Available online at www.sciencedirect.com
ScienceDirect
International Journal of Heat and Fluid Flow 27 (2006) 902–910
International Journal of HEAT AND FLUID FLOW
www.elsevier.com/locate/ijhff

Direct numerical simulation of large-eddy-break-up devices in a boundary layer

P.R. Spalart [a,*], M. Strelets [b], A. Travin [b]

[a] Boeing Commercial Airplanes, P.O. Box 3707, Seattle, WA 98124, USA
[b] Federal Scientific Center "Applied Chemistry", St. Petersburg 197198, Russia

Received 29 September 2005; received in revised form 18 January 2006; accepted 4 March 2006
Available online 15 June 2006

Abstract

Turbulent boundary layers create aerodynamic noise inside all vehicles, and especially inside jetliners. The objective of this project is to modify the turbulence in an attached boundary layer, not to reduce skin friction, but to weaken the wall pressure fluctuations, or to shift them to less damaging frequencies and wavelengths. In order to benefit an entire airliner windshield, the effect would be sustained over about 25 boundary-layer thicknesses, δ, which far exceeds the common rule that the relaxation of the turbulence takes about 10δ. Various flow-control devices are studied by Direct Numerical Simulation, which is free of modeling and provides the full details of the pressure field. The Reynolds number is far lower than in the real flow, but this limitation of DNS is tolerable, because the focus is on the larger eddies. The multi-block implicit numerical method can represent fairly complex devices at a manageable cost. Inflow turbulence is provided by a recycling procedure derived from that of Lund, Wu and Squires, but much simpler. It occupies less than 5δ in the streamwise direction. Flow visualizations, Reynolds stresses, and spectra are shown; the baseline spectra are within the experimental scatter. Co-rotating vortex generators are tried first. They reduce the turbulence intensity away from the wall, as hoped, but actually intensify the wall pressure fluctuations and were therefore abandoned. Large-eddy-break-up devices resembling a highway bridge are tried next, and succeed in reducing the wall fluctuations, but only over about 6δ. Thus, the technology is not successful yet for a windshield, but it might be applied to other windows or other vehicles, and the simulation methodology appears to be well developed and of some interest particularly regarding inflow conditions.

Keywords: Turbulence control; Boundary layer; Pressure fluctuations; Direct numerical simulation; Turbulent inflow condition; Vortex generators; Large-eddy-break-up device

1. Introduction

Aerodynamic noise is a factor for pilot and passenger comfort in most vehicles. Insulation and other measures add to the cost and weight of the system, and a reduction at the source would be most welcome. Obvious sources of aerodynamic or "airframe" noise are regions of massive separation, such as behind a rear-view mirror, a windshield wiper, a spoiler, or landing gear. Separation off a smooth surface, especially caused by a shock wave, is also troublesome, and has been controlled with vortex generators (Pearcey, 1961; ESDU, 1993; Ashill et al., 2005). Many cars have small devices on their mirrors or sun-roofs; these cannot suppress massive separation, but it appears that they alter the turbulence sufficiently to provide a noise benefit. Regarding airframe noise from airliners and particularly landing gear, solutions have been tested with some success (Dobrzynski et al., 2002), but often were impractical. No obvious steps have been taken in production, even on the most recent models. The industry remains secretive about other measures, for instance applied to the high-lift system.

* Corresponding author.
E-mail address: philippe.r.spalart@boeing.com (P.R. Spalart).

0142-727X/$ - see front matter © 2006 Elsevier Inc. All rights reserved.
doi:10.1016/j.ijheatfluidflow.2006.03.014

This state of the art makes the interior noise generated by an attached flow appear as the minimum that must be accepted, but our purpose here is precisely to manipulate an attached boundary layer and lower its impact, although only over a limited area, nominally a single but large window on an airliner. The applicability of a successful concept could be very wide, and include ground vehicles. On the other hand, the turbulent boundary layer (TBL) is rightly considered as a very "robust" flow type, one for which modifications beyond a few % are very difficult to effect with simple devices. A case in point is the large amount of work expended on riblets, which has yet to lead to any industrial application (although much of this is due to issues of cost and durability, which are separate from the issues of fluid mechanics). The same applies to large-eddy-break-up (LEBU) devices. It proves very difficult to defeat the turbulence in setting the skin-friction coefficient, short of obtaining laminar flow. In addition, the TBL is robust in the sense of returning to normal quite rapidly, in a distance of the order of 10δ, where δ is the full thickness of the TBL. An exception is that embedded streamwise vortices extend much farther than 10δ; this has made them common tools for TBL control especially when the area in need of control is not anchored (Pearcey, 1961; ESDU, 1993). However, they act in the direction of increasing skin friction, which at first sight would increase the pressure fluctuations since the pressure *rms* in a standard boundary layer is proportional to the skin friction, multiplied by a weak function of Reynolds number. Their effect on turbulence intensity will be considered later, and could be opposite. Other devices will also be considered, drawing on past attempts at TBL control, but with the new situation that drag reduction is not the objective. For instance, if we imagine completely suppressing the turbulence over a region of interest, and paying a penalty only downstream, this could work if only a small enough patch is of interest. However, the study is limited to simple static external devices (as opposed to, for example, a suction system), with some size constraints due to visibility, so that such a deep reduction is not expected.

