

# Simulation of a Turbulent Flow Subjected to Favorable and Adverse Pressure Gradients

Ali Uzun\*

*National Institute of Aerospace, Hampton, VA 23666*

Mujeeb R. Malik†

*NASA Langley Research Center, Hampton, VA 23681*

**This paper reports the results from a direct numerical simulation of an initially turbulent boundary layer passing over a wall-mounted “speed bump” geometry. The speed bump, represented in the form of a Gaussian distribution profile, generates a favorable pressure gradient region over the upstream half of the geometry, followed by an adverse pressure gradient over the downstream half. The boundary layer approaching the bump undergoes strong acceleration in the favorable pressure gradient region before experiencing incipient or very weak separation within the adverse pressure gradient region. These types of flows have proven to be particularly challenging to predict using lower-fidelity simulation tools based on various turbulence modeling approaches and warrant the use of the highest-fidelity simulation techniques. Simulation results are utilized to examine the key phenomena present in the flowfield, such as relaminarization/stabilization in the strong acceleration region succeeded by retransition to turbulence near the onset of adverse pressure gradient, incipient/weak separation, and development of internal layers where the sense of streamwise pressure gradient changes at the foot, apex and tail of the bump. The present direct numerical simulation is performed using a flow solver developed exclusively for graphics processing units, which is found to provide a significant speedup compared to an earlier solver optimized for central processing unit architectures.**

## I. Introduction


Turbulent boundary layers subjected to pressure gradients are commonly found in many practical applications. A particularly interesting application is found over the upper surface of subsonic and transonic airfoils, wherein the favorable pressure gradient over the leading-edge region is followed by an adverse pressure gradient further downstream, generated due to a change in body contour and/or the presence of a shock. The favorable pressure gradient generally leads to flow acceleration, which can be strong in certain cases, while the interaction with the adverse pressure gradient often leads to flow separation. Similar flow phenomena also exist in turbomachinery and wind turbine applications. Flows subjected to such conditions have proven to be particularly challenging to predict using lower-fidelity simulation tools based on various turbulence modeling approaches. A recent example demonstrating this deficiency of turbulence models can be found in the work of Spalart et al. [1], who employed wall-modeled large eddy simulation (WMLES) in the form of an improved delayed detached eddy simulation (IDDES) to simulate the strong flow acceleration over a transonic bump and the subsequent shock-induced flow separation. Despite the large number of grid points used in the WMLES (nearly 1.7 billion points), comparisons of the results with the direct numerical simulation data revealed that the WMLES failed to predict the flow development in the accelerating flow region prior to flow separation. This deficiency led to an incorrect prediction of shock location, thereby yielding an incorrect flow separation location.

In aeronautical applications, reliable lift and drag predictions are strongly dependent upon accurate computation of these complex flows. Accurate prediction of maximum aircraft lift could enable significant reduction in certification flight testing, saving hundreds of million dollars in aircraft development programs. To further investigate turbulent flows subjected to favorable/adverse pressure gradients and generate additional computational as well as experimental data, a new benchmark test case has been recently proposed [2]. This new test case, which will be referred to as the “speed bump” problem from this point onward, contains a wall-mounted geometry represented in the form of a Gaussian distribution profile, as depicted in Figure 1. The problem involves an initially turbulent boundary layer passing over the speed bump profile that generates a favorable pressure gradient region over the front portion, followed

---

\* Associate Principal Engineer, Senior Member AIAA.

† Senior Aerodynamicist, Computational AeroSciences Branch, MS 128, Fellow AIAA.



**Fig. 1** Upper, side view; lower, spanwise cross-section; side walls are at  $z/L = \pm 0.5$ , ceiling is at  $y/L = 0.5$  (courtesy of Dr. Philippe Spalart).

by an adverse pressure gradient in the aft section. The speed bump model has a uniform height along much of the span, and its height is smoothly tapered down to zero by an error function near the side edges, as can be seen in Figure 1.

A detailed experimental investigation campaign is planned for this test case. It is anticipated that several groups will study this problem using various turbulence modeling approaches, including direct numerical simulation (DNS). Comparisons of the results from a number of simulations should be of value to the turbulence research community. Our contribution to this community-wide exercise is in the form of a spanwise-periodic DNS, the results of which are presented in this paper. While this spanwise-periodic DNS obviously cannot duplicate the full experimental configuration, it can still serve as a useful reference against which the spanwise-periodic DNS results of other groups, as well as the results of lower-fidelity simulations, such as Reynolds-averaged Navier-Stokes (RANS) calculations and WMLES, can be compared. A wall-resolved LES (WRLES) or DNS of the full 3-D configuration at a realistic Reynolds number is planned as part of future work on a Department of Energy Leadership Computing Capability System, such as Summit (<https://www.olcf.ornl.gov/olcf-resources/compute-systems/summit/>\*) or the upcoming exascale systems, once the computational resources needed for such an extreme-scale simulation are secured.

## II. Computational Methodology

A direct numerical simulation (DNS) code, named **GTHORS** (GPU version of **T**urbulence with **H**igh-**O**rders **R**esolution **S**olver) has been specifically developed for graphics processing units (GPUs). This DNS code is based on explicit high-order finite-difference and explicit time integration schemes. We believe such schemes have the best potential to extract the full capability of GPUs since they are essentially made up of many independent multiply-and-add type operations, at which the GPU architecture excels. The code solves the compressible Navier-Stokes equations discretized on multiblock structured grids. The methodology employs the optimized fourth-order accurate central explicit finite difference scheme with a 13-point stencil, developed by Bogey and Bailly [3], to compute all spatial derivatives in the governing equations. For boundary and near-boundary points, matching one-sided and biased schemes, developed by Berland et al. [4], are used. To ensure numerical stability, the optimized 11-point sixth-order explicit selective filter, developed by Bogey and Bailly [5], is used. This optimized filter is derived from the standard explicit tenth-order filter. The optimized filter modifies the coefficients of the standard filter in order to improve the filter damping function in wavenumber space and constrain the numerical dissipation to waves discretized by fewer than four grid points. The frequency of filtering (in terms of time steps) and the amount of damping applied by the filter are fine tuned depending on the specific grid resolution used in a given problem. This is done in order to determine the minimal amount of numerical dissipation needed to keep the computation stable on a given grid resolution.

Time advancement can be performed using either a third-order, three-stage explicit Runge-Kutta scheme [6] or a second-order, single-stage explicit scheme developed by Versteppen and Veldman [7]. This second-order scheme is a modified Adams-Bashforth scheme with improved stability properties. The second-order scheme runs at half the time step of the third-order Runge-Kutta scheme but requires only one right-hand-side computation per time step as opposed to three right-hand-side evaluations needed for the Runge-Kutta scheme. It therefore provides a speedup factor of 3/2 over the Runge-Kutta scheme. The flow solver employs three levels of parallelism based on a hybrid combination of MPI+OpenMP+CUDA Fortran. Further details of the GPU flow solver, including its parallel scaling performance and validation for a turbulent channel flow problem are provided in the Appendix.

\*Website last accessed May 20, 2020.

### III. Test Case: Turbulent Flow over Speed Bump

The 3-D profile of the speed bump geometry [2] is given by

$$y(x, z) = \frac{h}{2} \left[ 1 + \operatorname{erf} \left( \left( \frac{L}{2} - 2z_0 - |z| \right) / z_0 \right) \right] \exp \left( - (x/x_0)^2 \right) \quad (1)$$

where  $x, y, z$  respectively, denote the axial, vertical and spanwise directions,  $L$  is the cross-section width,  $h = 0.085L$ ,  $x_0 = 0.195L$  and  $z_0 = 0.06L$ . The present DNS assumes spanwise periodicity and therefore a uniform profile along the span. For the present case, the 2-D profile shape is given by  $y(x) = h \exp \left( - (x/x_0)^2 \right)$ .

#### A. Simulation Details

The Reynolds number based on the upstream reference velocity,  $U_\infty$ , and  $L$  is  $Re_L = 1$  million. The freestream Mach number upstream of the bump is set to 0.2. Figure 2 depicts a schematic of the computational domain. The inflow boundary is placed at  $x/L = -0.8$  while the outflow boundary is placed at  $x/L \approx 2$ . The physical domain of interest ends at  $x/L = 1$ . The region from  $x/c = 1$  to 2 is the sponge zone, which is constructed by applying rapid grid stretching along the streamwise direction. This region contains fewer than 200 points because of the grid stretching applied. The sponge zone damps out the turbulence before it reaches the outflow boundary, on which characteristic outflow boundary conditions are imposed. Viscous isothermal boundary conditions are applied on the lower boundary, which contains the speed bump profile. The wall temperature is set the same as the reference freestream temperature. The outer boundary in the vertical direction is located at  $y/L = 0.45$ . On this boundary, a nonreflecting characteristic boundary condition formulation imposes the pressure and velocity components available from an incompressible RANS computation (performed using the Spalart-Allmaras model) and determines the density. To account for the effect of the ceiling in the experimental configuration, the RANS was performed with a viscous top wall at  $y/L = 0.5$ . The chosen  $y/L = 0.45$  location for the DNS outer boundary is located within the inviscid region outside the boundary layer developing on the top wall. Assuming that the RANS solution obtained in the vicinity of the top wall is reasonably accurate, this approach ensures that the corresponding effect is properly captured in the DNS. In any case, this is an approximation but direct simulation of the viscous upper wall is not afforded by the computational resources at hand.




Fig. 2 Domain schematic. Contours denote instantaneous pressure normalized by reference value.

The periodic domain span is set to  $0.04L$ . As will be seen, the domain span is rather narrow in the recovery region post separation, since the adverse pressure gradient in the aft section of the bump significantly increases the boundary layer thickness. We therefore expect that the flow development in that region will be affected by the narrow span to some extent. This particular span was primarily chosen to allow a number of groups to perform scale-resolving simulations with relative ease, the results from which will be used to perform code-to-code comparisons.

The domain is discretized using 12288 points along the streamwise direction, 320 points in the vertical direction and 576 uniform points along the spanwise direction. The total number of grid points is about 2.265 billion. In terms of wall units, the largest streamwise spacing varies from 6 to 8 units, while the largest spanwise spacing is about 4.5 units. The wall-normal grid spacing on the wall generally varies from about 0.4 to 0.95 units in much of the attached flow region and becomes smaller in the vicinity of the weakly-reversed flow. The largest vertical grid spacing in the vicinity of the boundary layer edge is around 10 units or smaller. Figure 3 provides the variation of the streamwise, spanwise and wall-normal grid resolutions in wall units along the speed bump. The third-order, three-stage explicit

Runge-Kutta scheme [6] is used for the time integration. It takes 750,000 time steps to compute a time interval of  $L/U_\infty$ . The solution is filtered [3, 5] at every time step with a filtering parameter of  $\sigma = 0.10$ . Note that  $\sigma = 0$  means that the filter is off, while  $\sigma = 1$  means that the filter is in full effect. For the DNS, 48 NVIDIA Tesla V100 GPUs, located at the NASA Advanced Supercomputing Division at Ames Research Center<sup>†</sup>, are used. Each GPU solves a grid block of  $512 \times 320 \times 288$  points. The DNS takes two months of run time to compute a time interval of  $20L/U_\infty$ . The initial numerical transients are driven out of the computational domain during the first  $5L/U_\infty$ . Statistical data are gathered over the remaining  $15L/U_\infty$ , which covers 8.33 physical domain flow-through times.

