Experiments in Fluids manuscript No.
(will be inserted by the editor)

R.T. Reynolds · P. Hayden · I.P. Castro · A.G. Robins

Spanwise variations in nominally two-dimensional rough-wall boundary layers

Received: date / Accepted: date

Abstract Laboratory experiments have been conducted in two separate boundary layer facilities to investigate steady spanwise variations in mean velocity discovered during studies of developing flows over regular arrays of large roughness elements. Regular spanwise variation was found with a steady wavelength, moderated by the growing boundary layer, which was an integer multiple of the repeating unit of roughness. Amplitude variations greater than $\pm 5\%$ in the mean were found over the roughness and greater than $\pm 10\%$ in turbulence quantities. Due to the dominating nature of this phenomena throughout the layer, care should be taken in undertaking local measurements aimed at identifying flow variations caused by roughness heterogeneity.

Keywords spanwise variation · secondary instability · rough-wall turbulent boundary layers

1 Introduction

In boundary layer flows over very rough surfaces, defined as surfaces whose roughness element height, h, is typically 5-20% of the boundary layer depth, δ, it might be anticipated that spanwise variabilities could be present, particularly if the roughness geometry is regular - i.e. a pattern which repeats in at least the spanwise direction. Such variabilities might occur even well above the roughness sublayer within which the spatial inhomogeneity reflects directly that of the roughness geometry (**?**). Intuitively, they *must* be present at very low δ/h (O(1), say) since then the 'boundary layer' is more like a free flow over a series of obstacles. If spanwise structure throughout the flow continues to large axial distances, by which time $\delta/h \gg 1$, detailed study of the nature of the inhomogeneities within the roughness sub-layer (e.g. **?**) could be inhibited, since variations would be partly caused by these larger-scale structures.

It has long been known that nominally two-dimensional smooth-wall boundary layers can exhibit spanwise variability in surface stress arising from imperfections in settling chamber screens. **?**, following the work of **?** and others, made measurements in nominally two dimensional smooth wall boundary layers and identified significant spanwise variabilities, finding variations in surface shear stress up to 10% or more created by spatial instabilities of the flow through the wind tunnel damping screens. He suggested that screens with open-area (porosity) ratios of more than 0.57 do not produce appreciable spanwise variation. Further, he showed theoretically the sensitivity of boundary layers to periodic variation in the lateral direction and suggested that either small variations in the open-area ratio or low porosity (high solidity) can lead to jet coalescence. This is a steady process documented with flow visualisation by a number of researchers and characterised by numbers of jets emerging from adjacent orifices grouping together, since they can only entrain fluid from each other, leaving fewer, larger jets shortly downstream of the orifices. **?** confirmed this idea, describing the bi-stable coalescing of jets from grids with an open-area ratio of 0.58, leading to persistent mean velocity variations and greatly reducing the decay rate of newly formed turbulence. They suggested that successful techniques to remedy this spanwise variation must eliminate the non-uniformities in the mean velocity early or prevent their formation in the first place.

The first order effect of the screens, as described by **?**, is to create a drop in the static pressure proportional to the square of the flow velocity and to refract the incoming flow in the direction normal to the screen. One factor affecting the pressure drop, the Reynolds number of the flow through the screen, based on wire diameter, is also an important consideration. **?** suggested that when tur-

R.T. Reynolds · I.P. Castro
School of Engineering Sciences, University of Southampton, Highfield, Southampton, SO17 1BJ, UK
Tel.: +44-2380-593394
Fax: +44-2380-593058
E-mail: i.castro@soton.ac.uk

P. Hayden · A.G. Robins
Fluids Research Centre, University of Surrey, Guildford, GU2 7XH, UK

bulence reduction is to be independent of the free-stream velocity, Reynolds numbers of fifty or less – corresponding to subcritical or critical Reynolds numbers – should be avoided. **?** also noted that though subcritical screens yield high reductions in turbulence, the pressure drop and sensitivity to clogging by dust in wind tunnel applications makes them undesirable. Supercritical screens shed vortices from the cylindrical wires whereas subcritical screens do not.

An extensive study of spanwise variation during laminar boundary layer testing was conducted by **?**. He described a series of modifications and tests done to reduce the amplitude of the variation. Small variations in the weave of the settling chamber screens were found to lead to non-uniformities throughout the flow. When interacting with the naturally unstable boundary layer, this generated streamwise vortices (thereby creating the variation in boundary layer thickness). Following work by **?**, the weak streamwise vortices within the boundary layer are defined as Klebanoff modes. In the study by **?**, all screens used were above the open-area ratio limit suggested by Bradshaw and operated at a range of subcritical and supercritical Reynolds numbers. However, Klebanoff modes were still present. Factors which reduced the spanwise variation incrementally included the use of new carefully tensioned screens, a new high quality axial fan to replace a blower fan, and ultimately a careful reordering of the tunnel screens based on laser scanning to identify weave variations. Changing the leading edge geometry of the roughness plate did not alter the wavelength or magnitude of the secondary flow.

Though normally encountered in transitional boundary layers (**?**), Klebanoff-type modes are also found elsewhere. They are essentially sets of longitudinal vortices whose spanwise spacing is of the order of the boundary layer thickness and whose amplitude can be significant. It turns out that one needs to take extraordinary measures to remove them.