For drag reduction, even a figure such as 3% is impressive, but noise is measured with much larger ratios, and a marginal alteration of the TBL would not be sufficient. Since aerodynamic noise is not the only source and is typically comparable with engine and ventilation noise, a reduction by 3 dB for this one source is a plausible minimum for what would make a modification worth implementing. This represents a 30% reduction of the pressure *rms*, or halving the energy, which must result from a major manipulation. Another practical aspect is that the wall pressure field is filtered by the response of the flexible structure, be it skin or window. Therefore, the pressure *rms* in itself is only a first indication of the eventual benefit inside the cabin. The dominant frequencies, propagation velocities, and spatial pattern (summarized by the two-point, two-time correlations) also matter, and need to be provided by the simulations or experiments (unless the device is tested in situ, in flight).

This new set of objectives opens a very wide design space and, also, rules out any simulation approach other than direct numerical simulation (DNS) or large-eddy simulation. The dimensions of the device are free within practical limits, as is the type of device, as is even the mechanism postulated to obtain the desired effect on the turbulence. DNS is very restrictive in terms of Reynolds number, but it is assumed that the manipulation will primarily target the largest eddies, especially since (at least in the windshield region on an airliner) the critical region of the spectrum is for a Strouhal number, St, below 1. Here, St is defined as $f\delta/U_e$, where f is frequency and U_e the edge velocity of the BL. These largest eddies are not Reynolds-number dependent, when normalized with the shear stress and δ. A large-eddy simulation with the same resolution in the outer-layer would allow a higher Reynolds number, which would be felt very near the wall, but would be somewhat more costly, and also introduce Sub-Grid-Scale modelling assumptions; these are always candidates for adjustments, and raise the issue of whether they control the findings. DNS has the great advantage of making no turbulence assumptions, which is essential when studying a new type of turbulence modification, and of providing detailed flow visualizations which would suggest why a candidate device is working or not, and in which direction to change its dimensions. A promising design found in simulations may be then tested in a wind tunnel, with a number of high-performance probes, or directly in flight, which is faster and faithfully includes all the transmission effects. However, the DNS of TBL with an embedded device is challenging, in that relatively complex geometries must be treated. In addition, the simulation must be "spatial", meaning that it has good turbulent inflow conditions, and covers a strong modification of the TBL, its relaxation over 25δ or so, and an outflow. The solutions applied to this challenge are presented in the next section, while the following section presents the results and the last one the outlook for this line of work. In general, the accuracy standards adopted in this study are not quite as high as in DNS studies conducted in support of turbulence theory or fine determination of universal constants (Spalart, 1988); the emphasis is on rapid testing of ideas, that is accurate enough for a reliable assessment. If and when a concept is found to be promising, simulations with somewhat finer grids (particularly in the wakes of the devices), wider domains, and higher Reynolds numbers will be in order, as will experimental and ultimately flight testing.

2. Simulation approach

2.1. Numerical aspects

The DNS of a TBL with embedded devices such as vortex generators appears to represent a step in complexity, and is possible with the NTS code (Strelets, 2001) thanks to its combination of multi-block capability, implicit numerics, and moderate numerical dissipation. The code

provides the capability for time-accurate computations with structured multi-block overset grids, and employs an implicit second order in time flux-difference splitting scheme similar to that of Rogers and Kwak (1988). The flow is treated as incompressible. The inviscid fluxes in the governing equations are approximated with centred 4th-order accurate differences everywhere, except for the close vicinity of the devices' stagnation points where 3rd-order upwind approximations are used. This is needed to suppress the odd–even pressure oscillations caused by strong pressure gradients in these areas. The viscous fluxes are approximated with centred 2nd-order differences. At every time step, the resulting finite-difference equations are solved with the use of Gauss–Seidel relaxations by planes and sub-iterations in "pseudo-time".

2.2. Turbulent inflow conditions

This is the more novel aspect of the method. It is desirable to have a standard TBL encountering the flow-control device, without a lengthy approach region. Lund et al. (1998, LWS) show that inflow conditions based on random numbers can easily take 20δ to recover; this would be quite wasteful when the useful region is of the order of 25δ. Their recycling method, in some sense, wasted about 8δ, which is much smaller, and could probably have been reduced further without damage. The approach here is strongly inspired by LWS, but simpler, and a few separate improvements were made.

LWS applied a different treatment to the inner- and outer-regions of the TBL, and blended the two formulas. This may appear optimal, but it requires re-scaling the friction velocity in addition to the BL thickness, and introduces several arbitrary parameters and relatively complex formulas. Such an approach fails to take advantage of two favorable facts. One is that the near-wall turbulence re-generates itself much faster than the outer-region turbulence, and therefore little damage is done by applying the outer-region scaling throughout. The other is that when the recycling station is taken quite close to the inflow, which is desirable in terms of computing cost, the conflict between inner-region and outer-region scaling essentially vanishes. A short recycling distance also reduces the shift in the time scales and spanwise length scales, neither of which is accounted for (therefore, recycling is better justified with shorter distances, until the distance becomes so short that the recycling and inflow planes are so close that no natural evolution of the eddies can take place). For these reasons, the present approach uses a single re-scaling. Similarly, corrections to the wall-normal velocity component v have very little effect, and are omitted. Finally, here there is no auxiliary simulation. The primary simulation provides its own recycling data.