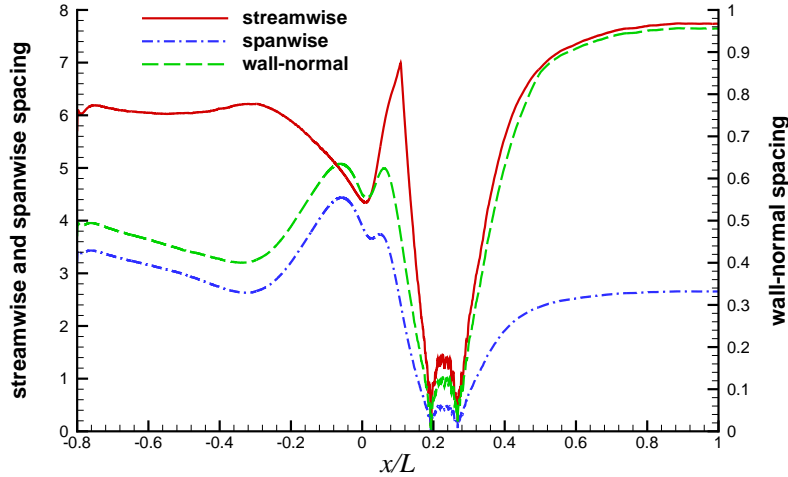


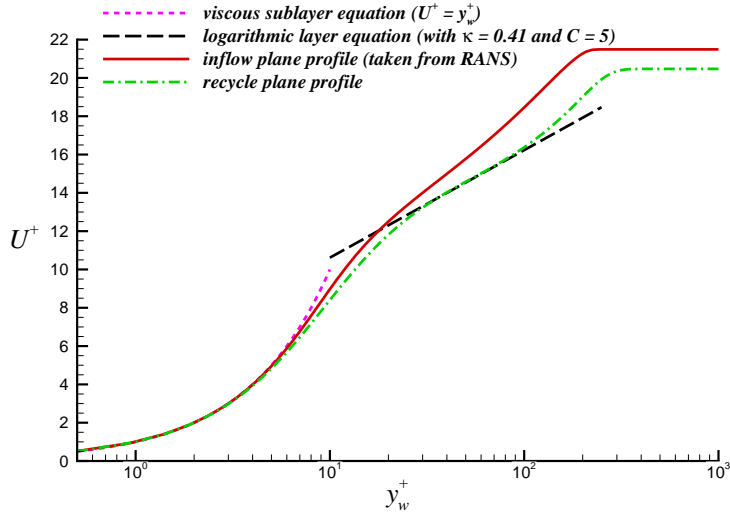
Fig. 3 Near-wall grid spacings in wall units.

For the turbulent inflow generation, we employ the rescaling-recycling technique, discussed in Uzun and Malik [8]. The mean flow imposed at the inflow boundary of the present DNS is taken from the aforementioned incompressible RANS solution. The mean boundary layer thickness on the inflow boundary is  $\delta_{in} \approx 0.0046L$ . The distance between the inflow and recycle planes is set to  $15\delta_{in}$ . The turbulent fluctuations extracted from the recycle plane are rescaled and reintroduced at the inflow after applying a spanwise scramble. See Uzun and Malik [8, 9] for further details. Figure 4 shows the inflow and recycle plane mean axial velocity profiles in wall units. As seen here, the inflow mean profile taken from the RANS does not possess a logarithmic layer that is well represented by a von Kármán constant of  $\kappa = 0.4\text{--}0.41$  and an intercept constant of  $C = 5.0\text{--}5.2$ . This is presumably due to the relatively low Reynolds number of the present test case. The logarithmic layer is defined as:  $U^+ = \kappa^{-1} \ln(y_w^+) + C$ , where  $y_w^+ = y_w u_\tau / \nu$ ,  $U^+ = U/u_\tau$ ,  $y_w$  is the wall-normal distance,  $U$  is the mean axial velocity,  $u_\tau = \sqrt{\tau_w/\rho}$  is the friction velocity,  $\tau_w$  is the wall shear stress, and  $\rho$  and  $\nu$ , respectively, are the density and kinematic viscosity on the wall. Despite this shortcoming of the RANS mean inflow profile, the turbulent boundary layer attains a more realistic mean state by the time it reaches the recycle station. The recycle plane mean profile has a logarithmic layer that is matched well by  $\kappa = 0.41$  and  $C = 5$ . The Reynolds number based on the local friction velocity and boundary layer thickness is  $Re_\tau \approx 350$  at the recycle station. We see that the logarithmic layer is rather small at this station because of the relatively low Reynolds number.

## B. Simulation Results

We now discuss the simulation results. The main features of the flowfield are presented first, followed by a more detailed investigation of the flow behavior over the upstream and downstream halves of the speed bump. Throughout the discussion, we refer to the portion of the speed bump upstream of the apex as the first or the upstream half, and the remaining portion as the second or the downstream half. The first half of the speed bump contains the strongly accelerated flow region preceded by a relatively weak upstream adverse pressure gradient that becomes stronger toward the foot of the bump, while the second half encompasses the deceleration due to adverse pressure gradient, incipient or very weak separation, and downstream recovery regions. We make use of the data gathered from the DNS to examine the various phenomena encountered in the flowfield.


<sup>†</sup>The GPU cluster contains 72 V100s total. Each V100 has a peak performance of 7.8 Tera-FLOPs per second in double precision.



**Fig. 4** Inflow and recycle plane mean axial velocity profiles.


*1. Main Features of the Flowfield*

Figure 5 depicts the total velocity magnitude contours near the wall, normalized by the reference freestream velocity, in the region where  $-0.5 < x/L < 0.1$ . The plane shown here is extracted at the tenth grid point in the wall-normal direction. On this plane, the wall-normal distance in wall units varies from about 4 to 6. For clarity, the region of interest is shown as two separate subfigures and the maximum value of the contour range differs between the two subfigures. Note that the bump apex is at  $x/L = 0$ . The top subfigure provides evidence of a fully turbulent upstream boundary layer, in the form of near-wall streaks, approaching the bump. However, the boundary layer accelerating over the front portion of bump displays a tendency toward relaminarization or stabilization, as can be seen in the bottom subfigure. We believe this phenomenon to arise from the relatively low Reynolds number of the incoming flow, which does not allow the turbulence to fully survive the acceleration caused by the favorable pressure gradient, and the further stabilizing effect of convex curvature. The flow retransitions to a turbulent state after passing over the apex.



**Fig. 5** Normalized total velocity magnitude contours near the wall on a plane where  $4 \lesssim y_w^+ \lesssim 6$ .

Figure 6 depicts the total velocity magnitude contours in the outer region, on a plane where the wall-normal distance in wall units varies from about 80 to 120. Note that for clarity, the vertical scale is enlarged by a factor of two relative to the horizontal scale in this figure. As seen here, the relaminarization/stabilization phenomenon is not immediately evident from solely examining the outer region of the flow; hence, one must look very close to the wall in order to observe it. This feature of the flowfield is investigated in more detail in section III.B.3.




**Fig. 6** Normalized total velocity magnitude contours in the outer region on a plane where  $80 \lesssim y_w^+ \lesssim 120$ .

The skin-friction distribution, depicted in Figure 7, provides further evidence of flow relaminarization/stabilization and also reveals significant differences between the DNS and RANS predictions. The data from the aforementioned incompressible RANS and a compressible RANS at Mach 0.176<sup>‡</sup>, also performed with a viscous top wall and the Spalart-Allmaras model, are shown. The skin-friction and surface pressure coefficients are given by

$$C_f = \frac{\tau_w}{\frac{1}{2}\rho_\infty U_\infty^2} \quad \text{and} \quad C_p = \frac{p - p_\infty}{\frac{1}{2}\rho_\infty U_\infty^2} \quad (2)$$


where  $\rho_\infty$ ,  $p_\infty$ ,  $U_\infty$ , respectively, are the reference freestream density, pressure and velocity,  $p$  is the mean surface pressure and  $\tau_w$  is the mean wall shear stress. To enable a future comparison of the present  $C_p$  data with the simulations performed by other investigators in smaller computational domains, the reference pressure for the  $C_p$  calculation is taken at  $(x/L, y/L) = (-0.6, 0.2)$ . The coefficients from the RANS calculations are computed similarly.



**Fig. 7** Skin-friction coefficient,  $C_f$ , distributions.

<sup>‡</sup>The experimental Mach number is expected to be around this value. The RANS calculation at this Mach number, as well as the incompressible RANS computation, were performed in separate independent studies by our colleagues.

We observe the primary  $C_f$  peak in the DNS at  $x/L \approx -0.055$ . The suction peak in the  $C_p$  distribution is positioned very slightly upstream of the bump apex at  $x/L = 0$ , as seen in Figure 8. Hence, the immediate drop following the primary  $C_f$  peak in the DNS begins while the flow is still accelerating upstream of the bump apex. The  $C_f$  drop continues as the flow enters the adverse pressure gradient region. The stabilizing effect of the favorable pressure gradient, generated upstream of the adverse pressure gradient region, vanishes once the flow goes past the apex. The adverse pressure gradient then starts to decelerate the flow; however, the flow is in the early stages of experiencing the adverse pressure gradient effect at that point. The flow still has enough momentum that allows the surviving instabilities to initiate a transition back to a turbulent state. This is verified by the presence of a secondary  $C_f$  peak at  $x/L \approx 0.05$  as well as the flow visualization depicted in Figure 5. The secondary  $C_f$  peak is quickly followed by another fast drop, due to the continued presence of adverse pressure gradient in the aft region, which continues to decelerate the flow and subsequently brings it to the brink of separation.




**Fig. 8** Surface pressure coefficient,  $C_p$ , distributions.

Comparison of the DNS prediction with the RANS results shows significant differences in the location and value of the peak  $C_f$ , as well as the absence of relaminarization/stabilization in the RANS results. The two RANS calculations provide similar  $C_f$  distributions. The peak RANS  $C_f$ , found at  $x/L \approx -0.028$ , has a much higher value than the primary DNS peak. Another important difference is the presence of more severe separation in the RANS computations, which is obvious from the reversed flow region in the RANS  $C_f$  distributions. The DNS, on the other hand, only shows incipient or very weak separation, as indicated by the small negative  $C_f$  values where  $0.195 \leq x/L \leq 0.268$ .

The incipient or very weak separation observed in the DNS is further verified by flow visualization. Figure 9 provides an instantaneous snapshot of the flowfield in the aft section of the bump, in terms of normalized axial velocity contours. For clarity, the maximum value of the contour range is set to a value lower than the largest value found in the flowfield. In the figure, localized regions containing negative axial velocity near the wall denote the presence of some reversed flow but, unlike the RANS results shown in Figure 7, we do not observe any regions of significant flow separation. This is indicative of incipient or very weak flow separation, as also verified by the DNS  $C_f$  distribution.

Further examination of Figure 7 also reveals important differences between the DNS and RANS predictions in the region upstream of the bump. Even though the DNS uses the mean inflow profile extracted from the incompressible RANS, the  $C_f$  distribution in the DNS displays a jump immediately downstream of the inflow plane. The DNS  $C_f$  then settles to a value higher than the RANS  $C_f$  in the region upstream of the bump. As seen in Figure 4, the absence of a logarithmic layer with well-accepted constants in the RANS result implies that the corresponding RANS  $C_f$  value at the inflow plane is likely erroneous. Even though the present Reynolds number is high for DNS, it is actually quite low for RANS. The observations made here about the RANS data are well-known issues associated with RANS computations performed at relatively low Reynolds numbers. The behavior seen in the DNS  $C_f$  variation immediately



**Fig. 9** Instantaneous snapshot of the flowfield in the aft section. Contours denote the instantaneous axial velocity normalized by reference velocity.

downstream of the inflow profile is related to the discrepancy identified in the RANS mean profile. Nevertheless, the DNS  $C_f$  attains a stable value at a short distance from the inflow plane. As already seen in Figure 4, the turbulent mean flow reaches a reasonable state at the recycle plane, which is positioned  $15\delta_{in}$  downstream of the inflow plane.

The  $C_p$  comparison, depicted in Figure 8, shows a reasonable agreement and much less dramatic differences between the DNS and RANS calculations. The RANS results show a slight strengthening of the suction peak with the Mach number. The DNS, performed at a freestream Mach number of 0.2, produces a bit stronger suction peak relative to the RANS computation at Mach 0.176. Some  $C_p$  differences are seen in the region where  $0.2 < x/L < 0.36$ , and the RANS computation at Mach 0.176 appears to be closer to the DNS in that range. This region encompasses part of the RANS separation bubble and the downstream recovery region. The  $C_f$  comparison in the same region shows more significant differences between the DNS and RANS computations.

## 2. Mean Boundary Layer Properties

Before we examine the variation of boundary layer properties along the speed bump, we briefly discuss how those quantities are computed. As noted by Spalart and Watmuff [10], in flows subjected to pressure gradients, the boundary layer edge velocity and the corresponding boundary layer thickness cannot be determined in a straightforward manner because the streamwise velocity does not reach a uniform value even at relatively large distances from the wall. Hence, the classical definitions of displacement and momentum thicknesses based on the integrals of streamwise velocity cannot be applied unambiguously because it is not clear where to truncate the integrals.