A detailed look at the mechanisms generating spanwise structure and scale growth in smooth-wall boundary layers was conducted by **?**, following the particle image velocimetry experiments conducted by **?**. Earlier, **?** discussed the role of inflectional velocity profiles in the development and growth of instabilities in smooth wall bounded flows. They illustrated the effect of large scale streamwise structure on the flow. **?** took this further, showing how an inflection instability mechanism in a plant canopy can lead to quasi-streamwise stretched eddy structures.

If the boundary layer develops over an aerodynamically rough surface, one might hope that such modes would be less important but, as noted earlier, if δ/h is low enough there must be spanwise variability throughout the flow arising from the spanwise structure of the roughness. Spanwise variation has been found in studies over rough hills. **?** conducted tests over mesh-roughened hills which produced periodic spanwise variation. It has been suggested that the variation present over the sinusoidal waves in this case was due to a Craik-Leibovich type instability creating Rayleigh waves whose amplitudes diminish with distance from the wall (**?**). However, additional tests over rough hills by **?** cast doubt on that idea; although spanwise variation with a wavelength of the order of the boundary layer thickness was again found over rough hills, it was also found over plane mesh roughness without the hills. This suggested that the mechanism creating the longitudinal vortices was not a direct function of the hill geometry and wavelength, a conclusion reinforced by the results found here. There is a suggestion (**?**) that the presence of hills simply encourages the spanwise behaviour but is not a prerequisite for it.

This paper discusses the essential question in the context of naturally grown rough wall boundary layers: to what extent does the roughness impose or 'tune' the longitudinal modes which might anyway arise in the normal way (i.e. the Klebanoff modes) or be present initially because of the roughness geometry (i.e. at δ/h=O(1))? The following Section outlines the experimental arrangements, major results are presented and discussed in Section 3 and the most important conclusions are highlighted in Section 4. Initial results of this collaborative study on spanwise variation were presented at the 10$^{\text{th}}$ European Turbulence Conference in July 2004 and published in the proceedings (**?**).

2 Experimental Setup

Experimental data in smooth-wall (flat-plate) and rough-wall boundary layers in subsonic wind tunnels have been obtained. Measurements were made in two 0.6×0.9 m subsonic, open circuit wind tunnels with different upstream flow conditioning. The testing characterised the extent of the inhomogeneity and the persistence of the induced periodic variation.

Measurements were conducted with single and crossed hot-wire anemometry at positions, x, up to 3.5 m downstream of the roughness leading edge and at a nominal free-stream speed of 10 m s^{-1}. This yields a Reynolds number, based on free-stream speed and boundary layer thickness at x=310h where h is the roughness element height, of 7.8×10^4 for 10 mm cube roughness – one of the surfaces used, see below. Alternatively, using a friction velocity, u_τ, see Cheng & Castro (2002), the same roughness and location give $\mathrm{Re}_\tau = 4520$. The freestream spanwise uniformity of both tunnels was $\pm 0.2\%$ for mean velocity with streamwise turbulence intensities in the freestream below about 0.1%. No clear, spanwise (periodic) spatial inhomogeneities at those amplitude levels (e.g. around $\pm 0.2\%$ on mean velocity) were found in the free stream of either wind tunnel working section.

In the primary facility, testing was conducted in a 'suck-down' wind tunnel at the University of Southamp-

ton. Prior to a 6:1 contraction, honeycomb and wire mesh screens condition the incoming flow. The working section has a rectangular cross-section of 0.6×0.9 m and is 4.5 m in length. A diffuser expands the flow to an axial fan at the tunnel exit. A 75 kW motor drives the variable speed, constant pitch fan blades controlling the air speed. Probes were mounted in the test section using a three axis traversing system sitting on top of the working section. Movement in the streamwise axis was manually controlled, while vertical and spanwise motion was computer controlled. The two traversing system stepper motors were operated by Parker PDX-15 stepper/indexers controlled via RS232 serial connections. The spanwise axis was manipulated by a traditional lead-screw traverse while the vertical axis was manipulated by a Parker Electro-Thrust pneumatic cylinder, which protruded into the tunnel from the roof. The accuracy of the traversing system depends on the thread type, backlash, and motor step angle. For this setup using backlash compensation, the repeatable accuracy was 0.02 mm, which is only 1% of the smallest 2 mm stepping distance used during testing.

There are three flow conditioning screens placed in the settling chamber of the Southampton tunnel; these have open-area ratios, β, of 0.69, 0.58, and 0.48, with the solidity increasing in the downstream direction. At a test section free-stream velocity of 10 m s^{-1}, corresponding to a convenient testing condition for urban roughness studies, the Reynolds numbers of the screens are 55, 55, and 22, respectively. Here the Reynolds number is defined as $Re_d = Ud/\nu$ with d, the wire diameter and U, the normal velocity entering the screen. Based on these numbers and the earlier discussion, it was therefore expected that some spatial inhomogeneities may occur in the settling chamber and thus that spanwise variations may be present in the working section wall boundary layers.

The second facility, at the University of Surrey, consisted of a 'blow-down' tunnel. A centrifugal 37 kW fan drives the flow through a wide-angle diffuser containing three screens and having a 48° included (double) angle and a 2.3 area ratio. Flow conditioning within the subsequent settling chamber is achieved with a honeycomb and four additional screens, each having $\beta = 0.59$. A 6.3-1 contraction then accelerates the flow into the test section, which has a 0.6×0.9 m cross-section and a 4.5 m length. The flow exits vertically into the laboratory *via* turning vanes located just downstream of the working section. A computer controlled traversing system was used to manipulate tunnel instrumentation through the test section ceiling via movable slats. At a working section velocity of 10 m/s the Reynolds number of the settling chamber screens was around 42. Again, one might expect perceptible spanwise variations in the working section boundary layers.

