Experiments in Fluids (2006) 41: 265–278
DOI 10.1007/s00348-006-0168-z

RESEARCH ARTICLE

Neelakantan Saikrishnan · Ivan Marusic
Ellen K. Longmire

Assessment of dual plane PIV measurements in wall turbulence using DNS data

Received: 19 December 2005 / Revised: 29 April 2006 / Accepted: 14 May 2006 / Published online: 27 June 2006
© Springer-Verlag 2006

Abstract Experimental dual plane particle image velocimetry (PIV) data are assessed using direct numerical simulation (DNS) data of a similar flow with the aim of studying the effect of averaging within the interrogation window. The primary reason for the use of dual plane PIV is that the entire velocity gradient tensor and hence the full vorticity vector can be obtained. One limitation of PIV is the limit on dynamic range, while DNS is typically limited by the Reynolds number of the flow. In this study, the DNS data are resolved more finely than the PIV data, and an averaging scheme is implemented on the DNS data of similar Reynolds number to compare the effects of averaging inherent to the present PIV technique. The effects of averaging on the RMS values of the velocity and vorticity are analyzed in order to estimate the percentage of turbulence intensity and enstrophy captured for a given PIV resolution in turbulent boundary layers. The focus is also to identify vortex core angle distributions, for which the two-dimensional and three-dimensional swirl strengths are used. The studies are performed in the logarithmic region of a turbulent boundary layer at $z^+ = 110$ from the wall. The dual plane PIV data are measured in a zero pressure gradient flow over a flat plate at $Re_\tau = 1{,}160$, while the DNS data are extracted from a channel flow at $Re_\tau = 934$. Representative plots at various wall-normal locations for the RMS values of velocity and vorticity indicate the attenuation of the variance with increasing filter size. Further, the effect of averaging on the vortex core angle statistics is negligible when compared with the raw DNS data. These results indicate that the present PIV technique is an accurate and reliable method for the purposes of statistical analysis and identification of vortex structures.

N. Saikrishnan (✉) · I. Marusic · E. K. Longmire
Department of Aerospace Engineering and Mechanics,
University of Minnesota, 107 Akerman Hall,
110 Union St. SE, Minneapolis, MN 55414, USA
E-mail: neela@aem.umn.edu

1 Introduction

A fundamental problem in the study of dynamics of turbulent boundary layers has been in understanding the relation between vortex structures and the Reynolds stresses. Wall turbulence studies, both experimental and computational, have yielded a variety of results and models. Various studies have been directed toward understanding the eddy structures within the turbulent boundary layer, with "hairpin vortices", first introduced by Theodorsen (1952), now widely accepted to be a primary feature of wall turbulence as described by Adrian et al. (2000a) and as representative average structures in attached eddy models, such as in Perry and Marusic (1995) and Marusic and Perry (1995). Adrian et al. (2000b) experimentally observed spatially coherent groups of hairpin vortices in the logarithmic and outer wake layers of turbulent boundary layers. The authors termed such a group of coherent hairpin vortices a "hairpin packet". Ganapathisubramani et al. (2003) showed with stereoscopic particle image velocimetry (PIV) data at a friction Reynolds number $Re_\tau = 1{,}160$ that these packets occupy a relatively small percentage of the total area but contribute significantly to the total Reynolds shear stress generated. Hence, hairpin packets play a very important role in the production of turbulence and transport of momentum and a wide range of studies, numerical and experimental, have been performed towards better understanding the shape, size, orientation and dynamics of these structures. A detailed discussion of these structures has been presented in Marusic (2001).

Dual plane PIV is a novel technique where the complete velocity gradient tensor can be determined experimentally. It has been successfully used for measurements

in the turbulent boundary layer by Kähler (2004) and Ganapathisubramani et al. (2005a). In these studies, three velocity components in one wall parallel plane were measured using stereo PIV and two or three velocity components in a parallel plane separated by a small distance were measured using conventional or stereo PIV. Either setup provides the complete velocity gradient tensor for incompressible flows, and hence the complete vorticity vector. This can be extremely useful in analyzing structures in the turbulent boundary layer.

In PIV, the velocity vector in each interrogation window is calculated by means of a cross-correlation algorithm, where the displacements of all the particles within the interrogation window are taken into consideration. Scales smaller than the size of the interrogation window are averaged within the window, and a representative velocity is associated with each window. The resolution of the velocity field depends on the experimental setup, but in general, this might be affected by practical issues such as seeding and zooming in on the interrogation volume. One main limitation of PIV in general is the minimum resolution attainable.

Direct numerical simulation (DNS) is a commonly used numerical technique to study structures in turbulent flows. For example, Kim et al. (1987) used DNS to obtain turbulent statistics in a fully developed channel flow. However, DNS is limited by the Reynolds number of flow, since the Reynolds number directly influences the range of scales in the problem, as discussed by Moin and Mahesh (1998). Recently, Hoyas and Jiménez (2006) analyzed intensities and correlations of velocity fluctuations in turbulent channels using DNS up to $Re_\tau = 2{,}003$. Moin and Mahesh (1998) have described the potential use of DNS data in order to assess the possible errors introduced in experimental measurements owing to lack of spatial resolution. Studies by Moin and Spalart (1987), Suzuki and Kasagi (1992) and Pompeo and Thomann (1993) focused on the use of DNS data to estimate the accuracy of hot-wire probes. The effect of averaging across the probe length was studied in detail, and correction methods were proposed for the same.

In the present study, the aim is to quantify and assess the effects of averaging that might be introduced as a result of the cross-correlation technique of PIV, using a DNS dataset of comparable Re_τ. The resolution of the DNS dataset is finer than the PIV dataset; hence, smaller scales can be resolved in the DNS. In order to understand the effects of averaging, the DNS dataset is averaged to the resolution of the PIV dataset, and ensemble averaged statistics and instantaneous velocity fields are compared between the raw and the averaged DNS datasets. The statistics from the different datasets yield vortex core angle distributions, which can be used to understand the characteristics of structures in the turbulent boundary layer and to understand the shortcomings, if any, of the experimental technique as compared to the numerical technique.

The comparison between the PIV and the DNS data is made in the logarithmic region of the turbulent boundary layer. In spite of the fact that both of these studies are conducted at similar Re_τ, they are strictly different flows. The PIV studies are made in the zero pressure gradient boundary layer over a flat plate, whereas the DNS is performed in a channel flow. However, it is expected that in the logarithmic region, the effect of the opposing wall of the channel will be negligible and hence the structures observed in the channel flow will approximate those in a developing boundary layer. This assumption may not hold at larger distances from the wall and care should be taken before extending this argument to the entire boundary layer.

The remainder of the paper is organized as follows. In Sect. 2, the experimental setup and the DNS dataset are briefly described. In Sect. 3, the averaging techniques are described, and its variations are discussed. In Sect. 4, the results obtained from the raw and the averaged DNS are described in comparison with the PIV. In Sect. 5, the vorticity vector is used to study the characteristics of structures in the boundary layer and conclusions are presented in Sect. 6.

2 Description of PIV and DNS datasets

The dual plane PIV dataset was obtained from experiments conducted in a suction-type boundary layer wind tunnel by Ganapathisubramani (2004). The measurement planes were located 3.3 m downstream of a trip wire in a zero pressure gradient flow with freestream velocity $U_\infty = 5.9$ m/s and $Re_\tau = 1{,}160$ ($Re_\tau = \delta U_\tau/\nu$, where δ is the boundary layer thickness, U_τ is the skin friction velocity and ν is the kinematic viscosity of the fluid). The Reynolds number based on the momentum thickness Re_θ was 2,808 and the value of δ in the region of the measurement planes was 70.8 mm. The streamwise, spanwise and wall-normal directions are along the x, y and z axes, respectively, and the fluctuating velocity components along those three directions are represented as u, v and w. All quantities are normalized using "wall" variables U_τ and ν and are denoted with a superscript $+$.

The experiments conducted used a three-camera polarization-based dual plane PIV system. The setup consisted of two independent PIV systems that captured images of olive oil droplets of size ~ 1 μm. The first system was a stereoscopic setup, which provided all three velocity components within a plane, while the second system was a conventional PIV system that measured the in-plane velocity components in a neighboring plane located 21 wall units above in the wall-normal direction. Simultaneous measurements were performed in the two planes by using the polarization property of laser light sheets to isolate one camera system from the other (Ganapathisubramani 2004; Kähler 2004). Further details of the experimental setup can be found in Ganapathisubramani et al. (2006) and a description of the measurement accuracy can be found

in Ganapathisubramani et al. (2005a). The experiments were conducted in the logarithmic region of the turbulent boundary layer, with the first plane located 110 viscous units from the wall.