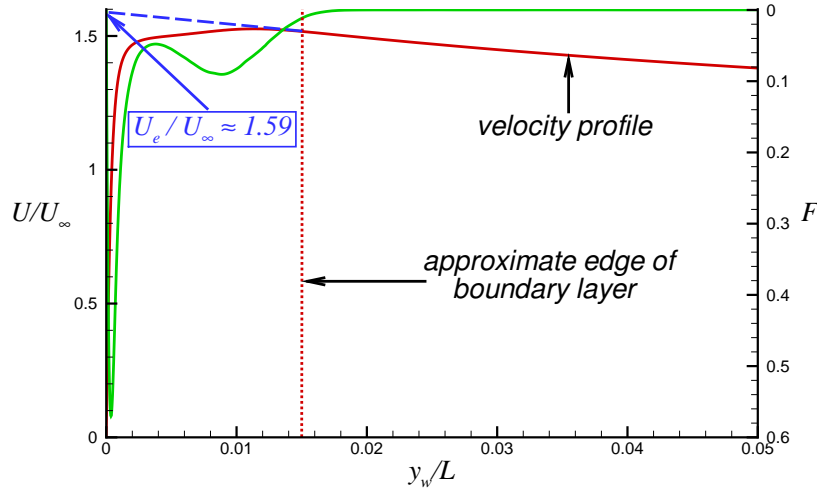
To get around this difficulty, Spalart and Watmuff [10] derived the following definitions of boundary layer edge velocity,  $U_e$ , and displacement and momentum thicknesses,  $\delta^*$  and  $\theta$ , respectively, based on the integrals of mean spanwise vorticity,  $\omega_z$ ,

$$U_e = \int_0^\infty -\omega_z(y_w) dy_w, \quad \delta^* = \frac{1}{U_e} \int_0^\infty -y_w \omega_z(y_w) dy_w, \quad \theta = \frac{2}{U_e^2} \int_0^\infty y_w \left[ \int_0^{y_w} \omega_z(n) dn \right] \omega_z(y_w) dy_w - \delta^* \quad (3)$$

where  $y_w$  denotes the local wall-normal distance at a given streamwise station and  $n$  is the variable of integration in the inner integral for  $\theta$ .  $U_e$  is also called the ‘‘pseudovelocity’’, and can be considered as an equivalent inviscid surface velocity. These relations assume incompressible flow and are based on the observation that the spanwise vorticity decays fast enough at large wall-normal distances for the above integrals to converge.



These integrals directly provide the boundary layer quantities  $\delta^*$  and  $\theta$ ; however, the issue of the determination of local boundary layer thickness,  $\delta$ , in flows subjected to pressure gradients remains to be addressed. A reasonable method to estimate the boundary layer thickness<sup>§</sup> is illustrated in Figure 10. This figure depicts the profiles of mean streamwise velocity and the quantity,  $F = -y_w \omega_z$ , at the bump apex. We observe that the peak region in the streamwise velocity is immediately followed by a linear drop that begins in the vicinity of  $y_w/L \approx 0.015$ . This region characterized by the linear drop in streamwise velocity is inviscid in nature. Taking the slope of the linear drop at  $y_w/L \approx 0.015$  and extrapolating the velocity at that location to the wall would essentially produce a good estimate of  $U_e$ , or an equivalent inviscid surface velocity. This extrapolation gives  $U_e \approx 1.59U_\infty$ . As will be seen, the corresponding vorticity integral at the bump apex gives  $U_e \approx 1.6U_\infty$ , which is very close to this extrapolated value. At  $y_w/L \approx 0.015$ , we see that the value of  $F$  is about 2% of the maximum  $F$ , found very close to the wall. Thus, the boundary layer thickness can be defined as the furthest point from the wall for which  $F > \varepsilon F_{max}$ , where  $\varepsilon = 0.02$ . This criterion is used to determine the local boundary layer thicknesses in this study.



**Fig. 10** Profiles of mean streamwise velocity and  $F = -y_w \omega_z$  at the bump apex.

We should also note that the upper limit of the vorticity integrals in Equation 3, which is infinity, is not straightforward to define in the present problem as the top boundary of our computational domain, located at  $y/L = 0.45$ , imposes the effect of a viscous wall, located at  $y/L = 0.5$  in the corresponding RANS computation. To eliminate any potential influence of the top wall on the integrals, they should be terminated as early as possible once sufficient convergence has been achieved. As seen in Figure 10, the vorticity (or  $F$ ) near the boundary layer edge is not exactly zero; hence, extending the integrals some distance beyond the boundary layer edge would ensure full convergence. Fortunately, the integrals performed over one and two local boundary layer thicknesses are found to show negligible differences and are therefore nearly identical. Thus, the quantities obtained from the integrals performed over two local boundary layer thicknesses will be shown in the upcoming discussion.


To compute the above integrals, the mean spanwise vorticity computed on the original DNS grid is first interpolated onto the local wall-normal line constructed for each streamwise station. The integrals are then performed along the local wall-normal lines using the interpolated data. A similar interpolation approach is taken for the mean velocity and Reynolds stress components to be examined later. The mean velocity and Reynolds stress components defined in the Cartesian coordinate system are first interpolated onto the local wall-normal line constructed for each streamwise station. Using the local wall-normal and surface tangent vectors, the interpolated values are then transformed to the corresponding quantities in the local orthogonal coordinate system.

Figure 11 provides the variation of  $\delta/L$ ,  $\delta^*/L$ ,  $\theta/L$ , shape factor,  $H = \delta^*/\theta$ , and  $U_e$ , as a function of  $x/L$ <sup>¶</sup>. The figure also includes the  $C_p$  and  $C_f$  distributions to aid in the discussion. As expected, all quantities display significant variations as the flow travels over the bump. We observe that  $U_e$  follows the pressure gradient consistently in the whole domain. It decreases in the adverse pressure gradient and increases in the favorable pressure gradient. At around the

<sup>§</sup>Credit is due Dr. Philippe Spalart for suggesting this method.

<sup>¶</sup>The same scale is used for  $\delta$ ,  $\delta^*$  and  $\theta$  because of the limit of 5 vertical axes imposed by the plotting software.

apex,  $U_e \approx 1.6U_\infty$ , which agrees well with the extrapolated value from the velocity profile, as discussed above.



**Fig. 11** Variation of mean boundary layer properties,  $C_p$  and  $C_f$  over the speed bump.

The incoming boundary layer, which starts off with a relatively low  $\delta$ ,  $\delta^*$  and  $\theta$  far upstream of the bump, encounters an initially mild adverse pressure gradient, which becomes stronger around the foot of the bump. The boundary layer growth, indicated by the increases in  $\delta$ ,  $\delta^*$  and  $\theta$  in the upstream adverse pressure gradient region, is accompanied by a corresponding drop in  $C_f$ . Interestingly, the growth in  $\delta$  continues until  $x/L \approx -0.25$ , while the drops in the other two quantities have already begun at around  $x/L \approx -0.29$  or  $-0.28$ . The shape factor,  $H = \delta^*/\theta$ , is about 1.57 near the inflow, and drops some once the turbulent boundary layer starts developing from the imposed inflow conditions. The adverse pressure gradient upstream of the bump slightly increases  $H$  to a local maximum of about 1.49 at  $x/L \approx -0.35$ . This value of  $H$  is a bit higher than the typical values of 1.3–1.4 expected for a zero pressure gradient turbulent boundary layer, and is a normal consequence of the adverse pressure gradient.

Once the flow goes past  $x/L \approx -0.3$ , it becomes subjected to the strong favorable pressure gradient effects, which persist until around the bump apex. These effects are manifested in a number of ways. We see that  $U_e$  rapidly increases and attains a maximum value of about 1.6 times the upstream reference velocity, at around the apex. The flow acceleration owing to the strong pressure gradient thins the boundary layer, as verified by the corresponding drops in  $\delta$ ,  $\delta^*$  and  $\theta$ . This thinning continues until slightly upstream of the apex.  $H$  drops to a local minimum of about 1.26 at  $x/L \approx -0.11$  before reaching a value of about 1.4 at around the apex, indicating that the reductions in  $\delta^*$  and  $\theta$  take place at different rates.  $H$  continues to increase past the apex and appears to briefly plateau in the vicinity of the secondary  $C_f$  peak. The flow acceleration also steepens the near-wall velocity gradient, which then leads to the rapid increase observed in  $C_f$  in the accelerating region.  $C_f$  reaches its primary peak some distance upstream of the apex and then starts to decrease while the flow is still under acceleration. A secondary  $C_f$  peak is reached shortly after the flow goes past the apex. This peculiar behavior in  $C_f$  is due to the relaminarization or stabilization very near the wall in the strong acceleration region and the transition back to turbulence in the initial deceleration region. Some discussion of this feature of the flowfield was provided earlier during the examination of Figures 5–7. Further investigation of this subject is provided in section III.B.3.

The strong acceleration effects due to the favorable pressure gradient over the upstream half of the bump are countered by the adverse pressure gradient effects in the downstream half. The adverse pressure gradient decelerates the flow, and substantially lowers  $U_e$  from its maximum value. The deceleration brings the flow to the brink of separation. As discussed earlier, the strongly accelerated flow that shows signs of relaminarization/stabilization upstream of the apex has enough momentum to transition back to turbulence in the early stages of the adverse pressure gradient. After its secondary peak, observed shortly after the apex due to the transition back to turbulence,  $C_f$  quickly drops down

to zero at around  $x/L \approx 0.195$  and remains slightly negative until  $x/L \approx 0.268$  because of incipient or very weak separation. We observe a substantial growth in  $\delta$ ,  $\delta^*$  and  $\theta$  as the flow decelerates. The growth rate of  $\delta^*$  is much faster than that of  $\theta$ .  $H$  consequently increases to a value of about 3.1 as  $C_f$  drops down to zero. As noted above, a brief plateau in  $H$ , that nearly coincides with the secondary  $C_f$  peak, appears during the boundary layer growth.  $H$  becomes as high as 3.4 in the short section where  $C_f$  goes slightly negative. That particular region contains a tiny patch of very weakly-reversed flow near the wall but we do not expect the very limited extent of the reversed flow to have a notable effect on the vorticity integrals.

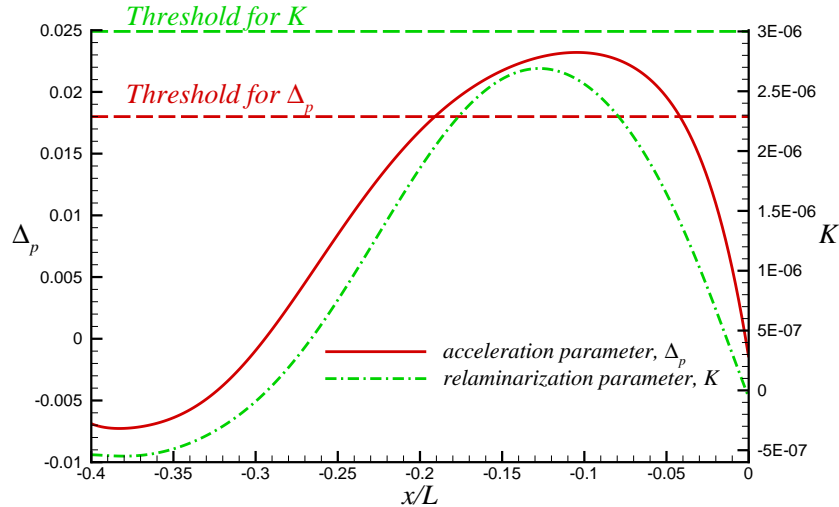
The flow recovery from the incipient or very weak separation, marked by the rise in  $C_f$ , starts while the adverse pressure gradient tapers down to a smaller value. Interestingly,  $\delta^*$  starts its decrease shortly after the flow recovery begins, at  $x/L \approx 0.28$ , while  $\theta$  continues its rise until  $x/L \approx 0.38$ , where the pressure gradient changes from adverse to mildly favorable. Both thicknesses level off further downstream and start growing again as the pressure gradient becomes nearly zero. The changes in  $\delta$  occur at locations close enough to these particular stations. The rise in  $\delta$  within the aft region continues until  $x/L \approx 0.3$ , followed by a slight decrease until  $x/L \approx 0.4$ , which is then succeeded by the growth in the recovery region. As will be seen later, the boundary layer in the decelerating region develops an internal layer that looks much like a free shear layer; thus, the behavior of boundary layer quantities relative to one another in the aft region may not precisely fit the norms of a more typical boundary layer.  $H$  drops down to a value of about 1.35 by the time flow reaches  $x/L = 1$ . This value is within the typical range expected for zero pressure gradient turbulent boundary layers. As stated earlier,  $U_e$  follows the pressure gradient consistently in the whole domain.  $U_e$  starts increasing in the downstream recovery region once the pressure gradient becomes favorable again at  $x/L \approx 0.38$ .

