Applied Scientific Research **49**: 217–244, 1992.

Decay of rotating turbulence: some particle tracking experiments

STUART B. DALZIEL

Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Silver Street, Cambridge, CB3 9EW, U.K.

Key words: rotating turbulence, particle tracking, experiments, inertial waves

Abstract. Recent development of measurement techniques based on particle image velocimetry (PIV) are enabling more detailed measurements to be made over extended regions of a flow than have been previously possible. These techniques are of particular value for turbulent flows where the structures present within such flows are incompletely understood and are not readily accessible to traditional measurement techniques. Unfortunately the considerable processing time and specialised equipment required with most PIV techniques limits their applicability when ensemble statistics are required for an evolving turbulent flow. This paper reports on the development and application of an efficient, fully automated particle tracking system. The system was developed as part of a study of the decay of turbulence in a rotating environment. Ensemble descriptions of the temporally evolving flow were required over an extended measurement domain. For each set of parameters particles were tracked with a sampling frequency of 12.5Hz over 60 seconds for 25 realisations. Typically 350 particles were identified and tracked at each time step. Processing speeds in the region ten to fifteen sample images per minute were achieved using a PC/AT compatible computer. The results of the experiments were found to be in broad agreement with previous investigations. However it was found that the method of generating the initial turbulent flow had a profound affect on the subsequent evolution due to the forcing of a strong, large scale systematic flow.

1. Introduction

Turbulence is widely considered one of the most difficult problems in fluid dynamics. Except in special cases, theoretical advancements are hampered by an inability to find exact solutions to the Navier-Stokes equations; approximations to these equations using closure models has met with only limited success. The range of scales involved in a turbulent flow means that only relatively low Reynolds number flows are accessible numerically to the present generation of computers. As a result experimental work continues to play a major role in developments, the scope being limited by our ability to measure and understand what is happening, rather than the use of approximate fluid dynamics. The quantity of primary interest in such experiments is the fluid velocity as a function of both space and time.

Over recent years the velocity measurement techniques available have advanced significantly. Starting from flow visualisation and hand-eye digitization, hot wire/film anenometry and then laser doppler anenometry were able to yield accurate point measurements of the flow. Fuelled by modern image processing technology there is renewed interest in flow visualisation as it now offers accurate quantitative information about a two or three dimensional region of the flow rather than just a single point. A wide range of image processing techniques have been developed to measure

the fluid velocity. The key feature in common with most these techniques is the seeding of the flow with particles: it is the motion of these particles which is actually measured. A thorough review of particle image velocimetry (PIV) techniques may be found in Adrian [1] and elsewhere in this volume.

In this paper we describe a simple, efficient method for measuring velocity by tracking a large number of individual particles. This method was developed as part of a study of the temporal decay of rotating turbulence, but has subsequently been used to measure a variety of other flows (turbulence near a stress-free surface, turbulent plumes and jets, two dimensional turbulence in a stratified environment, Rayleigh-Taylor instability, baroclinic instability and baroclinic vortex interaction). The particle tracking forms a part of a wide ranging image processing package, DigImage, developed for analysing fluid flows. In this paper we illustrate the particle tracking (PT) as it was applied to the rotating turbulence problem.

Considering its importance to a range of fields (e.g. meteorology, oceanography and turbo machinery), turbulence in rotating systems has been the subject of relatively few experimental studies. Part of this is due to experimental difficulties in generating and measuring such turbulence, and part due to the as yet incomplete understanding of the classical nonrotating problem. As a consequence there are very few data sets available for comparison with theoretical and numerical results. The results which are in the literature fall into two broad categories: inhomogeneous turbulence (typically produced by an oscillating grid – we shall comment on these briefly in the final section), and decaying homogeneous turbulence (typically produced by a single pass of a grid or flow past a grid). The work reported in this paper falls into the second category: the temporal decay of grid generated turbulence with a zero mean flow.

The majority of early work on classical turbulence has been undertaken in wind tunnels looking at the evolution of turbulence generated by a uniform flow past a grid. It is generally accepted that ensemble statistics may be replaced by their temporal equivalents in this set up. The space-time transformation and Taylor's frozen field hypothesis suggest point measurements are equivalent to a one dimensional transect in the direction of the mean flow. Such grid generated turbulence may be considered homogeneous and isotropic to a reasonable approximation. Unfortunately it is not practical to rotate a sufficiently large wind tunnel at adequate rotation rates to study the effect of rotation, particularly if we are interested in the later stages of decay. Instead we must consider an ensemble of experiments in which the turbulence exhibits a temporal evolution. The requirement for ensemble statistics means that a large number of experiments must be processed, and so the processing time was a fundamental design consideration during the development of the particle tracking.

Ibbetson and Tritton [11] were the first to make detailed measurements of the temporal decay of rotating turbulence. The turbulent flow was generated in air by pulling two grids apart vertically within a shallow rotating annulus, and measurements were made by sweeping a hot wire probe around the annulus. Their primary concern was to determine whether rotation produced an increase or decrease in the

decay rate: they found it to increase, but concluded this was due to inertial wave dissipation in the viscous bondary layers on the tank top and bottom.

The most recent work is by Jacquin et al. [12] who considered a tube rotating rapidly about its axis. Air was blown down this tube at up to $20\,\mathrm{ms}^{-1}$. The air passed through a suitable diffuser to bring it into solid body rotation and a grid to generate the turbulence. The decay downstream of the grid was then observed using hot wire probes. The main conclusion that rotation slows the decay is consistent with present thinking. As with Ibbetson and Tritton, Jacquin et al. found a strong anisotropy develop in the length scales parallel with and perpendicular to the rotation axis. In contrast very little change in anisotropy in the corresponding velocity fluctuations was observed.

In section 2 we detail the experimental set up, and in section 3 outline the particle tracking technique employed. We present classical and rotating results in sections 4 and 5 (respectively), before concluding in section 6.

2. Experimental apparatus

Figure 1 sketches the tank in which all the experiments were performed. The tank measures 1.2 m long (x direction), 0.25 m wide (y direction) and 0.6 m deep (z direction). The tank is positioned on a precision rotating table capable of rotation rates of up to 2π radians per second about the vertical z axis. The basic experiment

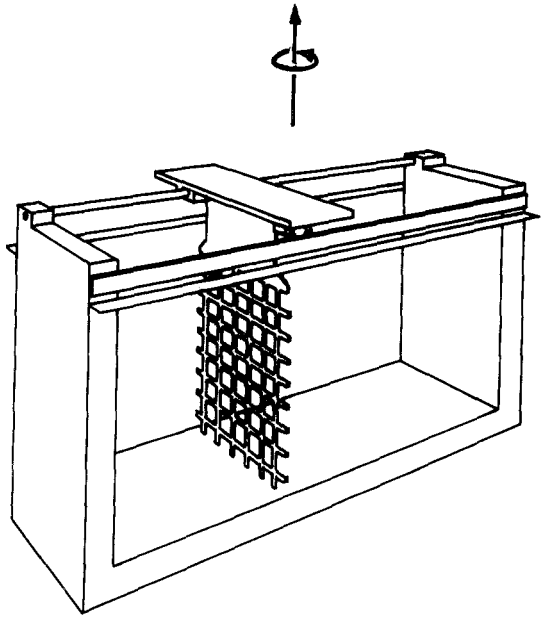

Fig. 1. Sketch of experimental setup showing tank, traversing mechanism and grid.

consists of a grid being towed down the length of the tank filled with water. The grid is machined from a single sheet of aluminium to give a square mesh consisting of 6.35 mm square bars at $M = 31.75$ mm centres. The grid is towed along the length of the tank at up to $U = 0.24\,\mathrm{ms}^{-1}$ by means of a traversing carriage located on the top of the tank. The grid Reynolds number ($\mathrm{Re}_M = UM/\nu$, where ν is the kinematic viscosity of water) is constant at $\mathrm{Re}_M \approx 7500$ over most of the length of the tank. The carriage and other associated equipment are controlled by means of a small microcomputer mounted on the rotating table, which in turn communicates with a terminal in the laboratory via a set of slip rings.

The design and construction of the tank and traversing mechanism included the ability to traverse a conventional hot film probe for velocity measurement. However the vibration in the traversing mechanism proved too great for reliable turbulence measurements. The development of the particle tracking system described in this paper grew out of the need to measure the flow field and the inability to utilise hot film anenometry (the project did not have access to an LDA system at the time).

The optical arrangement used in the particle tracking experiments is sketched in figure 2. The need to keep the effective diameter of the table relatively small (less than 1 m) forced all imaging to be indirected through a front silvered mirror. Reference points for image registration during the analysis were provided by fluorescent strip lights, covered in aluminium foil and painted matt black. The individual points were

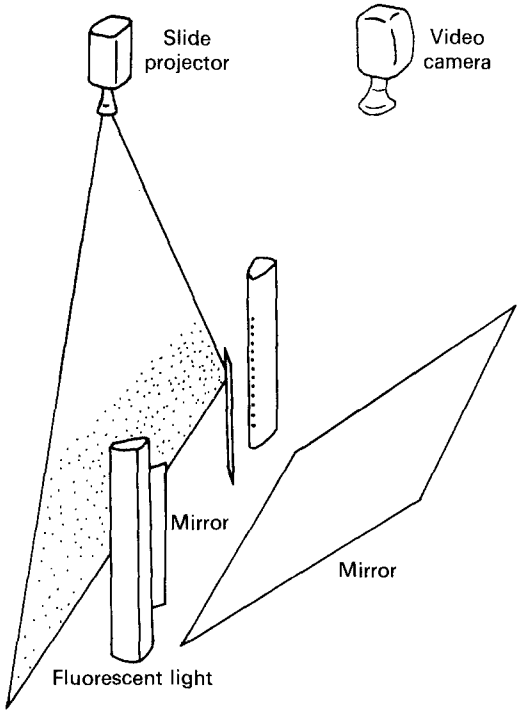

Fig. 2. Sketch of optical arrangement for experiments. The three mirrors are all front-silvered.

produced by piercing the foil against the glass tube with a sharp needle. This technique provides diffuse reference points (so the angle at which they are viewed is not critical) with an easily adjusted size. The points were superimposed as a row down each side of the image through the use of front silvered mirrors. The fluorescent lights were positioned so that the reference points lay in the in-focus plane of the camera.

The flow was illuminated using a sheet of light approximately 7 mm thick (δy) provided by a 250 W quartz halogen slide projector mounted above the tank. The light rays were approximately parallel to the rotation (z) axis with the sheet cutting along the length of the tank (x) at the centre line ($y = y_0$). The flow was imaged using a medium cost single CCD Super VHS camera (625 lines, 50 Hz field rate), fitted with a zoom lens (f1.6, focal length 10.5 mm to 126 mm), feeding through slip rings to a Super VHS video tape recorder. While the camera produced a Y/C colour signal, only the luminance (Y) component was passed through the slip rings to be recorded. This arrangement provided approximately 400 lines horizontal (intensity) resolution. Electronic shutter speeds of 1/50 s and 1/120 s were employed. At the higher shutter speed it was necessary to utilise the camera's automatic gain circuitry. The minimum useful illumination is quoted as 7 lux at f1.4 in the camera documentation: the camera was typically operating close to this limit. The signal noise due to the recording medium was generally less than that due to the camera at the relatively low light levels. The signal to noise ratio on playback remained adequate, allowing successful particle tracking even when the automatic gain was necessary. The equipment utilised to track the particles after the experiment was completed will be described in conjunction with the tracking process in the next section.