The approach consists in using the velocity field at a "recycling" station $x = x_r$ that is part-way inside the domain, suitably altered, to provide the velocity field at the inflow, $x = 0$. Since the BL is slightly thicker at x_r, a re-scaling in the wall-normal direction y is the minimum alteration that is needed. Let the thicknesses at $x = 0$ and $x = x_r$ be δ_0 and δ_r. Then the inflow condition for the velocity vector U is:

$$U(0, y, z, t) = U(x_r, y\delta_r/\delta_0, z + \Delta z, t - \Delta t),$$

where the spanwise shift, Δz, is introduced in order to keep turbulence at the inlet and recycling sections out of phase (in the simulations discussed below, the value of Δz was set equal to half the spanwise period). This is a minor step, taken to disorganize any durable spanwise variations of the mean flow, which would otherwise be recycled and possibly take much time to be damped by spanwise diffusion. Using the state from the last time step, $t - \Delta t$, is most convenient. The simplicity of the procedure is visible, compared with LWS, but the fundamental concept is theirs. The generalization to flows with pressure gradients and curved boundaries remains to be made.

The process is controlled by two parameters: the Reynolds number difference from inflow to recycling $\Delta R_x = x_r U_e/\nu$, and the ratio δ_0/δ_r. Assuming the approximate scaling given by Schlichting (1979) for the BL thickness, the flow should stabilize with the following value for the inflow thickness:

$$\delta = 0.37\left(\frac{\nu}{U_e}\right) \cdot \left[\frac{\Delta R_x}{(\delta_r/\delta_0)^{5/4} - 1}\right]^{4/5}$$

This gives fine control over the inflow thickness (even though the factor 0.37 is not exact, the slight difference in growth rates between the momentum thickness and the BL thickness is neglected, and the turbulence is not in a perfectly developed state), in that the response to small changes in ΔR_x and δ_r/δ_0 can be predicted (once the flow is well established).

The initial conditions for such a simulation are of some importance. For instance, random perturbations that are too weak or too inadequate in length scales can very well "die out." If so, the simulation will become laminar and produce a much thicker BL than expected with turbulence, because of the different scaling than in the equation above, which will make the failure obvious. This is more likely with a short recycling distance. Small values of ΔR_x and δ_r/δ_0 are cost-effective in the steady state, but impede the maturation of the turbulence that is needed initially. Therefore, it can be helpful to start with the recycling station father downstream, and then move it closer to the inflow by reducing ΔR_x and δ_r/δ_0 in harmony. Alternatively, if a flow-control device ends the normal BL and limits the position of the recycling station, the simulation could start with an extra length appended to the grid upstream, which would then be eliminated after the turbulence is well established, to reduce the cost of the long simulation needed for a sufficient time sample. Furthermore, the early part of the simulation does not need the entire length of the region that contains the device and the recovery, especially if the recycling parameters are being adjusted. In summary,

flexibility in restarting simulations from each other with different domains will be quite valuable, and several procedures are available to surmount the difficulties of the early regime, depending on the user's preference. Once the recycling procedure is well established, the solution loses memory of the initial-condition treatment.

The procedure used here benefits from another code, so that it is not universal. It is as follows. A RANS solution is calculated with the desired thickness, and perturbations added to it. These perturbations are obtained from a code that is used to initialize simulations of homogeneous turbulence. It yields "cubes" of turbulence of size 0.06^3, each with random phases, which are placed side-by-side on the wall after their velocity field is multiplied by a shape function, so that they are zero at the wall and in the freestream. This works well, although nothing is done to inject the correct Reynolds shear stress. A mature solution is obtained in a few flow-through times.

3. Results

3.1. Baseline flow

The simulation needs to be validated in terms of the standard TBL quantities, and of the wall pressure field. The unit length used below is roughly 1 m, for the boundary layer near the windshield of an airliner. The unit Reynolds number is 2.2×10^5 (which is far below full-size values) and the inflow thickness δ_0 is approximately 0.03, so that the momentum-thickness Reynolds number is 666. The length of the domain is 0.9, and the spanwise period is 0.12, which is over twice the outflow BL thickness and therefore quite sufficient. The recycling position x_r is 0.09, giving $\Delta R_x = 2 \times 10^4$ and $\delta_r/\delta_0 = 1.07$. The grid spacing in wall units is about $\Delta x^+ = 30$, streamwise, and $\Delta z^+ = 11$, spanwise, which is typical in DNS, and the first wall-normal spacing Δy^+ is about 1. This gives about 2.5 million grid points. The time step is 10^{-3}, normalized with edge velocity and unit length; the time sample contains 4500 steps.

Fig. 1(a) shows that the skin-friction coefficient C_f, although it first dips by about 7%, has stabilized by $x \sim 0.12$, which is only 4δ from the inflow. C_f is 0.0051 where $R_\theta = 740$, which is within 2% or 3% of the accepted value. The behaviour of the displacement and momentum thicknesses (not shown) is very smooth; the skin friction is not as smooth, but nothing of value to the technology selection would be learned by increasing the time sample until it is. Fig. 1(b) compares the four Reynolds stresses with past DNS studies, and is again very satisfactory except for a slight deficit in u' at the higher Reynolds number. The same applies to the mean velocity in Fig. 1(c). This suggests that the numerical method, grid spacing, time step, and recycling procedure are all capable of representing this TBL quite accurately.

Fig. 2 shows the lateral two-point correlation coefficient of wall pressure, which is of interest for two reasons. First, its behaviour for small Δz reflects the scales of the pressure-producing turbulent motion. The experimental comparison is quite satisfactory, and the appearance of a departure with finite slope illustrates the intensity of narrow structures. It is consistent with the physical basis for the log layer and $1/f$ behaviour of the spectra. Second, this curve reflects the adequacy of the spanwise period; ideally, it would fall to 0 well within the half-domain. Here, the larger domain used, of size 3.5δ, is satisfactory while the smaller one, 1.75δ, is marginal and better reserved for preliminary studies. The simulations shown below generally keep the ratio over 3.

