International Journal of Computational Fluid Dynamics

ISSN: 1061-8562 (Print) 1029-0257 (Online) Journal homepage: https://www.tandfonline.com/loi/gcfd20

A simple technique for the visualisation of eddy kinematics in turbulent flows

Geoffrey Lee , Carlo Scalo & Ugo Piomelli

To cite this article: Geoffrey Lee , Carlo Scalo & Ugo Piomelli (2012) A simple technique for the visualisation of eddy kinematics in turbulent flows, International Journal of Computational Fluid Dynamics, 26:4, 263-274, DOI: 10.1080/10618562.2012.693605

To link to this article: https://doi.org/10.1080/10618562.2012.693605

International Journal of Computational Fluid Dynamics
Vol. 26, No. 4, April–May 2012, 263–274

Taylor & Francis
Taylor & Francis Group

A simple technique for the visualisation of eddy kinematics in turbulent flows

Geoffrey Lee, Carlo Scalo* and Ugo Piomelli

Department of Mechanical and Materials Engineering, Queen's University, Kingston, ON, K7L 3N6, Canada

(*Received 13 January 2012; final version received 7 May 2012*)

We developed and tested a simple technique to predict, for flow visualisation purposes only, the evolution of coherent structures in between two given realisations of a turbulent flow. Classic coherent-structure eduction methods are adopted, such as the *Q*-criterion, pressure fluctuations and contours of velocity fluctuations. The kinematics of the evolving structures are reconstructed by means of an advection-based reconstruction technique and captured in a movie. The resulting quality of the animations has been assessed via the Structural Similarity Index (SSIM). The sensitivity to increasing spacing in time of the available flow realisations has been tested and several improvements implemented. The abrupt transition from reconstructed frames of the animation to the available realisations results in a noticeable lack of smoothness. The replacement of the available realisations with a similar advection-based average increases the perceived smoothness of the movies. This is confirmed by the reduced total variation of the SSIM index. The residual minor periodic variations of accuracy have been reduced by introducing a stochastic weighting function. The sensitivity of the results to changes in Reynolds number, resolution and structure representation methods has been tested.

Keywords: flow visualisation; large-eddy simulation; turbulent structures kinematics; Structural Similarity Index; advective interpolation; movie accuracy and smoothness; human visual system

1. Introduction

In turbulent flows, coherent motions (known as 'eddies') play a crucial role in the transport of mass, momentum and energy (Robinson 1991). Hence, a significant effort has concentrated, in the past, on the eduction, quantification and visualisation of coherent eddies.

Coherent structures were first discovered with the aid of experimental flow visualisation methods (Liepmann 1979). Early flow visualisations relied mostly on either line-of-sight (optical interferometry, schlieren and shadowgraphy) or light-scattering (smoke or dye visualisation) methods. Among the first examples of such discoveries were the near-wall streaks (Kline *et al.* 1967) and the spanwise rollers in mixing layers (Brown and Roshko 1974). Numerical simulations of turbulent flows (and, in particular, large-eddy and direct simulations) have given a major contribution to this subject. From a flow visualisation standpoint, they present significant advantages, compared to experimental results. First, since three-dimensional (3D) pressure and velocity fields are available, one can consider a variety of identifiers, and choose the one most appropriate to the flow under investigation. Also, the extraction of the information is local, instantaneous and non-intrusive. Only recently, experimental techniques that yield 3D, instantaneous flow fields such as, for instance, tomographic and time-resolved Particle-Image Velocimetry (Elsinga *et al.* 2006) have become available.

While early experimental studies only yielded particle paths, streamlines or streaklines, the availability of instantaneous data has made it possible to visualise a variety of quantities, highlighting particular aspects of the turbulent flow field. Among the most common choices are vorticity and pressure-related quantities and various invariants of the velocity gradient tensors. Most commonly, two-dimensional (2D) contours are shown in planes of interest; isosurfaces of various quantities are also very often shown, to highlight the 3D structure of the turbulent eddies. Comparisons of the advantages and disadvantages of the various methods can be found in Robinson *et al.* (1989) and Dubief and Delcayre (2000). Various vortex-identification methods have also been proposed (Kida and Miura 1998, Haller 2005).

In addition to the instantaneous snapshots, movies or animations highlighting the kinematic development of these eddies have become widespread. Early movies, obtained from experiments, typically used smoke or dye injection and proved to be particularly effective in analysing mechanisms of interaction between coherent

*Corresponding author. Email: cscalo.ca@gmail.com

ISSN 1061-8562 print/ISSN 1029-0257 online

http://dx.doi.org/10.1080/10618562.2012.693605
http://www.tandfonline.com

motions. For example, Head (1969) successfully captured the time evolution of the intermittent (turbulent/non-turbulent) edge of a smoke-filled turbulent boundary layer on a 16-mm film. With the same setup, Head and Graham (1969) looked at the interaction of trailing edge vortices and near-wall coherent eddies. A more detailed discussion on the implications of this structure visualisation method (with particular focus on hairpin/horseshoe vortices) is present in Head and Bandyopadhyay (1981). Recently, the data obtained from direct and large-eddy simulations are increasingly used to generate such movies, which are often part of refereed article submissions and form an integral part of scientific contributions.

Two types of costs are associated with the generation of animations: CPU time and storage. While generating 2D contours is reasonably inexpensive, the rendering and the encoding of the animation of the evolution of a turbulent field through isosurfaces, at a frame rate comparable to the time resolution of the computation, can take longer than the simulation itself. For example, one Navier–Stokes solver step for a Large Eddy Simulation (LES) with 1.6 million points requires, with the code used in this work, 1.83 s on a 2.8 GHz AMD Opteron®; a call to the Matlab® isosurface routine for the same flow field and on the same machine takes (including rendering time) 3.2 s for one frame; its cost scales linearly with the total number of grid points. Storage can also become an important issue. If in a numerical simulation only the three velocity components and the pressure field are stored (in double precision), one of the largest direct numerical simulations (DNS) computations of a boundary layer (by Schlatter and Örlü 2010) would require 96 GB for every instantaneous field. Possible alternatives to mitigate these issues are to decrease the spatial resolution of the data set, the frame rate or the temporal accuracy. The first approach leads to undesirable loss in spatial accuracy and unphysical smoothing of the data. The second approach can result in a poor kinematic description of the coherent motions and, also, in a visually unattractive movie. The third is the most commonly used one, and the one we will be considered here.

For a flow visualisation to be effective in terms of the accuracy with which the data are represented the interval between successive frames must be a fraction of the integral scale of the structure being observed. Taking a turbulent channel flow as an example, at a Reynolds number accessible by DNS, the integral time scale of eddies in the buffer layer is $T_{in} = \mathrm{o}(y/u_\tau)$ (where y is the distance from the wall and $u_\tau = (\tau_w/\rho)^{1/2}$ is the friction velocity, where τ_w is the wall-shear stress and ρ is the fluid density), while in the outer reaches of the logarithmic layer $T_{out} = \mathrm{o}(\delta/U_b)$ where δ is the channel half-width and U_b is the average flow velocity. The time-step, which is limited by the Courant-Friedrichs-Levy (CFL) condition in simulations of this type, is of the order of $\Delta t = \mathrm{o}(U_b/\Delta x)$. For a DNS of channel flow at $Re_\tau = 400$ (Moser *et al.* 1999), assuming a (maximum) interval between frames of one-tenth of the T_{out}, the number of time frames required to reproduce accurately the structure evolution in one flow-through time, $T_{fthr} = 2\pi\delta/U_b$, would be 10 T_{fthr}/T_{out}. The number of complete flow fields to be stored would be approximately 63, resulting in approximately 128 GB of storage.