The particles used were produced by crushing and seiving pliolite (a resin used in the manufacture of white paint) to obtain sizes in the range 150 μm to 250 μm. The very small particles and dust remaining after the seiving were removed by a decanting process after suspension in fresh water. Photographic wetting agent was used to prevent the particles adhering to the free surface, and around 3 wt% salt used to adjust the water density so that the particles were approximately neutrally buoyant. While the particles were of irregular geometry, they were sufficiently small and close to spherical for their precise geometry to be of little importance.

3. Particle tracking

In this section we describe the particle tracking approach to PIV used in this study. The underlying philosophy is one of simplicity and computational efficiency. Efficiency is of particular concern with the need for ensemble statistics in this evolving turbulent flow.

The particle tracking process may be divided into a number of distinct tasks: experimentation, image acquisition and enhancement, particle location, particle pairing and subsequent analysis. Generally, due to the volume of data involved, it is desirable to combine some of these tasks into a single phase. Here we shall consider

the tracking phase which consists of the acquisition, location and pairing tasks. The end product is six bytes of data per particle per time step, describing the particle positions, characteristics and relationships.

3.1. Image acquisition

Images are acquired from Super VHS video tape recorder (VTR) with a resolution of 512×512 pixels at 256 grey levels using a Data Translation DT-2862 arithmetic frame grabber. Due to the nature of video tape technology, it is not possible to acquire suitable high quality images from a "paused" VTR. A combination of the VTR's internal frame strobe, markers on the audio channel and interrupt driven routines on the host computer allowed efficient, reliable control of tape transport and acquisition of the exact frame required while in "play" mode.

During the tracking process, images are acquired as a sequence of four frames (not necessarily adjacent) from the playing VTR. The tape is then repositioned as a background task while the subsequent processing is undertaken. If desired, a *background* image is subtracted from the incoming signal to enable the subsequent application of uniform thresholds. If the particles are moving more than a small fraction of their diameter between each video field (1/50 s – each frame consists of two interlaced fields), it is necessary to remove half the information contained on the interlaced frame to ensure that a given particle appears at only one location on a given sample image. This also provides the possibility of sampling at the 50 Hz field rate, provided the small vertical displacement between the two fields is not important (the video camera may be modified so that it always produces only the odd or the even lines).

The fluorescent reference points are located and used to generate a mapping from the current pixel coordinate system to a reference pixel coordinate system using a least squares routine. In the majority of cases only a simple translation is required. The root mean square (rms) error in this mapping is calculated from the residuals and used as a check on the quality of the image. If the rms error exceeds a preset limit, then a further attempt will be made to acquire the image in the hope of improving its quality. The reference point mapping is used in conjunction with a pixel to world transformation to obtain the true world coordinates from the current pixel system.

3.2. Particle location

To locate a particle the image is first searched for a point satisfying a spatially uniform threshold condition (the optional subtraction of the background image during acquisition means that this threshold may be spatially varying relative to the raw image). The region of the image satisfying the threshold and connected to this first pixel is fathomed, and statistics are collected describing this region or *blob*. The shape, size and intensity characteristics of the blob are then checked against the allowable limits to determine if the blob represents a valid particle. The location of the particle is

determined from the volume (mass) centroid of the region satisfying the threshold and the position transformed to the world coordinate system. This strategy produces good estimates of the relative position of a given particle in a sequence of images, provided the particles are sufficiently close to spherical that any rotation of the particle is of no concern. As an estimate of their absolute position, the volume centroid is less accurate, but the small errors in this are much less important. More complex and computationally expensive models such as those assuming a Gaussian intensity profile (e.g. Perkins and Hunt [14]) are not appropriate for the current method of illumination and are not expected to yield significantly better results.

Particles less than one pixel in size may only be located with pixel accuracy. However, for particles occupying more than one pixel in each direction the volume centroid gives subpixel accuracy, the error being related to the area A occupied by the particle through

$$\sigma_x \leqslant a(\sigma_i \bar{i}) A^{-1/2}, \tag{1}$$

where σ_x is some measure of the error in estimating the position of the particle, $\sigma_i \bar{i}$ is a measure of the relative intensity variations within a pixel (which may include both the unresolved portion of the intensity signal and the noise introduced by camera, recording medium and the digitizing process), and a is some constant. Experience shows that particles extending at least three pixels in each direction may typically be located to an accuracy of better than 0.2 pixels (typically 0.1 pixels) relative to the same particle at some other position.

In practice we search for particles in an image twice. The second pass utilises less stringent threshold requirements than the first and is used to pick up blobs which we are less certain about. As will be explained in section 3.4, this second set of particles is handled somewhat more cautiously.

3.3. Particle pairing

A wide variety of methods have been used to pair the particle images between two successive image samples (see Adrian [1]). Here we develop an efficient pairing algorithm based on the *transportation algorithm* [9], a graph theory technique frequently used in operations research for determining optimal associations between two sets. With careful coding the computational time for this method increases only slightly more rapidly than the number of particles (in many cases the matching process is faster than the particle location).

The basic idea is to minimise some objective function ζ which is linear in the associations it includes between the two sets. Here one set is the particle images $P = \{p_{[i]} \text{ for } i = 1, M\}$ at the *old* time, $t_{[n]}$ (say), and the other the particle images $Q = \{q_{[j]} \text{ for } j = 1, N\}$ at the *new* time, $t_{[n+1]}$. We define a set of variables α_{ij} recording the set of particle pairings: if $\alpha_{ij} = 1$, then particle image $p_{[i]}$ at $t_{[n]}$ is considered to be the same particle as the image $q_{[j]}$ at $t_{[n+1]}$. If $p_{[i]}$ and $q_{[j]}$ are considered as different

physical particles, then $\alpha_{ij} = 0$. We express the objective function ζ as the linear combination

$$\zeta = \sum_i \sum_j \alpha_{ij} c_{ij}, \tag{2}$$

where c_{ij} is the *cost* for the association α_{ij} to be included in the pairings. The subscript i is summed from 0 to M, where M is the number of particle images at $t_{[n]}$. The value $i = 0$ is reserved for all particles outside the set of particles P found at $t_{[n]}$ (i.e. all particles outside the viewing region plus those obscured within the viewing region). Similarly, subscript j is summed from 0 to N, where N is the number of particle images at $t_{[n+1]}$ and $j = 0$ represents all particles outside the set of particles Q found at $t_{[n+1]}$. We shall discuss how to assign the costs c_{ij} in the next subsection. Before that we shall consider in more depth the structure of the problem and how to solve it.

The constraints on α_{ij} arise from a physical particle only being able to be in one location at a given time. Thus, for any $i = K \neq 0$, there must be one and only one j in $[0, N]$ for which $\alpha_{Kj} = 1$. Similarly, for any $j = L \neq 0$, there must be one and only one i in $[0, M]$ for which $\alpha_{iL} = 1$. An association of the form $\alpha_{ij} = 1$ for $i, j \neq 0$ represents a pairing between images in the two sets of particles, while $\alpha_{0j} = 1$ indicates that $q_{[j]}$ was not in the set P. Similarly $\alpha_{i0} = 1$ indicates that $p_{[i]}$ is no longer in the set Q of physical particles we have located at $t_{[n+1]}$. Note that there may be many values of i and j giving $\alpha_{0j} = 1$ or $\alpha_{i0} = 1$ as there can be many particles entering or leaving our viewing region and hence associated with the special index 0. In general α_{00} will not be zero as some of the particles not in the set P will also not be in the set Q. However the value of α_{00} is of no concern and so we assign it a zero cost.

The transportation problem differs from our present particle pairing problem in that the former does not have the possibility of associations outside the sets P and Q (i.e. $i = 1, M$ and $j = 1, N$). Nevertheless the structure of our present problem is very similar to the transportation problem. Both may be written as a linear program of the form $\mathbf{B}\boldsymbol{\alpha} = \boldsymbol{\xi}$ with $\zeta = \min \mathbf{c}^T\boldsymbol{\alpha}$, where all the coefficients of the $\mathbf{B}$ matrix are either 0 or 1 and are arranged in such a way that $\mathbf{B}$ is totally unimodular. For our present problem the right hand side $\boldsymbol{\xi}$ is integer, so that the only solutions for $\boldsymbol{\alpha} = \alpha_{ij}$ are also integer. The differences lie in that the present problem includes a set of constraints for the associations with $i = 0$ or $j = 0$ (particles external to P or Q, respectively). While these constraints follow the simple block identity structure of the transportation problem they are inequality constraints and so require the introduction of artificial variables (these artificial variables do not have the same block identity structure). With $\boldsymbol{\xi} = \mathbf{1}$ all the variables, except α_{00}, are constrained to be either zero or one. The variable α_{00} may be any nonnegative integer, depending on the total number of particles in the flow. However, as we have already stated, the value of α_{00} is of no consequence.

Solution of the particle pairing problem can be tackled in the same manner as the graph theory approach to the transportation problem. We consider α_{ij} and c_{ij} as two

dimensional tables and define the relational index $I(j)$ as the value of i, for given $j \geqslant 1$, for which $\alpha_{ij} = 1$. Similarly we define $J(i)$ as the value of j, for given $i \geqslant 1$, for which $\alpha_{ij} = 1$. Note that $I(J(i)) = i$ and $J(I(j)) = j$, provided the corresponding association is between particles in the two sets P and Q. In the transportation problem, this identity means that only one of $I(j)$ and $J(i)$ are needed. However, for the particle pairing problem, $I(0)$ and $J(0)$ can not be defined uniquely, hence the need for both relational indices.

The solution process starts by assigning some arbitrary set of associations which satisfy the physical constraints on particles. Careful choice of this initial assignment of $I(j)$ and $J(i)$ will lead to the optimal solution being attained much more rapidly. Suppose some association α_{kl} is currently not in the set of associations. If we were to allow the α_{kl} pairing to be made, then for the constraints on the $p_{[k]}$ and $q_{[l]}$ particle images to remain satisfied, the associations $\alpha_{I(l)l}$ and $\alpha_{kJ(k)}$ must leave the set of pairings. Further, to maintain the constraints on the $p_{[I(l)]}$ and $q_{[J(k)]}$ particle images, we must introduce the association $\alpha_{I(l)J(k)}$. Now the change in the overall cost (*reduced cost*) associated with this change over is

$$rc_{kl} = (c_{kl} + c_{I(l)J(k)}) - (c_{I(l)l} + c_{kJ(k)}). \tag{3}$$

If rc_{kl} is negative, then entry of α_{kl} into the set of pairings is favourable.