To obtain the vector fields, the images obtained from PIV were interrogated using a two-frame cross-correlation algorithm with discrete window offset. Firstly, using coarse 64×64 pixel windows, the mean displacement for each window was calculated. Next, with a 32×32 pixel window and the second window offset by the mean displacement, a vector field was obtained. This image was again interrogated using a 16×16 pixel window. At this step, the interrogation box in frame 1 was offset upstream and the frame 2 box was offset downstream by half the mean displacement calculated in the previous step. Thus, the final interrogation window size was 16×16 pixels, and a 50% overlap was used. The final vector fields obtained from the stereo PIV had a total size of $1.1\delta \times 1.1\delta$, and the spacing between the vectors was 12.3×12.3 wall units. The data from the two planes were used to compute the entire velocity gradient tensor in the lower plane. For the in-plane gradients, a second-order central difference scheme was used, while a first-order forward differencing scheme was used to calculate the streamwise and spanwise velocity gradients in the wall-normal direction. Finally, the wall-normal gradient of the wall-normal velocity was recovered from the continuity equation. Thus, the complete velocity gradient tensor was obtained using this experimental technique.

The DNS dataset under consideration here is a numerical simulation of a fully developed channel flow performed by del Álamo et al. (2004). The numerical technique involved the integration of the Navier–Stokes equations in the form of evolution problems for the wall-normal vorticity and the Laplacian of the wall-normal velocity. This is the same method as described by Kim et al. (1987). For spatial discretization, Chebychev polynomials were used in the wall-normal direction, while de-aliased Fourier expansions were utilized in wall-parallel planes. The temporal discretization used was a third-order semi-implicit Runge–Kutta scheme. Further details of the DNS can be found in del Álamo et al. (2004). The simulation of interest in the present study has $Re_\tau = 934$, which is referred to as L950. The computational domain in this simulation was $8\pi h$ units in the streamwise direction and $3\pi h$ units in the spanwise direction, where h is the channel half width. The sizes of the domain were set to try to account for all the energy containing structures in the flow, especially features with dimensions of the order of h. The spacing between the vectors was 11.46×5.73 wall units in the streamwise and spanwise directions, respectively, after de-aliasing.

3 Averaging techniques used

As described in Sect. 2, the resolution of the DNS dataset is higher than the PIV dataset. In order to obtain comparable resolutions in the two datasets, an averaging scheme was implemented in the DNS data. The averaging technique is described as follows. Each PIV field had a size of $1.1\delta \times 1.1\delta$ and contained a total of 100×100 interrogation windows. This yielded a field of 100×100 vectors, with a spacing of 12.3 wall units between vectors in both the streamwise and the spanwise directions. However, taking into account the overlap of the interrogation windows in the PIV processing, the effective resolution of the PIV data was 24.6×24.6 wall units. As described earlier, the spacing between the vectors in the DNS dataset was 11.46×5.73 wall units in the streamwise and the spanwise directions. Thus, the aim was to average the DNS values within the appropriate square interrogation windows to obtain a field of 100×100 vectors.

One time step of the DNS dataset contains $2{,}048 \times 1{,}536$ points in the streamwise and the spanwise directions, respectively. This field from the DNS dataset was divided into smaller fields of 107×214 points to obtain fields with a dimension of $1{,}226 \times 1{,}226$ wall units each, which was about the same size as the PIV field of $1{,}230 \times 1{,}230$ wall units. Then, the fields of 107×214 points were averaged down to 100×100 points to reduce the resolution of the DNS dataset. Thus, 133 fields of 107×214 points from a plane of one time step of the DNS data were obtained, and all fields from different time steps were processed for statistical analysis. In the current study, two time steps of data were used, and these fields provided good statistical convergence.

In the wall-normal direction, the laser sheets had finite thicknesses. Burn tests were used to estimate the thicknesses of the laser sheets used in the experiments, which were found to be 0.35 and 0.45 mm or 5.6 and 7.2 wall units for the stereoscopic and the single camera systems, respectively. Given this configuration of the

Table 1 Ensemble averaged mean and RMS statistics from dual plane PIV, raw DNS and 3D averaged DNS datasets. λ_{3D}^+ is the three-dimensional swirl and is defined in Sect. 4

	$\overline{U^+}$	σ_u^+	σ_v^+	σ_w^+	$\sigma_{\omega_x}^+$	$\sigma_{\omega_y}^+$	$\sigma_{\omega_z}^+$	$\overline{\partial U^+/\partial z^+}$	$\overline{\lambda_{3D}^+}$
PIV data	16.1	1.91	1.34	1.18	0.069	0.064	0.056	0.021	0.026
Raw DNS	16.6	1.91	1.35	1.07	0.073	0.066	0.071	0.022	0.025
Averaged DNS	16.6	1.86	1.30	1.00	0.054	0.054	0.056	0.019	0.021

PIV dataset, the aim was to extract data from the DNS at the corresponding wall-normal locations. In the DNS data, Chebychev polynomials were used in the wall-normal direction, due to which the spacing between the planes was not uniform. Hence, it was necessary to identify wall-normal planes corresponding to the locations of the PIV data planes accurately.

Two averaging techniques were used. In the 2D averaging technique, averaging was carried out over a two-dimensional window such that the resolution in a given plane was reduced only in the streamwise and spanwise directions. In the 3D averaging technique, the averaging in a given plane was performed in a three-dimensional box including the effect of averaging over the laser sheet thickness. The streamwise and spanwise out-of-plane gradients, the only quantities that required data from the upper PIV measurement plane, were computed using the same $\Delta z^+ = 21$ for both averaging techniques. In both methods, the streamwise–spanwise averaging was implemented as follows. A box filter of size 12.3×12.3 wall units was used to detect all data points lying within the corresponding interrogation window. This box was moved through the entire field of 107×214 points to reduce it to 100×100 points. Typically, this procedure resulted in about 2 points per box in the spanwise direction and 1 point per box in the streamwise direction. A simple average was performed to come up with a representative value for each box. This procedure was performed for all three velocity components U, V and W. Due to the discrete cutoffs used for detection of the points, the number of points in a box could be as high as 6, instead of the typical value of 2 per plane. Further, each point in the DNS dataset represents a finite area, given by the resolution of the dataset. Thus, the actual area of averaging is the area represented by the points and not the spacing between them. This grid of 100×100 points with a resolution of 12.3×12.3 wall units was averaged to a resolution of 24.6×24.6 wall units, using a window overlap to retain the number of points as 100×100. This averaged dataset is now of the same nominal resolution and size as the PIV measurements. This dataset is referred to as the averaged DNS dataset.

In order to understand the effects of varying the size of the streamwise–spanwise averaging window on the identification of vortex core statistics, additional tests

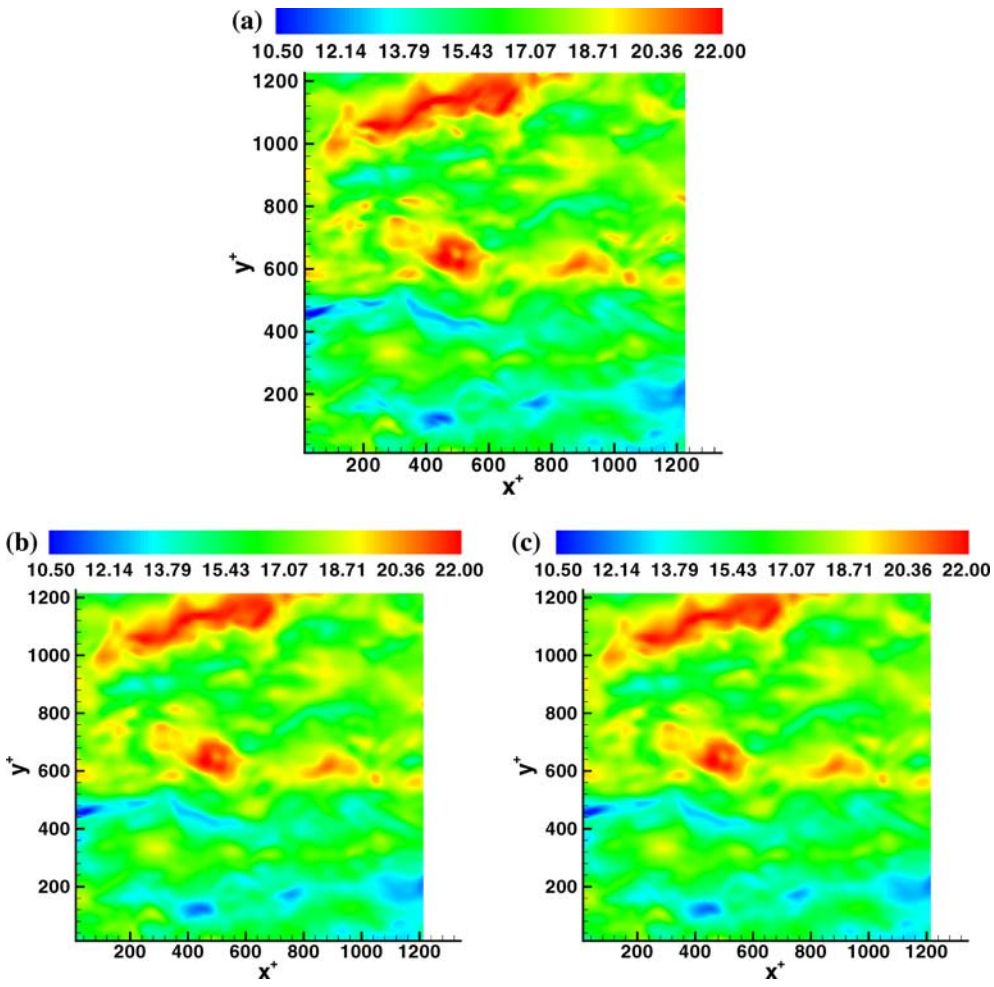

Fig. 1 Contour plots of U^+ extracted from **a** raw DNS data **b** 2D averaging **c** 3D averaging