Fig. 3 then compares the DNS wall-pressure spectrum with experiments and empirical laws. The Tupolev data and Efimtsov empirical laws were personal communications of R. Rackl of the Boeing Company, 2004. The Hodgson data were reported by Willmarth and Woolridge (1962) and by Willmarth (1975). Considering the emphasis on fairly low frequencies, an outer-layer scaling was chosen for frequency, based on δ and U_e. However, the inner-layer scaling based on wall units gives an even better comparison, in the appropriate frequency range. The pressure scaling uses the skin friction, instead of the freestream velocity; this is not supported by a full consensus in the community, but represents only a mild controversy for the present purpose. In the range of St between about 0.2 and 1, the DNS results are within 3 dB of the experiments; they are closest to Blake's (1970) once the Reynolds-number difference is considered. For higher St, they fall off due to the low Reynolds number, and do not have any of the St^{-1} range that is expected from theory. This is consistent with the absence

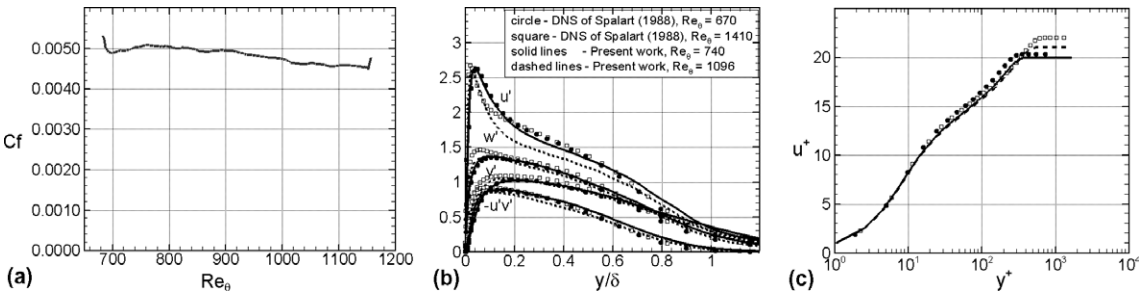

Fig. 1. Streamwise profile of the skin friction (a), wall-normal profiles of Reynolds stresses (b), and mean velocity (c).

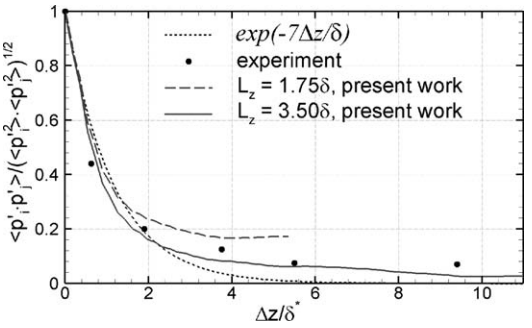

Fig. 2. Two-point correlation coefficient of wall pressure.

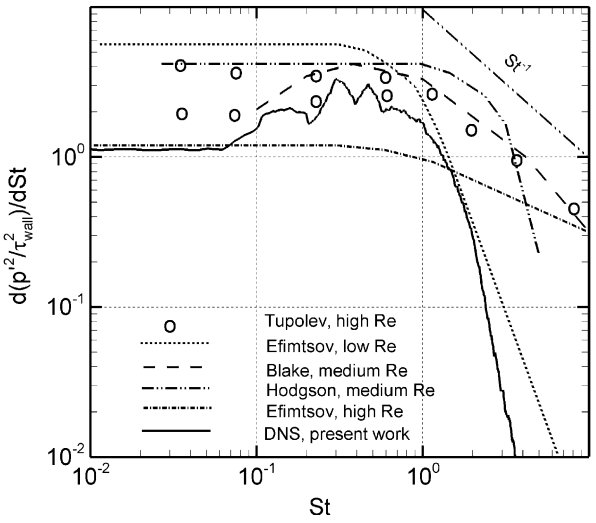

Fig. 3. Spectra of wall pressure fluctuations, baseline case. DNS, experiments and empirical laws.

of any extended logarithmic layer in the velocity profile (Fig. 1(c)). The region below $St \sim 0.2$ is more delicate. The time sample again appears marginally adequate, judging from the smoothness of the spectra, but the experimental scatter is more of an obstacle to a quantitative validation than the spectral roughness is. Also, the mechanics of the simulation system do not favour very low frequencies. Outer-layer eddies take about 0.13 time units to travel from inflow to recycling station, which corresponds to $St \sim 0.2$. Thus, frequencies much below $St = 0.2$ cannot be expected to have physical meaning. In experiments or in a global simulation, the spectrum cannot be expected to be unique to arbitrarily low frequencies either. At some point, the spectrum will show a memory of the transition process, and depend on whether transition was natural and contained turbulent spots, for instance. Wind-tunnel unsteadiness will also be felt. The empirical laws are level all the way to $St = 0$ more out of simplicity than based on conclusive data or theoretical reasons. In summary, care is needed in the region below 0.2, but it does not contain much of the energy. The bulk of the turbulence modification is expected in the region between 0.2 and 2, in which the DNS appears reliable and the Reynolds-number effect is weak. Therefore, meaningful simulations of the TBL manipulation are possible.

3.2. Vortex Generators (VG's)

The simulation shown fitted a grid around the VG's, but others were run using a body force to generate the vortices, and gave close results at a lower cost (Liu et al., 1996). The VG's are 0.025 tall and spaced by 0.12, and at 18° incidence; their planform is typical, and the circulation is expected to be of the order of 0.02. The grid is shown in Fig. 4; the total number of points is now near 3 million.