### 3. Further Examination of the Strong Acceleration Region

We now examine the strongly accelerated flow over the first half of the speed bump in more detail. Figure 12 shows the variation of two potentially useful parameters over this region. The first parameter is the acceleration parameter,  $\Delta_p$ , and the second one is the relaminarization parameter,  $K$ . These parameters are defined as

$$\Delta_p = -\frac{\nu}{\rho u_\tau^3} \frac{\partial p}{\partial s} \quad \text{and} \quad K = \frac{\nu}{U_e^2} \frac{\partial U_e}{\partial s} \quad (4)$$


where  $\nu$  is the kinematic viscosity,  $\rho$  is the density,  $u_\tau$  is the wall friction velocity,  $p$  is the surface pressure,  $U_e$  is the boundary layer edge velocity, all of which are determined from the mean flow and  $s$  is the surface distance. Note that  $\Delta_p$  accounts for the boundary layer wall region<sup>||</sup> because it incorporates the friction velocity, while  $K$  is solely based on the boundary layer edge velocity, and therefore, cannot account for the wall region.



**Fig. 12** Variation of acceleration and relaminarization parameters over the upstream half of speed bump.

<sup>||</sup>Note that nondimensionalization of  $\Delta_p$  would bring  $Re_L$  into the denominator. An increase in the Reynolds number would lower  $u_\tau$ , but the product of the larger  $Re_L$  and the smaller  $u_\tau^3$  in the denominator of  $\Delta_p$  would generally be greater than the corresponding product at the lower Reynolds number, and can be estimated more precisely using the available skin-friction correlations. Hence, an increase in the Reynolds number would require a stronger favorable pressure gradient, relative to the lower Reynolds number case, for  $\Delta_p$  to reach a certain threshold.

In the literature, there is disagreement about the critical value of  $K$  above which relaminarization can take place. A threshold value of about  $3 \times 10^{-6}$  has acceptance by some researchers [11–13]. The peak value of  $K$  in Figure 12, about  $2.7 \times 10^{-6}$ , is found to be smaller than this critical value. Hence, this particular threshold of the  $K$  parameter does not predict the relaminarization in the present problem. Regarding the acceleration parameter,  $\Delta_p$ , Patel and Head [14] concluded that “major departures” from the logarithmic layer occur when this parameter exceeds 0.018. They believe this major departure to be an indication of the onset of “reverse transition” or relaminarization. We observe that  $\Delta_p$  reaches its critical value at  $x/L \approx -0.19$  in the present flow. The corresponding value of  $K$  at that location is about  $2 \times 10^{-6}$ . To see whether this particular station marks a significant turn of events during the course of flow acceleration, we first examine the mean streamwise velocity profiles in wall units at several stations in the region of interest, shown in Figure 13. As seen earlier in Figure 4, the logarithmic layer intercept constant has a value of  $C = 5$  at the recycle plane station, which is positioned within the weak adverse pressure gradient upstream of the bump. The strengthening adverse pressure gradient encountered near the foot of the bump brings down  $C$  to a value of 4.8 by the time the flow reaches  $x/L = -0.4$ . The logarithmic layer shifts down some more as the flow travels to  $x/L = -0.3$ , around which the pressure gradient changes from adverse to favorable. We observe an upward shift and narrowing down in the logarithmic layer as the accelerated flow progresses to  $x/L = -0.2$ . The logarithmic layer at this location is rather narrow, indicating that it is about to disappear. This station is fairly close to  $x/L = -0.19$  where the acceleration parameter reaches Patel and Head’s critical value of  $\Delta_p = 0.018$  [14].



**Fig. 13 Mean streamwise velocity profiles in wall units in the region where  $-0.4 \leq x/L \leq 0$ .**

Figure 14 shows the mean streamwise velocity profile in wall units at  $x/L = -0.19$ , where  $Re_\theta \approx 1692$  based on the local  $U_e$ . With an intercept constant of  $C \approx 5.26$ , the logarithmic layer equation overlaps with the actual velocity profile over only a few wall units, indicating that the logarithmic layer has nearly disappeared. Hence, we find the  $\Delta_p = 0.018$  value suggested by Patel and Head [14] to be a reasonable predictor of the logarithmic layer breakdown in the accelerating region of the present flow. As the flow continues its acceleration toward the apex, the logarithmic layer completely disappears and the velocity profiles depicted in Figure 13 take shapes that resemble those of relaminarizing boundary layers observed in laboratory experiments [15].

Next, we examine the variation of peak turbulence intensities, as well as peak turbulent kinetic energy, with streamwise distance in the region of interest. The variations of peak turbulence intensities and peak turbulent kinetic energy in the range where  $-0.4 \leq x/L \leq 0.1$  are shown in Figure 15. In this figure,  $u'_p$ ,  $v'_p$  and  $w'_p$  denote the peak turbulence intensity in the local streamwise, wall-normal and spanwise directions, respectively, and  $k_p$  is the peak turbulent kinetic energy. These quantities are normalized separately using the freestream reference velocity,  $U_\infty$ , and the local edge velocity,  $U_e$ . Examination of Figure 15 reveals the interesting trends in the development of peak turbulence intensities. In Figure 15(a), we see that both  $u'_p/U_\infty$  and  $k_p/U_\infty^2$  continuously grow in the accelerating region bounded by  $-0.3 \lesssim x/L \lesssim 0$ . The rates of growth of  $u'_p/U_\infty$  and  $k_p/U_\infty^2$  appear to slow down momentarily just upstream of the apex, but the growth rates ramp up again once the flow goes past the apex. The brief leveling off in  $u'_p/U_\infty$  appears to take place near the primary  $C_f$  peak location observed at  $x/L \approx -0.055$ . Figure 15(b) shows

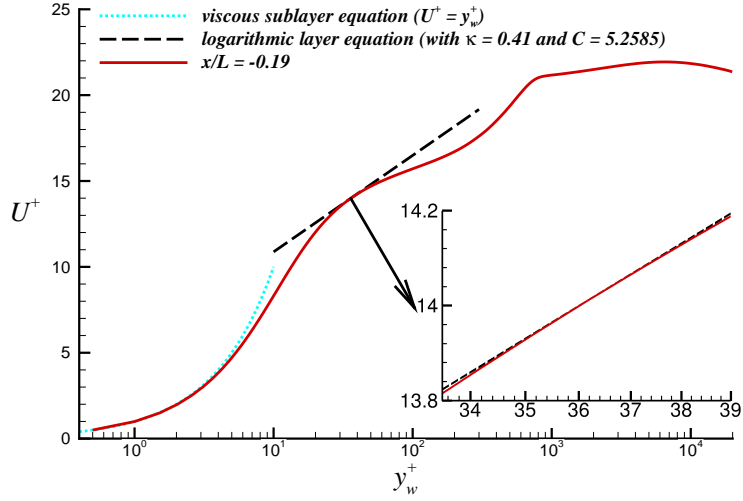




Fig. 14 Mean streamwise velocity profile in wall units at  $x/L = -0.19$ .



(a) Peak streamwise intensity and turbulent kinetic energy



(b) Peak wall-normal and spanwise intensities

Fig. 15 Streamwise variation of  $u_p'$ ,  $v_p'$ ,  $w_p'$  and  $k_p$  in the range where  $-0.4 \leq x/L \leq 0.1$ .

a bit of growth in  $v'_p/U_\infty$  and  $w'_p/U_\infty$  in the initial stages of flow acceleration, but this growth is relatively sedate compared to that of  $u'_p/U_\infty$ . Unlike  $u'_p/U_\infty$ ,  $v'_p/U_\infty$  starts a dip in the vicinity of  $x/L = -0.19$ , and  $w'_p/U_\infty$  follows suit shortly thereafter. Normalizing these turbulence intensities by  $U_e$  brings out the nature of these dips in a more dramatic fashion, as demonstrated by the variations in  $v'_p/U_e$  and  $w'_p/U_e$  curves. These quantities reveal that the drops in  $v'_p/U_e$  and  $w'_p/U_e$  have already begun some distance upstream of  $x/L = -0.19$  in the relatively early stages of flow acceleration, or in the vicinity of  $x/L \approx -0.25$  to be more precise. By the time the flow reaches just upstream of the apex, the peak  $v'_p/U_e$  and  $w'_p/U_e$  have gone down substantially by about 65% and 53%, respectively.

Scaling  $u'_p$  and  $k_p$  by  $U_e$  reveals where the corresponding dips in those quantities begin within the accelerating region, as seen in Figure 15(a). We observe that  $u'_p/U_e$  starts its dip shortly after  $x/L = -0.19$ , while  $k_p/U_e^2$  starts decreasing some short distance upstream, due to the fact that  $v'_p/U_e$  and  $w'_p/U_e$  have started dropping in approximately the same region. The overall drop in  $u'_p/U_e$ , however, is only about 15% and considerably less than that observed in the other components. This is because the flow still generates the streamwise fluctuation,  $u'/U_\infty$ , in the entire accelerating region, due to the presence of streamwise velocity gradient near the wall that continues to steepen until the primary  $C_f$  peak. It is well understood that in these types of shear flows (which are dominated by the velocity gradient along the wall-normal direction), the turbulent kinetic energy production mainly takes place in the production term of the streamwise Reynolds stress transport equation, and the fluctuating pressure redistributes the energy from the streamwise stress to the other two normal stresses. The dominant production term contains the product of the streamwise velocity gradient in the wall-normal direction and the Reynolds shear stress. The drops in  $v'_p/U_\infty$  and  $w'_p/U_\infty$ , despite the rise in  $u'_p/U_\infty$ , suggests that the strong favorable pressure gradient interferes with the transfer of the energy from the largest normal stress to the smaller ones. The rise in  $u'_p/U_\infty$  in the entire accelerating region is accompanied by the simultaneous increase in the local edge velocity, which then leads to the 15% drop observed in  $u'_p/U_e$ . The continuous growth of  $u'_p/U_\infty$  in the accelerating region, despite the dips in  $v'_p/U_\infty$  and  $w'_p/U_\infty$ , is sufficient to keep  $k_p/U_\infty^2$  continuously growing in the accelerating region as well, since  $u'_p/U_\infty$  is the biggest contributor to  $k_p/U_\infty^2$ . When scaled by  $U_e$ , the largest drop in  $k_p/U_e^2$ , observed just upstream of the apex, amounts to about 45%.

To recap, the strongly-accelerated flow experiences a relatively small decrease in  $u'_p/U_e$  but a much more substantial drop in  $v'_p/U_e$  and  $w'_p/U_e$  in the region that shows signs of stabilization or relaminarization. Such changes are indications of mixing suppression along wall-normal and spanwise directions along with a modest drop in streamwise fluctuations as the flow accelerates past  $x/L = -0.19$  and approaches the apex. To further examine these events, Figures 16 and 17 provide the normal Reynolds stress profiles along the local wall-normal direction at several stations in the region of interest. The corresponding Reynolds shear stress profiles are shown in Figure 18. In the upcoming discussion,  $\langle u'u' \rangle$ ,  $\langle v'v' \rangle$ ,  $\langle w'w' \rangle$  and  $\langle u'v' \rangle$ , respectively, are the streamwise, wall-normal, spanwise and shear components of Reynolds stress in the local orthogonal system at a given station, and the  $\langle \rangle$  operator denotes averaging in time and along the span. Reynolds stresses scaled by both  $U_\infty$  and  $U_e$  are shown in Figures 16, 17 and 18. The wall-normal distance is scaled by the local  $\delta$ . The mixing suppression along wall-normal and spanwise directions is verified by the lowering in the corresponding Reynolds stress profiles observed downstream of  $x/L = -0.2$ , which appears more severe in the scaling by  $U_e$ . The changes in the streamwise component within the same region are relatively less severe.