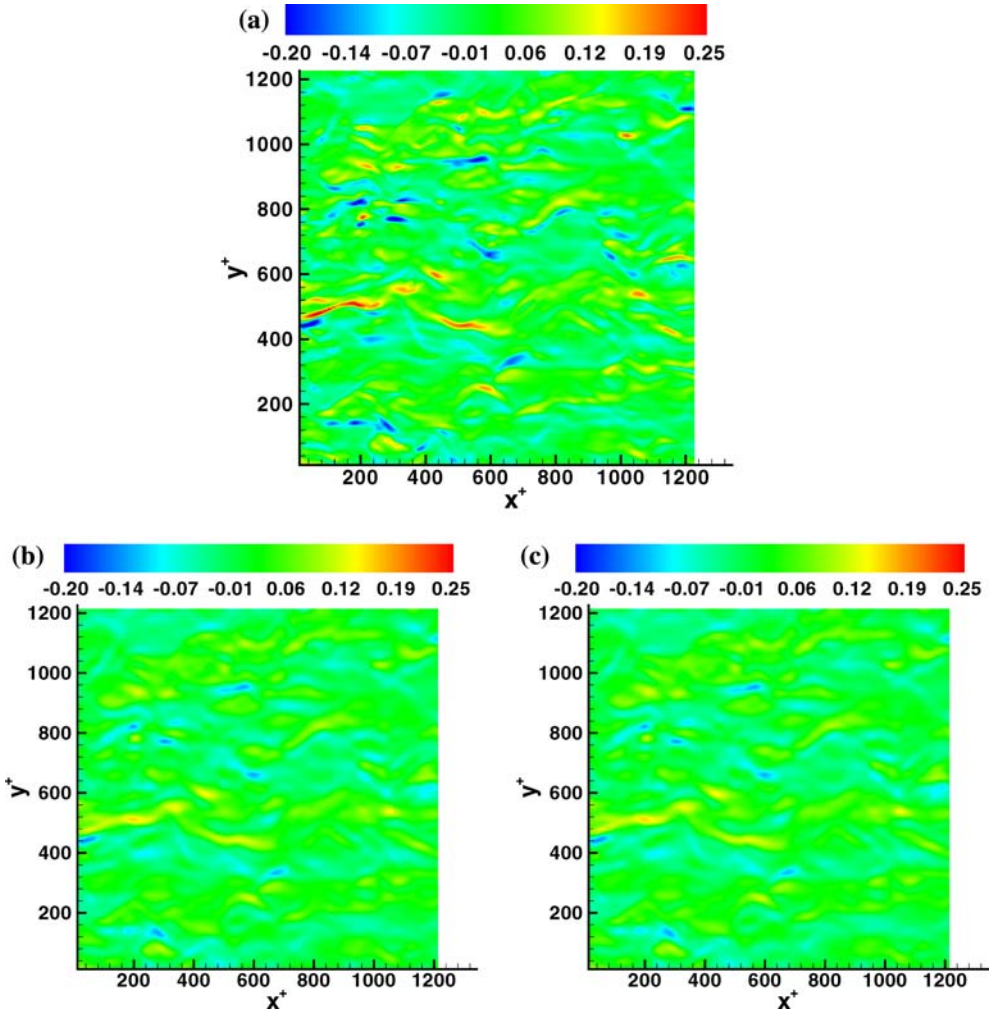

Fig. 2 Contour plots of $\partial U^+/\partial y^+$ extracted from **a** raw DNS data **b** 2D averaging **c** 3D averaging

were performed with larger filter sizes (L^+ equal to 49.0 and 75.6). Further, the window size was extended to even larger values to understand the effect of the spatial averaging and resolution on the RMS values of velocity and its derivatives.

In the DNS dataset, the plane numbers 60, 61, 62, 66 and 67 correspond to planes at 106.738, 110.313, 113.943, 129.012 and 132.914 viscous units from the wall, respectively. For 3D averaging, the planes used were 60, 61 and 62 for plane 1 and planes 66 and 67 for plane 2, while for 2D averaging, only planes 61 and 67 were used.

To calculate the gradients in the raw DNS dataset, a second-order finite difference scheme was implemented in all three directions. For wall-normal gradients, information from planes 60, 61 and 62 was used. On the other hand, in the case of the averaged DNS obtained from either averaging technique, the appropriate averaging as described earlier was performed to yield averaged planes 61 and 67. These averaged planes were used to calculate in-plane gradients to second-order accuracy and wall-normal gradients to first-order accuracy using a forward difference scheme. This essentially mimics what happens with the PIV dataset.

4 Instantaneous and ensemble averaged results

The instantaneous and statistical results from the dual plane PIV, the raw DNS and the averaged DNS datasets are presented in this section. Some instantaneous velocity and velocity gradient fields are also presented to provide a view of the structure of the turbulent boundary layer. The mean and RMS statistics of the velocity and vorticity components computed from the dual plane PIV, raw DNS dataset and the 3D averaged DNS dataset are listed in Table 1. In the table, quantities with overbars are mean quantities and σ represents the variance of the fluctuating part of the various quantities.

As expected, the averaging does not affect the mean value of the streamwise velocity. The probability density functions of the velocity components and their derivatives were plotted and were found to follow a

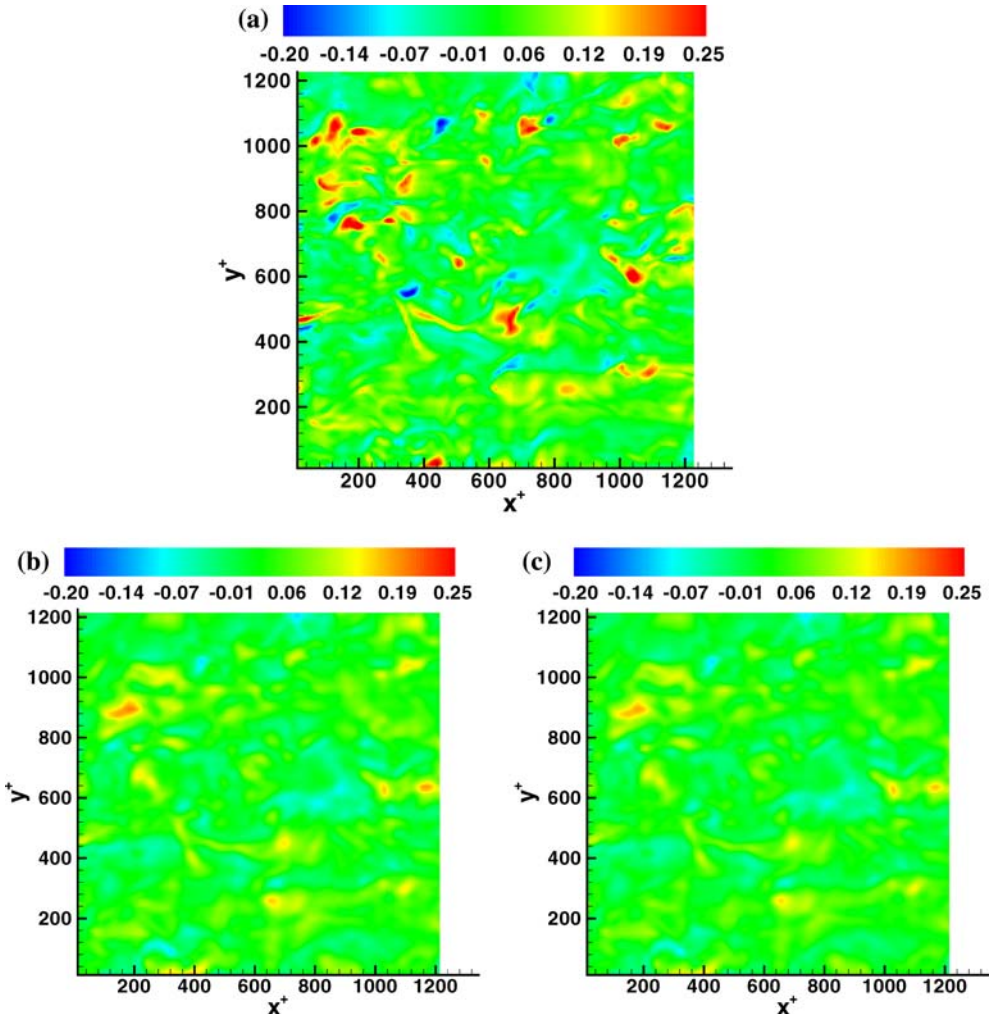

Fig. 3 Contour plots of $\partial U^+/\partial z^+$ extracted from **a** raw DNS data **b** 2D averaging **c** 3D averaging