The reduced cost is calculated for all $\alpha_{kl} = 0$ with $k \geqslant 1$ and $l \geqslant 1$, although if it has already been evaluated for $rc_{I(l)J(k)}$ it need not be evaluated again for rc_{kl}. Note that associations to the $i = 0$ and $j = 0$ sets may still be made through the $\alpha_{I(l)J(k)}$ variable in such a way that does not violate the inequality constraints on these variables. The k, l pair which yields the most negative value of rc_{kl} will be allowed to enter the set of pairings along with $\alpha_{I(l)J(k)}$. The $\alpha_{I(l)l}$ and $\alpha_{kJ(k)}$ associations must leave the set of pairings. These changes may be reflected by simultaneously updating the relational indices

$$I(l) := k,\ J(k) := l,\ I(J(k)) := J(I(l)),\ J(I(l)) := I(J(k)). \tag{4}$$

This process is repeated until no association has a negative value of rc_{kl}. The linear nature of ζ ensures that this solution has the minimum value for the objective function and corresponds to the optimal set of associations. (In principle this optimal set as defined by c_{ij} may not be unique, but in practice this situation rarely arises.)

3.4. Costing associations

The only algorithmic restriction on the assignment of costs c_{ij} is that they may not include information on the state of α_{ij} for the current time step. If information about α_{ij} were used in c_{ij} then the objective function would be nonlinear and a more sophisticated algorithm (e.g. simulated annealing) would be required to determine the optimal solution. Each cost may include information about the positions of $p_{[i]}$ and

$q_{[j]}$, the velocity history (normally only of $p_{[i]}$, but may include other particles at $t_{[n]}$), their geometric properties (e.g. shape), intensity, colour, fluid dynamics or a wide variety of other possible factors. However, experience has shown that a relatively simple costing strategy is all that is required in the majority of cases. In this subsection we outline the strategy we have employed for the current set of experiments. While we have utilised more sophisticated strategies in other flows (e.g. flows containing dye as a tracer), they are beyond the scope of this paper.

The basic method of costing an association α_{ij} is a modification to looking for the nearest neighbour. We define the basic cost of the association as

$$b_{ij} = |\mathbf{x}_{[i]} + \gamma \mathbf{u}_{[i]} \delta t - \mathbf{x}_{[j]}|^m, \tag{5}$$

where $\mathbf{x}_{[i]}$ and $\mathbf{x}_{[j]}$ are the locations of $p_{[i]}$ and $q_{[j]}$ (respectively) projected onto the $y = y_0$ plane, $\mathbf{u}_{[i]}$ is some estimate of the velocity (in the y_0 plane) of $p_{[i]}$ at $t_{[n]}$, γ is some weighting function (typically unity), $\delta t = t_{[n+1]} - t_{[n]}$ is the time step, and m is some positive value (typically $m = 2$). If $\gamma = 0$ and $m = 1$, then b_{ij} forms the basis for a nearest neighbour solution. For $\gamma = 1$ and $m = 1$ we are applying a minimum acceleration criterion, while $\gamma = 1$ and $m = 2$ may be viewed as minimising the rate of work on the fluid elements containing the particles.

To overcome difficulties which arise when considering associations with $i = 0$ or $j = 0$, the possibility of multiple particles casting a single image, or with other anomalous behaviour, we utilise a number of functions or modifiers in conjunction with the basic cost. The cost of an association α_{ij} is written as

$$c_{ij} = C(b_{ij} \eta_i \varepsilon_j \mathscr{D}_j \tau_j), \tag{6}$$

where the modifier functions η_i, ε_j, $\mathscr{D}_j$ and τ_j are typically step functions acting as price premiums or discounts. The purpose of these functions is explained below.

The function $C(\cdot)$ controls the maximum cost, $c_{\max}$ (say), an association is allowed to have. If the expression inside $C(\cdot)$ yields a result larger than $c_{\max}$, then $C(\cdot)$ will set the resultant value of c_{ij} to infinity. The pairing algorithm outlined in the previous section may be modified to ensure that associations with infinite costs are never in the solution. The associations α_{i0} and α_{0j} are costed either on the cost of associating with a particle lying just outside the physical boundaries of the image, or at the maximum cost $c_{\max}$, whichever is smaller. Thus the maximum cost effectively places a limit on how far a particle may be from its predicted position; this is equivalent to a bound on the apparent acceleration or rate of work.

The use of infinite costs in place of $c_{ij} > c_{\max}$ greatly improves computational efficiency as we need not evaluate costs for associations between well separated particles, nor worry about such associations in the pairing algorithm. The number of finite values of c_{ij} will be greater than or equal to $\mathrm{Max}(M, N)$ but much less than $(M + 1)(N + 1)$.

One of the difficulties with basic cost function b_{ij} is the need to have an estimate for the velocity $\mathbf{u}_{[i]}$. If a pairing was made for $p_{[i]}$ at the previous time step, then we may use this to estimate $\mathbf{u}_{[i]}$. However, if we have no history for $p_{[i]}$ then we must make a reasonable guess at $\mathbf{u}_{[i]}$ but make an allowance for the guess being wrong. This guess could be based on a knowledge of the flow or the behaviour of other particles in the neighbourhood of $p_{[i]}$. For thin light sheets or flows where there are no appreciable velocity gradients parallel to the viewing direction within the illuminated region, the local mean velocity is a good predictor of $\mathbf{u}_{[i]}$. For thick light sheets in three dimensional flows we may be able to do no better than predict $\mathbf{u}_{[i]}$ as some constant. Regardless of how we predict $\mathbf{u}_{[i]}$ where it is not known from the particle history, we shall increase the maximum distance between the predicted position of $p_{[i]}$ and any paired image $q_{[j]}$ by reducing the cost of any such association by multiplying by some factor $\eta_i \leqslant 1$ (we may wish to prevent matches on large, small or elliptical particles when $\mathbf{u}_{[i]}$ is not known by setting $\eta_i = \infty$ for such particles). Typically this discount factor will be derived from some estimate of the maximum error in the estimated value of $\mathbf{u}_{[i]}$. If η_i is set too large, then only particles conforming very closely to the predicted $\mathbf{u}_{[i]}$ will be able to enter the tracking process. If too small, then initial matchings may be made over unaceptable distances. However, provided $c_{\max}$ is sufficiently small, any erroneous pairings introduced by a value of η_i which is too small will not persist past the frame pair at which they entered the problem: $\mathbf{u}_{[i]}$ at the next step will not predict the particle's position with sufficient accuracy. For particles $p_{[i]}$ which were matched at the previous time step we know $\mathbf{u}_{[i]}$ so offer no discount and set $\eta_i = 1$.

We utilise a number of ways of detecting and using particles which appear at almost the same location at $t_{[n+1]}$. The first is based on the shape of a particle: if a blob appears very elliptical, then it may represent two particles close together. We assign to the elliptical blob one ordinary particle with the modifier $\varepsilon_j = 1$, and a possible second particle with ε_j set to some constant greater than one. Associations with the ordinary particle proceed as normal. Associations with the second particle are costed higher by the factor ε_j. If ε_j is sufficiently large, and $c_{\max}$ sufficiently small, then an association with this second particle will only be made if we predict two particles should be there on the basis of the information at $t_{[n]}$. Note that particles with no velocity history are prohibited from matching with such particles.

The modifier $\mathscr{D}_j$ operates in much the same manner as ε_j except that it is based on the overall area of the particle rather than the shape. Particles smaller than a preset limit have $\mathscr{D}_j = 1$, while those larger are treated as two, the second with a price premium of $\mathscr{D}_j \geqslant 1$.

As noted in subsection 3.2, we locate particles using two different thresholds. Particles satisfying the more stringent threshold criteria are assigned a modifier value of $\tau_j = 1$, while those satisfying only the less stringent threshold suffer a premium set by $\tau_j \geqslant 1$, restricting the size of the error between their predicted and actual locations at $t_{[n+1]}$.

Extensive use of this costing and pairing strategy in a variety of different flows suggests only a very weak dependence on the ε_j, $\mathscr{D}_j$ and τ_j modifiers. It is for this

reason, combined with simplicity, that we have adopted step rather than continuous functions. Only η_i and $c_{\max}$ have a marked effect on the tracking process.

3.5. Implementation

The tracking algorithm we have outlined in the previous four subsections has been implimented on a PC/AT compatible micro computer using a Data Translation DT-2862 arithmetic frame grabber. All software has been written by the author in a combination of assembler and Fortran in such a manner as to allow fully automated tracking. To improve computational efficiency the association costs c_{ij} are evaluated only once and stored after rescaling to a single byte integer value. The cost matrix c_{ij} requires $(M + 1)(N + 1)$ bytes of memory limiting this simple implementation to 511 particles in each sample (c_{ij} is 256 Kbytes). An alternative implementation of the algorithm making use of the fact that many c_{ij} are infinite allows up to 4095 particles to be tracked by the PC/AT. Practical considerations limit the number of particles to around 2000. (When the experimental work reported in this paper was undertaken only 511 particle could be tracked.)

The speed of operation depends on the number of particles and quality of the images. When tracking a relatively small number of particles, the execution speed is limited to around 25 frame pairs per minute by the need to acquire images in groups of four samples from a playing video tape. When tracking more than around 200 particles the computations for the location and pairing of particles becomes important. With approximately 400 particles the processing rate varies from six (20 MHz 80386) to fifteen (25 MHz 80486) frame pairs per minute depending on the speed of the computer.

For each time step the pairing algorithm is executed twice. The first pass operates on $p_{[i]}$ and $q_{[j]}$ at $t_{[n]}$ and $t_{[n+1]}$, respectively. Any particles at $t_{[n]}$ which are not matched are stored away for later use. Particles at $t_{[n+1]}$ which were not matched are then compared with particles a $t_{[n-1]}$ which were also left unmatched, to determine if any further pairings may be made. This second pairing operation increases the number of matches made by between 1% and 10%. The number of matches made at any time step depends primarily on the velocity component normal to the light sheet and the thickness of the light sheet.

3.6. Velocities

There are a number of different methods of estimating the velocities from the positions of a particle along a path. The simplest and arguably most effective of these fit a surface to the positions of a given particle in $\mathbf{x}, t$ space. In practice we fit a straight line to the x, t and z, t data sets over some time interval $s\delta t$ during which the change in velocity is relatively small. The slope of the fitted lines may then be used to estimate the velocity. In the limit of $s = 1$ this degenerates to the simple finite difference approach. The fitting process should utilise all the information in $s\delta t$, suggesting either

a least squares fit or a Chebeshev polynomial. Both these strategies are in some sense optimal and in practice yield very similar results. The results presented in this paper employ the least squares approach.

With the least squares fit the error in the estimated velocity σ_u is given by the standard error in the slope of the line which may be expressed as

$$\sigma_u = \left\{\frac{12}{(s+2)(s+1)s}\right\}^{1/2} \frac{\sigma_x}{\delta t}. \tag{7}$$

With the values used in this paper ($\sigma_x < 0.2$, $s = 5$, $\delta t = 0.08s$) we have a velocity error $\sigma_u < 0.24\,\mathrm{mms}^{-1}$. Typically the error is less than half this value.