The need for visual smoothness (the difference between successive frames must be small enough not to be detectable by the human eye) may increase this requirement. Zooming on a small region of the flow, for instance, would amplify the relative motion of the structures, and require that more finely spaced flow realisations be stored.

CPU and storage issues are both addressed in the work by Jobard *et al.* (2002) who propose a method for advecting textures with an assigned 2D flow field. Images are seeded with particles that are then tracked at the pixel level with a hybrid Lagrangian and Eulerian scheme. The velocity field is contaminated by a low percentage of noise in order to compensate for pixel duplication and spurious flow divergence effects. Noise has two distinct roles: it is advected (in parallel with the textures) and re-inserted in the Lagrangian coordinate integration step and, just for post-processing purposes, is used in the blending of frames in order to achieve high temporal correlations, therefore, visual pleasantness of the animations. The issues with rendering time are not present since advection is carried out at the pixel level. This method results in savings of both computational resources (since the rendering is only carried out on the coarsely spaced fields) and storage, however, its effectiveness in turbulent flows is unknown.

In this article, we propose an advection-based method inspired by the work of Jobard *et al.* (2002). We choose, however, to advect physical quantities (the variables to be plotted) rather than the pixels. Thus, our approach does not decrease the CPU time required by a visualisation, but only the storage required. We hope, however, that the physical knowledge on the structure of turbulence that can be used with our approach will compensate the higher cost of our method. Our goal is to develop a method that allows the generation of visually pleasing animations, using time-coarsened data sets. Given a series of flow realisations that are sufficiently closely spaced (in time) to represent accurately the flow-field evolution, intermediate fields can be reconstructed that result in a visually smooth evolution of the structures. We

emphasise again that the coarsening can only be carried to the level that allows the accurate description of the flow physics; the aim of this method is to separate this requirement from that of visual smoothness.

Some difficulties may arise when trying to assess the perceptual quality, rather than the physical accuracy, of animations capturing the evolution of turbulent structures. An appropriate metric is needed to address the aforementioned quality issues. Simple point-to-point comparison between a reference (distortion-free) field and its approximate reconstruction can be misleading as shown in Figure 1. The original frame (a) has been distorted in five different ways (Figure 1b–f), yet achieving the same root-mean square (rms) difference (of the 8-bit greyscale field) with respect to the original image. Images (b) and (c) clearly have a higher perceived quality than others. The loss of information in the compressed image (d) and the blurring in image (e) diminishes the perceived quality. An average quality score should be given to image (f), which was perturbed with scattered black and white pixels, but still shows a good structural integrity. Image defects, therefore, have to be weighted based on their visibility to the human visual system (HSV) in the global quality assessment criteria.

The human eye is highly adapted to extract structural information from the viewing field and the Structural Similarity Index (SSIM) has been designed to serve specifically for this purpose (Wang *et al*. 2004). To improve the relevance of our quality assessment of the generated animation, the first step is to focus the analysis on the 2D bitmap image representation of the 3D data, via the SSIM index, rather than the reconstructed 3D data itself.

In the following, we will describe the numerical method used to generate the data, and the metrics used to evaluate image quality and smoothness. We will then apply the proposed method to turbulent channel flow and discuss the results, and draw concluding remarks.

2. Problem formulation

2.1. Data set description

The available set of instantaneous velocity fields is taken from a large-eddy simulation of a turbulent channel flow at $Re_B = U_B\,\delta/\nu = 5000$ and 11,000 based on the channel half-width δ and the streamwise average velocity U_B, corresponding to $Re_\tau = u_\tau\delta/\nu = 300$ and 600, respectively. The filtered equations of conservation of mass and momentum are:

$$\frac{\partial u_i}{\partial x_i} = 0 \tag{1}$$

$$\frac{\partial u_j}{\partial t} + \frac{\partial u_i u_j}{\partial x_i} = -\frac{\partial p}{\partial x_j} - \frac{\partial \tau_{ij}}{\partial x_i} + \frac{1}{Re_B}\frac{\partial^2 u_j}{\partial x_i\,\partial x_i} \tag{2}$$

Figure 1. Comparison of 'Boat' images with different types of distortions, all with MSE = 210. (a) Original image (256 × 256 resolution at 8 bits/pixel), (b) contrast-stretched image, SSIM = 0.9168, (c) mean-shifted image, MSSIM = 0.9900, (d) JPEG-compressed image, SSIM = 0.6949, (e) blurred image, SSIM = 0.7052, (f) salt–pepper impulsive noise contaminated image, SSIM = 0.7748. Reproduced with permission from Wang *et al*. (2004).

where x_1, x_2 and x_3 (or x, y and z) are, respectively, the streamwise, wall-normal and spanwise directions, and u_i the filtered velocity components in those directions. The equations are solved in discretised form using a well-validated second-order-accurate finite difference code (Keating *et al.* 2004), based on a staggered grid arrangement. The code is parallelised using the Message Passing Interface (MPI) protocol. The sub-grid scale model used is the classic Plane-Averaged Dynamic Eddy-Viscosity model (Germano *et al.* 1991, Lilly 1992). Periodic boundary conditions are applied to all quantities in the streamwise and spanwise direction. The velocity field obeys no-slip conditions at the walls.

At $Re_B = 5000$, two-grid resolutions are used, with $64 \times 64 \times 64$ and $128 \times 96 \times 128$ points, respectively, in the streamwise, normal-to-wall and spanwise directions. Only the finer mesh is used at the higher Re_B. The resolution in wall units of the fine $Re_B = 5000$ calculation is $\Delta x = 14^+$, $\Delta z = 7^+$ and $\Delta y_{min} = 0.7^+$; the coarse Re_B 5000 and $Re_B = 11{,}000$ calculations share the same resolution of $\Delta x = 28^+$, $\Delta z = 14^+$ and $\Delta y_{min} = 1.4^+$.

The equations of motion were integrated in time until a statistical steady state is reached; statistics are then accumulated for approximately 10 large-eddy turnover times (LETOTs) δ/u_τ. The mean velocity and Reynolds stresses were found to be in good agreement with the reference data (Moser *et al.* 1999), with maximum differences of 6%. Realisations of the flow field were then stored, spaced by 0.1 δ/U_B, equal to 0.005 LETOTs, or $1.5\,\nu/u_\tau^2$ (viscous time units) at $Re_B = 5000$ and $3\,\nu/u_\tau^2$ at $Re_B = 11{,}000$.

2.2. *Advective-based interpolation technique*

We intend to reproduce the evolution of turbulent structures using only a subset of all of the available closely spaced realisations with $T = 0.1\ \delta/U_B$. We will refer generally to ϕ as any scalar field used for visualisation purposes. Each of these quantities calculated from the ith realisation is indicated with $\phi^i = \phi(x,t_i)$, where $t_i = (i-1)T$. We will coarsen the set of realisations by an (integer) factor n so that, in the coarsened space, the interval between two consecutive realisations is $t_{i+n} - t_i = nT$, and the corresponding scalar fields are ϕ^i and ϕ^{i+n}. For $n = 6$, the spacing in time is equal to $1/30^{\text{th}}$ of an LETOT.