Gaussian distribution. The averaging has the effect of reducing the spread of the points about the mean value in a given distribution, and hence reduces the RMS values. The averaging reduces the RMS values of the velocity components by less than 7%, while the RMS values of the vorticity components are reduced by 20–26%. The greater reduction in vorticity values are most likely due to faster drop off in spatial correlations as compared with velocity. In general, the PIV velocity and vorticity variances fall closer to the raw DNS values than to the averaged DNS values. In addition, the values of the streamwise and spanwise velocity gradients are somewhat "smoothed" in the averaged DNS due to the larger spacing between the planes used for differencing. The effect of this can be seen in the term $\overline{\partial U^+/\partial z^+}$, where the differencing technique reduces the mean value of the quantity. The value of λ^+_{3D}, the mean of the three-dimensional swirl, matches within 4% between the PIV and raw DNS datasets, while the averaged DNS value is about 16% less.

In order to further investigate the effect of averaging on the RMS values, the 3D averaging procedure described earlier was implemented on the raw DNS dataset, with different filter sizes, ranging from 12.3 to 490 wall units. The results for the different filter sizes are shown in Fig. 14a, b in the Appendix. In these figures, the RMS values are normalized with the values for the raw DNS. As can be seen from the plots, the effect of spatial averaging is to reduce the RMS values as described earlier. In Fig. 14a, the RMS values fall off slowest for u^+, probably due to long correlations in the streamwise direction, and fastest for w^+ where the spatial correlations are shortest (see Ganapathisubramani et al. 2005b). The vorticity values fall off much more steeply, consistent with the concept of even more limited spatial correlations. It can be seen that averaging with a window size of 24.6×24.6 wall units reduces the RMS values by less than 7% for velocities and 20–26% for vorticity, as observed earlier from the numerical values in Table 1. The plots also show that with increasing window size, the effect of filtering increases greatly, and with a window size of 100×100 wall units, about 60–70% of the vorticity fluctuations are damped out. This would likely be significant in the identification

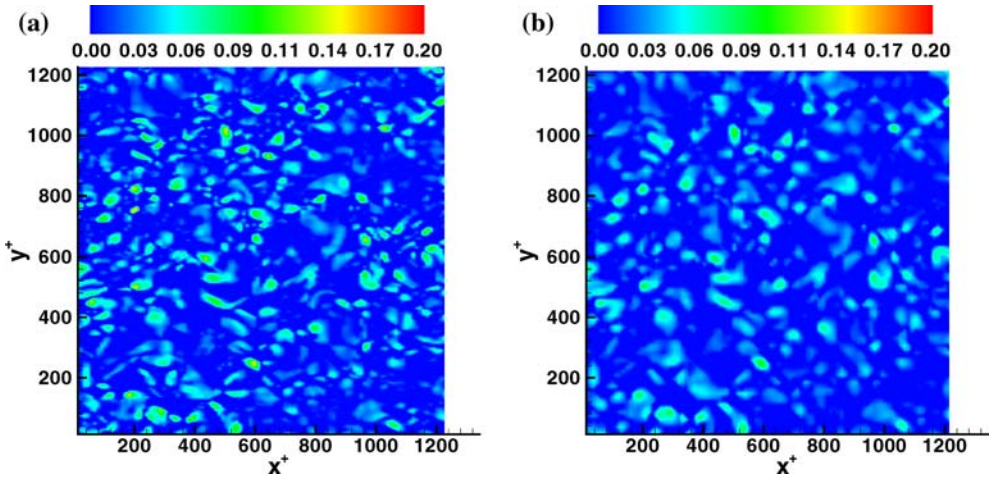

Fig. 4 Contour plots of λ_{2D}^{+} extracted from **a** raw DNS data **b** averaged DNS data

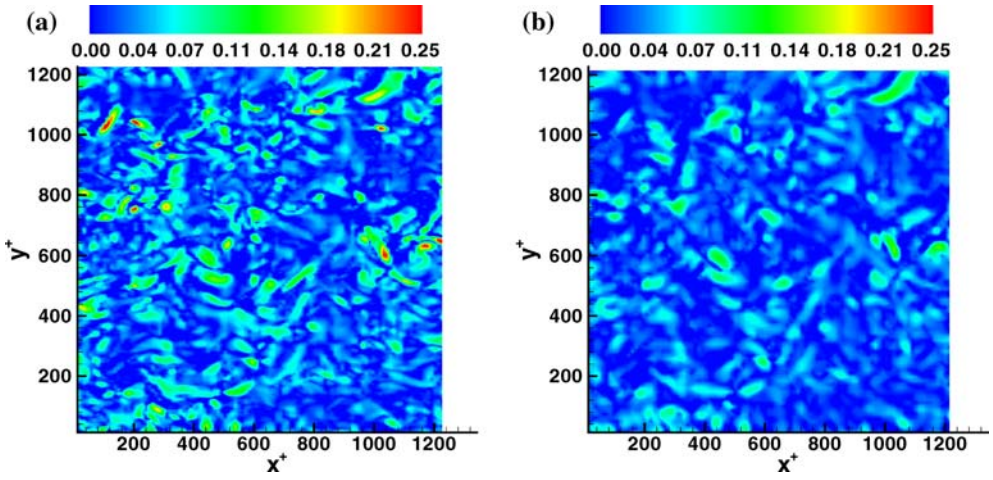

Fig. 5 Contour plots of λ_{3D}^{+} extracted from **a** raw DNS data **b** averaged DNS data

of vortex cores based on vorticity and quantities derived from the velocity derivatives.

It should be noted that the result in Fig. 14 is valid strictly for the wall-normal location of $z^+ = 110$. It is expected that closer to the wall, where small-scale fluctuations are relatively more important, the attenuation of the RMS values would be greater as a result of filtering at a given scale. Representative results for wall-normal locations in the viscous buffer region and the outer region are also shown in the Appendix for completeness. The results here are similar to those reported in Suzuki and Kasagi (1992) and Ligrani and Bradshaw (1987) that analyzed the effectiveness and spatial resolution of hot wires based on their length. The current plots however relate directly to the effectiveness of a wide range of possible PIV windows at capturing various turbulence statistics.

Instantaneous contour plots of the streamwise velocity U^+ are shown in Fig. 1. The plots show the effects of averaging introduced by implementation of the techniques described earlier, with the averaged plots showing some smoothing of the velocity contours. It is also apparent that features that are more distinct and sharp in the raw DNS plots appear blurred in the averaged fields, which is as expected from averaging of any kind. An example of this may be seen in the top left corner of the plots in Fig. 1, where we observe that the velocity values are closer to the mean in the averaged plots than in the raw DNS. In Fig. 2, it is noted that the averaging acts to reduce the level of variations for the quantity $\partial U^+/\partial y^+$. However, the basic structures observable in the raw DNS can be seen in both the 2D and 3D averaged plots. The effect of averaging on the wall-normal gradient shown in Fig. 3 is more appreciable due to the difference in methods used to obtain the wall-normal gradients between the raw and averaged data as described earlier.