We should note that the particle tracking process and the method of calculating velocities over paths with $s > 1$ leads to a sampling bias in favour of regions of the flow in which the velocity normal to the light sheet is relatively small. This arises through the need for particles to remain in the light sheet over $s + 1$ samples. The value $s = 5$ has been chosen so that relatively few particles will ever cross the 7 mm light sheet in less than $s\delta t = 0.4$ s. In the present experiments this bias is only significant in the very early stages of the decay process and will not affect the discussion in the following sections.

4. Nonrotating turbulence

We shall consider first the nonrotating ($f = 0$) limit as the behaviour of this well known flow will show any major problems with our experimental apparatus or measurement method. There are a wide range of methods of characterising such a flow. For the present we will confine our attention to relatively simple velocity statistics.

We define the spatial-ensemble mean velocity for the illuminated region as

$$\bar{\mathbf{u}}_{\mathrm{se}}(t_{[n]}) = \sum_r \sum_i \mathbf{u}_{[i][r]}(t_{[n]}) \Big/ \sum_r M_{[r]}(t_{[n]}), \tag{8}$$

where $u_{[i][r]}(t_{[n]})$ is the velocity for particle i in realisation r at time $t_{[n]}$ in the light sheet. There are $M_{[r]}(t_{[n]})$ particles for which the velocity is known in realisation r at time $t_{[n]}$. The summation is over all particles located in each realisation (approximately 350 particles) and all realisations (typically $R = 25$ realisations) at the specified time. As our tank represents a closed system the flow must have a zero mean velocity. Therefore we might expect (8), which is a sample of the global mean velocity, to evaluate to zero. However, as we can see from Fig. 3, $\bar{\mathbf{u}}_{\mathrm{se}}$ does not vanish as expected. Two features are readily apparent in the horizontal component of $\bar{\mathbf{u}}_{\mathrm{se}}$: an oscillation with a period of approximately 1.25 s ($Ut/M \approx 9.5$), and an exponentially decaying mean value. The first of these has the correct frequency for deep water waves in the

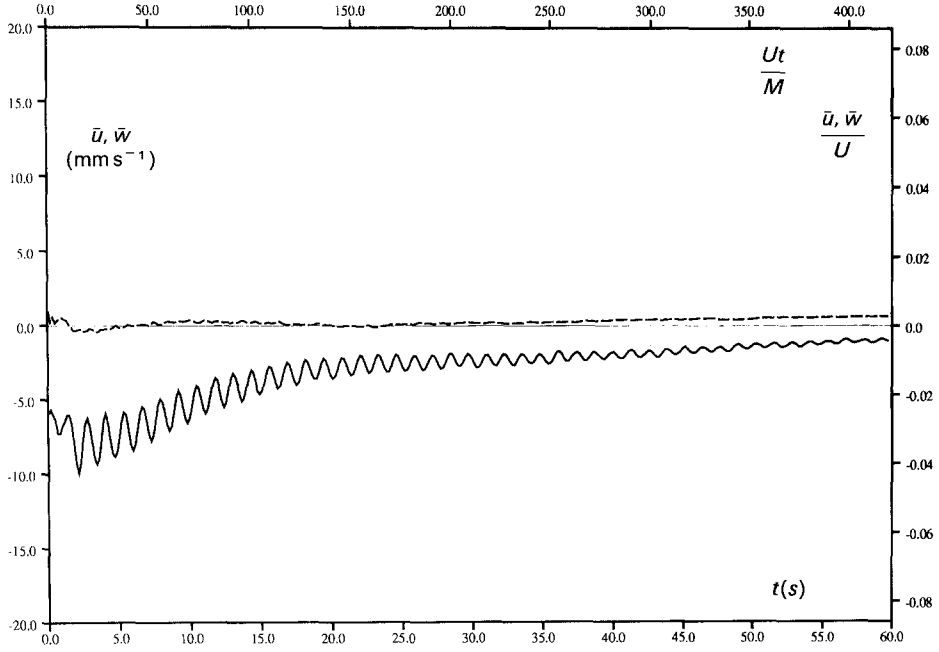

Fig. 3. Evolution of the mean velocity within the light sheet for $f = 0$. The solid line is for the horizontal component, and the dashed line for the vertical component.

tank. These waves are initiated by the passage of the grid. The mechanics of driving the grid prevented a rigid lid from being used. While slats were fitted between the bars of the grid to reduce the effect of the free surface, a significant wave component still exists.

The origin of the decaying flow in Fig. 3 is also due to the passage of the grid. Careful observations show that the passage of the grid sets up a significant large scale circulation within the tank. This motion is predominantly in the direction of the grid near the centre of the tank, with return flow near the boundaries. Increasing the solidity of the grid near the walls and making a double pass with the grid reduced the size of this circulation, but it did not prove possible to eliminate it. The embarrassing existence of the systematic flow is not unique to the apparatus used in these experiments (Maxworthy, personal communication) although much longer tanks do not seem to suffer this problem (e.g. Britter et al. [4]). However, most studies do not address the issue, partly due to the difficulty with measuring and allowing for such a flow, especially using traditional techniques. The majority of experiments on homogeneous turbulence have been in wind tunnels where the problem is much less severe with a much large ratio of mesh length to tunnel width (giving a much better separation of scales). A similar phenomenon is known to exist in oscillating grid experiments with the production of a larger scale systematic flow (e.g. McDougall [13]). In both cases the scale of the systematic flow is set by the overall tank geometry and not the mesh size.

Due to the existence of spatial structure in the surface waves and in the systematic circulation, statistics should be evaluated relative to a spatially varying ensemble mean, $\bar{\mathbf{u}}_e$, say. In practice the mean is evaluated for each time step by distributing the particle velocities onto a regular grid (using a local weighted least squares fit) and averaging the grid over all realisations. Figure 4a plots the evolution of the turbulent fluctuations relative to this mean. If the initial turbulence field is approximately homogeneous and isotropic then during the initial stages of decay, starting some time $t > t_0$ after the passage of the grid, and extending until Ut/M is greater than around 200, the turbulence intensities are then found to follow the power law

$$\frac{\bar{u}^2}{U^2} = A\left[\frac{U(t - t_0)}{M}\right]^{-\beta}, \tag{9}$$

where A and β are constants and t_0 is the virtual origin at which the fluctuations would be infinite. Batchelor's [2] initial period decay law gives $\beta = 1$ for the type of grid used in these experiments. More extensive data sets suggest a value of β between 1.2 and 1.4 (e.g. Comte-Bellot and Corrsin [5]). As is frequently the situation with power laws, the exponent β for the optimal fit is sensitive to location of the virtual origin t_0. Typically the virtual origin is found to be located at Ut_0/M between 1 and 6, with the initial period, during which the power law applies, starting some ten to twenty grid time scales (M/U) later.

Figure 4b reproduces the fluctuation data in a log-log format. Between approximately 10 and 75 grid time scales after the passage of the grid the energy decays following a power law with $\beta \approx 1$. The size of the ensemble is not sufficiently large to pin down β more precisely as the value of t_0 is also unknown. The initial much flatter curve is due to a combination of the behaviour immediately behind the grid and the conditional sampling provided by the particle tracking ($v < \delta y/(s\delta t)$). The bias due to the conditional sampling is significant only in the very early stages of the decay. The other notable feature of this plot is the marked decrease in the decay rate at times larger than $Ut/M \approx 75$ ($t \approx 10$ s). The decay does not follow a power law after this point and cannot be attributed to the normally increased decay during the final period. Rather it would appear as though the largest turbulent scales are starting to interact with the large scale systematic flow set up by the initial passage of the grid. Moreover, as this circulation is on the same scale as the dimensions of the tank, the overall experimental geometry will be playing a significant role. A power-law decay appears to establish again after $Ut/M \approx 250$ once the scales of the systematic flow are fully incorporated into the decay process. The decay rate after this point will be influenced by the geometry of the tank but is again with $\beta \approx 1$.

Clearly the large scale circulation imposed by the initial passage of the grid imposes severe experimental difficulties. A variety of methods were tried to minimise its generation, the results presented in this section representing the best alternative. As will be seen in the next section, rotation compounds the difficulties with the generation of inertial waves.