One way to reconstruct the evolution of ϕ is an advection-based interpolation method:

$$\frac{\partial \phi}{\partial t} + \mathbf{V}_t \cdot \nabla \phi = \frac{D_t \phi}{Dt} = 0 \tag{3}$$

where $\mathbf{V}_t$ is an advection velocity that needs to be specified (numerical experiments have shown that the best choice for $\mathbf{V}_t$ is the local mean velocity). The basic physical assumption underlying Equation (3) is that between one realisation and the other the evolution of ϕ is governed only by mean flow transport, neglecting, among other things, transport by turbulent fluctuations and diffusion.

It is now possible to predict ϕ at time $t_i < t < t_{i+n}$, based on the available information at times t_i and t_{i+n}, by advancing Equation (3) in time first forward, from t_i to t, then backwards, from t_{i+n} to t, and taking a weighted average of the two solutions:

$$\phi^*(\mathbf{x}, t) = (1 - W_i(t))\phi^+ + W_i(t)\phi^- \tag{4}$$

where

$$W_i(t) = \frac{t - t_i}{nT} \tag{5}$$

is a weighting function. The forward and backward solutions of the advection equations

$$\phi^+ = \int_{t_i}^{t} \frac{D_t \phi}{Dt} dt + \phi^i, \qquad \phi^- = \int_{t_{i+n}}^{t} \frac{D_t \phi}{Dt} dt + \phi^{i+n} \tag{6}$$

can be approximated, to first order, by

$$\phi^+ \simeq \phi^i(\mathbf{x} - (t - t_i)\mathbf{V}_t) \tag{7}$$

$$\phi^- \simeq \phi^{i+n}(\mathbf{x} + (t_{i+n} - t)\mathbf{V}_t) \tag{8}$$

where the right-hand-sides of Equations (7) and (8) can be conveniently computed numerically using tri-linear interpolation.

2.3. *Performance metrics*

Given that the focus of this article is on the flow visualisation, the quality assessment (accuracy and smoothness evaluation) of the resulting animations must be based on the comparison between the original and reconstructed images (respectively, indicated as unstarred and starred quantities), rather than actual 3D data. We will indicate the projection operation of the 3D original, ϕ, and reconstructed, ϕ^*, fields onto a 2D image as follows:

$$\phi(\mathbf{x}, t) \to \mathcal{I}(t), \qquad \phi^*(\mathbf{x}, t) \to \mathcal{I}^*(t) \tag{9}$$

where by cursive letters, such as $\mathcal{I}(t)$, we will concisely indicate the collection of red, green and blue (RGB) values for each pixel (three monochromatic images) composing the final colour frame of the animation at time t. The accuracy of the animations will, therefore, be determined by comparing the reconstructed image $\mathcal{I}^*(t)$ with the corresponding frame from the original

data $\mathcal{I}(t)$, that is to say, by evaluating some appropriate metric

$$A = A(\mathcal{I}^*(t), \mathcal{I}(t)). \tag{10}$$

The smoothness of the animations, on the other hand, is a self-referenced quality measurement, and it can be evaluated by comparing two successive frames, at t and $t + \Delta t$, of the reconstructed animations

$$S = S(\mathcal{I}^*(t), \mathcal{I}^*(t + \Delta t)). \tag{11}$$

As mentioned in Section 1, the development of appropriate quality measurements is critical, since the pixel-to-pixel difference may be misleading. In this work, the norms in Equations (10) and (11) are evaluated by a locally weighted and image-averaged version of the SSIM (Wang *et al*. 2004). The SSIM was developed to assess the quality of an image, with respect to a reference one, based on the degradation of structural information. Pixels can exhibit strong correlations, especially when spatially proximate. Apart from detecting simple differences between two images, such as different levels of luminance and contrast, the SSIM index compares normalised distribution patterns of pixel intensities in order to extract purely structural information. This parameter is very useful in our case because we want to focus on structural information change in the image rather than a simple norm of a pixel-to-pixel difference. This makes the SSIM index ideal for assessing both the perceived quality and the accuracy of the reproduced animations of coherent structures.

The SSIM index can be defined by considering two monochromatic images, $\mathcal{X}$ and $\mathcal{Y}$ of the same size. The symbols represent the arrays of pixels each with its own luminance intensity

$$\mathcal{X} = \{x_1, x_2, \ldots, x_N\}; \qquad \mathcal{Y} = \{y_1, y_2, \ldots, y_N\} \tag{12}$$

where N is the total number of pixels in the images. For example, for 8-bit greyscale images, values of x_k, y_k range from 0 to 255. The statistical (unbiased) estimates of the mean intensity, the signal contrast and the correlation between the two images $\mathcal{X}$ and $\mathcal{Y}$ are, respectively

$$\mu_x = \frac{1}{N}\sum_{k=1}^{N} x_k \tag{13}$$

$$\sigma_x = \left[\frac{1}{N-1}\sum_{k=1}^{N}(x_k - \mu_x)^2\right]^{1/2} \tag{14}$$

$$\sigma_{xy} = \frac{1}{N-1}\sum_{k=1}^{N}(x_k - \mu_x)\big(y_k - \mu_y\big). \tag{15}$$

Taking into account the differences in luminance, contrast and image structure (the former measured by the normalised correlation σ_{xy}), we obtain the final form of the SSIM index (Wang *et al*. 2004)

$$\mathrm{SSIM}(\mathcal{X}, \mathcal{Y}) = \frac{\big(2\mu_x\mu_y + C_1\big)\big(2\sigma_{xy} + C_2\big)}{\Big(\sigma_x^2\sigma_y^2 + C_1\Big)\Big(\mu_x^2\mu_y^2 + C_2\Big)} \tag{16}$$

where C_1 and C_2 are stabilising constants (to avoid divisions of zero by zero) taken as $C_1 = (0.01L)^2$ and $C_2 = (0.03L)^2$, where L is the dynamic range of the pixel values (255 for 8-bit greyscale images). For image-quality assessment, it is useful to apply the SSIM index locally rather than globally by adopting a windowing approach across the whole image. The image statistics Equations (13), (14) and (15) are locally estimated using a 11 × 11 circular-symmetric Gaussian weighting function instead of simple arithmetic average. This leads to the definition of a mean SSIM index (MSSIM) used to evaluate the overall image quality (Wang *et al*. 2004)

$$\mathrm{MSSIM}(\mathcal{X}, \mathcal{Y}) = \frac{1}{M}\sum_{j}^{M}\mathrm{SSIM}\Big(\{\mathcal{X}\}_j, \{\mathcal{Y}\}_j\Big) \tag{17}$$

where $\{\mathcal{X}\}_j$ and $\{\mathcal{Y}\}_j$ are the image contents at the local j-th window of the reference and the distorted image, respectively, with Gaussian weighting. For colour images, it will be implied that such mean index is averaged over the three RGB channels, separately treated as monochromatic images. The accuracy and smoothness indices (10) and (11) will, therefore, be calculated as

$$A = \mathrm{MSSIM}(\mathcal{I}^*(t), \mathcal{I}(t)) \tag{18}$$

$$S = \mathrm{MSSIM}(\mathcal{I}^*(t), \mathcal{I}^*(t + \Delta t)) \tag{19}$$

3. Results

We will now perform a quantitative assessment of the accuracy and smoothness of the reconstructed animations; two additional corrections are found to be necessary to improve the animation quality; they will be presented and their effects are discussed. The effects of changes in Reynolds number, spatial resolution and visualisation methods are discussed as well as results from a brief cost analysis.