The urban-type roughness consisted of an array 10 mm cubic roughness elements with a 25% area density coverage, as shown in figure 1. The array is in a staggered configuration, though an aligned array configuration is possible by rotating the roughness. Additional testing included measurements on 20 mm cube roughness in the same configurations and also 10 mm square, random height roughness elements which had the same density and total volume. Note that, for the staggered configuration shown, the 'wavelength' of the first row of roughness seen by the flow is $2h$, which is in fact the spanwise wavelength of the repeating unit – defined as the smallest roughness patch which when tessellated in both directions produces the entire surface; we term this latter wavelength Δ. The axial wavelength of the repeating unit is $4h$ and this obviously becomes the spanwise wavelength for the aligned cube case. Testing also included wire mesh roughness and a smooth (flat-plate) surface. More detailed information about the cube roughness configurations can be found in **?**. A leading edge ramp was used at the entrance to the working section to sweep the flow up and onto the roughness panels, which for these studies were placed on the floor of the wind tunnel. No device was used to trip the boundary layer as the roughness surface created turbulent flow and separation regions between elements immediately.

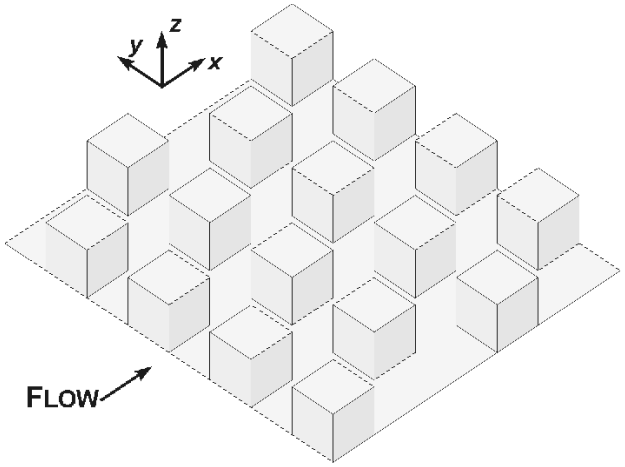

Fig. 1 The roughness array. Shown with flow direction for staggered array. Aligned array obtained with 90° flow rotation.

Flow diagnostics at both Surrey and Southampton were obtained with hot wire anemometry (HWA). Additional measurements of dynamic pressure, reference pressure, and reference temperature were also made. All data were collected and processed *via* computer with software designed specifically for the corresponding measurements. The dynamic pressure measurements were made using pitot-static tubes connected to Furness Controls Limited micromanometers and reference atmospheric pressure and temperature were collected manually (automatically at Surrey) from a digital instrument located in the laboratory.

The acquisition system used for sampling analog voltage signals (from any source, but typically the micromanometer and hot wire anemometry output) consisted of a 16-bit IOTECH ADC488/8SA with eight channel simultaneous sampling at up to 100 kHz. The digital signals were sent via IEEE 488 cable to an Apple G4 computer with a National Instruments PCI-GPIB data acquisition card.

Constant temperature HWA was used to obtain high frequency velocity data with hot wire probes (from Dantec or Auspex) driven by a hot wire bridges manufactured by the University of Newcastle, New South Wales. An overheat ratio of 1.5 for the wires reduced the wire response to temperature fluctuations. Typical wire frequency response was at least 20 kHz. The hot wire bridge signals were passed through signal conditioning circuitry mounted in the Newcastle Rack system. Electronic filtering removed high frequency content to prevent aliasing, and voltage offset and gain were applied to manage the output voltage signal to the A/DC so as to maximise resolution.

Data acquisition was controlled using the EnFlo software created by the University of Surrey using National Instruments Labview software. The comprehensive acquisition and processing package was able to operate the traverse, acquire signals from the A/DC, and process and output the results in a variety of ways. The software system was run on an Apple G4. Very similar, but not always identical, systems to those described above were used to obtain the data in the Surrey experiments.

3 Results and discussion

Measurements were made at a number of downstream stations in a number of different flows; Table 1 summarises the major cases and locations studied. Note that for the rough-wall cases z (and δ) is measured from the base of the roughness elements and in all cases $y = 0$ is the spanwise centreline of the working section. Spanwise variation was found, to some degree, in all the boundary layers tested and existed throughout the boundary layer depth, not just, for example, in a roughness sublayer. For cases of small roughness ($h/\delta < 3\%$) the wavelengths of the spatial oscillations, when detectable, were of the order of the boundary layer depth and had amplitudes typically less than $\pm 2\%$. The cube roughness arrays, in contrast, produced very regular spanwise wavelengths. The amplitude of the oscillation was also more significant, generating variation in some locations of over $\pm 5\%$ in mean streamwise velocity, $\overline{U}$, and $\pm 10\%$ in streamwise turbulence intensity, $\sqrt{\overline{u'^2}}$. For the cube roughness, a wavelength of $8 \times h$ (i.e. 2Δ) was dominant over much of the fetch.