A comparison of the 2D and 3D averaging plots in Fig. 1b, c also shows that the effect of averaging over effective sheet thickness in the wall-normal direction is not appreciable. This difference was checked on individual in-plane and out-of-plane derivatives as well as

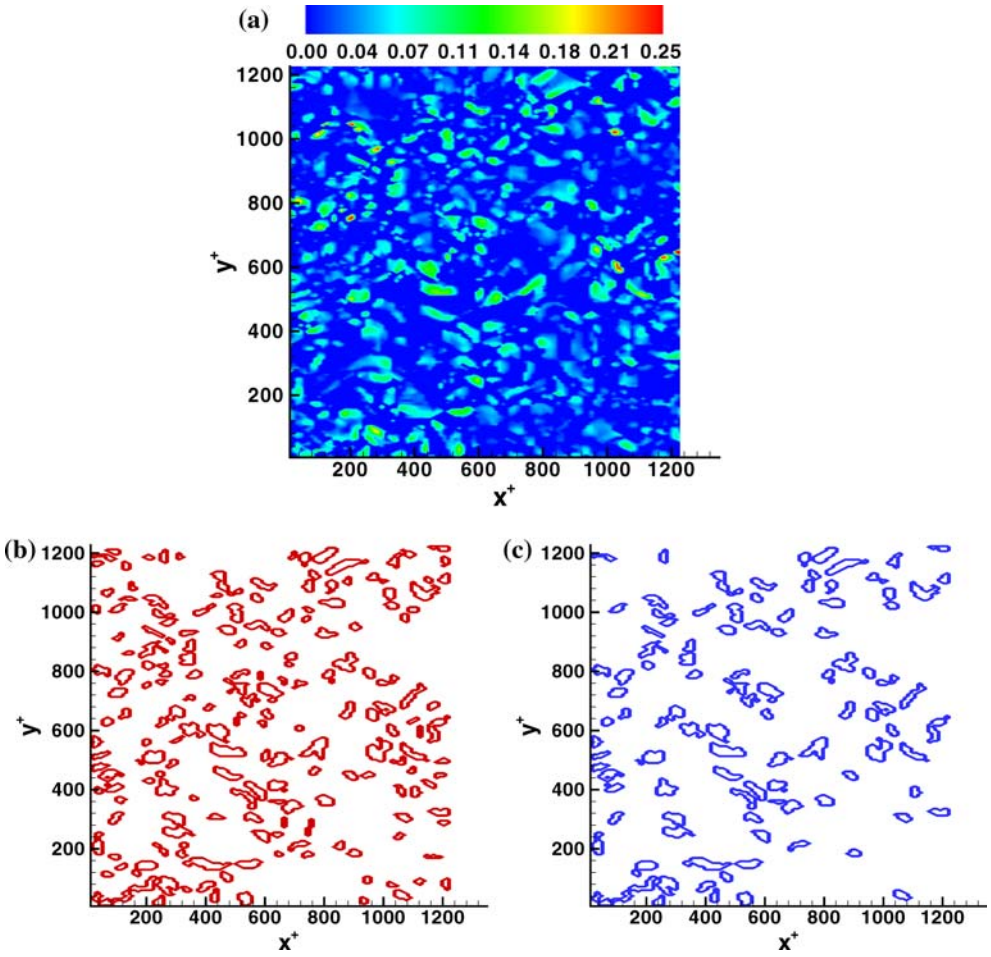

Fig. 6 Performance of the vortex core identification algorithm for raw DNS data. **a** λ_{3D}^{+}. Cores identified with **b** 5 point threshold **c** 12 point threshold

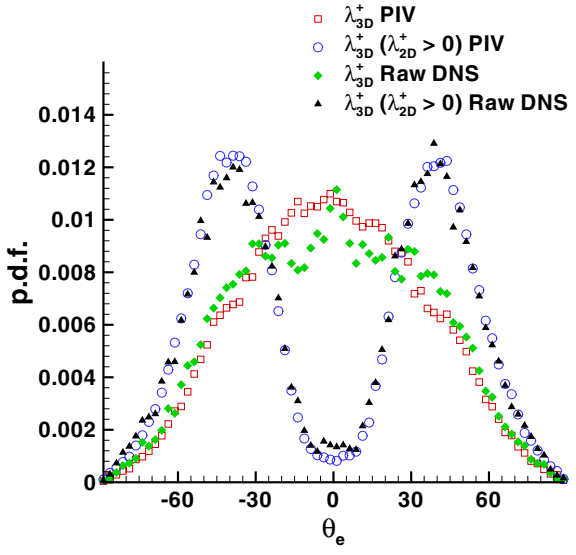

Fig. 7 Probability density function (p.d.f.) of the elevation angle θ_e for the raw DNS and the PIV datasets

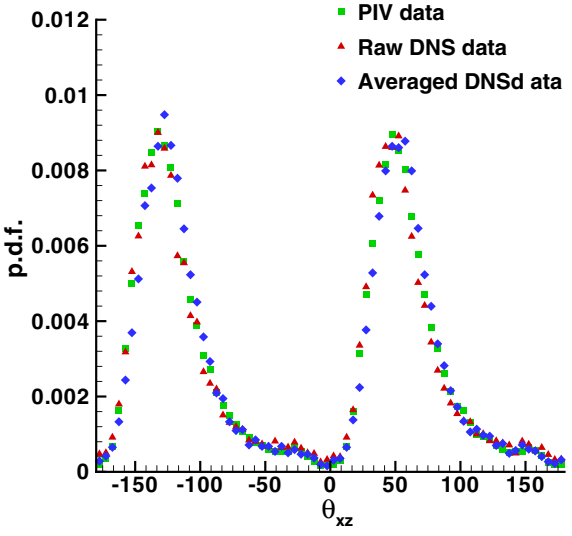

Fig. 8 Probability density function of θ_{xz} from raw DNS, average DNS and PIV

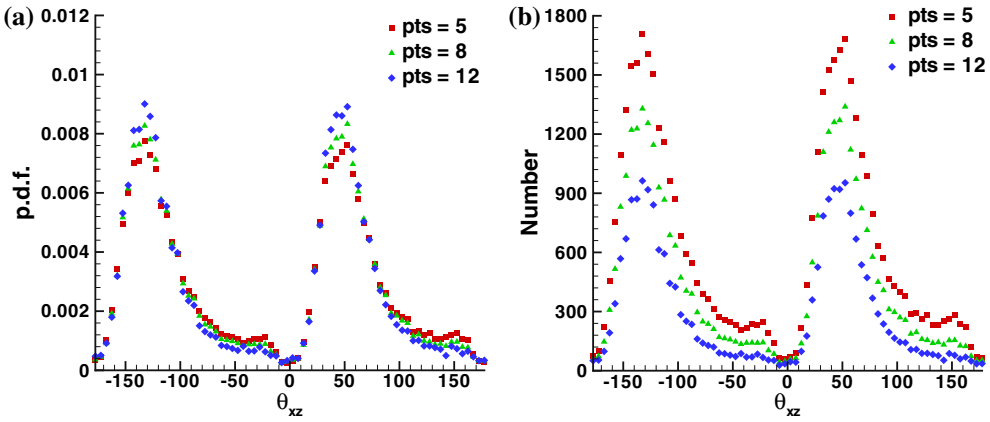

Fig. 9 Effect of area threshold for the raw DNS data on **a** p.d.f. of θ_{xz} **b** absolute number of cores identified

on derived quantities such as vorticity and swirl. Figure 2b, c illustrates this for the spanwise gradient of the streamwise velocity, while Fig. 3b, c shows this for the wall-normal gradient. Since there is no appreciable difference in all these cases, we can safely conclude that the effect of averaging over the sheet thickness is minimal. Further, raw fields were averaged in the wall-normal direction over the sheet thickness with no streamwise–spanwise averaging and the results were found to be statistically identical to those obtained from the raw field. For further analysis, we will use only the 3D averaging to represent "averaged" DNS data.

Plots of the two-dimensional and three-dimensional swirl strengths are shown in Figs. 4 and 5, respectively, for raw and averaged data. The three-dimensional swirl strength λ^{+}_{3D}, defined as the imaginary part of the complex eigen value (λ_{ci}) of the full velocity gradient tensor, can be used to visualize vortices as suggested by Zhou et al. (1999). λ^{+}_{3D} is an indicator of the presence of regions of fluid swirling about an axis. λ^{+}_{2D}, based on the two-dimensional velocity gradient tensor, as described by Adrian et al. (2000a), is an indicator of the presence of regions of fluid swirling about an axis aligned with the wall-normal direction. Mathematically, λ^{+}_{3D} is greater than or equal to λ^{+}_{2D}, depending on the deviation of the swirling motion from the wall-normal direction.

By comparing the plots from the raw and averaged DNS data in Figs. 4 and 5, we can see that structures that appear sharper in the raw DNS are blurred by averaging and that a lower number of cores are observed. This could be due to a reduction of values of the local maxima by averaging and due to the merging of structures upon lowering the resolution. As shown in Table 1, the mean of average swirl strength is seen to reduce upon averaging. This would be expected as in the calculation of the mean value of the swirl, points having zero swirl are also used. Upon averaging, the number of points containing zero swirl values is seen to increase, and hence the corresponding mean value decreases.

The following section will discuss statistical distributions of vortex core angles obtained from the instantaneous swirl fields.

5 Vortex structure deduced from instantaneous results

The complete velocity gradient tensor can be used to compute the inclination angle of any individual vortex structure by determining the orientation of the vorticity vector averaged over the region of the vortex core. It is emphasized here that our interest is to determine the characteristics of vortex core structures responsible for the bulk of the kinetic energy and not the bulk of the enstrophy, which is dominated by small-scale motions approaching the Kolmogorov length scale.