In Figure 2, three animation frames reconstructed with the advection-based technique (Equation (7)) are shown for different values of the spacing n for the fine $Re_B = 5000$ calculation. The frames are all at a time, t_f, at the midpoint of the coarsened interval, which extends from the $t_{f-n/2}$ to $t_{f+n/2}$. We consider

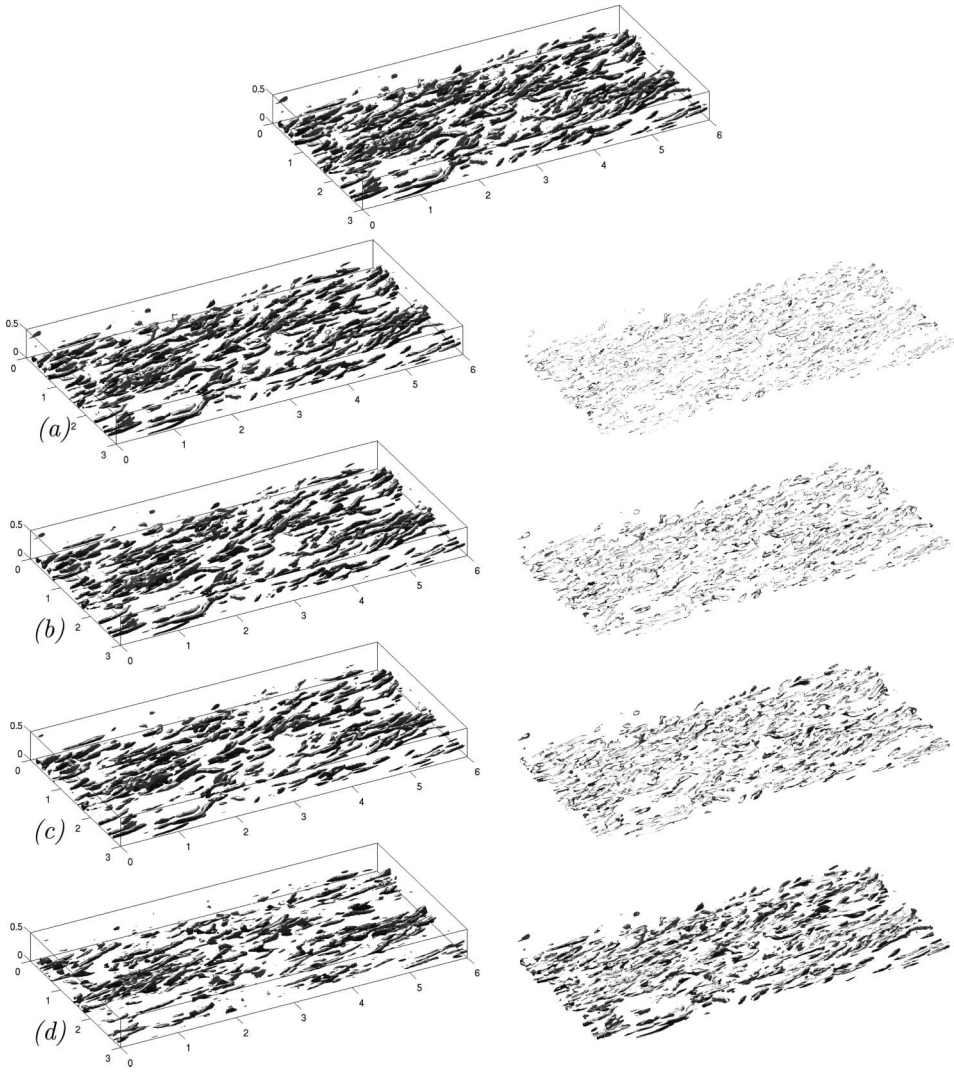

Figure 2. Reconstructed frame (left) and corresponding pixel-to-pixel difference (right) with respect to original frame for highly resolved $Re_B = 5000$ case and $Q = 3.0\,U_B^2/\delta^2$. Top: reference image. Reconstructed image for (a) $n = 2$ ($A = 0.90$), (b) $n = 4$ ($A = 0.85$), (c) $n = 6$ ($A = 0.80$) (d) $n = 20$ ($A = 0.70$).

isosurfaces of the second invariant of the velocity gradient tensor,[1]

$$Q = -\frac{1}{2}\frac{\partial u_i}{\partial x_j}\frac{\partial u_j}{\partial x_i}. \tag{20}$$

The sensitivity of the reconstruction to the realisation spacing n is highlighted by the pixel-to-pixel image difference between the reconstructed and original frames, which grows in complexity as n increases (more detail of the original frame is loss). The accuracy index decreases from $A = 0.90$ to $A = 0.70$ for n increasing from 2 to 20. Major differences are evident mostly around the edges of structures distorted by turbulent straining (which is not included in the model). The prediction of the location of larger structures remains rather accurate even for large values of n, as they are less sensitive to turbulent diffusion, and are dissipated at a slower rate, which results in qualitative agreement between the reconstructed and reference frame. For all but the coarsest ($n = 20$) reconstruction, however, the general agreement between the fields is reasonable, and the first three images convey very similar information. For visualisation purposes, an accuracy index $A > 0.8$ appears to be acceptable.

The movie accuracy and smoothness indices as function of time are shown in Figure 3. The accuracy index returns to unity every n frames, when an actual flow realisation is used, and there is no difference between $\mathcal{I}$ and $\mathcal{I}^*$. The accuracy decreases for reconstructed frames that are more distant (in time)

from the reference ones, reaching the lowest value at the centre of the time interval $t_{i+n} - t_i$. The overall accuracy decreases with n as expected. As mentioned above, for $n \leq 6$, the reconstructed frame contains at least 80% of the information contained in the original one. The regular and abrupt variations in the accuracy index have a striking effect to an observer of the correspondent animation, as movie smoothness becomes an issue. Note that, for $n = 6$, the structures visualised here (which are typically located at $y^+ = 50-150$) move by approximately one-tenth of the channel length (or approximately 10–13 grid points) between two actual realisation. If the advective interpolation were not used, the smoothness of the movie would be appreciably decreased, to a value of approximately $S = 0.68$, 0.660, 0.655, 0.65 for $n = 2$, 4, 6 and 20, respectively. This results in an unpleasant movie where the under-resolved evolution of structures cannot be tracked by the human eye. Also note that, even when a very short interval between frames is used, smoothness levels very close to 1 cannot be achieved. If, for instance, we consider two realisations separated by one-tenth of the separation used here (i.e. 3×10^{-3} LETOTs), the smoothness index remains at 0.88. This rapid decrease is due to many factors, such as the interpolation required to calculate velocity gradients at the cell centres from their staggered locations, and the details of the rendering algorithm, which may be affected by small changes in the underlying field. From a visual inspection of the movies, the smoothness indices greater than approximately 0.7 are desirable.

