prasad 2000). We opt to use a dual-plane target herein, and so, by imaging the calibration target within the experimental setup, 2D mapping functions $\mathbf{G}_i^{k\pm}$ can be calculated. The mappings $\mathbf{G}_i^{k\pm}$ map from the pixel coordinates of each camera

$$\mathbf{X}_i = \begin{pmatrix} X_i \\ Y_i \end{pmatrix} \quad \text{with } \mathbf{X}_i \in \mathbb{P}^2, \quad (\text{B.1})$$

to world coordinates

$$\mathbf{x} = \begin{pmatrix} x \\ y \\ z_k \pm \frac{\Delta z}{2} \end{pmatrix} \quad \text{with } \mathbf{x} \in \mathbb{W}^3. \quad (\text{B.2})$$

Here, $\mathbb{P}^2 \subset \mathbb{R}^2$ is the pixel coordinate space, on the two z planes of the target $z_k \pm \frac{\Delta z}{2}$, where z_k is the mid-plane of the target and Δz is the spacing between the two planes, and where $\mathbb{W}^3 \subset \mathbb{R}^3$ is the world coordinate space. Specifically,

$$\mathbf{x} = \mathbf{G}_i^{k\pm}(\mathbf{X}_i), \quad (\text{B.3})$$

where i is the camera identifier and $k\pm$ distinguishes between the two planes of the calibration target positioned at $z = z_k \pm \frac{\Delta z}{2}$. Thus there are two mappings, \mathbf{G}_i^{k+} and \mathbf{G}_i^{k-} , for each camera. In general, the mapping functions $\mathbf{G}_i^{k\pm}$ are not known *a priori* and are commonly determined by least squares fitting of a polynomial function between the known world coordinates \mathbf{x} of the features on the calibration target and the corresponding pixel coordinates of each camera \mathbf{X}_i .

For a non-scanning system, with a light sheet centred at z_k , the geometric mapping \mathbf{G}_i^k at z_k can be inferred by linear interpolation between the two mappings at $z_k \pm \frac{\Delta z}{2}$:

$$\mathbf{G}_i^k(\mathbf{X}_i) = \frac{1}{2} (\mathbf{G}_i^{k+}(\mathbf{X}_i) + \mathbf{G}_i^{k-}(\mathbf{X}_i)). \quad (\text{B.4})$$

To calculate the world displacements $\Delta \mathbf{x}$, we need a 3D mapping. We opt to use an approach similar to Soloff *et al* (1997) and start by imaging the dual-plane calibration target with both PIV cameras. Then, by identifying a common point on the calibration target, the four pixel coordinates (two from each camera) are used to determine the 3D mapping

$$\mathbf{x} = \mathbf{F}^k(\mathbf{X}_A, \mathbf{X}_B), \quad (\text{B.5})$$

where the subscripts A and B distinguish between the two PIV cameras, and k defines the mid-plane location of the calibration target z_k where the mapping is calculated. Given the two planes (separated by Δz) of the dual-plane target, a polynomial function can be fitted, for a given target position z_k , with high order dependence in x and y (dependent on the number of features of the target), and linear dependence in z . Given the thin light sheet, higher order z dependence is generally not required, but could be achieved by including more planes on the multi-plane target or by accurately traversing a planar target in z . The world velocities $\mathbf{u} = (u, v, w)$ on a plane positioned at $z = z_k$ can then be found from

$$\mathbf{u} \simeq \mathbf{J}_A^k \mathbf{U}_A + \mathbf{J}_B^k \mathbf{U}_B, \quad (\text{B.6})$$

where $\mathbf{U}_i(\mathbf{X}_i, z_k) = \frac{\Delta \mathbf{X}_i}{\Delta t}$ are the 2D (pixel) velocities for camera i and $\mathbf{J}_i^k = \frac{\partial \mathbf{F}^k}{\partial \mathbf{X}_i}$ is the (3×2) Jacobian matrix (with units of world/pixel) associated with the mappings at the z_k location of camera i .

For the scanning system, the mappings and Jacobian matrices are required for all z in the scan. To this end, the dual-plane target is positioned at a number of z_k locations within the volume of interest. This enables the mappings \mathbf{G}_i^k and Jacobians \mathbf{J}_i^k (plus inverses of both) to be calculated at several z locations spanning the volume. A least squares fit in z of the polynomial coefficients of the mapping functions is then used to generate z -dependent 2D mappings $\mathbf{G}_i(\mathbf{X}_i, z)$ and z -dependent Jacobians $\mathbf{J}_i(\mathbf{X}_i, z)$ over the whole volume to be scanned.