Fig. 5 illustrates the turbulence in its baseline state, and after the manipulation. The resolution of small eddies is quite fine, which supports the choice of grid spacing and algorithm. The typical random bulges are first seen, and indicate that the spanwise period is adequate; two-point correlations of wall pressure were studied and although they do not firmly reach 0 within the domain, they fall sufficiently low for evaluating new devices. They also agree well with experiment, roughly following the law $\exp(-7z/\delta)$; this will be essential when addressing the transmission through structure or glass. The dominant propagation velocity is near $2/3U_e$, and is also essential; it would be slightly higher at full-size Reynolds numbers. The figure then shows how, farther downstream, the single vortex per period gradually overturns the TBL, without strongly suppressing the smaller eddies. This creates a large bulge, but one that does not wander in time, as in the study by Liu et al. (1996).

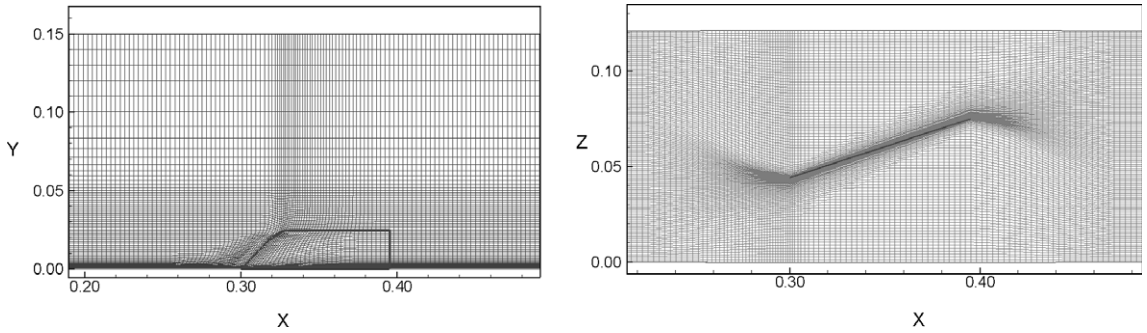

Fig. 4. Grid near Vortex Generator.

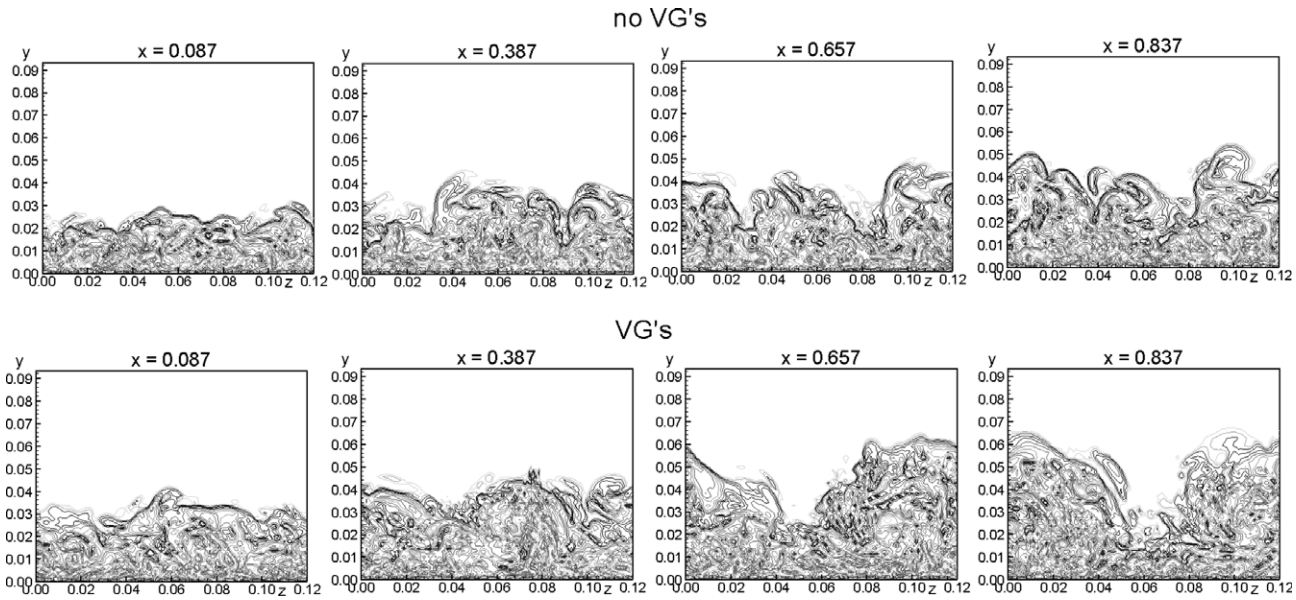

Fig. 5. Visualization of the boundary layer by vorticity, with and without VG's.

The rationale behind VG's for the present purpose was the following. Rapid rotation is known to suppress turbulence, giving hope that the wall pressures would also be calmed. Another sketchy argument is that the TBL may maintain about the same level of skin friction and "total" turbulence energy, but the Reynolds stresses are now split into two contributions, τ_1 and τ_2. τ_1 is created by the deviation from the time average, whereas τ_2 is created by the deviation from the spanwise average. Potentially, the mixing due to the vortices would take over to some extent, so that the second component would deplete the first component. Only that first component contributes to noise. This idea is supported by Fig. 6. The baseline case has zero τ_2, because it is homogeneous spanwise. The VG case has a strong τ_2, but τ_1 is indeed tangibly weaker over the lower half of the boundary layer, near $y = 0.01$, for all x beyond about 0.4. This could be a good sign.