For given n, at times coinciding with one of the reference snapshots (at time t_i or t_{i+n}), the interpolation method (Equation 4) returns the original, non-interpolated frame. This results in an abrupt increase of detail in the animation, given the presence of small structures being rendered. To remove this problem, we introduced a frame-substitution method in which the two interpolated fields $\phi^*(x,t_i - \Delta t)$ and $\phi^*(x,t_i + \Delta t)$ are used as reference frames to recalculate $\phi^*(x,t_i)$ via Equation (4) resulting in $\phi^*(\mathrm{x},t_i) \neq \phi^i(\mathrm{x},t_i)$. The total variation of the accuracy history is, therefore, reduced (Figure 4).

The periodic rise in accuracy and smoothness between the reference realisations can, however, still be seen in Figure 4, which may be appreciable by a human observer in the animations, especially for large n: the evolution of the isosurfaces appears to change

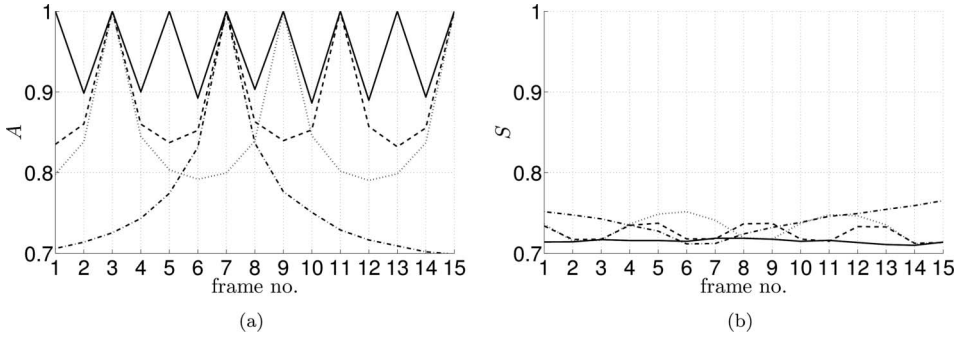

Figure 3. Reconstructed-movie accuracy (a) and smoothness (b). Highly resolved $Re_B = 5000$ case with isovalue of $Q = 3.0\, U_B^2/\delta^2$, for $n = 2$ (——), $n = 4$ (– – –), $n = 6$ (. . . . .) and $n = 20$ (– · –).

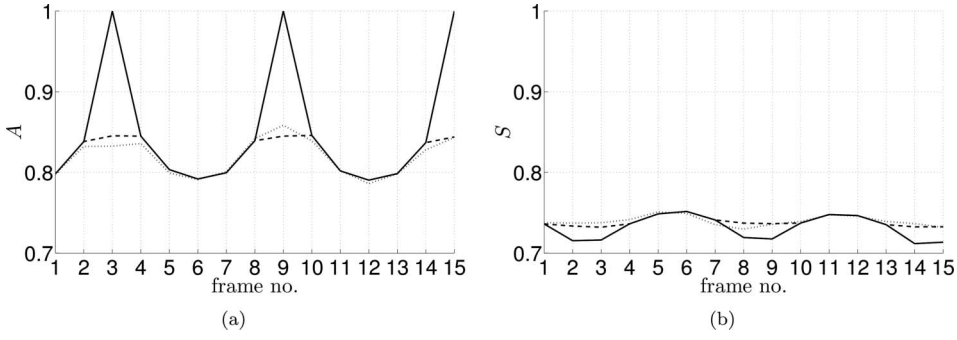

Figure 4. Reconstructed-movie accuracy (a) and smoothness (b). Highly resolved $Re_B = 5000$ case with isovalue of $Q = 3.0\, U_B^2/\delta^2$, for $n = 6$ without frame-substitution and $\varepsilon = 0$ (——), with frame-substitution and $\varepsilon = 0$ (– – –) and with frame-substitution and $\varepsilon = 0.05$ (. . . . .).

with a regular, periodic, pattern between the two reference realisations. To compensate for this unwanted effect, a stochastic component, similar to the noise-blending technique devised by Jobard *et al.* (2002), was added to the weighting function

$$W_i^*(t) = \frac{\frac{t-t_i}{nT} + \varepsilon\mathcal{N}(0,1)}{1+\varepsilon} \tag{21}$$

where $\mathcal{N}(0,1)$ is a normally distributed Gaussian variable and $\varepsilon = 0.050$. This value was chosen as the optimal value, since this noise amplitude level was strong enough to eliminate the spurious periodicity of variations in smoothness and accuracy but not to affect significantly the overall accuracy of the animation (Figure 4a). An observer of the correspondent animations can appreciate the absence of unpleasant periodic variations in the movie quality as well as the (inevitable) slight deterioration of the overall accuracy due to the introduction of noise. Determining the optimal value for ε results, once again, in a trade-off between accuracy and visual pleasantness of the animation. Overall, the introduction of noise can be considered as a second-order adjustment with respect to the frame-substitution method.

It is to be expected that the reconstruction method proposed here will depend greatly on the variable being plotted, and on the resolution. Turbulent quantities that are representative of the large-scale motions will be represented more accurately by the simple advection-based model. Figure 5 shows contours of u' and v' fluctuations and isosurfaces of $Q = 2$ and $Q = 3$ for the same flow realisation. First, we observe the larger scale of the u' fluctuations, compared to v'. In Figure 5(c) and (d), on the other hand, we note how for lower levels of Q more eddies are observed, especially in the buffer layer. The proposed model cannot be expected to be as accurate when reconstructing near-wall structures, as it is for outer-layer eddies; similarly, the smaller eddies present in the contours of v' cannot be reconstructed as well as the larger streaky structures. A quantitative assessment of the sensitivity of the proposed method to the structure size is shown in Figures 6 and 7, in which we see that a reduced structure size has a beneficial effect on both the accuracy (Figure 6a) and the smoothness (Figure 7a). The same test has been repeated for pressure isosurfaces (Figures 6b and 7b), which highlight bigger structures. In this case, increasing the threshold also results in a sensible reduction of the size of the visualised structures. Also note that, in the case of Q isosurfaces, the proposed interpolation method will naturally decrease the population of visualised structures (Figure 2); therefore, in order to have a reconstructed movie with a similar number of visualised turbulent structures, a lower level of the threshold need to be chosen with respect to the original movie. All of the considerations made above remain unchanged when introducing noise in the weighting function (Equation 21).