Despite the mixing suppression along wall-normal and spanwise directions that begins fairly early in the accelerating region,  $C_f$  continues its rise until its primary peak at  $x/L \approx -0.055$ , indicating that the near-wall velocity gradient continues to increase in the accelerating region until this point. Mixing suppression would normally tend to lessen the near-wall velocity gradient; however, the favorable pressure gradient simultaneously works to oppose to the mixing suppression effect. The rise in  $C_f$  until  $x/L \approx -0.055$  implies that the favorable pressure gradient effect is strong enough to counteract the mixing suppression effect and continue increasing the near-wall velocity gradient until that point. The negative impact of suppressed mixing on  $C_f$  is felt deep into the accelerating region, over a relatively short distance near the apex as the favorable pressure gradient begins to ease. The stabilizing effect of convex surface curvature, which becomes stronger as the flow approaches the apex, also plays a role on  $C_f$  reduction. The convex curvature effect is discussed in section III.B.4. The apex roughly marks the end of the favorable pressure gradient and the beginning of the adverse pressure gradient. All peak turbulence intensities, regardless of the scaling by  $U_\infty$  or  $U_e$ , display a rapid increase once the stabilizing effect of the favorable pressure gradient disappears past the apex, within the early stages of the adverse pressure gradient. Transition of the stabilized flow back to a turbulent state generates the secondary  $C_f$  peak location at  $x/L \approx 0.05$  with associated peaks in  $u'_p$  and  $k'_p$  curves slightly upstream, at  $x/L \approx 0.04$  or so. The corresponding peaks in  $v'_p/U_\infty$  and  $w'_p/U_\infty$  appear some distance in the wake of the secondary  $C_f$  peak, at  $x/L \approx 0.07$ , while the  $v'_p/U_e$  and  $w'_p/U_e$  curves display only a brief slowdown in their streamwise growth at that location. Further examination of the flow features in the downstream half of the bump is provided in section III.B.5.






Fig. 16 Normal Reynolds stress profiles scaled by  $U_\infty$ , where  $-0.4 \leq x/L \leq -0.025$ .

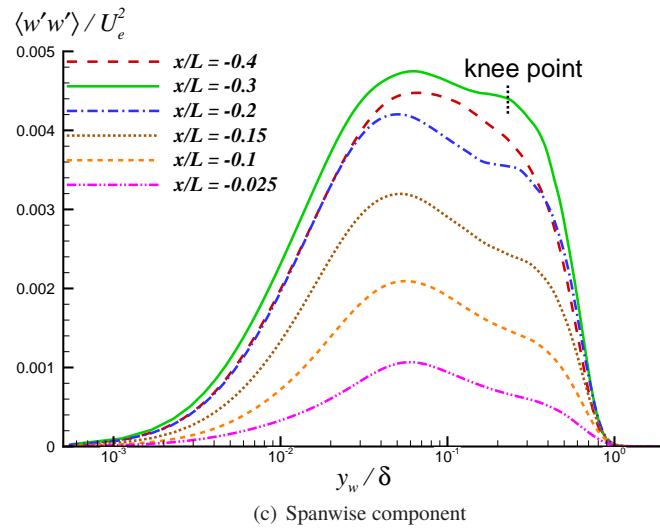
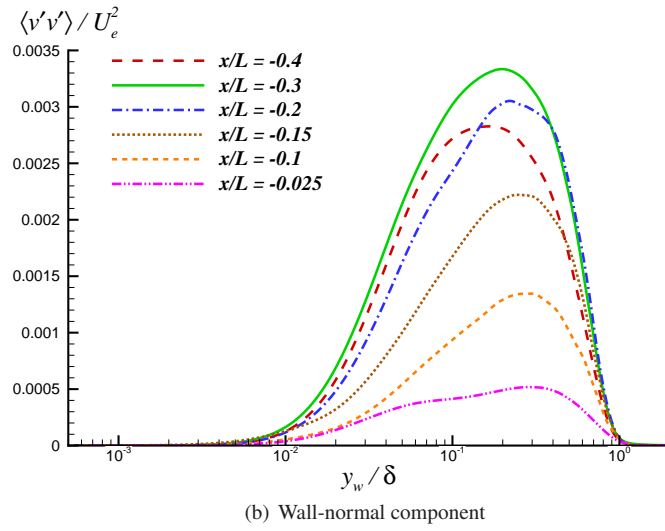





Fig. 17 Normal Reynolds stress profiles scaled by  $U_e$ , where  $-0.4 \leq x/L \leq -0.025$ .






(a) Shear stress profiles scaled by  $U_\infty$



(b) Shear stress profiles scaled by  $U_e$




(c) Shear stress and mean streamwise velocity profiles in wall units

**Fig. 18 Reynolds shear stress and mean streamwise velocity profiles in the accelerating region.**

Another noteworthy observation from Figures 16 and 17 is discussed next. The streamwise station of  $x/L = -0.3$  shows the emergence of a so-called “knee point” in the streamwise and spanwise Reynolds stresses at  $y_w/\delta \approx 0.23$ , marked in Figure 17. The knee point is where the brief slowdown in the outward decrease of these quantities is succeeded by a much faster decay rate. This is an indication of the formation of an internal layer in the accelerating region. Such internal layers are known to be triggered by sudden changes in local wall boundary conditions, discontinuity in surface curvature or by a change in the sense of streamwise pressure gradient. Note that the present geometry is defined by a smooth analytical function, thus it does not contain any discontinuity in its surface curvature. The problem formulation also does not contain sudden changes in local wall boundary conditions anywhere. Hence, the internal layer is believed to be triggered by the switch from the initially adverse to favorable pressure gradient in the vicinity of  $x/L = -0.3$ .

Figure 19 shows the location of the knee point identified in the streamwise and spanwise Reynolds stress profiles at  $x/L = -0.3$ , and its position relative to the peaks in wall-normal and shear components at the same station. The knee point appears to coincide with where the Reynolds shear stress profile peaks. The wall-normal stress peak is also positioned fairly close by. As seen in Figures 16 and 17, once formed, the knee point shifts outward with streamwise distance, indicating that the internal layer grows from the wall. The peaks in Reynolds shear and wall-normal stresses generally follow the knee point, while the peaks in the other two components are positioned closer to the wall within the growing internal layer. At the apex, the knee point is found at  $y_w/\delta \approx 0.44$ . As will be seen, the internal layer is further characterized by the development of an inner peak in the Reynolds shear stress profiles.



**Fig. 19 Reynolds stress profiles at  $x/L = -0.3$ .**

On a related note, wall-normal stresses shown in Figures 16(b) and 17(b) also provide the first hint for the formation of another internal layer as the flow nears the apex, where the pressure gradient changes from favorable to adverse. The wall-normal stress profile at  $x/L = -0.025$  shows that the rate of increase in the wall-normal stress in the outward direction momentarily slows down in the vicinity of  $y_w/\delta \approx 0.08$  before picking up pace again. As will be seen, this particular portion of the wall-normal stress profile will develop an inner peak at the apex. Further examinations of the second internal layer and the flow development downstream of the apex are provided in section III.B.5.

The evolution of Reynolds shear stress profiles within the accelerating region, depicted in Figures 18(a) and (b), reveals the emergence of an inner peak within the first internal layer identified above. These profiles are scaled by both  $U_\infty$  and  $U_e$ . The scaling by  $U_e$  brings out the weakening of the inner peak due to relaminarization/stabilization in a more dramatic fashion. Note that the production of Reynolds shear stress depends on the streamwise velocity gradient in the wall-normal direction and wall-normal Reynolds stress. Despite the steepening near-wall velocity gradient due to the favorable pressure gradient in much of the accelerating region, mixing suppression reduces the wall-normal stress level, which then impacts the production of Reynolds shear stress near the wall. Downstream of  $x/L \approx -0.12$ , the inner peak becomes relatively stronger than the outer peak, which decays at a higher rate deep into the accelerating region. Figure 18(c) depicts the Reynolds shear stress and mean streamwise velocity profiles in wall units in the region where  $-0.2 \leq x/L \leq -0.025$ . We observe that as the flow approaches the apex, the mean streamwise velocity gradient in the vicinity of the outer stress peak is gradually flattened by the acceleration; this naturally has a detrimental effect on the production of normal and shear stresses around the edge of the internal layer in the accelerating region.

#### 4. Convex Surface Curvature Effect

Previous research related to surface curvature effects on turbulent boundary layers has provided the understanding that convex surface curvature has a stabilizing effect on turbulence, and leads to a reduction in Reynolds stresses and skin friction. The ratio of the local boundary thickness to the radius of curvature is an important parameter in quantifying the significance of curvature. For example, Muck et al. [16] state that “mild convex curvature, with a radius of curvature of the order of 100 times the boundary-layer thickness, tends to attenuate the preexisting turbulence”.

Figure 20 shows the radius of curvature,  $R$ , variation in the range where  $-0.4 \leq x/L \leq 0.1$ . Because of the geometrical symmetry with respect to  $x/L = 0$ , the variation of  $R$  is also symmetric with respect to  $x/L = 0$ . The bump geometry has convex curvature in the range where  $-0.137886 \leq x/L \leq 0.137886$  and concave curvature elsewhere. At  $x/L \approx \mp 0.137886$ ,  $R$  goes to infinity as the curvature changes from concave to convex or vice versa. Figure 21 depicts the variation of  $\delta/R$  over the convex curvature portion of the bump, and  $s_c/\delta$ , where  $s_c$  is the surface distance measured from the start of convex curvature. We observe that  $\delta/R$  increases as the accelerated flow moves toward the apex. In the vicinity of the primary  $C_f$  peak, where  $x/L \approx -0.055$ ,  $\delta/R$  approaches 0.05. The largest  $\delta/R$ , found at  $x/L \approx 0.009$ , is about 0.0676. We also see that  $s_c/\delta \approx 10$  at around the apex, meaning that by the time the flow reaches the apex, it has been subjected to the stabilizing convex curvature effect over a surface distance that is about 10 times the  $\delta$  at the apex. These  $\delta/R$  values are considerably greater than the 0.01 value that represents “mild” convex curvature [16]. Thus, there is a nonnegligible stabilizing effect of the convex curvature that comes into play as the accelerated flow moves toward the apex. The convex curvature effect extends into the adverse pressure gradient region as well. The effect of convex curvature is intermingled with the pressure gradient effect for the configuration under investigation, and it is therefore hard to separate one effect from the other. Nevertheless, these findings reinforce the understanding that the relaminarization/stabilization phenomenon observed deep into the accelerating region is the combined outcome of the favorable pressure gradient and the convex curvature.

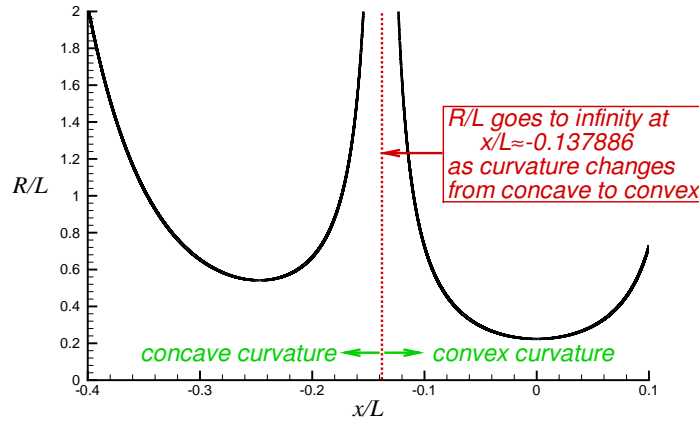


Fig. 20 Radius of curvature variation in the range where  $-0.4 \leq x/L \leq 0.1$ .