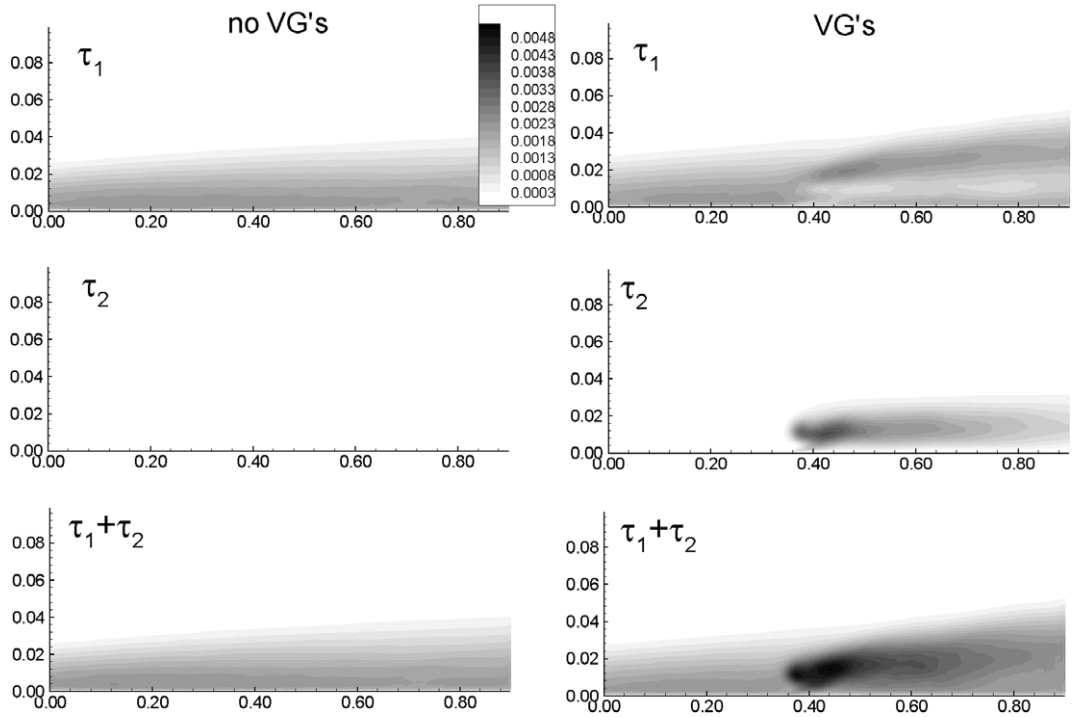

Fig. 6. The two components of Reynolds shear stress, and the total, in the baseline and the VG case.

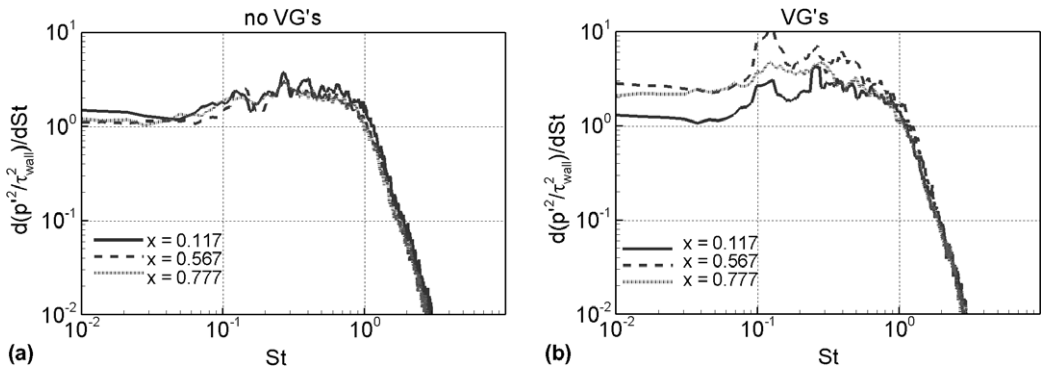

Fig. 7. Spectra of wall pressure fluctuations with and without VG's.

The unfortunate finding is that the alteration of the turbulence energy does not result in any benefit in terms of wall pressure, as seen in Fig. 7. In fact, the spectrum is everywhere higher, by up to 2 dB. The conclusion must be that this device is of no value for the purpose of interior-noise reduction. Some experimentation with VG's of smaller size and spacing failed to produce any promising designs. The attention then turned to completely different devices.

3.3. Large-eddy-break-up device

This device is closer to traditional thinking in turbulence reduction, but recall that the objective here is different, so that the failure of LEBUs to find drag-reduction applications does not rule out success here. The design, seen in Fig. 8, is sized to interfere with the large outer-layer eddies, and also to create drag in the inner-layer, which should depress the skin friction and consequently the pressure activity. Designs were tried with the main blade at the BL edge, and with the blade well inside the BL. Results are presented for the latter type, which is somewhat more effective. It extends roughly from $x = 0.33$ to $x = 0.38$. The entry region had to be lengthened by 0.05, because the pressure gradient due to the LEBU interfered with the recycling. Thus, moderate adjustments are needed to the recycling procedure depending on the cases. The figure illustrates grid refinement next to the blade (unlike in Immersed Boundary approaches). On the other hand, the refinement is not extended to the wake of the blade, so that small eddies in the wake are probably not supported; however, such eddies cannot influence the wall pressures.