Vertical profiles of mean longitudinal velocity were required to determine the variation of boundary layer depth with distance downstream. Figure 2 shows typical examples of mean velocity and turbulence intensity obtained on the spanwise centre-line and plotted against z/δ, where δ is taken as the distance from the bottom of the roughness elements to the point where the mean velocity was 99% of the free-stream value. Note that in the case of the rough-wall flows the zero plane displacement, d, has been subtracted from both z and δ. d was taken to be the value found after extensive work over the same surfaces by Cheng & Castro (2002). The figure shows that there is a small development in both U and u' with fetch. Estimated absolute errors in U and u' are below$\pm 1\%$ and $\pm 3\%$, respectively; the experimental scatter and repeatability in the data is usually significantly better than that.

10 mm cubes		20 mm cubes		Smooth	
x, mm	δ, mm	x, mm	δ, mm	x, mm	δ, mm
105	30.3	220	54.1	300	10.0
185	33.4	370	61.7	750	16.2
345	39.4	700	71.8	1650	28.4
655	50.6	1050	84.7	2590	40.0
1015	61.4	1600	108	3430	50.0
1245	68.5	1840	115.5		
2180	93.8	2240	127.5		
3130	117	3040	148		
2.0-7.8		3.6-9.9		2.0-7.8	

Table 1 Major x-locations of data collection, with corresponding boundary layer thickness, for the three surfaces. The final row gives the ranges of Re x 10^{-4} for each surface.

Figures 3 and 4 show examples of the time-averaged streamwise mean velocity and turbulence intensity variations, expressed as percent changes from the spanwise average values, at z/δ=0.6 in the boundary layer developing over the 10 mm (staggered array, uniform height) cube roughness. An 80 mm wavelength is visible in both cases and the mean velocity and turbulence intensity vary by over $\pm 5\%$ and $\pm 8\%$, respectively, even at this height in the boundary layer. The presence of such large-amplitude structure in the flow, in what would be considered the outer layer of the boundary layer, was initially unexpected. No such spanwise variations were present in smooth wall boundary layers developed in the same wind tunnels; indeed, typical spanwise 'scatter' in mean velocity (with no dominant wavelength) was below about $\pm 1\%$ at all axial locations. Likewise, as noted earlier, we found no analogous spanwise behaviour over rough surfaces having much smaller h/δ (e.g. the 'mesh' roughness). It must be concluded that the variations seen in figs.3 and 4 are a result of the roughness geometry and the relatively large h/δ.

Given the work described above, tests in the two separate facilities were undertaken in an attempt to isolate the mechanisms associated with the development of the large magnitude spanwise variations. Although the flow conditioning and contraction sections of the two facilities were similar overall (as described earlier), they differed in

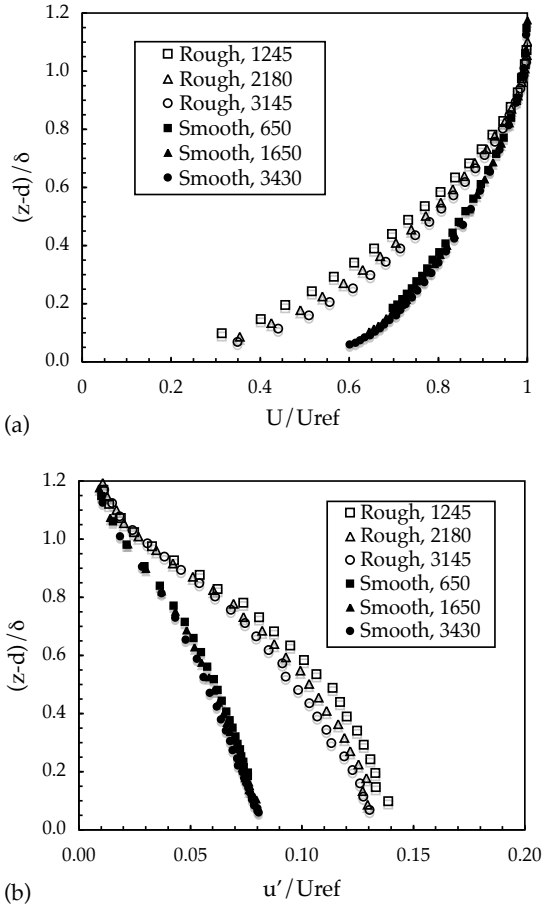

Fig. 2 Profiles of mean velocity (a) and rms longitudinal turbulence (b) for the smooth wall and the staggered 10 mm cube surface; values in legends refer to distances in mm from the leading edge. $d = 0$ and 7.5 mm for smooth and 10 mm cube surfaces, respectively.

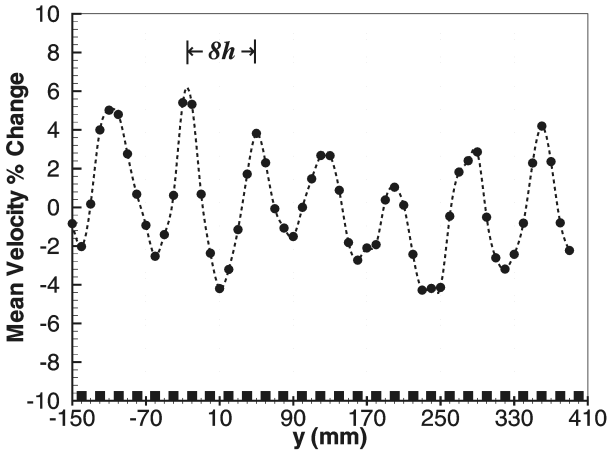

Fig. 3 Percentage variations in streamwise mean velocity with spanwise location. 10mm staggered cubes, z/δ=0.6, x/h=125.