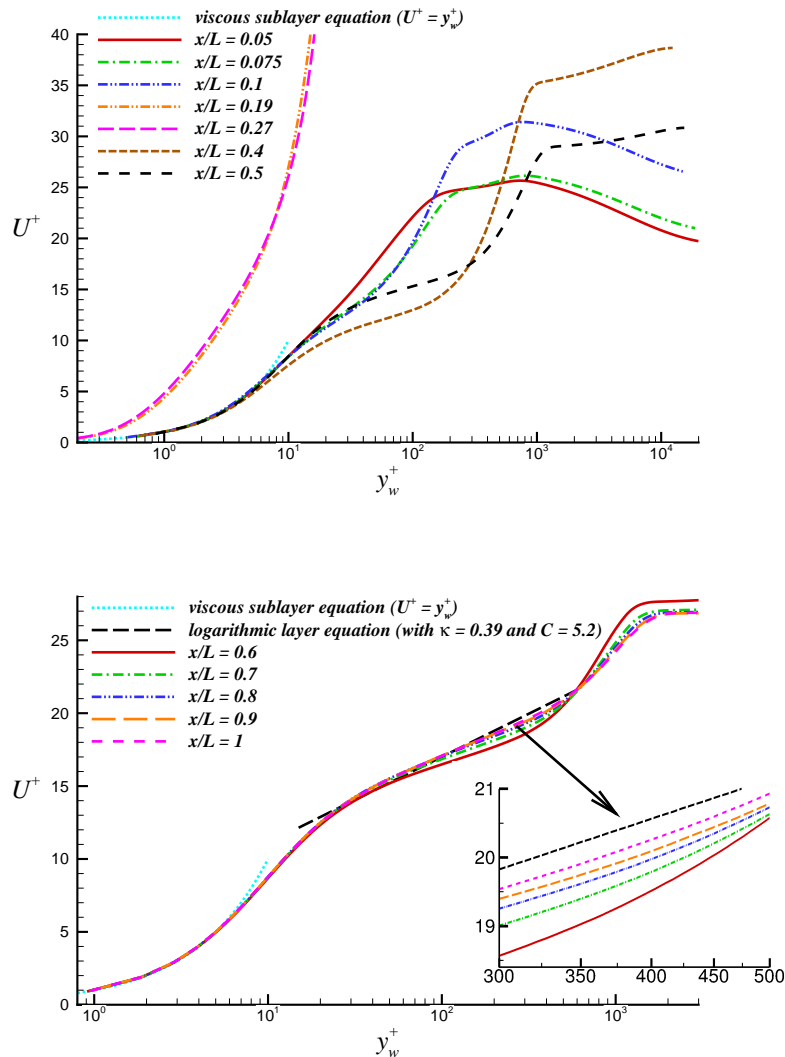


Fig. 21 Variation of  $\delta/R$  and  $s_c/\delta$  over the convex curvature portion of the bump.

### 5. Further Examination of the Deceleration, Separation and Downstream Recovery Regions

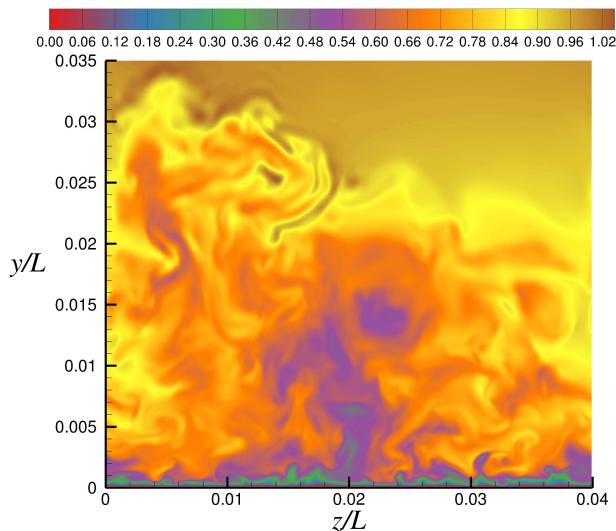
We now shift our attention to the second half of the speed bump, which contains the deceleration, incipient or very weak separation, and downstream recovery regions. Figure 22 shows the mean streamwise velocity profiles at several stations in that region. The profile at  $x/L = 0.05$ , corresponding to where the secondary  $C_f$  peak is located, and those in the near wake of the secondary  $C_f$  peak (i.e., at  $x/L = 0.075$  and  $0.1$ ), reveal that despite the transition back to turbulence indicated by the flow visualization in that region, there is not enough development length to generate a logarithmic layer in the velocity profiles. The continued presence of an adverse pressure gradient slows down the flow and brings it to the brink of separation. As seen earlier, the  $C_f$  distribution becomes slightly negative in the range where  $0.195 \lesssim x/L \lesssim 0.268$ , indicating that the flow undergoes incipient or very weak separation in that region. The profiles at  $x/L = 0.19$  and  $0.27$ , which are positioned near the separation and reattachment points, respectively, show the markedly different nature of these profiles compared to those in the attached region. This is not surprising since the wall unit scaling is based on a nearly-zero friction velocity at these stations.



**Fig. 22 Mean streamwise velocity profiles in the downstream half of speed bump.**

The flow recovery from separation begins in the wake of the reattachment point. Several stations taken within the downstream recovery region, from  $x/L = 0.5$  to  $1$ , show the emergence of a logarithmic layer in the turbulent boundary layer. As can be seen in the zoomed-in view included in the bottom subfigure of Figure 22, the shifting portions of velocity profiles suggest that the logarithmic layer is still in the process of adjustment as the flow reaches the end of the physical domain at  $x/L = 1$ . Examination of Reynolds stress profiles in the recovery region later in this section will

provide further evidence of the continuing boundary layer development. A reasonable fit to the emerging logarithmic layer is given by a von Kármán constant of  $\kappa = 0.39$  and an intercept constant of  $C = 5.2$ . As seen in Figure 11,  $\delta$  in the early recovery region following the incipient/weak separation is about the same as the domain span, which is  $0.04L$  in the present DNS. The boundary layer thickness continues to increase further downstream. Hence, the domain span is rather narrow when compared to the boundary layer thicknesses in the recovery region. A snapshot of the flow at  $x/L = 0.5$ , depicted in Figure 23, confirms the presence of large-scale structures that are comparable in size to the span. This observation raises the issue of the constraint imposed by the narrow span, which could skew the flow development and affect logarithmic layer properties in the recovery region to some extent. Further investigation of the span effect will be made when additional computational resources become available.



**Fig. 23** Normalized total velocity magnitude contours at  $x/L = 0.5$ .

Figures 24 and 25 depict the evolution of normal Reynolds stress profiles in the region from the bump apex until the end of the physical domain. Figure 26 shows the Reynolds shear stress profiles in the same region. For brevity, only the Reynolds stresses scaled by  $U_\infty$  are shown here. Figure 27 shows the variation of peak turbulence intensities in the downstream half of the flowfield. This figure provides the quantities scaled by both  $U_\infty$  and  $U_e$ . As seen in Figure 27, the transition of the stabilized flow back to turbulence in the early stages of an adverse pressure gradient generates the highest peak streamwise intensity found in the flowfield, slightly upstream of the secondary  $C_f$  peak. Here, we see values of  $u'_p/U_\infty \approx 0.232$  and  $u'_p/U_e \approx 0.156$ . The overshoot in  $u'_p$  due to the transition back to turbulence is rather short-lived. This is also verified by the decrease in the streamwise Reynolds stress profiles, observed downstream of  $x/L = 0.05$  in Figure 24(a). According to Figure 27, the peak streamwise intensity rapidly drops until  $x/L \approx 0.1$ , after which it starts increasing again. As seen in Figure 24(a), this switch is due to the emergence of an outer streamwise Reynolds stress peak.

We clearly see the evolution of the outer peak in the streamwise Reynolds stress profiles starting from  $x/L = 0.065$ , which is believed to be associated with the formation of a new internal layer, triggered by the switch from favorable to adverse pressure gradient at the bump apex. As will be seen, this new internal layer forms as a free shear layer embedded within the decelerating boundary layer. The first hint for this new internal layer was seen earlier in the wall-normal Reynolds stress profiles depicted in Figure 16(b) or Figure 17(b). Figure 24(b) shows that an inner peak has formed in the wall-normal Reynolds stress profile by the time the flow reaches the bump apex. The inner peak is found at  $y_w/\delta \approx 0.054$ . Note that the outer peak in the wall-normal profile at that station is the remnants of the outer edge of the first internal layer. At the bump apex, we also find the primary peak of the spanwise Reynolds stress profile at  $y_w/\delta \approx 0.05$ , which is very close to where the inner peak in the wall-normal profile forms. This inner peak in the spanwise stress profile is the surviving portion of its primary peak engulfed within the first internal layer, after the flow experiences mixing suppression due to acceleration. The Reynolds shear stress profile at the bump apex has its inner peak at  $y_w/\delta \approx 0.033$ , which is not too far off from the inner peak emerging in the wall-normal profile at  $y_w/\delta \approx 0.054$ . This inner peak in the Reynolds shear stress profile starts its development within the first internal layer further upstream within the accelerating region, as seen earlier in Figure 18. These observations suggest a natural link between the first and second internal layers, which is discussed further below.


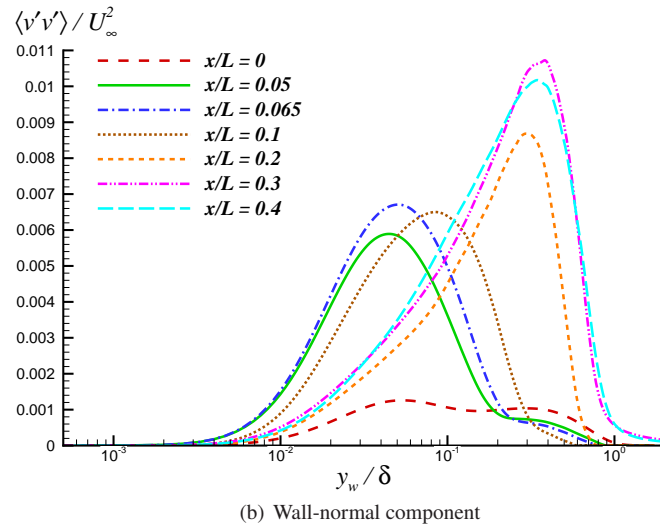



Fig. 24 Normal Reynolds stress profiles scaled by  $U_\infty$ , where  $0 \leq x/L \leq 0.4$ .






Fig. 25 Normal Reynolds stress profiles scaled by  $U_\infty$ , where  $0.4 \leq x/L \leq 1$ .

As the flow going past the apex transitions back to turbulence, we observe the strengthening of the primary peak in the wall-normal and spanwise Reynolds stresses in Figures 24(b) and (c). This rise is briefly interrupted by a leveling off or a slight drop at around  $x/L \approx 0.07-0.1$ , as indicated by the respective turbulent intensity components plotted in Figure 27. The peaks in wall-normal and spanwise Reynolds stress profiles move toward the wall as the flow passing the apex transitions back to a turbulent state, and then begin to shift outward in the post transition region. At  $x/L = 0.1$ , the peak in wall-normal and spanwise Reynolds stress profiles is at  $y_w/\delta \approx 0.085$  and  $0.081$ , respectively. The corresponding peak in the streamwise Reynolds stress at the same station is found at  $y_w/\delta \approx 0.107$ . Note that the inner peak of the streamwise Reynolds stress profile at the same station, found at  $y_w/\delta \approx 0.017$ , weakens because of the reduction in the near-wall velocity gradient in the decelerating region, which has a detrimental effect on the streamwise Reynolds stress production. The inner peak in the Reynolds shear stress profiles also becomes stronger following the transition back to turbulence, and is positioned at  $y_w/\theta \approx 0.091$  for the streamwise station of  $x/L = 0.1$ . This peak is in the vicinity of the peaks observed in the other Reynolds stress components at the same station. As the flow reaches  $x/L = 0.2-0.3$ , the profiles of all normal Reynolds stresses as well as the shear stress peak fairly close to one another, and the peaks shift away from the wall to where  $y_w/\delta \approx 0.3-0.4$ . These peaks generally stay in the vicinity of  $y_w/\delta \approx 0.35$  as the flow enters the recovery region downstream of the very-weakly reversed region.

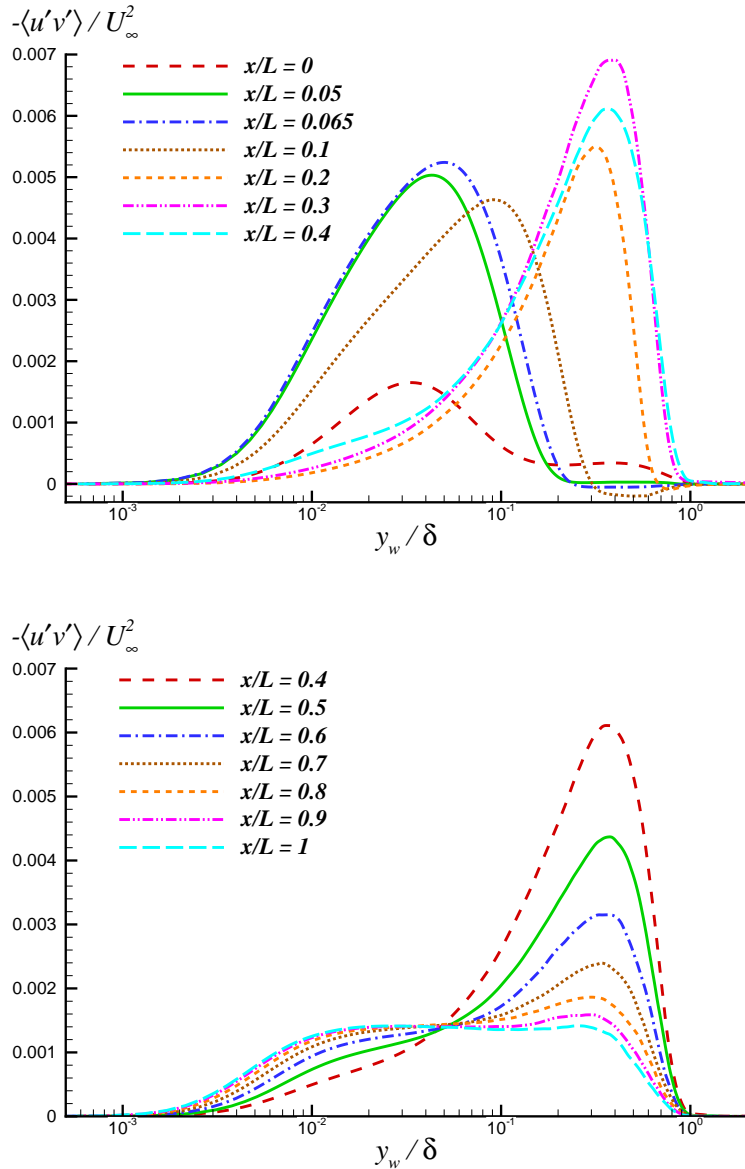


Fig. 26 Reynolds shear stress profiles scaled by  $U_\infty$ , where  $0 \leq x/L \leq 1$ .