Fig. 9 shows that downstream of the LEBU, the turbulence has lost its largest bulges, and generally created smaller eddies in the outer-layer, which can be beneficial in terms of de-energizing larger eddies. Thus, the turbulence alteration appears consistent with the postulated mechanism, without being dramatic. This is confirmed by Fig. 10, where the turbulent kinetic energy integrated across the BL has excursions at the LEBU, then drops as

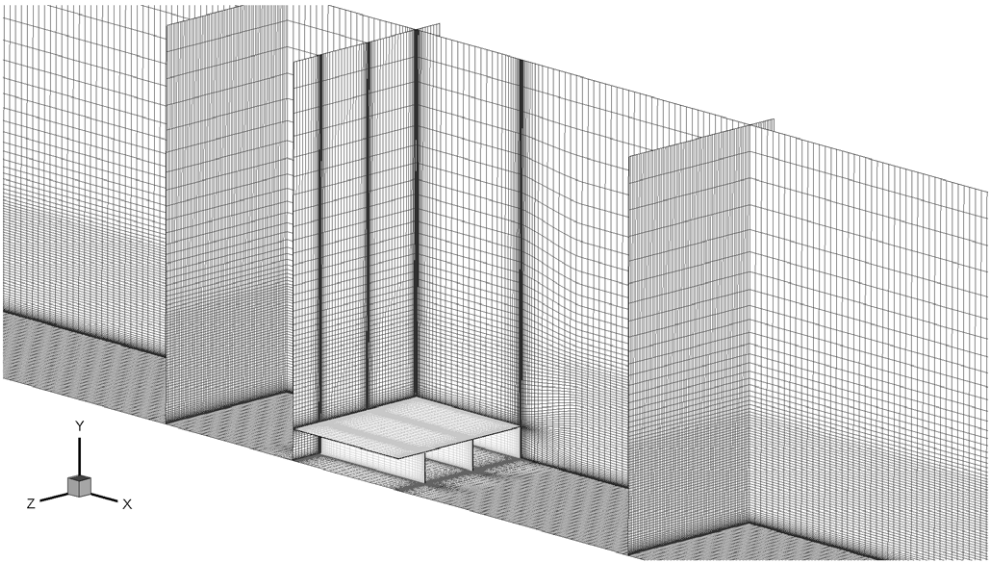

Fig. 8. Large eddy break up device grid.

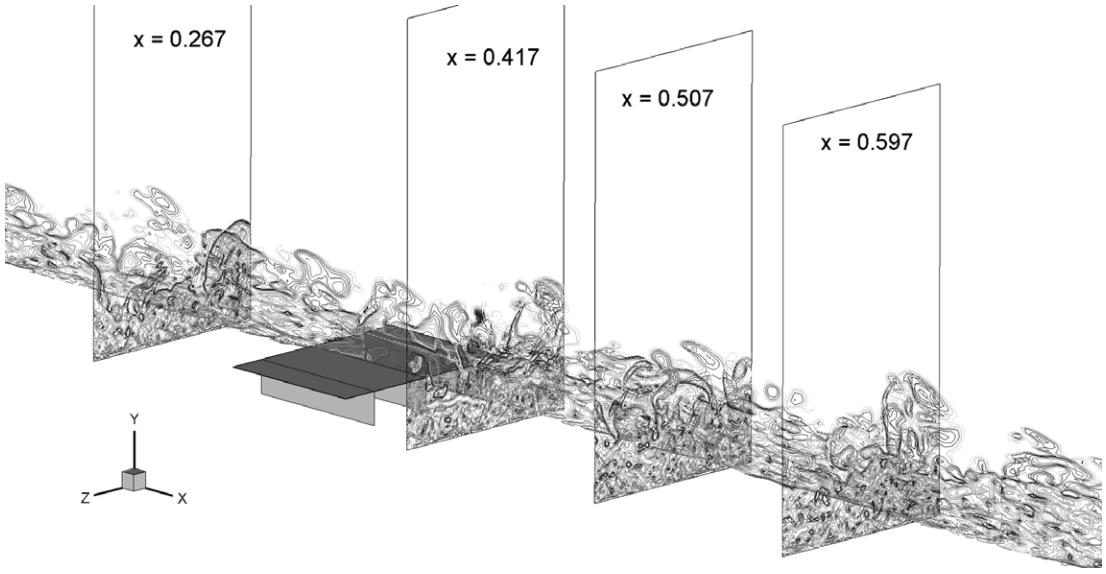

Fig. 9. Visualization of the flow with LEBU. Vorticity contours, flow from left to right.

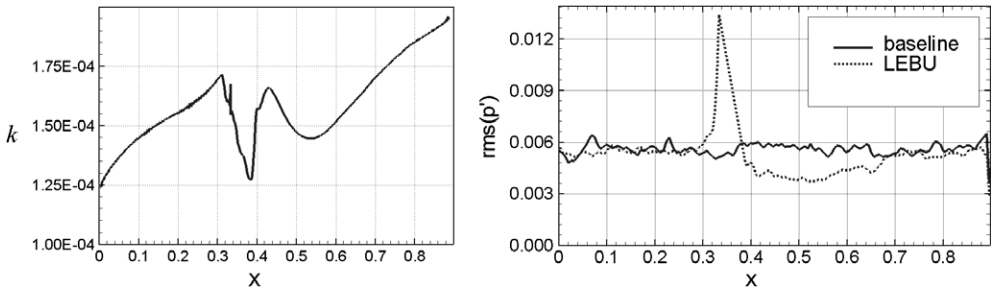

Fig. 10. Left, turbulence kinetic energy, integrated in y, with LEBU. Right, wall pressure *rms* with and without LEBU, normalized with U_e.