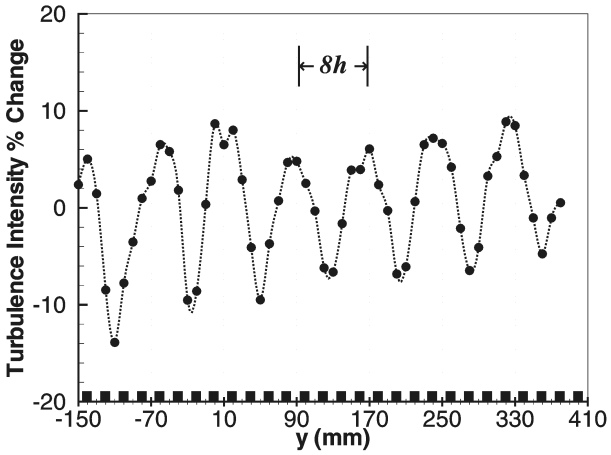

Fig. 4 Percentage variation in streamwise turbulence intensity with spanwise location. 10mm staggered cubes, z/δ=0.6, x/h=125.

fine detail so it was thought that the amplitude (at least) of any spanwise variations might differ appreciably. The test sections were both 0.6 x 0.9 x 4.0 m, accommodating the same exact panel mounted roughness sections. Figure 5 shows a comparison of spanwise profiles taken at similar locations in each facility. Though the collapse is not perfect, the same pattern is evident with identical wavelengths and very comparable amplitudes.

During the course of the study, the flow conditioning screens in the Southampton facility were removed and thoroughly cleaned. A considerable build up of dust had accumulated on the screens; the finest ($\beta = 0.48$), most downstream screen was largely clogged. After cleaning the screens, measurements over the 10 mm cube staggered array roughness were repeated. Further tests were conducted after removing the finest screen; $\beta = 0.48$ is well below the 0.57 lower limit suggested by **?** and might thus be expected to lead to jet coalescence. Cleaning the screens reduced the magnitude of the mean streamwise velocity variation somewhat, but removing the finest screen had little effect on the velocity pattern. This may indicate that the porosity of the second screen (0.57) is also too low, or that porosity at these levels is not, in fact, a significant variable for controlling the behaviour within the working section boundary layers.

The amplitude of the spanwise variation decreased with height in the boundary layer over the cube surfaces. This is seen, for example, in Figure 6 over a roughness comprising square random height roughness obstacles with an average height of 10 mm, and an 80 mm repeating unit. At this x=1560 mm location the wavelength remained constant with height. In general, over any of the cube surface arrangements, one wavelength dominated the flow at a given streamwise location, independent of wall-normal location. Only during wavelength transition, as discussed below, was the dominant wavelength a function of height over the cubes.

At a given relative height in the boundary layer, the decay in the oscillation amplitude with fetch can be seen

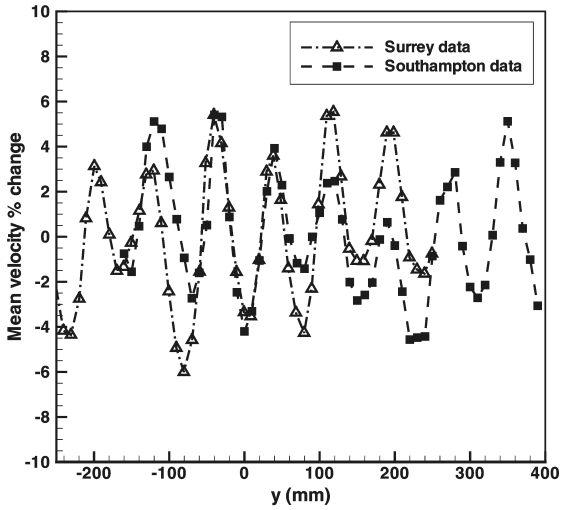

Fig. 5 Comparison of spanwise variation in two facilities. 10 mm staggered cube, z/δ=0.6, x/h=125.

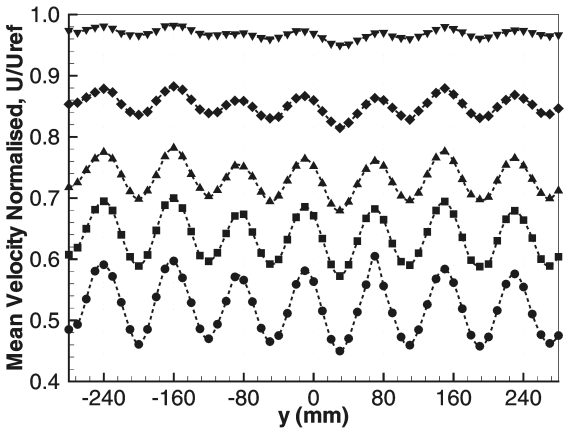

Fig. 6 Spanwise variation of mean streamwise velocity at (from top) z/δ=0.89, 0.67, 0.50, 0.39, 0.28. 10 mm random height roughness.

in Figure 7. In this case, the amplitude is characterized by its average peak-to-peak variation across the span, normalized by the mean value itself. The results for the 20 mm and 10 mm staggered cube are compared with smooth wall data taken in the same facility over a flat plate. The initially high amplitude variations over the cubes decay with fetch towards the smooth wall values, which remain relatively constant with downstream location. A noticeable distinction exists for the 10 mm cube,compared to the larger 20 mm cube, in that the peak-to-peak amplitude remains above 5% for a largeportion of the streamwise distance downstream. This discrepancy will be discussed further below. Notethat the streamwise direction, x, is not scaled with roughness dimension, h, in fig.7 to facilitate comparison with a smooth surface.