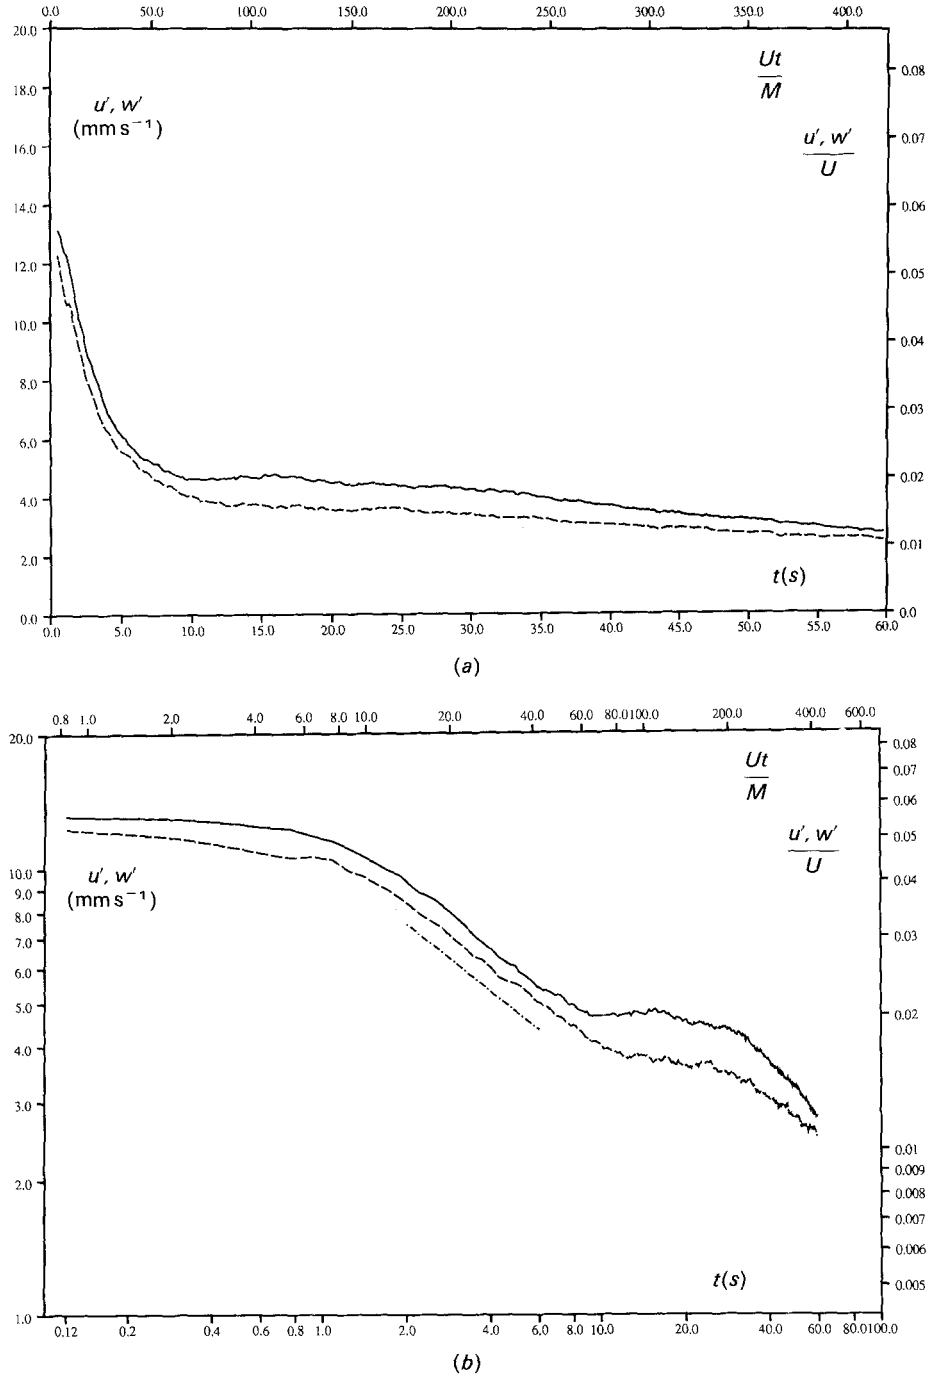

Fig. 4. Evolution of the turbulent velocity fluctuations about the spatially varying ensemble mean flow for $f = 0$. Plot (a) has a linear scale while (b) a logarithmic scale showing the power law behaviour of the decay ($t_0 = 0.4s$). Solid lines indicate the along-tank horizontal component, while dashed lines denote the vertical component. The initial period decay law with $\beta = 1$ is shown as a dot-dash line in (b).

5. The effect of rotation

In the interior of the tank, rotation affects the motion through the Coriolis force acting primarily on large scales and low velocities (i.e. small Rossby number). The systematic flow found for nonrotating turbulence has both these features and will hence be affected by the rotation from a very early stage. Figure 5 shows the evolution of the $\bar{\mathbf{u}}_{se}$ mean velocity at a rotation rate of $\Omega = f/2 = 0.5\,\text{rad.}\,\text{s}^{-1}$ (for the sake of brevity we shall present results only for this rotation rate). The relatively short period (~ 1.25 s) free surface waves (now Kelvin waves) are still present. In addition, we are able to see two (or more) relatively large amplitude inertial wave modes in place of the simple mean flow found in the previous section.

Plane inertial waves in a rotating system may be shown to have the velocities obeying

$$u = \pm(k_2\lambda_3 - k_3\lambda_2)\cos(\kappa(k_1x + k_3z) \mp k_3ft) - \lambda_1 \sin(\kappa(k_1x + k_3z) \mp k_3ft),$$
$$w = \pm(k_1\lambda_2 - k_2\lambda_1)\cos(\kappa(k_1x + k_3z) \mp k_3ft) - \lambda_3 \sin(\kappa(k_1x + k_3z) \mp k_3ft), \tag{10}$$

in the plane of the light sheet (e.g. Greenspan [8]). The magnitude of the wave number is given by κ and the orientation by the unit vector $\mathbf{k}$. The amplitude vector $\boldsymbol{\lambda}$ may be

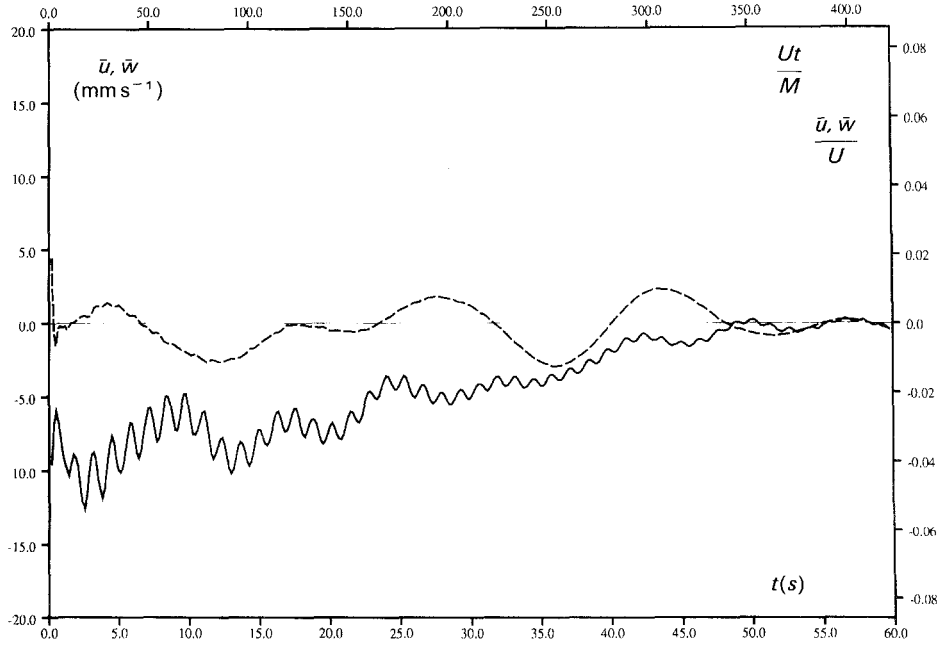

Fig. 5. Evolution of the mean velocity within the light sheet for $f = 1$ rad/s. The solid line is for the horizontal component and the dashed line for the vertical.

any vector perpendicular to **k**. Subscripts indicate the components in the three axis directions. The frequency ω of these plane waves is related to the angle between the rotation axis and the unit wave number vector by

$$\omega = k_3 f. \tag{11}$$

The solution given by (10) satisfies both the linear and nonlinear intertial wave equations. Note, however, that superposition of modes is not possible in the latter case.

While the inertial waves in this flow are likely to be nonliner, we may make some progress by assuming a degree of superposition is possible. As noted above, Fig. 5 shows evidence of two inertial wave modes. Experiments at different rotation rates have shown the relative amplitude of these modes appears to depend on the rotation rate, presumably the result of variations in the ratio of the time to traverse the grid and the rotation period. By considering the vertical mean velocity we can see a mode with $\omega \approx 0.4f$, implying the wave number vector is inclined at an angle of approximately 66° to the rotation axis. In addition to this $\omega \approx 0.4f$ mode, the horizontal mean velocity shows a response at $\omega = f$. The frequency and absence of vertical motion at this frequency are consistent with the wave number vector parallel to the rotation axis.

As with the nonrotating flow in the previous section, we may determine the spatio-temporal structure of the mean flow and utilise it when evaluating the statistics of the turbulence. Figure 6a plots the evolution of the fluctuations about the varying ensemble mean flow. Figure 6b plots the same data on logarithmic axes from which we see the existence of a power law giving a somewhat lower decay rate than observed for $f = 0$. While the value of β is not known for either case, overlay in the two plots allows a direct comparison to be made if we assume t_0 is the same for both cases. The decay process does not show the abrupt change in behaviour found for $f = 0$, suggesting that interaction between the inertial waves is generating turbulence at relatively large scales preventing an initial scale separation between the inertial waves and the turbulence (we had such a scale separation in the nonrotating flow in the previous section). Note that the energy contained in the inertial wave field is significantly larger than the systematic flow in the $f = 0$ limit.

A series of four plots showing the spatial structure and temporal evolution of the ensemble mean velocity field is given in Fig. 7. This series has been chosen to show the spatial structure of the inertial waves and the relatively rapid changes it undergoes. Note that the standard error in this spatially varying ensemble mean is proportional to $\bar{\mathbf{u}}'_e/R^{1/2}$, where $\bar{\mathbf{u}}'_e$ are the fluctuations about $\bar{\mathbf{u}}_e$ and R is the number of realisations (25 for the present experiments). While this represents a relatively small degree of uncertainty, statistics for a much larger ensemble would be desirable.

The degree of anisotropy between the horizontal and vertical components is illustrated by the structure function

$$K = (\bar{u}'_e - \bar{w}'_e)/(\bar{u}'_e + \bar{w}'_e), \tag{12}$$

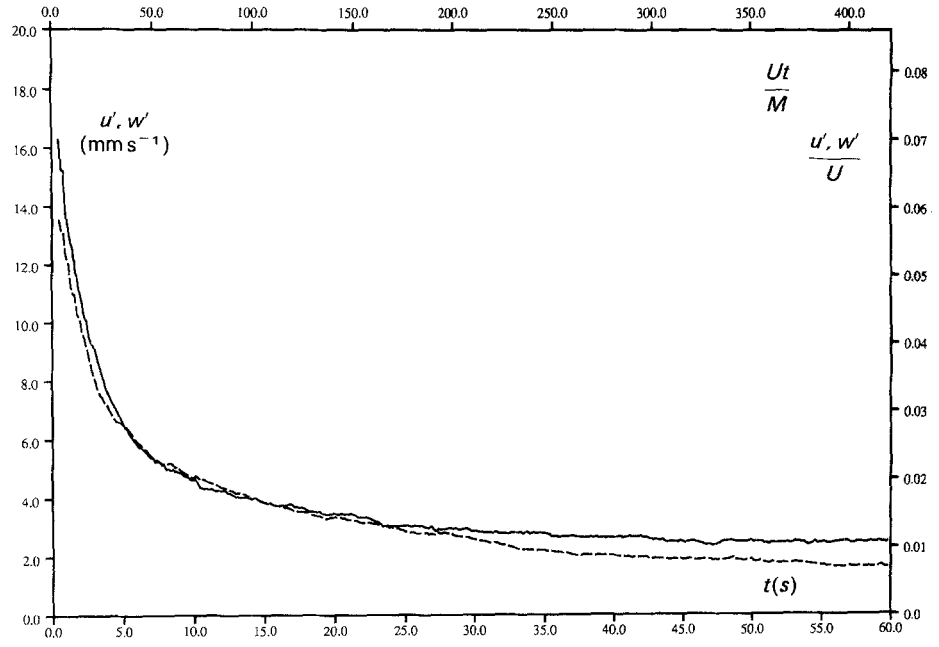

(*a*)