The region-growing algorithm used to identify vortex cores is discussed in detail in Ganapathisubramani et al. (2006) and Ganapathisubramani (2004), and is also described briefly below. The aim is to identify connected points of swirl strength greater than a certain threshold

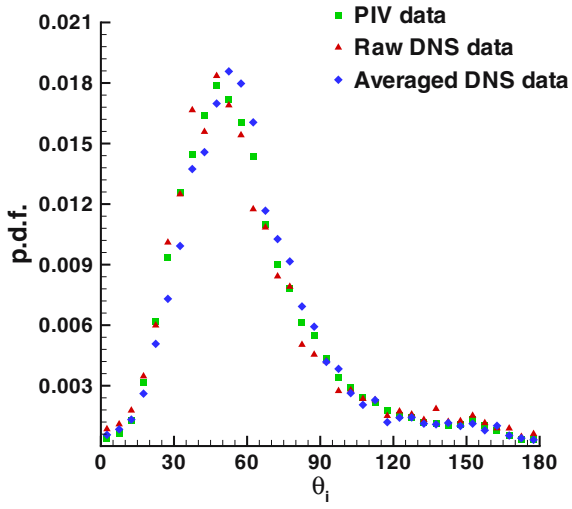

Fig. 10 Probability density function of eddy inclination angle θ_i

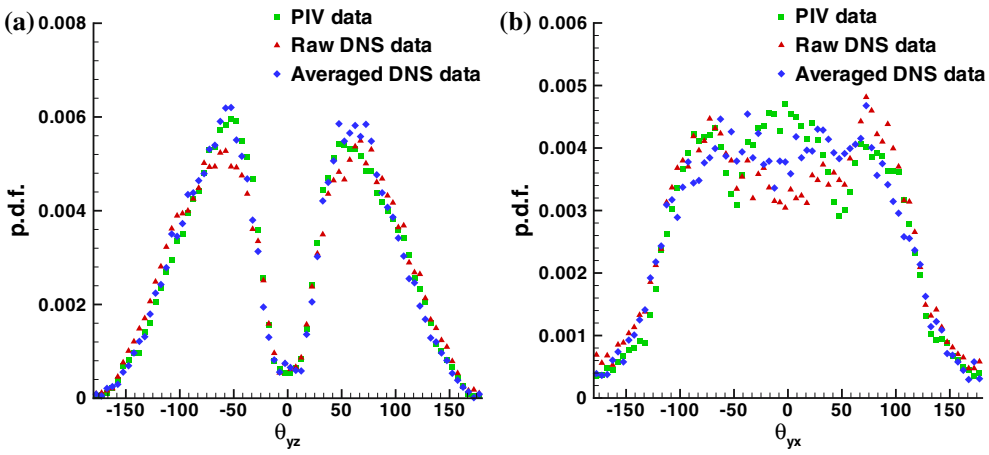

Fig. 11 Probability density function of **a** θ_{yz} **b** θ_{yx}

magnitude. The algorithm first identifies all local maxima in the field greater than the threshold. This threshold was fixed at 10% of the maximum value in the raw DNS dataset in order to filter out contributions from weak vortex structures and measurement noise. These points were then used as seed points for a region growing algorithm that identifies a connected region of swirl with values greater than the threshold. In order to reduce contributions from measurement uncertainty and very small scales, a lower limit on the minimum area or number of points included in a connected region is also imposed. The choice of this value is crucial to the functioning of the algorithm, and its effect on the statistical distribution is discussed in detail later.

Once the core regions have been identified, the average values of the three components of vorticity are computed in each connected region. An additional constraint imposed is that the average value of at least one vorticity component must be greater than the standard deviation of that component for that core to be included in the distribution of a related core angle. Also, a given projection angle is computed only if the average vorticity value of at least one component used in its computation is greater than its standard deviation. These criteria essentially eliminate contributions of weak vorticity values to the probability density function (p.d.f.) of specific projection angles.

Figure 6 shows a sample swirl field and the results of the vortex core identification algorithm. Two thresholds on the number of points are used in the results shown. In general, the sizes and shapes of the identified cores are seen to agree well with locations of significant swirl. A higher threshold on the number of points ensures that only stronger cores with greater circulation are identified since the average vorticity of the core must exceed the vorticity threshold. Thus, an increase in the area threshold decreases the number of cores identified by eliminating weaker and smaller cores.

Figure 7 shows the p.d.f. of the elevation angle θ_e made by the vortex cores for the raw DNS and PIV datasets. The elevation angle, which is defined as the angle made by the vorticity vector with the $x-y$ plane, can vary from $-90°$ to $+90°$. The distribution, depicted by square symbols for PIV and diamond symbols for DNS, shows that a large percentage of structures have very low elevation angles. This suggests that these contributions are made by structures that lie in the plane of

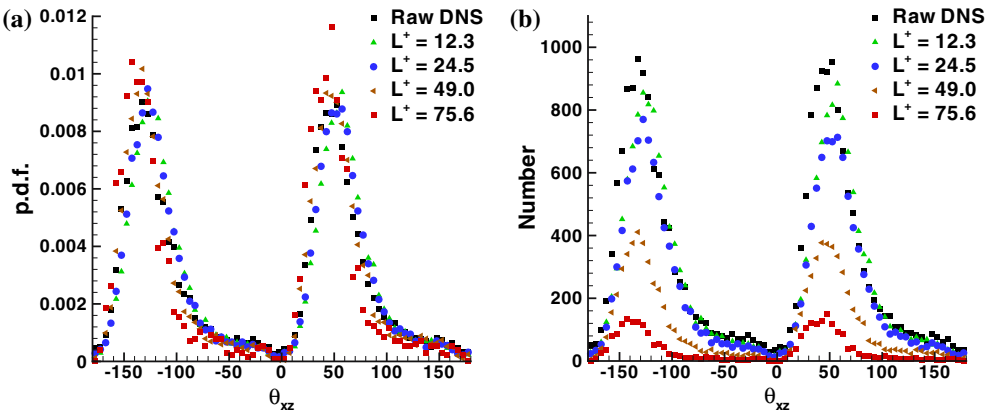

Fig. 12 Effect of variation of box filter sizes on **a** p.d.f. of θ_{xz} **b** absolute number of cores identified

measurement. In order to obtain the inclination angles of cores that are oriented with some angle to the streamwise–spanwise planes, we need to filter out the in-plane structures. It is here that the two-dimensional swirl is useful. λ_{2D}^{+} can be used to identify structures swirling about an axis perpendicular to the wall parallel plane. Thus, the use of $\lambda_{2D}^{+} > 0$ as an additional constraint provides us with a means of identifying only out-of-plane structures. The resulting curves of λ_{3D}^{+} are also plotted in Fig. 7. The DNS and the PIV datasets seem to provide consistent results, and in both the cases the peaks of the curves of λ_{3D}^{+} ($\lambda_{2D}^{+} > 0$) lie at about $\pm 40°$. However, these peaks are broad, suggesting that most of the structures are inclined at angles between $\pm 30°$ and $\pm 50°$ to the wall parallel plane. In all further p.d.f. plots of vortex angles where λ_{3D}^{+} is used, the criterion $\lambda_{2D}^{+} > 0$ is enforced.

The projection angles of the vortex structures in the $x-z$, $x-y$ and $y-z$ planes are calculated using the projections of the vorticity vector in the respective planes. Firstly, the probability distribution of the inclination angle θ_{xz}, which is the angle made by the projection of the vorticity vector in the $x-z$ plane with the positive x-axis, is shown in Fig. 8. The plots from the raw DNS, the averaged DNS and the PIV datasets match extremely well indicating peaks close at $\theta_{xz} = 45°$ and $-135°$, but again with broad distributions. The peaks are consistent with the presence of positive and negative legs of forward-leaning hairpin vortices. Another interesting observation is that there is a finite probability of θ_{xz} lying in the ranges $-90°$–$0°$ and $90°$–$180°$, indicating the presence of backward leaning cores. A detailed physical explanation of this can be found in Ganapathisubramani et al. (2006).

Figure 9 illustrates the effect of the threshold on core area or the minimum number of points per identified core. The effect of this constraint is that smaller and, in general, weaker cores are filtered out by increasing the minimum number of points. The other advantage of using a larger threshold area is that it reduces identification of cores due to measurement noise. On the other extreme, a very large threshold area prevents identification of a significant number of genuine vortex cores. So a balance has to be struck in the choice of this threshold. Figure 9a shows the p.d.f. of θ_{xz} plotted for the different thresholds applied to the raw DNS data. It can be seen that increasing the threshold area does not significantly change the peaks of the plot, but the value at each peak is higher. This can be seen in Fig. 9b, where it is observed that the number of structures filtered for different angles is different due to which the p.d.f. in Fig. 9a changes. For all angle distributions shown in Figs. 8, 10 and 11, the threshold represents about 780 square wall units, which corresponds to 12 points for the raw DNS and 5 points in the average DNS and PIV datasets. These numbers were chosen based on the spatial grid sizes in the different datasets in order to apply the same threshold on in-plane area occupied by the core.