was hoped, but recovers earlier than would have been needed. The pressure *rms* also shown in Fig. 10 first takes the accepted value in this Reynolds-number range, namely, about two times the skin friction. It then presents a mixed result. On one hand, the target 30% reduction is achieved over the x-interval [0.4, 0.65] (coupled with a similar reduction in skin friction); thus, this device is much more successful than the VG's. On the other hand, this interval of length 0.25 is only about six times the undisturbed BL thickness at the location of the LEBU; the recovery is quite complete after 10 δ. While this behaviour is reassuring in terms of placing the DNS results close to the accepted estimate, it represents a failure in terms of seeking an invention. The "robustness" of boundary-layer turbulence evoked earlier is in full force.

Fig. 11 shows the spectra, with energy integrated over a running window, 1/3 octave wide. This gives a smoother curve and also helps identify the frequency range which contains most of the energy, when the frequency axis is logarithmic. The decrease is quite uniform over the spectrum, and close to the expected factor of 1/2, or 3 dB (the baseline value of skin friction is used everywhere for normalisation). The new peak near $St = 0.25$ may be spurious, which can be tested by altering the inflow length or otherwise altering the possible feedback mechanism from $x = x_r$ to $x = 0$, but the energy it contains is modest. The tendency

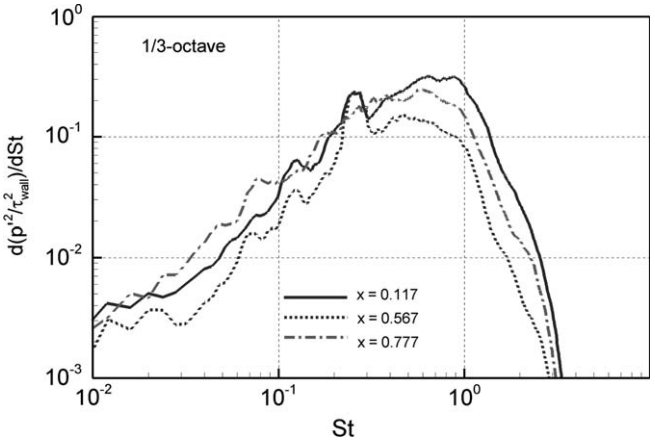

Fig. 11. Wall pressure spectra with LEBU (centred at $x = 0.35$).

to lower the dominant frequency range is probably not favorable, based on the properties of the glass in airliners, but the principal issue is the recovery length.

4. Outlook

This study has, so far, failed to produce a viable invention, although one device did reduce the pressure fluctuations. However, the simulations have screened out two concepts that appeared to have potential, much more economically than wind-tunnel or flight tests would have. The DNS-based numerical system appears quite capable of producing extensive quantitative information in this relatively complex physical situation, with proper caution towards the low end of the frequency spectrum and a sound strategy for scaling with Reynolds number. Feedback within the recycling region might be more noticeable with embedded devices than in a simple TBL, but several tools are available to detect such a feedback. Direct Numerical Simulations of a TBL, with several million points and with obstacles properly gridded, are now possible on Personal Computers. The simulations not only indicate the potential value of a device (pending studies of the transmission through the surface), but provide visualizations to suggest the next generation or at least the size of the next candidate. The search for a successful device continues.

References

Ashill, P.R., Fulker, J.L., Hackett, K.C., 2005. A review of some recent developments in flow control. Aero. J. 109 (5), 205–232.

Blake, W.K., 1970. Turbulent boundary-layer wall-pressure fluctuations on smooth and rough walls. J. Fluid Mech. 44 (4), 637–660.

Dobrzynski, W., Chow, L.C., Guion, P., Shiells, D., 2002. Research into landing gear airframe noise reduction. AIAA, 2002–2409.

ESDU, 1993. Vortex generators for control of shock-induced separation. Data Item 93024. ESDU Int. plc, London.

Liu, J., Piomelli, U., Spalart, P.R., 1996. Interaction between a spatially-growing turbulent boundary layer and embedded streamwise vortices. J. Fluid Mech. 326, 151–180.

Lund, T.S., Wu, X., Squires, K.D., 1998. Generation of turbulent inflow data for spatially-developing boundary-layer simulations. J. Comp. Phys. 140, 233–258.

Pearcey, H.H., 1961. Shock-induced separation and its prevention by design and boundary layer control. In: Lachmann, G.V. (Ed.), Boundary Layer and Flow Control. Pergamon, Oxford.

Rogers, S.E., Kwak, D., 1988. An upwind differencing scheme for the time-accurate incompressible Navier-Stokes equations. AIAA Paper 88-2583-CP.

Schlichting, H., 1979. Boundary-Layer Theory. McGraw-Hill, New York.

Spalart, P.R., 1988. Direct simulation of a turbulent boundary layer up to $R_\theta = 1410$. J. Fluid Mech. 187, 61–98.

Strelets, M., 2001. Detached Eddy Simulation of Massively Separated Flows. AIAA Paper 2001-0879.

Willmarth, W.W., 1975. Pressure fluctuations beneath turbulent boundary layers. Ann. Rev. Fluid Mech. 7, 13–38.

Willmarth, W.W., Woolridge, C.E., 1962. Measurements of the fluctuating pressures at the wall beneath a thick turbulent boundary layer. J. Fluid Mech. 14, 187.