Other parameters that affect the quality of the proposed reconstruction are the grid resolution and Reynolds number that directly determine the visualised structure size. The reconstruction of a flow animation

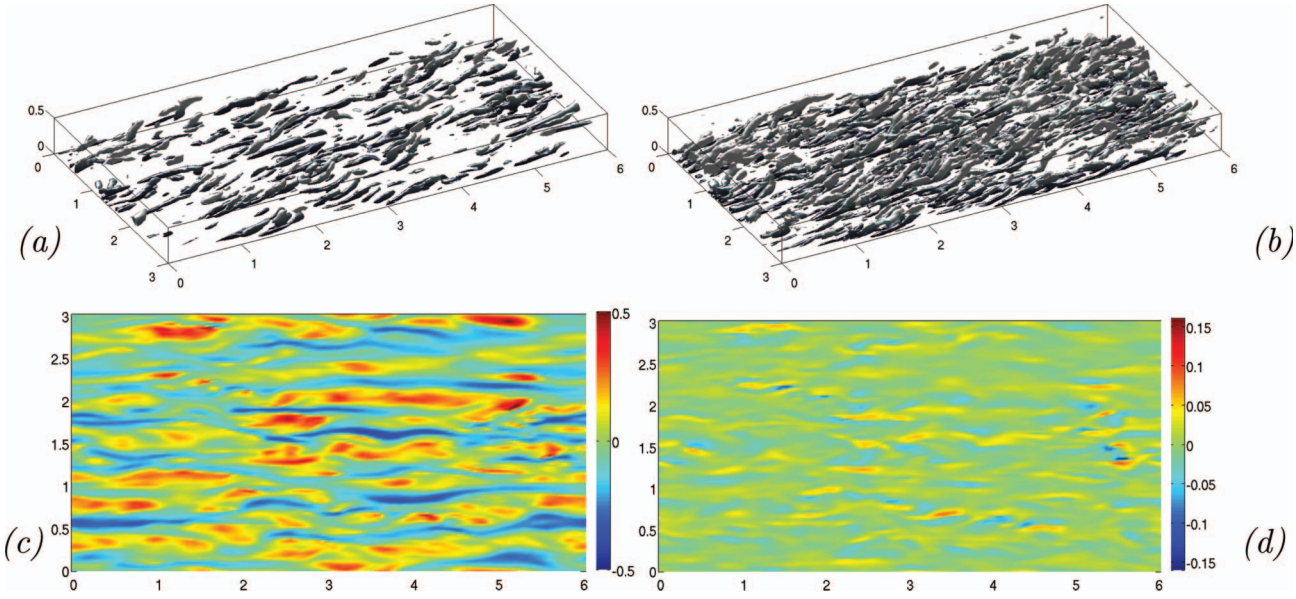

Figure 5. Highly resolved $Re_B = 5000$ case; isosurfaces of $Q = 3.0\,U_B^2/\delta^2$ (a) and $Q = 2.0\,U_B^2/\delta^2$ (b), contours of streamwise (c) and wall-normal (d) fluctuating velocity.

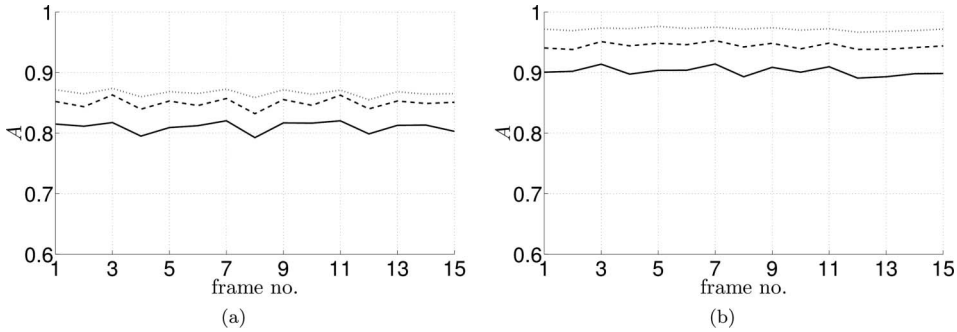

Figure 6. Comparison of accuracy for Q-isosurfaces (a) and p' -isosurfaces (b) of advection-based interpolated movie for $\varepsilon = 0$ and $n = 4$ and frame replacement active. In (a) $Q = 2.0\,U_B^2/\delta^2$ (——), $Q = 3.0\,U_B^2/\delta^2$ (– – –) and $Q = 4.0\,U_B^2/\delta^2$ (. . . . .); in (b) $p\prime = -0.010\,\rho U_B^2$ (——), $p\prime = -0.015\,\rho U_B^2$ (– – –) and $p\prime = -0.020\,\rho U_B^2$ (. . . . .).

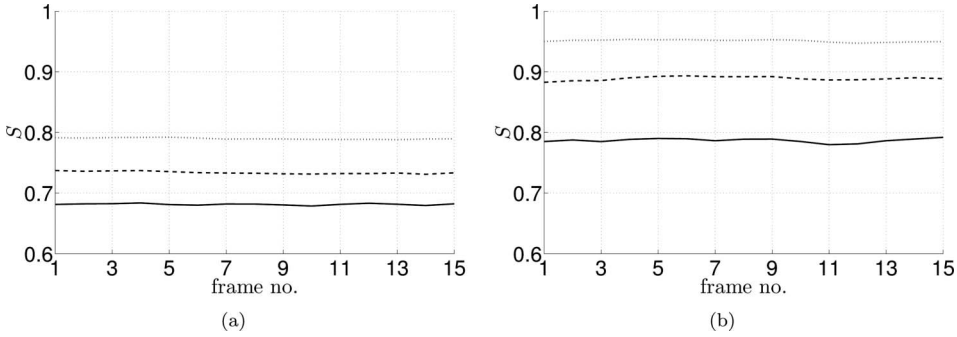

Figure 7. Comparison of smoothness for Q-isosurfaces (a) and p' -isosurfaces (b) of advection-based interpolated movie for $\varepsilon = 0$ and $n = 4$ and frame replacement active. In (a) $Q = 2.0\,U_B^2/\delta^2$ (——), $Q = 3.0\,U_B^2/\delta^2$ (– – –) and $Q = 4.0\,U_B^2/\delta^2$ (. . . . .); in (b) $p\prime = -0.010\,\rho U_B^2$ (——), $p\prime = -0.015\,\rho U_B^2$ (– – –) and $p\prime = -0.020\,\rho U_B^2$ (. . . . .).