The first step in calculating stereo velocities is calculating the 2D velocities in pixel coordinate space \mathbb{P}^2 for each of the PIV cameras at a given z location $\mathbf{U}_i(\mathbf{X}_i, z)$. The velocities $\mathbf{U}_i(\mathbf{X}_i, z)$ in pixel coordinate space \mathbb{P}^2 at a known z location are mapped to a common grid in world coordinate space \mathbb{W}^3 . The z position of the velocity fields is used so that the pixel velocities of the two cameras in pixel coordinate space \mathbb{P}^2 can be mapped to the corresponding world coordinate space \mathbb{W}^3 using the z dependent 2D mappings generated during the calibration:

$$\mathbf{U}_i(\mathbf{X}_i, z) \xrightarrow{\mathbf{G}_i(\mathbf{X}_i, z)} \hat{\mathbf{U}}_i(\mathbf{x}). \quad (\text{B.7})$$

The Jacobian \mathbf{J}_i , initially in \mathbb{P}^2 space separately for each camera, is also mapped to world coordinate space \mathbb{W}^3 in the same manner:

$$\mathbf{J}_i(\mathbf{X}_i, z) \xrightarrow{\mathbf{G}_i(\mathbf{X}_i, z)} \hat{\mathbf{J}}_i(\mathbf{x}). \quad (\text{B.8})$$

Note that the values of $\hat{\mathbf{J}}_i$ still represent the world units per pixel. From $\hat{\mathbf{U}}_A$ and $\hat{\mathbf{U}}_B$, the two sets of pixel velocities now in \mathbb{W}^3 , we can calculate the velocity field in world units and in \mathbb{W}^3 using

$$\hat{\mathbf{u}} \simeq \hat{\mathbf{J}}_A \hat{\mathbf{U}}_A + \hat{\mathbf{J}}_B \hat{\mathbf{U}}_B. \quad (\text{B.9})$$

We do this for each pair of 2D2C planes produced by the 2D PIV algorithm for cameras A and B to construct a sequence of 2D3C planes at different z locations in the scan. These 2D3C planes are finally combined to construct volumetric 3D3C measurements.

B.2. Error

For each z location of the scan, an estimate of the error in the stereo reconstruction can be determined by back-projecting the world velocities $\hat{\mathbf{u}}$ onto the velocities (in pixels/frame) of the two cameras. For convenience, we do this calculation on the common world grid to obtain $\hat{\mathbf{U}}_i^*$ as follows:

$$\begin{pmatrix} \hat{\mathbf{U}}_A^* \\ \hat{\mathbf{U}}_B^* \end{pmatrix} = \hat{\mathbf{J}}^{-1} \hat{\mathbf{u}} \quad (\text{B.10})$$

where $\hat{\mathbf{J}}^{-1}(\mathbf{x})$ is the inverse Jacobian in world coordinate space \mathbb{W}^3 found from

$$\mathbf{J}^{-1}(\mathbf{X}_i, z) \xrightarrow{\mathbf{G}_i(\mathbf{X}_i, z)} \hat{\mathbf{J}}^{-1}(\mathbf{x}), \quad (\text{B.11})$$

where $\mathbf{J}^{-1} = \frac{\partial \mathbf{X}_i}{\partial \mathbf{F}_i^{-1}}$ is the (4×3) Jacobian matrix (with units of pixel/world) of the inverse 3D mapping \mathbf{F}_i^{-1} . A measure of the error can be obtained by comparing the back-projected velocities with the velocities calculated by the PIV algorithm mapped to \mathbb{W}^3 . Assuming equal weighting among each of the velocity components, we construct the error field from the magnitude of the vector difference:

$$E(\mathbf{x}) = \frac{1}{4} \left(\|\hat{\mathbf{U}}_A^* - \hat{\mathbf{U}}_A\| + \|\hat{\mathbf{U}}_B^* - \hat{\mathbf{U}}_B\| \right). \quad (\text{B.12})$$

This error estimate can be used as an additional quality check of the velocity fields by removing vectors where the

reconstruction error in $E(\mathbf{x})$ exceeds some threshold value (typically chosen to be 0.5 pixels/frame as in Wieneke (2005)). An estimate of the error introduced by the mappings can also be found using the forward mapping on the back-projected fields $\hat{\mathbf{U}}_i^*$ to obtain $\hat{\mathbf{u}}_i^*$ and then the inverse mappings on this field to obtain a second set of back-projected mappings. These back-projected mappings can then be compared to $\hat{\mathbf{U}}_i^*$ to highlight possible errors associated with the mappings.

B.3. Calibration refinement

An important step in calculating the velocities using stereo PIV is the calibration refinement technique. The stereo reconstruction of velocity, as outlined in appendix B.1, relies on the calibration between each camera and the world coordinates being accurate. This ensures that a velocity perceived from one camera is correctly reconstructed with a velocity from the other camera, i.e. the two velocities are coincident in world coordinates. Although, in practice, the calibration procedure is performed carefully, it is difficult to perfectly align the calibration target with the position of the light sheet and even small discrepancies can yield large errors in the stereo reconstruction (Wieneke 2005). With the scanning system, this is complicated further as any small misalignment between the laser beam and the optics on the traverse carriage can lead to a systematic positional error in the light sheet. To correct for this, we opt to use an iterative two-step calibration refinement technique using a methodology similar to that of Willert (1997) for a fixed light sheet.