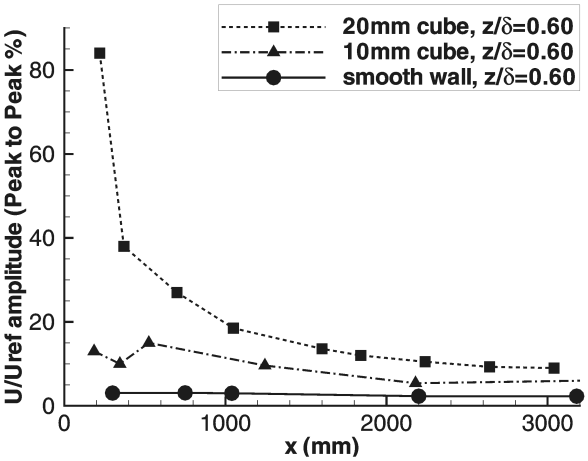

Fig. 7 Spanwise peak-to-peak amplitude decay with streamwise fetch for three surfaces.

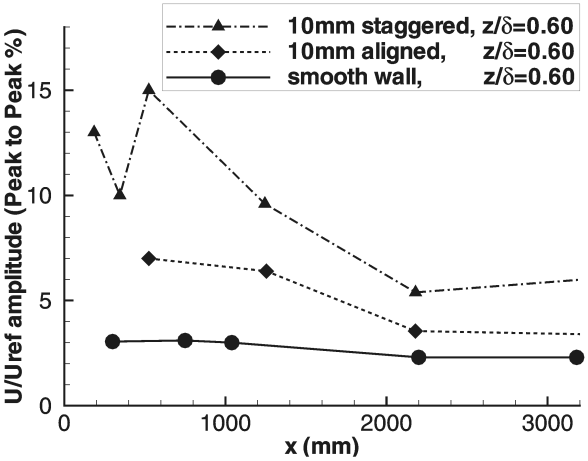

Fig. 8 Spanwise peak-peak amplitude decay with cube arrangement.

It was found that altering the arrangement of the roughness (and thus the wavelength of the repeating unit) affected both the amplitude and the wavelength of the spanwise variation. Figure 8, for example, shows a comparison of amplitude variation for the 10 mm cube roughness in staggered and aligned configurations. Although intuitively one might expect the aligned arrangement to yield larger-amplitude spanwise variation due to the regularly spaced 'street canyons' aligned in the streamwise direction (see Fig.1), it is the staggered array which actually creates larger variations having a more consistent wavelength. The aligned configuration amplitude eventually tends toward the smooth wall value of just under 2%, whereas the staggered array amplitude remains higher. Similar behaviour was apparent at other heights in the boundary layer.

Additional properties of the spanwise variations were identified using crossed hot wires, which yielded measurements of the vertical mean velocity component. Fig-

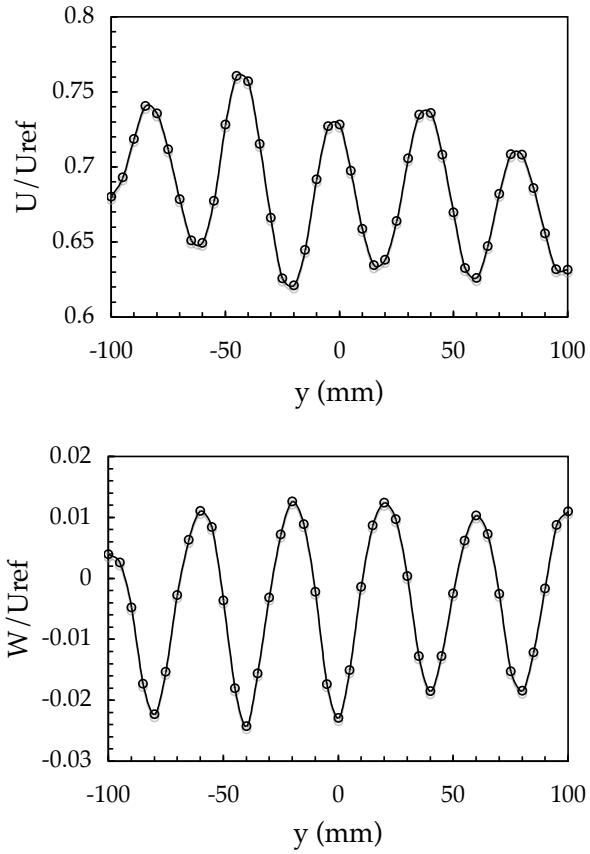

Fig. 9 Phase shift between streamwise and wall normal velocity fluctuations. 20 mm staggered cube, z/δ=0.6, x/h=11.5.