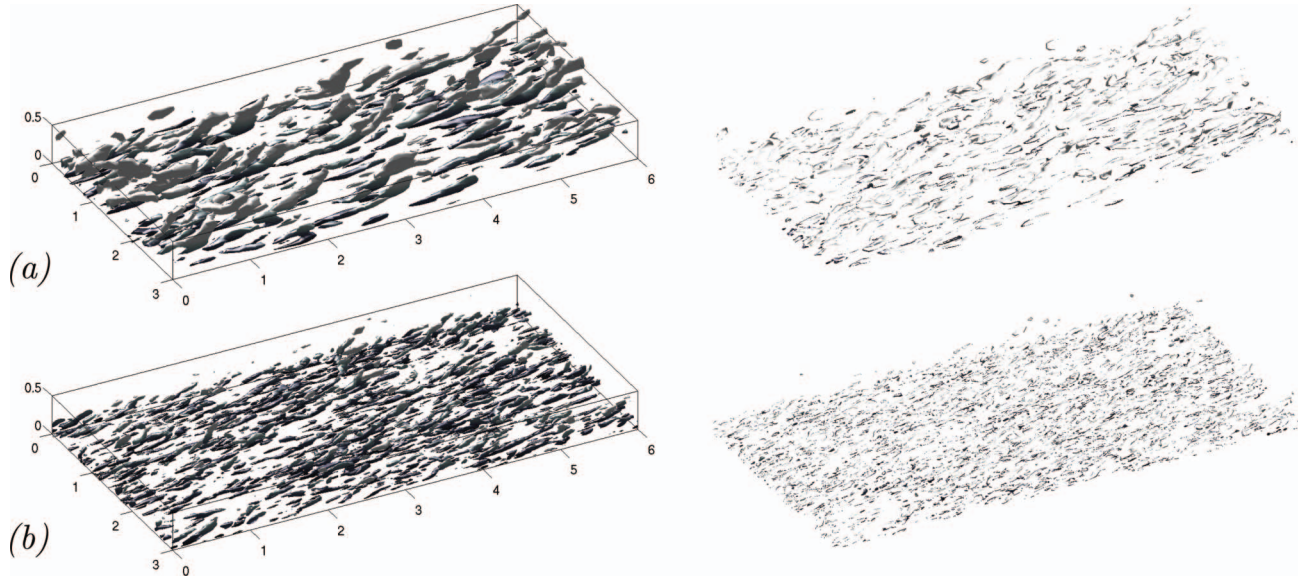

Figure 8. Reconstructed frame (left) and corresponding pixel-to-pixel difference with respect to original frame (right) showing effects of Reynolds number on image accuracy. Low resolution $Re_B = 5000$ calculation (a) $A = 0.842$, $n = 4$ for $Q = 1.0\,U_B^2/\delta^2$ and $\varepsilon = 0$; $Re_B = 11{,}000$, $A = 0.845$, $n = 2$ (b) for $Q = 4.0\,U_B^2/\delta^2$ and $\varepsilon = 0$.

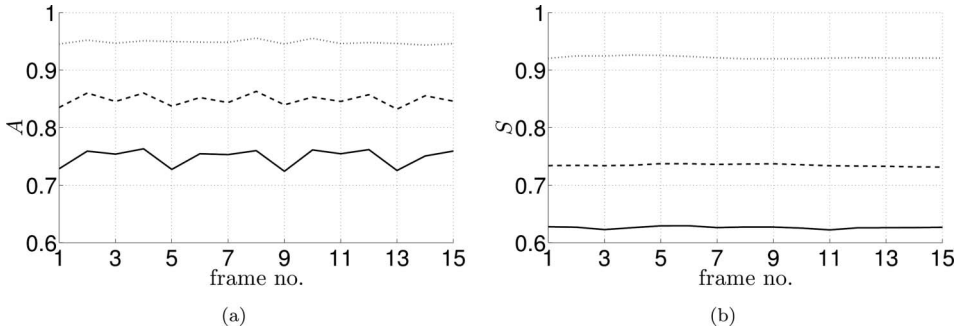

Figure 9. Effect of Reynolds number and grid resolution on accuracy (a) and smoothness (b). Coarsely resolved $Re_B = 5000$ (. . . . .), finely resolved $Re_B = 5000$ (– – –), low resolution $Re_B = 11{,}000$ (——) for $\varepsilon = 0.0$, $n = 4$ and $Q = 3.0\,U_B^2/\delta^2$.

at higher Reynolds numbers, or with higher spatial resolution, is more challenging since the evolution of a wider range of structure sizes must be predicted all at once. For example, in order for the $Re_B = 11{,}000$ animation to have an accuracy comparable to the $Re_B = 5000$ one (at the same resolution in wall units), the distance between the realisations must be reduced and/or, in case of Q-isosurfaces, the threshold for the higher Reynolds number data has to be increased (reducing the population of structures). For example, in order for a reconstructed frame from the higher Reynolds number case to retain the same accuracy as the lower Reynolds number counterpart, the value of n had to be halved and the value of Q increased from 1 to 4 to reduce the population of structures (Figure 8). This shows that the method performs better with larger resolved eddies in the flow that, also, evolve over longer time scales, therefore, exhibit more predictable kinematics.

Figure 9 quantifies the aforementioned effects allowing an assessment of Reynolds number effects (for a fixed resolution) and resolution effects (for a fixed Reynolds number). Increasing the grid resolution results in a decrease in the accuracy (Figure 9a), which is expected because the domain becomes more densely populated. Increasing the Reynolds number, keeping the resolution (in wall units) fixed, causes an even larger drop in accuracy. In both cases, the visualised resolved structures are much finer and a high correlation between the reconstructed and reference images is more difficult to obtain. The method fails to correctly reproduce both the actual and perceived evolution of the smaller structures. The smoothness index exhibits the same qualitative dependency from Reynolds number and resolution (Figure 9b). Lower values of smoothness for the $Re_B = 11{,}000$ case (or highly resolved $Re_B = 5000$ case) are related to the higher density of structures that will inevitably decrease the correlation between two consecutive frames. Lower values of smoothness are consistent with the higher total variation of the accuracy time history.

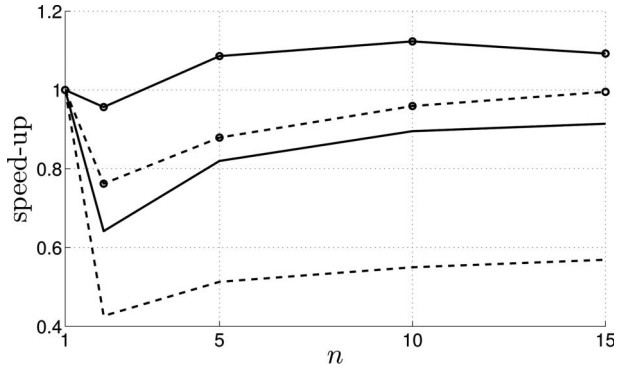

Figure 10. Movie generation speed-up for coarsely resolved $Re_B = 5000$ (– – –) and finely resolved $Re_B = 5000$ (——) cases, $Q = 1.0\,U_B^2/\delta^2$ (with circles), $Q = 3.0\,U_B^2/\delta^2$ (without circles), for frame substitution and stochastic weighting with $\varepsilon = 0.05$. Speed-up is defined as the overall time necessary to generate a non-interpolated movie over the time required to generate the same movie by interpolating every n frames.

The distance from the wall is also an important factor affecting the visualised structure size and characteristic time scale, not only for the isosurfaces but especially for contour plots. Accuracy can drop by 7% from the channel's centreline to the rms peak location for u' and v' animations (not shown). This is due to the error made when approximating the structure transporting velocity $\mathbf{V}_t$ with local average one in Equation (3) in the regions of the flow where the velocity variance is higher. Also, visualising smaller structures, with respect to u' contours, such as contours of v' or u' v' can cause accuracy to drop by 2 and 9%, respectively.

3.1. Cost analysis

The main advantage of the proposed frame-reconstruction technique is to reduce the data storage requirements while preserving, within acceptable

levels, the visual pleasantness and accuracy of the generated animations. However, while reductions in storage requirements, in the proposed technique, increase linearly with the snapshot spacing n, the total computational cost may increase or decrease depending on the data size, the structure eduction method and other tuneable parameters of the technique.

The total movie-generation time is primarily taken up by the advection-based interpolation, image rendering and I/O activity. Figure 10 shows the speed-up in generating Q-isosurface animations for the complete advection-based frame-reconstruction technique, including noise and frame substitution. The latter introduces a moderate computational overhead (not shown), vanishing with n, since one extra interpolation step is needed every n frames, for visual pleasantness purposes. High values of Q yield animations with fewer isosurfaces and faster rendering times. In this case, the total movie generation time is taken up by the advection-based interpolation time which increases with n, causing an overall loss in performance (speed-up is lower than 1). However, the computational overhead of the reconstruction method stays below 20% for $n > 3$ for the fine-grid case, being the one that benefits the most from the reduced I/O activity. The generation time for the coarse data set animation at high Q is also limited by the 3D interpolation, having very little I/O activity, resulting in a significant performance loss. In this case, however, very few structures are visualised.