Effectively, an algorithm similar to PIV is run on synchronised images from the different cameras all projected on to world coordinates. In the first step, the mean disparity between the images is used to refine the light sheet position, i.e. the position of each pair of light sheets is updated until the mean disparity between the images is minimised. The second step performs finer adjustment across the whole image, to effectively remove any residual skewing of the light sheet relative to the coordinate system. This ensures that the 2D velocity information from each of the PIV cameras overlap. In general, this two-step approach should be iterated but typically we find that any residual error in position is much smaller than δz after one pass.

B.4. PLIF calculation

To determine the density field, the signal from the fluorescent dye added to the flow has to be calibrated. Assuming negligible attenuation and a linear camera response (see Crimaldi (2008) for more details), the calibration of the experimental images is simply

$$C(\mathbf{X}, z) = \frac{I(\mathbf{X}, z) - I_0(\mathbf{X}, z)}{I_1(\mathbf{X}, z) - I_0(\mathbf{X}, z)} C_1, \quad (\text{B.13})$$

where $C(\mathbf{X}, z)$ is the calculated concentration field, $I_1(\mathbf{X}, z)$ is a reference image containing a known homogeneous

mixture of dye that can be used to remove any spatial variation present in the light sheet and establish the concentration of dye in the images, and C_1 is the dye concentration for which $I(\mathbf{X}, z) = I_1(\mathbf{X}, z)$. These two reference images are determined as follows. First, a scan-position-dependent background reference image $I_0(\mathbf{X}, z)$ is calculated at each z position of the scan by recording $O(10)$ scans of the volume containing no fluorescent dye and taking the mean at each z_i location. To minimise errors, no fitting is used for the background images and a distinct image is found for every z location in the forward and backward scan sequences. In a similar manner, a second reference image $I_1(\mathbf{X}, z)$ is calculated by averaging over $O(10)$ scans of the volume containing a homogeneous mixture of the highest concentration of dye in the experiment C_1 , again producing $2N$ images.

Finally, the concentration measurements found from equation (B.13) are mapped to world coordinate space \mathbb{W}^3 . As discussed in appendix B.1, this is achieved in the scanning system using a least squares mapping $\mathbf{G}_C(\mathbf{X}_C, z)$ calculated from the images recorded of the calibration target for the PLIF camera (Cam C). The concentration \hat{C} in world coordinate space \mathbb{W}^3 can then be found from

$$C(\mathbf{X}, z) \xrightarrow{\mathbf{G}_C(\mathbf{X}_C, z)} \hat{C}(\mathbf{x}). \quad (\text{B.14})$$

As for the PIV cameras (Cam A and B), the coordinate system of Cam C can be refined using the strategy discussed in appendix B.3. In this case, only the second step of refinement is needed, and the coordinate system of Cam C can be refined using simultaneously acquired particle images from Cam C, with its longpass filter removed, and one (or both) of the PIV cameras. To ensure all of the coordinates systems lie in the same plane, the disparity map is calculated using the refined coordinate system of Cam A or B (or both) with the full disparity map (or the average disparity map if the refinement is conducted with both PIV cameras) applied to Cam C images only.

To relate the concentration to the density field ρ we have

$$\rho = \frac{\Delta\rho}{C_1} \hat{C} + \rho_{\min}, \quad (\text{B.15})$$

where $\Delta\rho = \rho_{\max} - \rho_{\min}$ and ρ_{\max} , ρ_{\min} are the densities of the fluid with the maximum concentration of dye and the fluid containing no dye, respectively.

It is worth noting that, in the scanning system, we have twice as many ρ fields as velocity fields because the PLIF calculation only requires a single image (compared to the two needed to calculate a velocity field using PIV). That being said, there is a choice in how to make use of the extra information. We could use either a single ρ field, corresponding to the first or second raw particle image used in the PIV calculation, or the mean of the two ρ fields could be taken to help remove random noise from the data. Alternatively, the velocity information calculated for the corresponding pair of ρ fields could be used to advect the fields forward and backward in time by $\pm \frac{1}{2} \mathbf{u} \Delta t$ before taking the mean of the two fields.