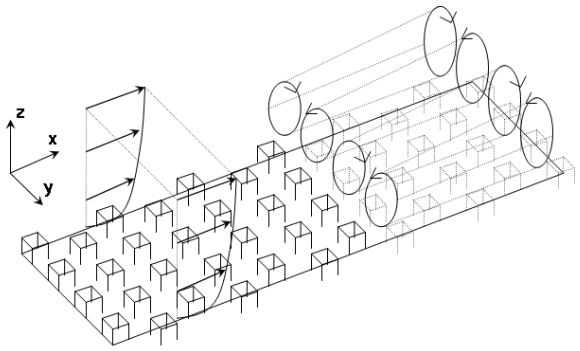

Fig. 10 Visualization of weak streamwise vortex organization in boundary layer over cubes.

ure 9 shows the streamwise and wall normal velocity variations with spanwise location. A phase shift of 180° is visible so that the wall normal velocity reaches local maxima where minima in the streamwise velocity occur. This immediately suggests that the spanwise variations are caused by streamwise (longitudinal) vortex pairs. Figure 10 shows a sketch of a possible arrangement of a secondary flow structure which would be consistent with the kind of data in figure 9 (although at twice the wavelength). The presence of this weak vortex system within a strong shear layer also explains the amplitude decrease with height. The vortex system is constrained by the wind tunnel side walls, which may explain their steady nature (this is a subject of further work). The vortex pairs are also stretched by the growing boundary layer. Ultimately, one could speculate that the stretching is unsustainable and the vortex pairs merge, leading to a doubling of the wavelength, as discussed further below.

Further insights can be gained by looking at the development of the spanwise wavelength with distance downstream. At some axial locations, the spanwise variations clearly contained more than one spatial frequency. To identify the dominant mode, (spatial) spectra of the data were obtained *via* 'autocorrelations' formed from the spanwise 'signals' typified by those shown in the earlier figures. Examples of the spectral development with fetch are given in Figure 11, for the 10 mm cube roughness. All the spectra shown in the figure were obtained at $z/\delta = 0.6$. Notice first that near the leading edge ($x/h \leq 20$, say) the spectra have sharp peaks at a wavelength, which we call hereafter λ, around 20 mm, i.e. $2h$. This is perhaps not surprising, for it reflects the likely dominance of the flow over the first row of cubes. Likewise, for the 20 mm cube roughness (not shown), the initial peaks had $\lambda = 40$ mm. In both cases these early peaks disappear further downstream. It is clear from fig.11 that until $x/h \leq 200$ the dominant peak has $\lambda \approx 80$ mm, i.e. 4Δ. At $x/h = 218$ two peaks are evident but further downstream the lower wavelength peak disappears, leaving a peak around $\lambda = 8\Delta$. The behaviour is illustrated in Figure 12, for both the 10 mm and the 20 mm cube surfaces. In regions of fixed λ, λ/δ naturally falls as the boundary layer grows downstream but at certain axial locations secondary, doubled wavelength peaks appear and the lower ones eventually disappear. The vertical lines are merely illustrative of approximate locations where the higher peak becomes dominant..

Similar results were obtained at other z/δ, but the axial fetch at which the large λ peaks appear while lower ones disappear was found to differ. This is illustrated in the summary sketch of Figure 13, where λ values deduced from spectra like those in fig.11, at $z/\delta = 0.35$ and 0.6, are superimposed on a plot of the boundary layer growth with fetch, at the locations where the spectra were obtained. The diagonal dashed lines are approximate demarcations between regions in which λ was fixed, at the four integral multiples of Δ and without evidence of any other significant spectral peak. Notice the position on the $z/\delta = 0.35$ line (at $x/h \approx 300$) where two spectral peaks, at $\lambda = 8h$ and $16h$ were evident. Of course, the presence of two peaks does not suggest that two vortex systems co-exist - merely that they probably exist alternately but randomly. It appears that once the boundary layer depth is large enough to allow a vortex doubling process to be-

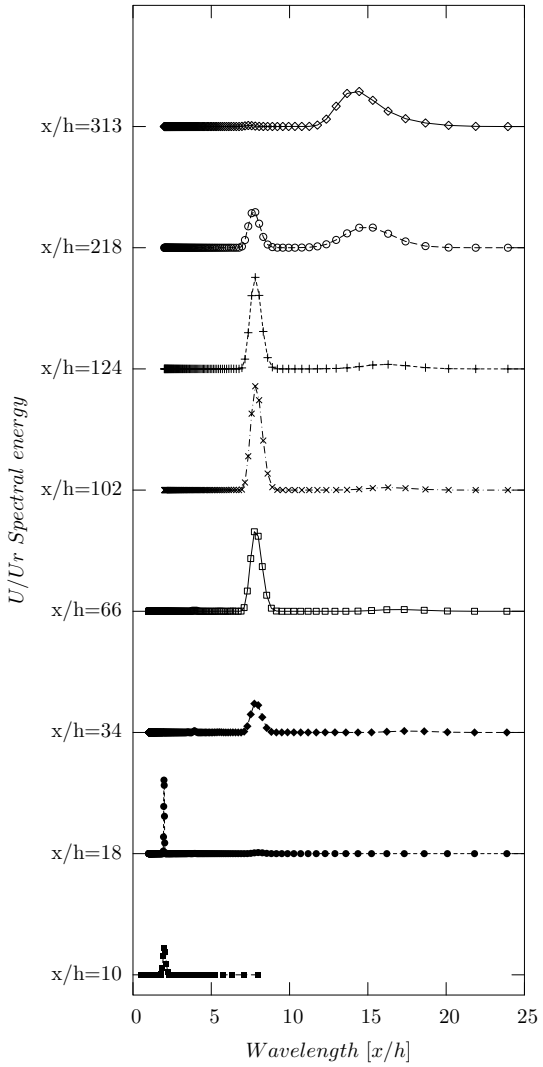

Fig. 11 Spectra of the spanwise variations in mean velocity over the 10 mm staggered cubes at $x/\delta = 0.6$ and for various distances downstream.