Fig. 27 Streamwise variation of  $u_p', v_p', w_p'$  in the range where  $0 \leq x/L \leq 1$ .

Figure 28 shows the Reynolds shear stress profiles and the corresponding mean streamwise velocity profiles at three stations in the decelerating region. As expected, the flow in the immediate vicinity of the wall is decelerated considerably by the adverse pressure gradient. The flow in the outer region slows down as well, however, deep into the decelerating region, the outer flow, located at a wall-distance of  $y_w/\delta \approx 1$  and beyond, still moves at a speed comparable to the freestream velocity. The buffer zone between the high and low speed regions, in which the peaks of Reynolds shear stress profiles are generated, is characterized by a prominent velocity-gradient profile resembling that of a free shear layer. As seen above, the peaks of normal stress components in the decelerating region are located in close proximity to the shear stress peak. These observations suggest that there is an intimate connection between the internal layers generated in the accelerating and decelerating regions. The high-gradient portion of the first internal layer that encompasses the near-wall Reynolds shear stress peak moves away from the wall and evolves into the free shear layer embedded within the decelerating boundary layer. The mixing in the free shear layer generates the peaks in all Reynolds stress component profiles examined earlier.

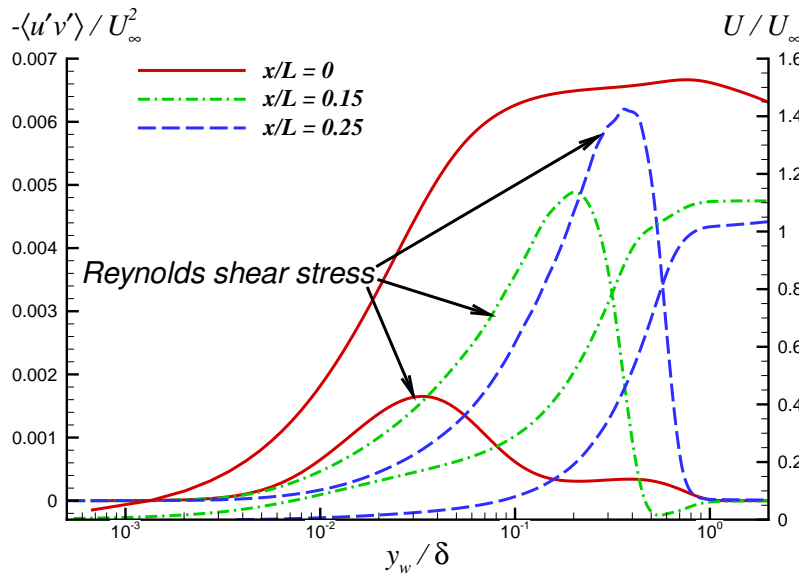



Fig. 28 Reynolds shear stress and mean streamwise velocity profiles, where  $0 \leq x/L \leq 0.25$ .

Figure 27 shows that the spanwise intensity,  $w'_p/U_\infty$ , reaches its global peak within the very-weakly reversed flow region, while the wall-normal intensity,  $v'_p/U_\infty$ , reaches its global peak in the early stages of recovery from separation. The secondary peak in  $u'_p/U_\infty$  downstream of the bump is also found to occur within the very-weakly reversed flow region. These peaks are generated within the free shear layer region of the flow identified above. Scaling the wall-normal and spanwise intensities by  $U_e$  shows a continuous rise from the apex until they reach their respective peak. The peak in  $w'_p/U_e$  is broad; it begins in the reversed region and extends into the early stages of recovery, while that in  $v'_p/U_e$  is narrower and positioned within the early recovery region, at  $x/L \approx 0.36$ . The local peaks observed in the  $v'_p/U_\infty$  and  $w'_p/U_\infty$  curves at  $x/L \approx 0.07$  appear as a brief slowdown in the corresponding curves scaled by  $U_e$ . The secondary peak in  $u'_p/U_e$  appears to be more of a twin peak that covers part of the reversed and the early recovery region. The secondary peak value in terms of  $u'_p/U_\infty$  is considerably lower than the corresponding primary peak value, whereas the scaling by  $U_e$  shows that the secondary peak in terms of  $u'_p/U_e$  is nearly as strong as the corresponding primary peak. The drop in  $u'_p$  after the secondary peak continues until  $x/L \approx 0.54$ , after which it shows signs of leveling off in both scalings. As seen in Figure 25, this is because of the development of an inner streamwise Reynolds stress peak near the wall once the flow reattaches and becomes subjected to a mildly favorable pressure gradient. The strength of the inner peak surpasses that of the outer peak, starting at  $x/L \approx 0.54$ , as the reattached flow keeps developing in the recovery region. Further into the recovery region, both  $v'_p$  and  $w'_p$  decay monotonically regardless of the scaling, and appear to be leveling off toward the end of the physical domain.

The evolution of normal Reynolds stress profiles in the recovery region can be seen in Figure 25. The bottom subfigure of Figure 26 shows the corresponding Reynolds shear profiles in the same region. As noted earlier, the pressure gradient changes from adverse to mildly favorable at  $x/L \approx 0.38$ . The first indication of a near-wall stress peak in the recovery region can be found in the streamwise and spanwise Reynolds stress profiles shown at  $x/L = 0.3$ – $0.4$ . The stations positioned downstream depict how the inner peaks develop further. All Reynolds stress components generate their respective inner peak near the wall, although the domain length is not sufficient for the wall-normal and shear component peaks to fully develop by the time flow reaches the end of the physical domain at  $x/L = 1$ . This resurgence of Reynolds stress peaks near the wall immediately following the very-weakly reversed flow region can be viewed as the development of yet another internal layer triggered by the change in the sense of pressure gradient. The outer stress peak generated by the mixing in the free shear layer portion of the flow begins to decay in all Reynolds stress profiles with streamwise distance. This is because the velocity gradient within the free shear layer is decreased by the mildly favorable pressure gradient, and this causes the free shear layer to gradually vanish in the recovery region, as depicted by the profiles shown in Figure 29. This figure also shows that the boundary layer is still continuing to develop by the time the end of the physical domain at  $x/L = 1$  is reached. With a longer domain length, we would expect to see the outer peaks in the Reynolds stresses to completely disappear, and the mean velocity as well as Reynolds stress profile shapes approach those of a fully developed turbulent boundary layer under zero pressure gradient conditions.



**Fig. 29** Reynolds shear stress and mean streamwise velocity profiles, where  $0.5 \leq x/L \leq 1$ .

## IV. Conclusions

A direct numerical simulation of an initially turbulent boundary layer passing over a wall-mounted speed bump geometry has been performed using a flow solver developed exclusively for graphics processing units. The turbulent boundary layer upstream of the bump is initially subjected to a mild adverse pressure gradient that becomes stronger as the bump is approached. The pressure gradient turns favorable at the bump foot and subjects the turbulent boundary layer to a strong acceleration over the upstream half of the bump. This acceleration is strong enough to cause a stabilization or relaminarization of the relatively low Reynolds number turbulent boundary layer as it approaches the apex. The convex surface curvature of the bump is believed to be an additional stabilizing factor for the turbulent boundary layer over the bump. The ratio,  $\delta/R$ , increases as the flow approaches the apex and reaches a maximum value of 0.0676, which significantly exceeds the value suggested in Muck et al. [16] for the stabilizing effect of surface curvature on turbulence. The critical value of the acceleration parameter, which is 0.018 according to Patel and Head [14], is found to reasonably predict the location where the logarithmic layer in the mean streamwise profile is about to disappear in the accelerating region. That particular station is located at  $x/L \approx -0.19$ . The mean velocity profiles located between this station and the bump apex at  $x/L = 0$  are found to resemble those of relaminarizing boundary layers observed in laboratory experiments [15]. The surviving instabilities manage to initiate a transition back to a turbulent state shortly after the flow passes over the apex and enters the adverse pressure gradient.

An internal layer is found to develop over the front half of the speed bump within the accelerating region. This internal layer is believed to be triggered by the switch from adverse to favorable pressure gradient at the foot of the bump. The initial outer edge of this internal layer is approximately marked by the peaks in the corresponding Reynolds shear and wall-normal stress profiles at the beginning of the accelerating region. The internal layer grows from the wall as the flow travels toward the apex. It is not particularly thin; its thickness at  $x/L = -0.3$  is about  $4.4 \times 10^{-3}L$  and increases to about  $6.6 \times 10^{-3}L$  at around the apex. An inner Reynolds shear stress peak also emerges within this internal layer as it develops, but its growth gets hampered by the mixing suppression arising from the flow relaminarization/stabilization deep into the accelerating region. The peaks in the corresponding streamwise and spanwise Reynolds stresses at the beginning of the accelerating region, which are located closer to the wall, remain engulfed within the internal layer as it develops. The spanwise stress peak is significantly weakened by the mixing suppression while the streamwise stress peak is affected to a lesser extent.

The adverse pressure gradient generated over the second half of the bump decelerates the flow and brings it to the brink of separation. Flow visualizations as well as the statistical data indicate that the flow only experiences incipient or very weak separation. The decelerating region is further characterized by the emergence of a new internal layer, in the form of a free shear layer embedded within the decelerating boundary layer. Evidence suggests that there is an intimate connection between this free shear layer and the internal layer generated in the accelerating region. The mixing in the free shear layer generates the peaks in all Reynolds stress components observed in the aft region, generally at about 0.3 to 0.4 times the local boundary layer thickness away from the wall. As the recovery from incipient or very weak separation begins, the pressure gradient changes from adverse to mildly favorable, and yet another internal layer begins to develop near the wall. This internal layer develops toward a fully turbulent boundary layer beneath the free shear layer positioned further above the wall, which gradually vanishes with streamwise distance.

Finally, comparison of the present DNS results with the RANS  $C_f$  and  $C_p$  data shows that even though there is reasonable overall agreement in the  $C_p$  predictions, RANS does not detect the flow relaminarization observed in the accelerating region. It also predicts more severe separation in the aft (or decelerating) region. The inability of RANS to detect relaminarization in the present problem is hardly surprising, and its prediction of more severe separation relative to what is seen in the present DNS is another example showing its deficiency in adverse pressure gradient flows involving separation. A spanwise-periodic DNS of the same geometry at a higher Reynolds number that would preclude relaminarization/stabilization in the accelerating region is planned as part of the ongoing efforts related to this test case. At the moment, it is not clear whether the higher Reynolds number would create more severe separation in the aft region of the speed bump geometry. That remains to be seen. Nevertheless, the present test case has clearly demonstrated that the flowfield is already quite complex even without massive flow separation. Hence, even if the higher Reynolds number case does not generate massive separation, the DNS results from that case should still be of value to turbulence modeling efforts. The performance of turbulence models is hardly fully satisfactory in complex attached flows subjected to pressure gradients, let alone separated flows; hence, the results from the planned higher Reynolds number DNS, with a larger span, will facilitate comparisons with RANS results and those of other wall-modeled simulation approaches, which should be useful for further assessing the true predictive capability of such lower-fidelity simulation tools. Those findings will be reported in our future publications.