Figure 10 shows the eddy inclination angle θ_i, obtained by accumulating all the forward leaning eddies and backward leaning eddies together from the distribution in Fig. 8. This is the angle of inclination of an average representative eddy with respect to the $x-y$ plane. These eddies (most likely hairpins) are statistically averaged and hence must be symmetric about the $x-z$ plane (since these flows are two-dimensional in the mean). The PIV and raw DNS data are seen to match very well, while the averaging of the DNS tends to shift the peak of the curve toward a slightly higher angle. However, the quantitative agreement in this case is also very close, suggesting that the predictions of the PIV and the DNS datasets are almost identical.

The projections of the vorticity vector on the other planes can also be used to understand the typical structures in the turbulent boundary layer. The angle made by the projection of the vorticity vector on the $y-z$ plane with the positive y-axis is referred to as θ_{yz}, and the angle made by the projection on the $x-y$ plane with the positive y-axis is referred to as θ_{yx}. Figure 11a shows the p.d.f. of θ_{yz} for the raw DNS, the averaged DNS and the PIV data. The plot of averaged data appears to be closer to the PIV data, but overall good qualitative and quantitative agreement is seen among all the datasets. Further analysis of the PIV data and interpretation based on the hairpin shaped structure has been discussed in Ganapathisubramani et al. (2006).

Figure 11b shows the p.d.f. of θ_{yx}. As seen with earlier quantities, good qualitative agreement between the different datasets is observed, but some quantitative differences are noted, particularly for low angles where considerable scatter is observed for all three cases. In the PIV data, the p.d.f. seems to contain three peaks. The main peak occurs at 0°. If we assume that the field is dominated by hairpin structures, symmetric or otherwise, this could be caused by the regions near and including the heads of hairpin loops. The other two peaks, located at $\theta_{yx} = \pm 75°$, could be caused by the necks of hairpin vortices. The peak at 0° is not visible in the raw DNS, but the other two peaks are present. The averaging seems to bring the distribution closer to the PIV result.

In order to come up with a useful limit for the averaging window size, larger streamwise–spanwise boxes were tested on the vortex core statistics. Figure 12a replots the p.d.f. of θ_{xz} for the raw data and earlier "averaged" data ($L^+ = 24.5$), together with two larger box sizes. It can be seen from the plot that averaging with a small box size does not change the distribution of the angles significantly for $L^+ = 24.5$ and $L^+ = 49.0$. However, when the size is increased to $L^+ = 75.6$, there is a clear shift in the distribution toward smaller angles. A further increase in the box size is expected to yield even poorer results. Figure 12b shows the number density of cores identified for the different filter sizes. It is clear from this plot that the coarser filtering greatly reduces the number of identified cores.

A great deal of care should be taken in obtaining statistics for larger filter sizes because the constraints like the vorticity threshold and the area threshold must be considered in relation to the interrogation area size in order to understand and interpret the results. For example, once the filter size matches or exceeds the core threshold area, a large percentage of cores will go undetected. In addition, once the in-plane filter size significantly exceeds the out-of-plane resolution (in our case 21 wall units for normal gradients of streamwise and spanwise velocity), quantities dependent on direction should become biased. For example, vortex angle distributions should become biased away from the wall-normal direction due to greater smoothing of streamwise and spanwise gradients.

From the figures described in this section, it can be seen that, in general, the results from the PIV appear to match those from the raw and averaged DNS datasets within limits of statistical uncertainty. This indicates that the current PIV technique appears to be a very good technique at the current resolution for the identification of vortex cores and for calculation of vortex core angle statistics.

6 Conclusions

The vortex core angle distributions in wall turbulence were analyzed using statistical tools applied on an experimental dataset from a dual plane PIV system and a DNS dataset in the logarithmic region of a channel of similar Reynolds number. The effects of averaging in the DNS data were studied and a comparison with PIV data was made. The RMS values of all components of velocity and vorticity were quantified for a wide range of effective measurement resolutions. These results can be used to estimate the percentage of turbulence intensity or enstrophy captured or lost for a given PIV interrogation resolution across turbulent boundary layers.

The averaging of the raw DNS data filters out the small-scale structures from the field. Instantaneous fields of velocity and velocity gradients show smoothing of the velocity contours. It is also apparent that features which are more distinct and sharp in the raw DNS plots appear blurred in the averaged fields. The effect of averaging over an estimated sheet thickness does not seem to be appreciable. The effect of averaging over the spacing between the two sheets (to obtain wall-normal gradients) is small but noticeable.

The averaged data give us a good idea of stronger structures expected to drive the main energy-containing interactions within the boundary layer. Based on the comparisons with raw and filtered DNS data, the dual plane PIV data were sufficiently resolved to characterize vortical structures. Therefore, this study serves as a strong validation for the experimental technique. It clearly illustrates that the spatial averaging inherent to PIV measurements did not affect the vortex core angle statistics significantly because of the close match between the experimental and numerical datasets.

Acknowledgements The authors would like to express sincere thanks to Prof. Robert Moser for providing the DNS data for the channel flow. We are indebted to Dr. Bharathram Ganapathisubramani, Pramod Subbareddy and Dr. Nicholas Hutchins for their help in using the PIV and DNS datasets and many discussions during the course of this study. Support from the National Science Foundation through Grant CTS-0324898 and from the David and Lucile Packard Foundation is gratefully acknowledged.

7 Appendix: Effect of filter size on RMS statistics

The following section provides plots of the RMS values of velocity, vorticity and the Reynolds shear stress in the logarithmic, viscous buffer and outer regions of a

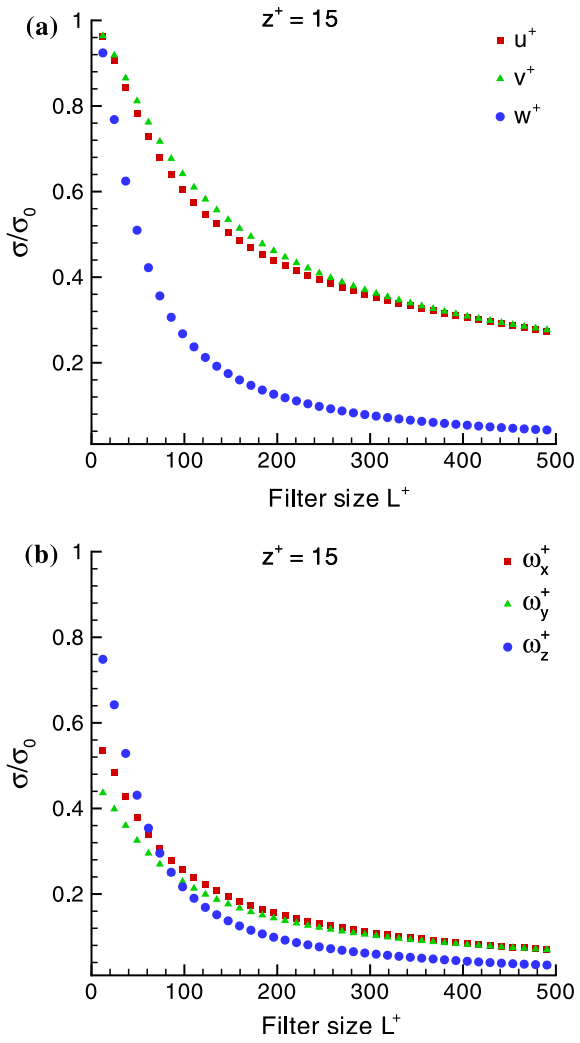

Fig. 13 Variation of RMS statistics in the viscous buffer region with filter size of **a** velocity **b** vorticity. Quantities are normalized by raw DNS values

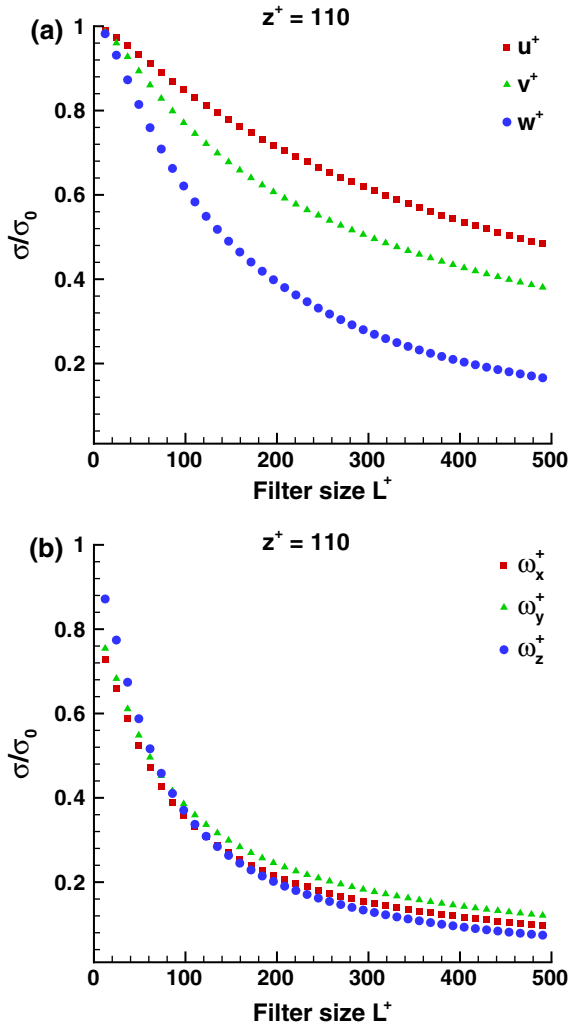

Fig. 14 Variation of RMS statistics in the logarithmic region with filter size of **a** velocity **b** vorticity. Quantities are normalized by raw DNS values