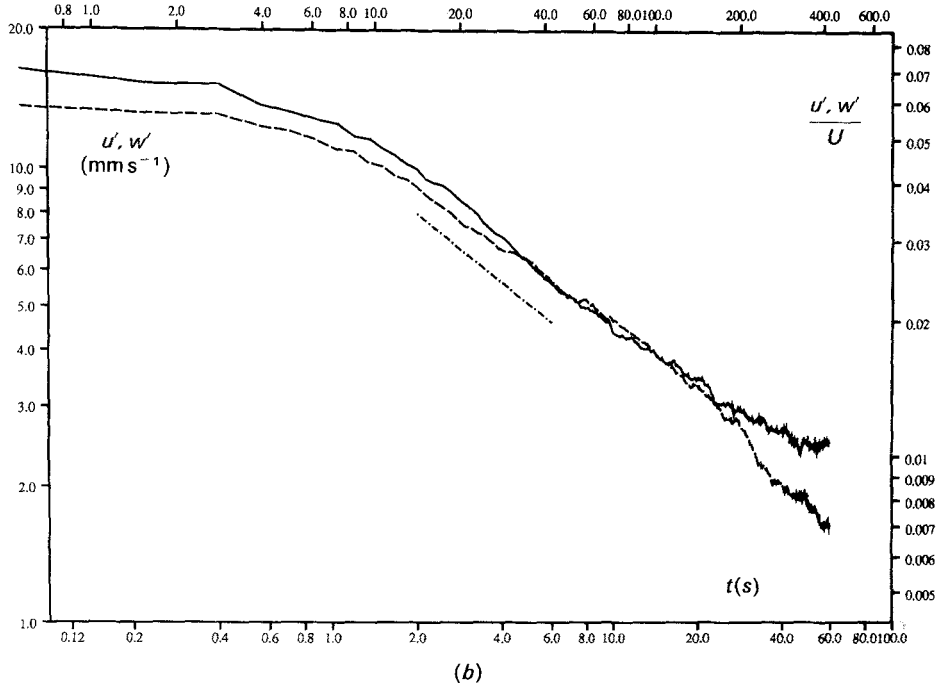

(*b*)

Fig. 6. Evolution of the turbulent velocity fluctuations about the spatially varying ensemble mean flow for $f = 1$ rad/s. Plot (a) has a linear scale while (b) a logarithmic scale showing the power law behaviour of the decay ($t_0 = 0.4$s). Solid lines indicate the along-tank horizontal component, while dashed lines denote the vertical component. The initial period decay law with $\beta = 1$ is shown as a dot-dash line in (b).

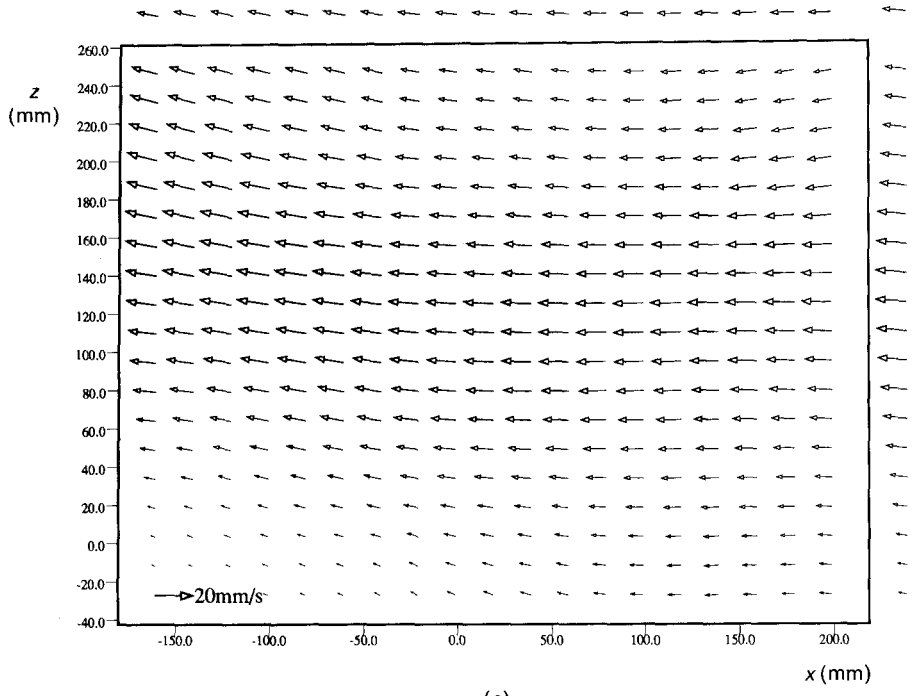
z
(mm)
260.0
240.0
220.0
200.0
180.0
160.0
140.0
120.0
100.0
80.0
60.0
40.0
20.0
0.0
-20.0
-40.0
20mm/s
-150.0
-100.0
-50.0
0.0
50.0
100.0
150.0
200.0
x (mm)

(*a*)

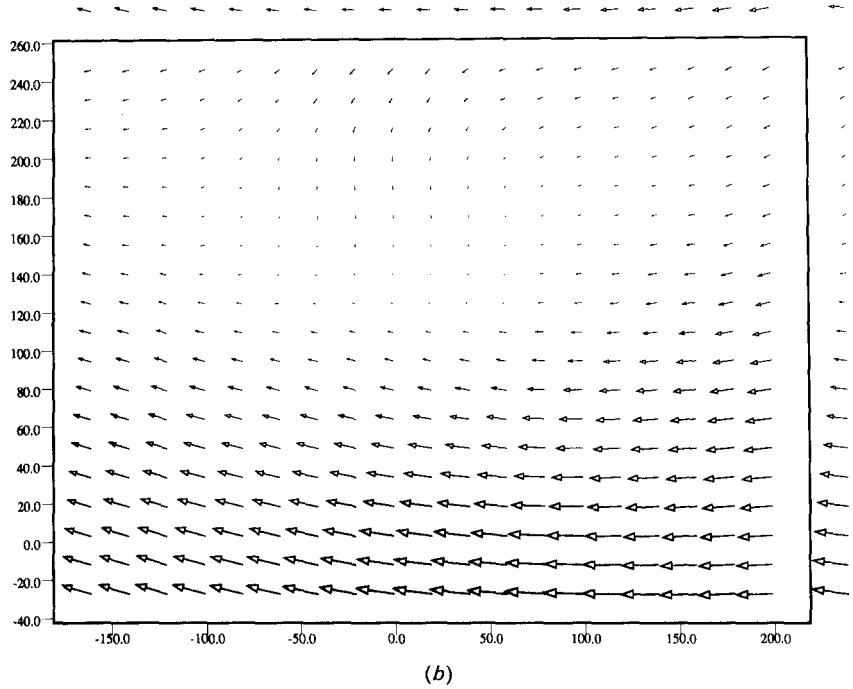
260.0
240.0
220.0
200.0
180.0
160.0
140.0
120.0
100.0
80.0
60.0
40.0
20.0
0.0
-20.0
-40.0
-150.0
-100.0
-50.0
0.0
50.0
100.0
150.0
200.0

(*b*)

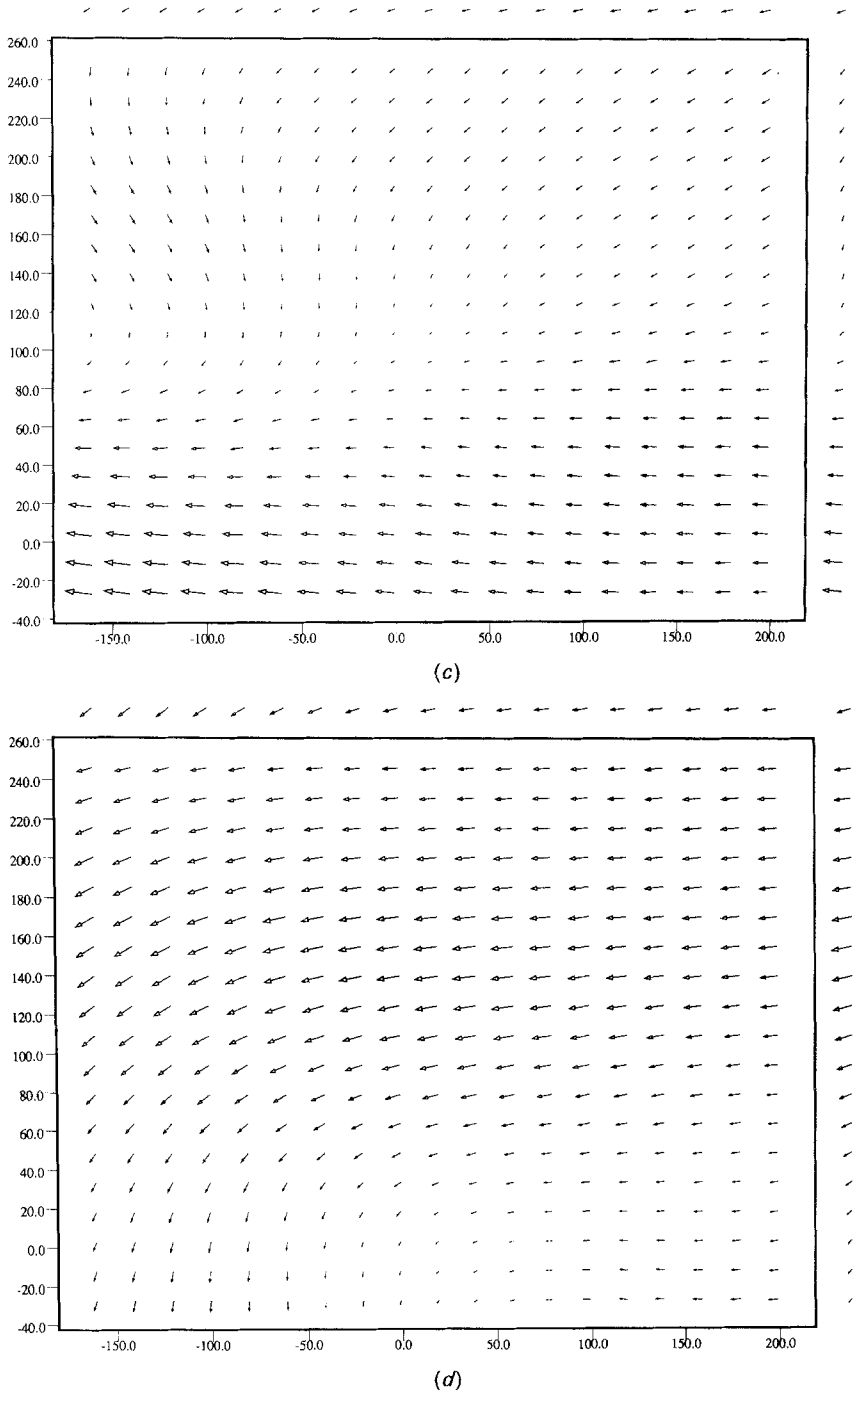

Fig. 7. Series of plots showing the evolution of the spatially varying ensemble mean motion for $f = 1$ rad/s. The passage of a wave crest through the series is clearly visible. The velocity scale is indicated on the first figure. The times corresponding to these plots are (a) $4s(Ut/M = 29)$, (b) $6.24s(Ut/M = 46)$, (c) $8.48s(Ut/M = 62)$, (d) $10.72s(Ut/M = 79)$. The arrows at the top and right-hand extremes of the plots are the averages for the corresponding column or row.

which is plotted in Fig. 8 for both the $f = 0$ rad/s and $f = 1$ rad/s cases. In the absence of rotation we see a persistent 10% anisotropy, a figure consistent with previous investigators for classical turbulence. For $f = 1$ rad/s contamination from the systematic flow yields a smaller degree of anisotropy at early stages. As the decay rate of vertical fluctuations is greater than for horizontal fluctuations, there is a gradual increase in the anisotropy.