ORCID iDs

J L Partridge  <https://orcid.org/0000-0003-2940-7341>
 Adrien Lefauve  <https://orcid.org/0000-0003-3692-2886>
 Stuart B Dalziel  <https://orcid.org/0000-0002-8487-2038>

References

- Brücker C 1997 3D scanning piv applied to an air flow in a motored engine using digital high-speed video *Meas. Sci. Technol.* **8** 1480
- Crimaldi J P 2008 Planar laser induced fluorescence in aqueous flows *Exp. Fluids* **44** 851–63
- Dahm W J A, Southerland K B and Buch K A 1991 Direct, high resolution, four-dimensional measurements of the fine scale structure of Sc 1 molecular mixing in turbulent flows *Phys. Fluids A* **3** 1115–27
- Dalziel S B, Carr M, Sveen J K and Davies P A 2007 Simultaneous synthetic schlieren and PIV measurements for internal solitary waves *Meas. Sci. Technol.* **18** 533
- Dalziel S B, Hughes G O and Sutherland B R 2000 Whole-field density measurements by synthetic schlieren *Exp. Fluids* **28** 322–35
- Dalziel S B, Linden P F and Youngs D L 1999 Self-similarity and internal structure of turbulence induced by Rayleigh–Taylor instability *J. Fluid Mech.* **399** 1–48
- Discetti S and Coletti F 2018 Volumetric velocimetry for fluid flows *Meas. Sci. Technol.* **29** 042001
- Holt S E, Koseff J R and Ferziger J H 1992 A numerical study of the evolution and structure of homogeneous stably stratified sheared turbulence *J. Fluid Mech.* **237** 499–539
- Jiménez J 2018 Coherent structures in wall-bounded turbulence *J. Fluid Mech.* **842** P1
- Krug D, Holzner M, Lüthi B, Wolf M, Tsinober A and Kinzelbach W 2014 A combined scanning PTV/LIF technique to simultaneously measure the full velocity gradient tensor and the 3D density field *Meas. Sci. Technol.* **25** P1
- Lawson J M and Dawson J R 2014 A scanning piv method for fine-scale turbulence measurements *Exp. Fluids* **55** 1857
- Lefauve A 2018 Waves and turbulence in sustained stratified shear flows *PhD Thesis* University of Cambridge
- Lefauve A, Partridge J L, Zhou Q, Dalziel S B, Caulfield C P and Linden P F 2018 The structure and origin of confined Holmboe waves *J. Fluid Mech.* **848** 508544
- Lucas D, Caulfield C P and Kerswell R R 2017 Layer formation in horizontally forced stratified turbulence: connecting exact coherent structures to linear instabilities *J. Fluid Mech.* **832** 409–37
- McDougall T J 1979 On the elimination of refractive-index variations in turbulent density-stratified liquid flows *J. Fluid Mech.* **93** 83–96
- Meyer C R and Linden P F 2014 Stratified shear flow: experiments in an inclined duct *J. Fluid Mech.* **753** 242–53
- Münch B, Trtik P, Marone F and Stampanoni M 2009 Stripe and ring artifact removal with combined wavelet-fourier filtering *Opt. Express* **17** 8567–91
- Olsthoorn J and Dalziel S B 2017 Three-dimensional visualization of the interaction of a vortex ring with a stratified interface *J. Fluid Mech.* **820** 549–79
- Prasad A K 2000 Stereoscopic particle image velocimetry *Exp. Fluids* **29** 103–16
- Prasad A K and Jensen K 1995 Scheimpflug stereocamera for particle image velocimetry in liquid flows *Appl. Opt.* **34** 7092–9
- Schanz D, Gesemann S and Schröder A 2016 Shake-the-box: Lagrangian particle tracking at high particle image densities *Exp. Fluids* **57** 70
- Soloff S M, Adrian R J and Liu Z 1997 Distortion compensation for generalized stereoscopic particle image velocimetry *Meas. Sci. Technol.* **8** 1441
- Tian X and Roberts P J W 2003 A 3D LIF system for turbulent buoyant jet flows *Exp. Fluids* **35** 636–47
- Verso L, van Reeuwijk M, Gurka R, Diamessis P J, Taylor Z J and Liberzon A 2017 Experimental study of the initial growth of a localized turbulent patch in a stably stratified fluid *Int. J. Heat Fluid Flow* **66** 127–36
- Wang C, Gao Q, Wei R, Li T and Wang J 2017 Weighted divergence correction scheme and its fast implementation *Exp. Fluids* **58** 44
- Webster D R, Roberts P J W and Ra'ad L 2001 Simultaneous DPTV/PLIF measurements of a turbulent jet *Exp. Fluids* **30** 65–72
- Wieneke B 2005 Stereo-PIV using self-calibration on particle images *Exp. Fluids* **39** 267–80
- Willert C 1997 Stereoscopic digital particle image velocimetry for application in wind tunnel flows *Meas. Sci. Technol.* **8** 1465
- Zhou Q, Taylor J R, Caulfield C P and Linden P F 2017 Diapycnal mixing in layered stratified plane couette flow quantified in a tracer-based coordinate *J. Fluid Mech.* **823** 198–229