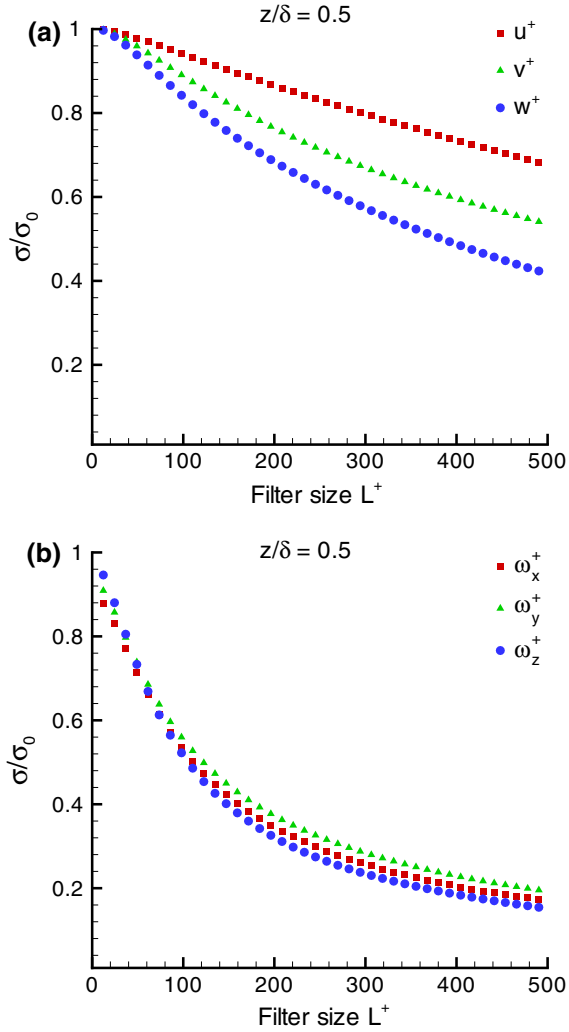

Fig. 15 Variation of RMS statistics in the outer region with filter size of **a** velocity **b** vorticity. Quantities are normalized by raw DNS values

turbulent boundary layer (Figs. 13, 14, 15, 16). These results can be used to estimate the percentage of turbulence intensity or enstrophy captured or lost for a given PIV interrogation resolution across turbulent boundary layers.

In all of the following plots, the normalization is done using the corresponding RMS value computed from the fully resolved DNS data set. RMS values of the various quantities at different wall-normal locations from the fully resolved DNS data are shown in Table 2. For the calculation of streamwise and spanwise derivatives in the wall-normal direction, a first-order approximation using two planes separated by $\Delta z^+ = 21$ is used, consistent with the description in Sect. 3. It must be noted that at the wall-normal location $z^+ = 15$, the effect of spacing between the planes can be seen clearly from the large drop in values of the vorticity components containing wall-normal gradients.

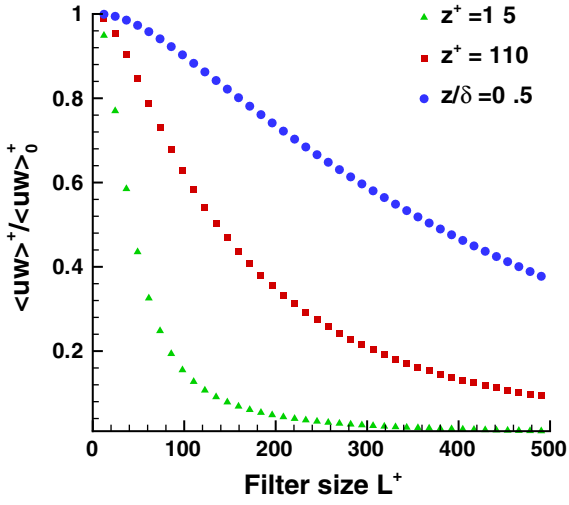

Fig. 16 Variation of Reynolds shear stress with filter size, normalized by raw DNS values

Table 2 RMS statistics at various wall-normal locations for the fully resolved DNS dataset

	σ_u^+	σ_v^+	σ_w^+	$\sigma_{\omega_x}^+$	$\sigma_{\omega_y}^+$	$\sigma_{\omega_z}^+$	$\langle -uw \rangle^+$
$z^+ = 15$	2.80	1.26	0.55	0.164	0.235	0.173	0.641
$z^+ = 110$	1.91	1.35	1.07	0.073	0.066	0.071	0.856
$z/\delta = 0.5$	1.27	0.88	0.81	0.029	0.026	0.026	0.384

References

Adrian RJ, Christensen KT, Liu Z-C (2000a) Analysis and interpretation of instantaneous turbulent velocity fields. Exp Fluids 29:275–290

Adrian RJ, Meinhart CD, Tomkins CD (2000b) Vortex organization in the outer region of the turbulent boundary layer. J Fluid Mech 422:1–54

del Álamo JC, Jiménez J, Zandonade P, Moser RD (2004) Scaling of the energy spectra of turbulent channels. J Fluid Mech 500:135–144

Ganapathisubramani B (2004) Investigation of turbulent boundary layer structure using stereoscopic particle image velocimetry. Ph.D. thesis, University of Minnesota

Ganapathisubramani B, Longmire EK, Marusic I (2003) Characteristics of vortex packets in turbulent boundary layers. J Fluid Mech 478:35–46

Ganapathisubramani B, Longmire EK, Marusic I, Pothos S (2005a) Dual-plane PIV technique to determine the complete velocity gradient tensor in a turbulent boundary layer. Exp Fluids 39:222–231

Ganapathisubramani B, Hutchins N, Hambleton WT, Longmire EK, Marusic I (2005b) Investigation of large-scale coherence in a turbulent boundary layer using two-point correlations. J Fluid Mech 524:57–80

Ganapathisubramani B, Longmire EK, Marusic I (2006) Experimental investigation of vortex properties in a turbulent boundary layer. Phys Fluids 18:055105

Hoyas S, Jiménez J (2006) Scaling of the velocity fluctuations in turbulent channels up to $Re_\tau = 2003$. Phys Fluids 18:011702

Kähler CJ (2004) Investigation of the spatiotemporal flow structure in the buffer region of a turbulent boundary layer by means of multiple plane stereo PIV. Exp Fluids 36:114–130

Kim J, Moin P, Moser RD (1987) Turbulence statistics in fully developed channel flow at low Reynolds number. J Fluid Mech 177:133–166

Ligrani PM, Bradshaw P (1987) Spatial resolution and measurement of turbulence in the viscous sublayer using subminiature hot-wire probes. Exp Fluids 5:407–417

Marusic I (2001) On the role of large-scale structures in wall turbulence. Phys Fluids 13(3):735–743

Marusic I, Perry AE (1995) A wall wake model for the turbulent structure of boundary layers. Part 2. Further experimental support. J Fluid Mech 298:389–407

Moin P, Mahesh K (1998) Direct numerical simulation: a tool in turbulence research. Annu Rev Fluid Mech 30:539–578

Moin P, Spalart P (1987) Contributions of numerical simulation databases to the physics, modeling, and measurement of turbulence. Technical report TM10022 NASA

Perry AE, Marusic I (1995) A wall wake model for the turbulent structure of boundary layers. Part 1. Extension of the attached eddy hypothesis. J Fluid Mech 298:361–388

Pompeo L, Thomann H (1993) Quadruple hotwire probes in a simulated wall flow. Exp Fluids 14:145–152

Suzuki Y, Kasagi N (1992) Evaluation of hot-wire measurements in wall shear turbulence using a direct numerical simulation database. Exp Therm Fluid Sci 5:69–77

Theodorsen T (1952) Mechanism of turbulence. In: Proceedings of the second Midwestern conference on fluid mechanics, March 17–19

Zhou J, Adrian RJ, Balachandar S, Kendall TM (1999) Mechanisms for generating coherent packets of hairpin vortices in channel flow. J Fluid Mech 387:353–396