The dominant effect of rotation on the flow is in the growth of the associated length scales. We define the two-point correlation functions

$$R_{ij}(\Delta\mathbf{x}, t = t_c) = \frac{\Sigma\Sigma\,(u_i'(\mathbf{x})u_j'(\mathbf{x} + \Delta\mathbf{x}) - \bar{u}_i'\bar{u}_j')}{\bar{u}_i'\bar{u}_j'}, \tag{13}$$

where $\Delta\mathbf{x}$ is the separation (in the $y = y_0$ plane) between the points at which the two velocity components (projected onto the $y = y_0$ plane) are evaluated. Subscripts are used here to represent different velocity components rather than different times. The summation is over all particle pairs with the appropriate separation, falling in the specified time interval $(t_c \pm \Delta t_c)$, and over all realisations. Note that the finite thickness of the light sheet (δy) means that the correlation functions defined by R_{ij} will in general be slightly smaller (numerically) than the true correlation functions for velocities lying in the $y = y_0$ plane. This will not be significant, however, for $|\Delta\mathbf{x}| \gg \delta y$.

Figure 9 shows the $R_{uu} = R_{11}$ and $R_{ww} = R_{33}$ two-point correlation functions

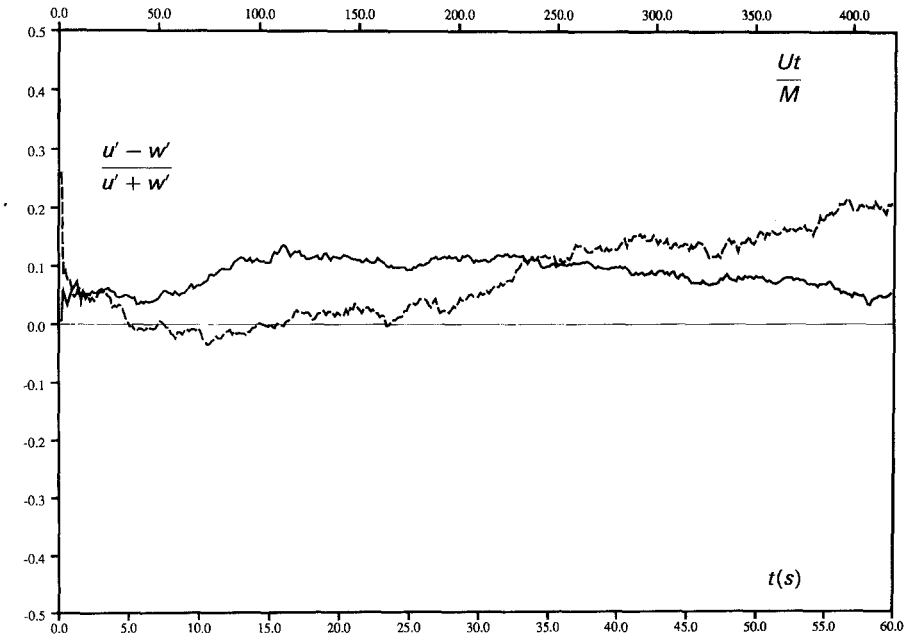

Fig. 8. Evolution of the structure function $K = (\bar{u}_e - \bar{w}_e)/(\bar{u}_e + \bar{w}_e)$ for the classical (solid line) and rotating (dashed line) flows.

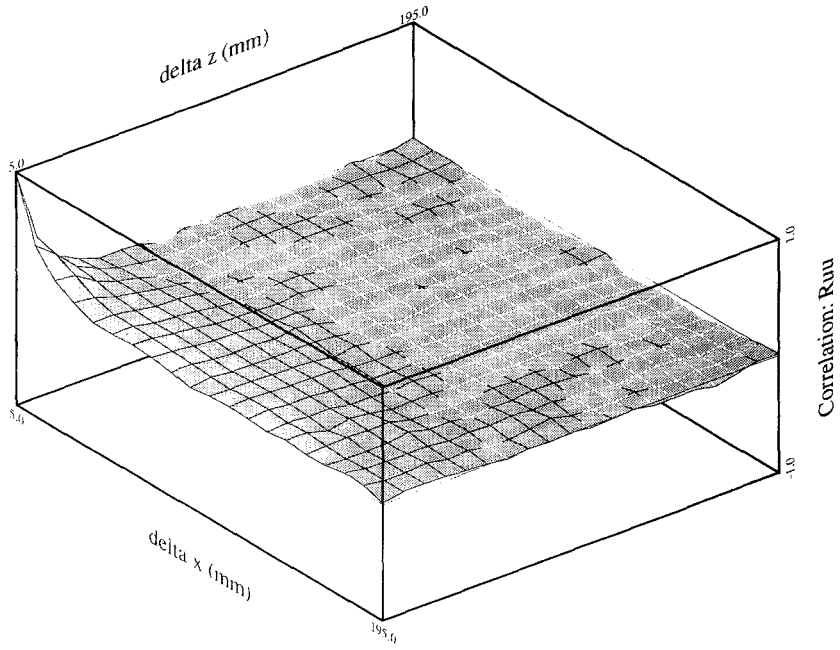

(*a*)

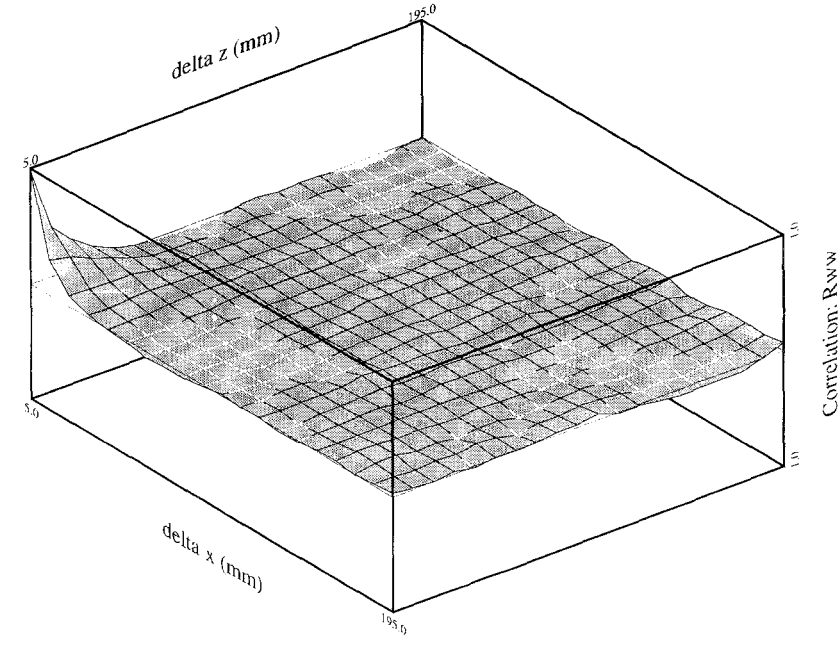

(*b*)

Fig. 9. Two-point correlation functions for the $f = 1$ rad/s ensemble. (a) $R_{uu} = R_{11}$ and (b) $R_{ww} = R_{33}$ at $t = 3s$ ($Ut/M = 22$); (c) $R_{uu} = R_{11}$ and (d) $R_{ww} = R_{33}$ at $t = 20s$ ($Ut/M = 147$).

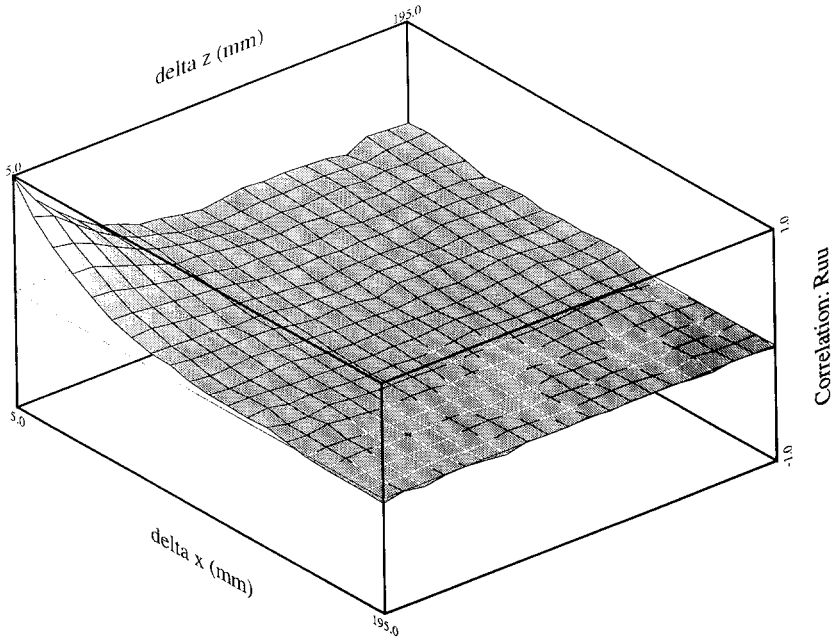

(*c*)

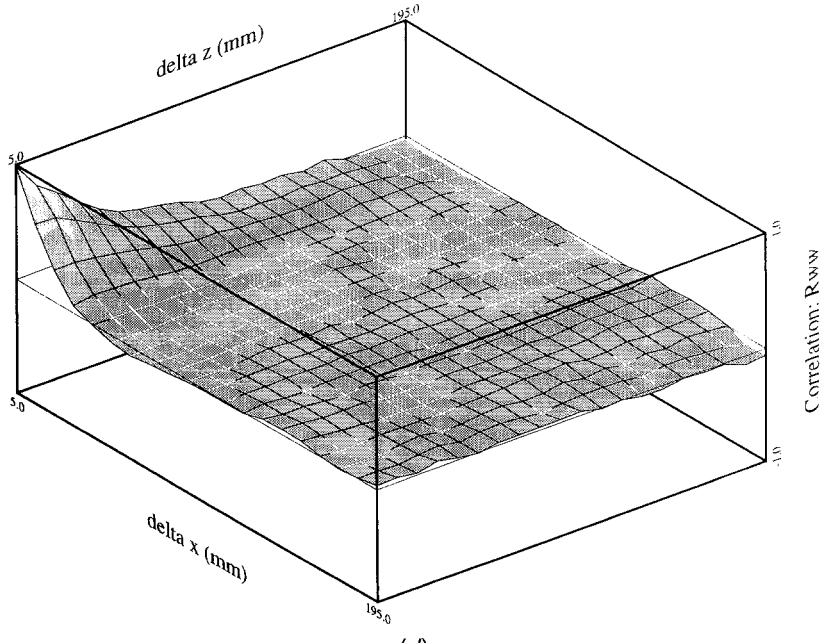

(*d*)

Fig. 9 (Continued)

evaluated at two different times. The first pair of plots is evaluated relatively early at $t = 3s$ ($Ut/M = 22$) before the effects of rotation are dominant. The behaviour as a function of spatial separation follows the normal form of correlations between velocities parallel to the separation remaining positive, while those perpendicular pass through zero. The $R_{uv} = R_{13}$ correlation function (not plotted) is approximately zero for all separations, as would be expected. The integral length scales evaluated from these correlations are $L_{uux} = 0.69\,M = 22\,\text{mm}$ for horizontal velocities and separations, and $L_{wwz} = 0.50\,M = 15\,\text{mm}$ for vertical velocities and separations. The inequality of these two length scales reflects the growing anisotropy introduced by rotation in addition to the weak anisotropy inherent in grid turbulence and contamination from the systematic flow. For isotropic turbulence, the length scales with the velocity and separation perpendicular should be half the value of that when the velocity and separation are parallel. Here we find they slightly exceed this relationship ($L_{uuz} = 12\,\text{mm}$ and $L_{wwx} = 10\,\text{mm}$), an affect probably due to the anisotropy introduced by rotation and the systematic flow.