On the other hand, low values of Q result in rendering-intensive movies (flow field is highly populated with structures) which benefit from the method primarily due to the smoothing of the Q-field, caused by the interpolation, which decreases the number of visualised structures, and, therefore, the rendering time. Speed-up finally becomes greater than one when I/O activity becomes important for larger data sets.

In all cases, the movie-generation speed-up increases with the data set size suggesting that the proposed method can be cost-effective for large-scale simulations where I/O time becomes the limiting factor for movie-generation time.

4. Conclusions

We have implemented and tested a simple technique to predict the evolution of coherent structures in between two given realisations of a turbulent flow. The present work brings together, for the first time, frame reconstruction methods used in texture advection techniques (extending them to 3D instantaneous turbulent flow fields) and image quality metrics, based on the receptivity of the HSV, into one simple visualisation technique designed with the aid of basic knowledge of turbulent flows. The overall objective was to generate high-quality animations of turbulent flows at a reduced cost especially for large scale computations.

The structures were identified as isosurfaces of positive values of Q, pressure isosurfaces and with flooded contour plots of velocity fluctuations. The prediction method used is a modified linear interpolation technique based on a mean velocity advection model. The accuracy of the resulting animations is assessed via the SSIM index by direct comparison with the original data set. The degree of similarity of adjacent frames was used as a quantitative measure of the perceived smoothness of the resulting animations. A frame-substitution method was implemented to avoid the abrupt transition between an interpolated and non-interpolated frame, which compromises the smoothness of the visualisation. The latter can be increased by increasing the spacing between reference realisations, n. This results, however, in an unrealistic streamwise stretching of the structures occurring periodically (with period n) in the animations. To rectify this, the weighting function was perturbed with random noise. Both adjustments result in a decrease of accuracy which is, on the other hand, beneficial towards obtaining higher level of pleasantness of the animations. The required amplitude of the random noise perturbation to the linear weighting was found to be acceptable within the range of $\varepsilon = 0.01 - 0.1$.

It was found that with the proposed method, visualising fewer and bigger structures is beneficial towards the overall quality of the animations. For example, at higher Reynolds numbers, where the structures being visualised are finer, n needs to be much smaller in order to retain the same accuracy as a lower Reynolds number case. Changing the threshold value for Q and p' isosurfaces has the same effect on the resulting animations. The contours of velocity fluctuations achieve, in general, higher values of accuracy than isosurfaces of pressure or Q. The former are, however, sensitive to the distance from the wall. Reconstructed animations in the buffer layer have lower accuracy given the higher rms of the instantaneous transporting velocity than in the centreline. Finally, the introduction of noise is effective in reducing the periodic variations of accuracy and smoothness which may still be present with the frame-substitution method. However, its effects are of secondary importance and do not affect significantly the overall performance of the proposed method.

The primary limitations of the adopted reconstruction method lie in its sensitivity to the visualised structure size. Tracking the evolution of smaller structures, either due to high-Reynolds-number data

or higher resolution, requires finer spacing in time between the reference snapshots, reducing the cost-effectiveness of the method. Also, testing is required in more complex flows, with 2D and 3D mean velocity fields, where artificial diffusion effects are expected to be stronger. However, even in simple flows, such as the one investigated, due to the inherent inhomogeneity in the wall-normal direction, the accuracy of the reconstruction of the 3D Q field, for example, will vary with y. Therefore, improvement, such as making the noisy weighting function dependent of the local intensity of the turbulent activity is one option worth investigating. Finally, we have shown that the proposed technique is beneficial both in memory and in computational requirements in the case of rendering-intensive movies processed from large data sets. A computational overhead is, however, present when the data set size (I/O activity) is not a limiting factor.

Acknowledgements

The authors thank the High Performance Computing Virtual Laboratory (HPCVL), Queen's University site, for the computational support. GL acknowledges the support of NSERC under the USRA program. UP acknowledges the support of the Canada Research Chair program, and of NSERC, under the Discovery Grant program.

Note

1. From now on, all quantities reported are normalised by the constant density ρ, U_b and δ, unless otherwise stated.

References

Brown, G.L. and Roshko, A., 1974. On density effects and large structure in turbulent mixing layers. *Journal of Fluid Mechanics*, 64 (04), 775–816.

Dubief, Y. and Delcayre, F., 2000. On coherent vortex identification in turbulence. *Journal of Turbulence*, 1, 1–22.

Elsinga, G., *et al.*, 2006. Tomographic particle image velocimetry. *Experiments in Fluids*, 41, 933–947. doi: 10.1007/s00348-006-0212-z.

Germano, M., *et al.*, 1991. A dynamic subgrid-scale eddy viscosity model. *Physics of Fluids A*, 3, 1760–1765.

Haller, G., 2005. An objective definition of a vortex. *Journal of Fluid Mechanics*, 525, 1–26.

Head, M.R., 1969. *Flow visualization in the turbulent boundary layer (16 mm cine film)*. Cambridge: Cambridge University Engineering Department, Technical report.

Head, M.R. and Bandyopadhyay, P., 1981. New aspects of turbulent boundary-layer structure. *Journal of Fluid Mechanics*, 107, 297–338.

Head, M.R. and Graham, J.M.R., 1969. *Flow visualization of transition (16 mm cine film)*. Cambridge: Cambridge University Engineering Department, Technical report.

Jobard, B., Erlebacher, G., and Hussaini, M.Y., 2002. Lagrangian-Eulerian advection of noise and dye textures for unsteady flow visualization. *IEEE Transactions on Visualization and Computer Graphics*, 8 (3), 211–222.

Keating, A., *et al.*, 2004. Large-eddy simulation of heat transfer downstream of a backward-facing step. *Journal of Turbulence*, 5 (20), 1–27.

Kida, S. and Miura, H., 1998. Swirl condition in low-pressure vortices. *Journal of Physical Society of Japan*, 67 (7), 2166–2169.

Kline, S.J., *et al.*, 1967. The structure of turbulent boundary layers. *Journal of Fluid Mechanics*, 30 (4), 741–773.

Liepmann, H., 1979. The rise and fall of ideas in turbulence. *American Scientist*, 67, 221–228.

Lilly, D.K., 1992. A proposed modification of the Germano subgrid-scale closure method. *Physics of Fluids A*, 4, 633–635.

Moser, R.D., Kim, J., and Mansour, N.N., 1999. Direct numerical simulation of turbulent channel flow up to $\mathrm{Re}_\tau = 590$. *Physics of Fluids*, 11 (4), 943–945.

Robinson, S.K., 1991. Coherent motions in the turbulent boundary layer. *Annual Review of Fluid Mechanics*, 23 (1), 601–639.

Robinson, S.K., Kline, S.J., and Spalart, P.R., 1989. *A review of quasi-coherent structures in a numerically simulated turbulent boundary layer*. Washington, D.C.: NASA, TM 102191.

Schlatter, P. and Örlü, R., 2010. Assessment of direct numerical simulation data of turbulent boundary layers. *Journal of Fluid Mechanics*, 659, 116–126.

Wang, Z., *et al.*, 2004. Image quality assessment: from error measurement to structural similarity. *IEEE Transactions on Image Processing*, 13 (4), 600–612.