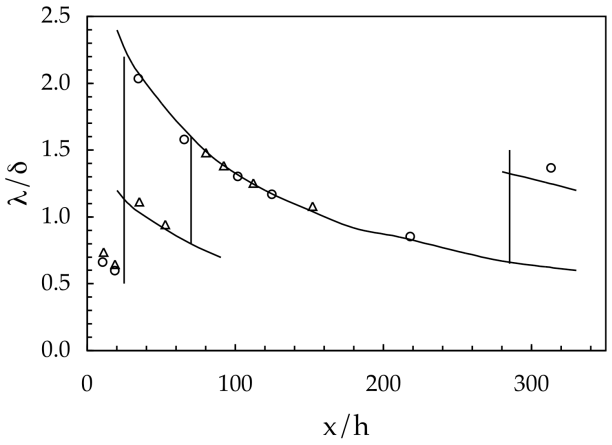

Fig. 12 Variation with streamwise fetch of the dominant wavelength, deduced from spatial spectra. z/δ=0.6, staggered cube arrays.

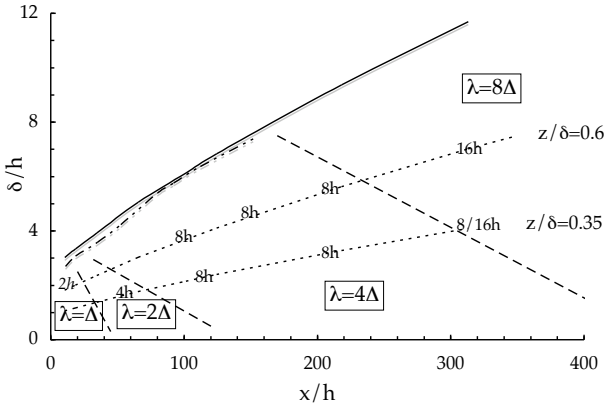

Fig. 13 Scaled sketch of wavelength doubling with boundary layer growth. Wavelength, λ, as multiple of spanwise repeating unit, Δ.

gin – even spasmodically – this is evident first in the outer region, becoming more regular with fetch and then eventually seen nearer the surface. Note that the vortex pattern sketch in fig.10 refers to the $\lambda = 2\Delta$ region in fig.13. There is some uncertainty about the extent of this particular region; the 10 mm cube spectral data contained only a hint of its presence at $z/\delta = 0.6$ and no data are available at $\delta = 0.35$ for $x/h \leq 100$. Nonetheless, the overall agreement between these two staggered cube array surfaces is striking.

4 Concluding remarks

Overall, the results suggest that particular roughness geometries can amplify and organize spanwise variations in turbulent flow which seem similar to those resulting from the Klebanoff modes reported over smooth walls in transitional regimes. These variations have magnitudes which are significant and can exhibit well-defined wavelengths under specific arrangements. The amplitude variation increases with decreasing height in the boundary layer and decreases with fetch, but remains evident even beyond 300 roughness element heights downstream from the leading edge. The wavelength initiated by the leading edge of the roughness disappears relatively rapidly and, thereafter, the wavelength initially matches the wavelength of the roughness repeating unit but then periodically doubles as the boundary layer grows. Our data are consistent with the speculation that these spanwise variations are signatures of longitudinal vortices whose size is 'tuned' by fixed geometry of the roughness pattern

and the growing boundary layer depth. They seem not to occur (or have amplitudes too small to be detectable) once δ/h is sufficiently large and at least O(10).

Additional testing is now underway to investigate the influence of the upstream ramp on the initiation of a vortex structure. Work is also planned to alter the spanwise domain of the wind tunnel under the assumption that the vortex pairing is controlled not only by the boundary layer depth but also the available spanwise width. Other results to be reported include the effect of roughness packing density and the upstream boundary layer depth approaching the roughness. The influence of free-stream velocity has not been explored. In separate PIV studies (Reynolds, 2005), spatial correlation analysis supports the presence of spanwise variations above the roughness sublayer, with streamwise bands of negative correlation running alongside positive correlation results.

The existence and character of spanwise flow variations in rough wall boundary layers seem not to have been reported previously and should be considered when planning wind tunnel experiments over rough surfaces, particularly those having element height to boundary layer depth ratios in excess of a few percent. It is likely that similar behaviour would occur over any rough surface of sufficient amplitude and containing identifiable units of repeating roughness pattern.

Acknowledgements Partial funding for this work was from the Urban Meteorology Programme under the Universities Weather Research Network (Natural Environment Research Council Grant DST/26/39). RTR is grateful for support from the School of Engineering Sciences. We thank Mr. T. Lawton at the University of Surrey for consultation.