The correlation functions at $t = 20s (Ut/M = 147)$ show clearly the enhanced vertical correlation of the velocities, and so will yield larger values for the vertical integral length scale. However, the systematic flow set up by the initial passage of the grid contaminates these figures to a significant degree, with the correlation functions *increasing* with vertical separation for larger separations. This makes evaluation of the integral length scales from these statistics very difficult. Further discussion on these length scales and how they are affected by rotation and the systematic flow are beyond the scope of this paper.

One of the strengths of particle tracking is the ability to obtain Lagrangian statistics much more efficiently than the manual digitization. Figure 10 plots the Lagrangian autocorrelation functions for particles in the light sheet at $t = 5s (Ut/M = 37)$. Due to the relatively thin light sheet we are considering there is a strong bias at longer times to particles with small velocities normal to the sheet. The number of particles used in this calculation is also plotted in the figure, and can be seen to decline approximately exponentially, as would be expected for a random distribution of cross-sheet velocities. Caution should be used when interpreting this plot as the turbulence is decaying with time; this is also a feature of comparable Lagrangian statistics in wind tunnel experiments (e.g. Snyder and Lumley [15]). The present results, however, are also contaminated by the systematic flow. The fluctuations in the trace apparent for separations greater than around 2 seconds are due to the free surface Kelvin waves: our Eulerian method of correcting the velocities for the systematic flow is not appropriate for Lagrangian autocorrelations (further discussion is beyond the scope of this paper). Note also that there are too few particles in the sample for separations exceeding 4 seconds.

6. Conclusions

In this paper we have outlined a simple, efficient PIV method of automatically tracking individual particles in a fluid flow. The method has been implemented on a

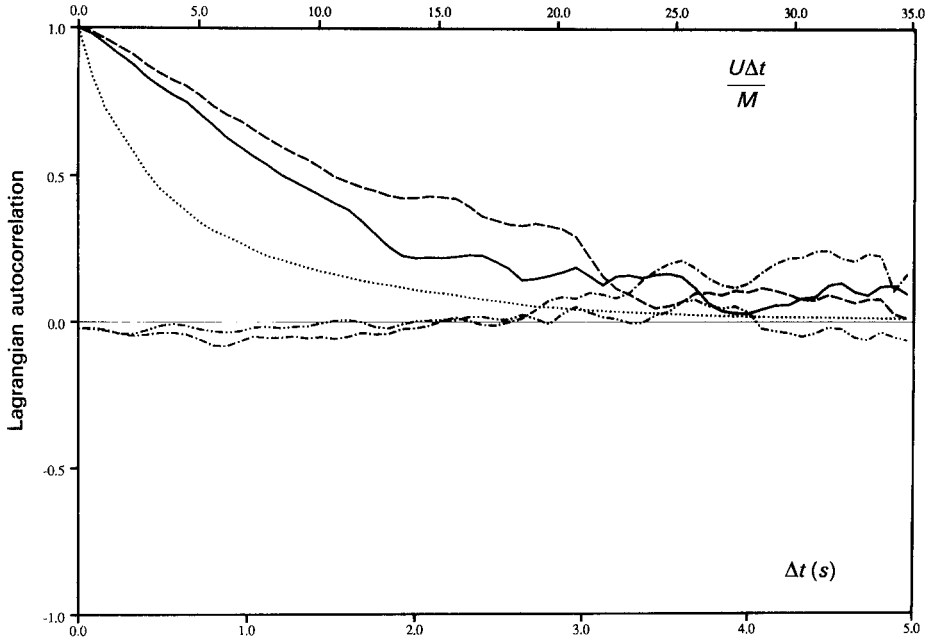

Fig. 10. Autocorrelation functions with $f = 1$ rad/s for particles tracked from $t = 5s$ ($Ut/M = 37$). The $u - u$ component is plotted as a solid line, $w - w$ as dashes, $u - w$ as dot-dash, and $w - u$ as dot-dot-dash. The relative number of particles belonging to the original set still illuminated by the light sheet is also plotted as a dotted line. Note that at large time separations the number of particles becomes too low for the autocorrelations to be meaningful.

PC/AT compatible micro computer utilising a medium cost frame grabber and Super VHS video tape recorder. While the algorithm described in this paper has been implemented to allow up to 4095 particles to be tracked simultaneously, practical considerations impose a limit of approximately 2000 particles. The experiments described in this paper utilised an early version of the particle tracking system following typically 350 particles. The high throughput (up to 25 frame pairs per minute in some situations) makes this method ideal when ensemble statistics of a temporally varying flow are required. Both Eulerian and Lagrangian views of the flow may be obtained from the particle paths.

The particle tracking technique has been applied to a wide range of flows and is illustrated here as part of a study of the decay of rotating turbulence. In addition to being a convenient method for measuring the velocities in this zero mean flow experiment, the ability to gather information over a two dimensional region has proved invaluable in analysing the systematic flow contaminating the experiment.

Severe limitations in the experimental set up have been discovered. The initial passage of the grid is found to set up a large scale systematic flow. In the absence of rotation this flow is predominantly in the direction of the grid near the centre of the tank, with the return flow near the boundaries. Multiple passes of the grid and increased solidity near the walls reduced this effect but did not prove possible to

eliminate it. When the system is rotating, the systematic flow breaks up into a large scale inertial wave field in the order of an inertial period. The wave field appears to consist of the superposition of two modes, one with the wave number vector parallel to the rotation axis, and the other with it inclined at approximately 66°. This wave field has a high degree of coherency throughout an ensemble of 25 realisations. The spatially and temporally varying structure of this wave field is established from the ensemble flow, and used as the basis of a correction to establish the turbulent component of the flow. The separation in length scales did not prove sufficient for this technique to be entirely effective.

The generation of such a strong systematic flow during the passage of the grid is of deep concern in the rotating system. Earlier investigators have not reported such difficulties, primarily as a result of such information being less accessible to traditional measurement techniques. The problem may be expected to be less severe with the symmetric grid arrangement of Ibbetson and Tritton [11] as the velocity component of the large scale flow introduced by the passage by the grid will be parallel to the rotation axis over most of the tank. However, at the initial mid-height position of the two grids and at the top and bottom of their annulus we would expect a significant radial component which may be strongly affected by the Coriolis force. The experiments of Jacquin et al. [12] will be less affected due to the mean flow parallel to the rotation axis and the absence of boundaries perpendicular to that axis. Moreover, as our present difficulties are the result of initial transients, the statistically steady flow in their tube does not have the same form of forcing function.

The oscillating grid experiments of Bretherton and Turner [3], Hopfinger et al. [10], Dickenson and Long [6] and Fluery et al. [7] may also be adversely influenced by the existence of a large scale systematic flow. Experiments by McDougall [13] show that in the absence of rotation such oscillating grids generate a systematic flow with an amplitude comparable to the fluctuations. This flow will interact with the Coriolis force in regions of the tank remote from the grid. This feature is being investigated further by Drayton (personal communication) using the present particle tracking technique. The advection of smaller scale turbulence by this flow may well lead to the formation of the intense vortices and rapid spatial change from three dimensional to nearly two dimensional motion characteristic of such experiments.

Despite the experimental limitations, rotation has been found to decrease slightly the rate of decay of turbulent fluctuations. In addition to a small increase in the anisotropy as the decay progresses, rotation leads to a marked increase in the integral length scales. These findings are in broad agreement with those of other authors.

References

1. Adrian, R. J., Particle-image techniques for experimental fluid mechanics. *Annu. Rev. Fluid Mech.* 23 (1991) 261–304.
2. Batchelor, G. K.: *The Theory of Homogeneous Turbulence.* Cambridge University Press (1953).

3. Bretherton, F. P. and Turner, J. S., On the mixing of angular momentum in a stirred rotating fluid. *J. Fluid Mech.* 32 (1968) 449–464.
4. Britter, R. E., Hunt, J. C. R., Marsh, G. L. and Snyder, W. H., The effects of stable stratification on turbulent diffusion and the decay of grid turbulence. *J. Fluid Mech.* 127 (1983) 27–44.
5. Comte-Bellot, G. and Corrsin, S., The use of a contraction to improve the isotropy of grid-generated turbulence. *J. Fluid Mech.* 25 (1966) 657–682.
6. Dickinson, S. C. and Long, R. R., Oscillating-grid turbulence including effects of rotation. *J. Fluid Mech.* 126 (1983) 315–333.
7. Fluery, M., Mory, M., Hopfinger, E. J. and Auchere, D., Effects of rotation on turbulent mixing across a density interface. *J. Fluid Mech.* 223 (1990) 165–191.
8. Greenspan, H. P., *The Theory of Rotating Fluids.* Cambridge University Press (1968).
9. Hichcock, F. L., The distribution of a product from several sources to numerous localities. *J. Math. Phys.* 20 (1941) 224.
10. Hopfinger, E. J., Browand, F. K. and Gagne, Y., Turbulence and waves in a rotating tank. *J. Fluid Mech.* 125 (1982) 505–534.
11. Ibbetson, A. and Tritton, D. J., Experiments on turbulence in a rotating fluid. *J. Fluid Mech.* 68 (1975) 639–672.
12. Jacquin, L., Leuchter, O., Cambon, C. and Mathieu, J., Homogeneous turbulence in the presence of rotation. *J. Fluid Mech.* 220 (1990) 1–52.
13. McDougall, T. J., Measurements of turbulence in a zero-mean-shear mixed layer. *J. Fluid Mech.* 94 (1979) 409–431.
14. Perkins, R. J. and Hunt, J. C. R., Particle tracking in turbulent flows. In Fernholz, H.-H. and Fielder, H. E. (eds), *Advances in Turbulence*, Vol. 2. Springer-Verlag (1979) pp. 286–291.
15. Snyder, W. H. and Lumley, J. L., Some measurements of particle velocity autocorrelation functions in a turbulent flow. *J. Fluid Mech.* 48 (1971) 41–71.