## Appendix: GPU Code Performance and Validation

### A. GPU Code Scaling Performance

The flow solver implements strategies that overlap communication with computation, which are crucial for achieving good performance on GPUs. This overlap is achieved as follows. The GPUs first compute the near-boundary information needed by their neighbors and copy the data to their host central processing units (CPUs) for MPI communication. The GPUs then compute the interior points while the host CPUs handle the MPI communication in parallel. The host CPUs run several OpenMP threads (one dedicated thread per face of a grid block) to perform the MPI data exchange. Assuming a large enough GPU workload, by the time the GPUs have finished computing interior points, the MPI communication among the host CPUs has been completed and the host CPUs have copied the exchanged MPI data back to their corresponding GPUs. The GPUs can then update the points near their block interface using the data received from their neighbors and can proceed further. The data copies between the GPU and the host CPU take place asynchronously, meaning that the GPU can do the computing (as long as there is sufficient work) and the data exchange with the host CPU simultaneously. Moreover, the GPU can run several independent tasks in parallel, as long as the resources needed by the computational kernels are available. The GPU operations are synchronized at critical points.


During the course of the developmental work, it was realized that the chosen explicit algorithms run extremely fast on the GPU since they are essentially made up of many independent multiply-add operations, at which the GPU excels. However, the host-assisted MPI communication among the GPUs was found to be a major bottleneck, as already known from others' similar experience in GPU code development. We observed that, unless we assign a large workload to the GPU and overlap computation with communication, MPI communication over the network cannot keep up with GPU computation. When the workload per GPU is too small, an inefficiency arises because the GPU finishes its work way too quickly and then sits idle while waiting for the communication to catch up. This idle time is nothing but a wasted opportunity to have the GPU do useful work. Hence, in order to have the GPU run at full efficiency and make the best use of limited computational resources, the present strategy is to assign the maximum possible workload to the GPU and overlap computation with communication as much as possible. To achieve the maximum possible workload, we adjust the grid block size per GPU in order to nearly or fully max out the available GPU global memory, which is 16 Gigabytes per GPU on the Summit system. This corresponds to about 24-25 million points per GPU on Summit. The GPUs at the NASA Advanced Supercomputing Division (NAS) facility have 32 Gigabytes of global memory, which allows nearly 50 million points per GPU on that system.

In this work, the host-assisted MPI communication approach is chosen over the so-called "GPU-aware MPI" option, which accomplishes direct communication among the GPUs without the complication of sending the data through the host CPUs. This choice was made because the performance of the GPU-aware MPI implementation available at the time of code development was found to be rather poor and unacceptable. Although the implementation of host-assisted MPI communication is more involved, assigning a large workload to the GPUs and having the host CPUs handle the MPI communication overlaps computation with communication effectively and hides the communication cost, as noted above. It is therefore well worth the additional complication since it improves the overall code performance.

Because of the slow communication issue relative to computation, we do not expect to see good strong scaling performance on the current systems. Communication simply would not be able to keep up with computation if the GPUs were assigned too little work. This issue can only be resolved by a much faster communication network, which presently does not exist. Hence, the only way to make the best use of limited resources at present is to run the code in its maximum efficiency mode. Although the maximum efficiency mode may not minimize the wall-clock run time, it will certainly minimize the total number of node-hours needed for a given problem size, which is the more relevant performance metric given the fact that the available resources are limited.

A weak scaling study has been performed on the Summit system, in which the total workload per GPU is kept fixed while increasing the total number of GPUs. Each node on Summit contains 6 NVIDIA Tesla V100 GPUs, with 16 Gigabytes of global memory per GPU. Our test case for the scaling study involves the flow over a flat plate. For the weak scaling study runs, the total number of GPUs is varied from 150 (25 nodes) up to 6000 (1000 nodes). Note that 1000 nodes make up 21.7% of the entire system. For each case, we assign a grid block of  $256 \times 320 \times 288$  points to each GPU. Test runs for each node-count case were performed to measure the total time taken to run 1000 time steps and compute the corresponding average time per step. The third-order, three-stage explicit Runge-Kutta scheme is used for time advancement. Figure 30 shows how the average time taken per step, in milliseconds, varies as the number of Summit nodes are increased. Note that the horizontal axis is plotted in logarithmic scale. Depending upon how the nodes assigned to a particular job are distributed within the system and the network route among the nodes, as well as the network traffic due to other jobs running on the system, code execution speeds may vary from one run

to another. We therefore see a slight variation in the average time per step as the number of nodes changes. The average time per step is nearly constant and remains between 149 and 150 milliseconds. This is the payoff for utilizing a strategy that assigns a large workload to the GPU and overlaps GPU computation with MPI communication as much as possible. The “mean” value of the time taken per step (i.e., the value averaged over all node-count runs) is about 149.5 milliseconds for the given grid block size per GPU. This corresponds to roughly one nanosecond per grid point per Summit node per time step.



**Fig. 30 Weak scaling performance on the Summit system.**

The overall performance of the GPU code was measured to be 10% of the peak theoretical double-precision performance of the V100 GPU, which is rated at 7.8 Tera-floating-point-operations (TFLOPs) per second. The most compute-intensive individual kernels (or subroutines) were found to achieve as high as 17% of the peak. To determine the performance limiting factor, the code was profiled using the available performance analysis tools. For a grid block of about 24 million points per GPU, it was found that, at each time step, the code moves about 71 Gigabytes worth of total data from the global memory to the registers of the processing units, and about another 39 Gigabytes worth of total computed data from the registers back to the global memory. Thus, there is a data movement of about 110 Gigabytes at each time step, between the global memory and the registers. Now, in order to illustrate what limits the code performance, suppose the code only moves this much data between the main memory and registers, but does not do any computing at all. How long would this data movement alone take? To answer this question, we first note that the V100 GPU architecture provides 900 Gigabytes/second of peak memory bandwidth. For simplicity, let us assume that the data movement takes place at the peak bandwidth. That would mean that at least  $110/900 \approx 0.122$  seconds would be needed to move that amount of data. Our profiling measurements show that the most compute-intensive kernels generally achieve about 80 to 90% of the peak memory bandwidth, so the actual data movement takes a bit longer than this estimate. With an overall average rate of 85% of the peak bandwidth, the data movement would take about 0.144 seconds. We also know that for the given grid block size per GPU, the code takes about 0.15 seconds to perform all operations and advance the simulation for one time step. The profiling measurements show that the actual compute time is around 19% of the total elapsed time. Even though the actual computations overlap with the memory operations, the data movement between the memory and registers still constitutes a significant chunk of the time taken per computational time step according to the above timings. These observations lead us to the conclusion that our code performance is bound by the available memory bandwidth. In other words, the memory bandwidth is not sufficient to transfer data into and out of the registers at the rate demanded by the processing units. Thus, the well-known adage of the computer science world, “The FLOPs are free, you are paying for the memory bandwidth!” is still very much valid in our case.

As noted above, with about 80 to 90% of the peak memory bandwidth achieved by the most compute-intensive kernels, our memory-bound GPU code is not far off from its maximum possible performance; hence, reaching a much greater percentage of the peak theoretical FLOPs per second performance of the V100 architecture is not feasible. We anticipate that potential memory bandwidth improvements in future-generation GPU architectures should enable our code to achieve higher FLOP counts per second on those systems.

## B. Performance Comparison to CPU Flow Solver

We now provide the performance comparison between this new GPU code and our previous CPU code, which was most recently used in the simulation of flow separation problems [8, 9]. The speedup factors depend on how exactly the comparisons are made. For example, the “node-to-node” speedup factor, which is derived from the performance

comparison based on one Summit node with 6 GPUs versus one dual-socket Intel Skylake CPU node with 40 cores, comes out to about 75 $\times$ . This comparison is for the version of the CPU code based on the third-order, three-stage explicit Runge-Kutta scheme, same as that employed in the GPU code. Note that the CPU code uses high-order compact finite-difference and filtering schemes, whose implementations are well-optimized for the CPU. We estimate that the implementation of explicit finite-difference and filtering schemes, employed in the GPU code, into the CPU solver would result in a performance improvement of only about 25%. Even in such a case, the node-to-node speedup factor would still be about 60 $\times$ . Basing the comparison on the number of GPUs versus the number of CPU cores would lead to an equivalent speedup factor of  $60 \times 40/6 = 400\times$ , meaning that one GPU is worth 400 CPU cores. We also note that one dual-socket Intel Skylake CPU, which contains 40 cores total, is comparable in price and power consumption to a V100 GPU. If we were to base the comparison on one V100 GPU versus one dual-socket Intel Skylake CPU, the corresponding speedup factor would be  $400/40 = 10\times$ .

The CPU code also has a second-order implicit time integration version. It is about 2.5 times more costly per time step than the explicit integration scheme. The explicit time integration scheme is normally run at a CFL number of around 0.8. Because of time accuracy concerns, the second-order implicit time integration scheme should not normally be run at a CFL number greater than 5 or so. This corresponds to a factor of about 6.25 increase in the time step with the implicit scheme. The increased computational cost per time step of the implicit scheme gives a speedup factor of about  $6.25/2.5 = 2.5$  over the explicit scheme. Thus, comparing the GPU code to the implicit version of CPU code *as is*, we obtain a node-to-node speedup factor of  $75/2.5 = 30\times$ . The corresponding comparison based on one V100 GPU versus one dual-socket Intel Skylake CPU would yield a speedup factor of  $30/6 = 5\times$ .

As noted earlier, it is possible to further accelerate the GPU code by switching to a second-order, single-stage explicit scheme developed by Verstappen and Veldman [7]. The second-order scheme runs at half the time step of the third-order Runge-Kutta scheme but requires only one right-hand-side computation per time step as opposed to three right-hand-side evaluations needed for the Runge-Kutta scheme. It therefore provides an acceleration factor of  $3/2 = 1.5\times$  over the Runge-Kutta scheme. With the second-order explicit time advancement scheme implemented in the GPU code, the performance comparison to the implicit version of CPU code (which is also second-order accurate in time), based on one V100 GPU versus one dual-socket Intel Skylake CPU, would provide a speedup factor of 7.5 $\times$ .

We should note here that a significant effort was also put into optimizing the implicit time integration version of the CPU code discussed above. The explicit time integration version of the CPU code is essentially the conversion from the optimized implicit version. Hence, the above performance comparisons are between the fastest-running versions of the GPU and CPU codes, and are as fair as currently possible.

We also note that the connection speed between the GPU and the host CPU has some effect on the overall code performance. Because of the slower GPU to CPU connection on the NVIDIA Tesla V100 GPU nodes installed at NAS (which is PCI-Express, while Summit nodes have the faster NVLink\*\*), the GPU flow solver runs about 5% slower on the NAS GPU nodes relative to Summit. The block interface data to be exchanged among the GPUs during the computation are first copied to their respective host CPUs, which then perform the actual data exchange via MPI communication and copy the exchanged data back to the GPUs. Although the code overlaps communication with computation as much as possible, the GPU to CPU connection speed and specific MPI implementation details still play a role on the overall performance.

### C. GPU Code Validation


For code validation, a turbulent channel flow problem was considered. The Reynolds number of the fully-developed turbulent channel flow is  $Re_\tau = \rho_{\text{bulk}} u_\tau h / \mu_{\text{wall}} = 590$ , where  $\rho_{\text{bulk}}$  is the bulk density,  $u_\tau$  is the wall friction velocity,  $h$  is the channel half-height and  $\mu_{\text{wall}}$  is the viscosity on the wall. The domain size is  $2\pi h$  in the streamwise direction,  $x$ ,  $2h$  in the wall-normal direction,  $y$ , and  $\pi h$  in the spanwise direction,  $z$ . The flow is periodic both in the streamwise and spanwise directions and is bounded by solid walls at  $y = 0$  and  $2h$ . Because of the imposed streamwise periodicity, a source term is added to the streamwise momentum and energy equations to drive the flow at a constant mass flow rate. The Mach number based on bulk velocity and sound speed on the wall is set to 0.2 in order to facilitate a comparison with incompressible DNS results available in the literature.

The grid used in the DNS contains  $768 \times 512 \times 768$  points along  $x$ ,  $y$  and  $z$  directions, respectively. The total number of points is about 302 million. The grid resolution in wall units is  $\Delta x^+ \approx 4.8$  in the streamwise direction and  $\Delta z^+ \approx 2.4$  in the spanwise direction. In the vertical direction,  $\Delta y^+ \approx 0.23$  on the wall and  $\Delta y^+ \approx 5.4$  at the channel centerline. For

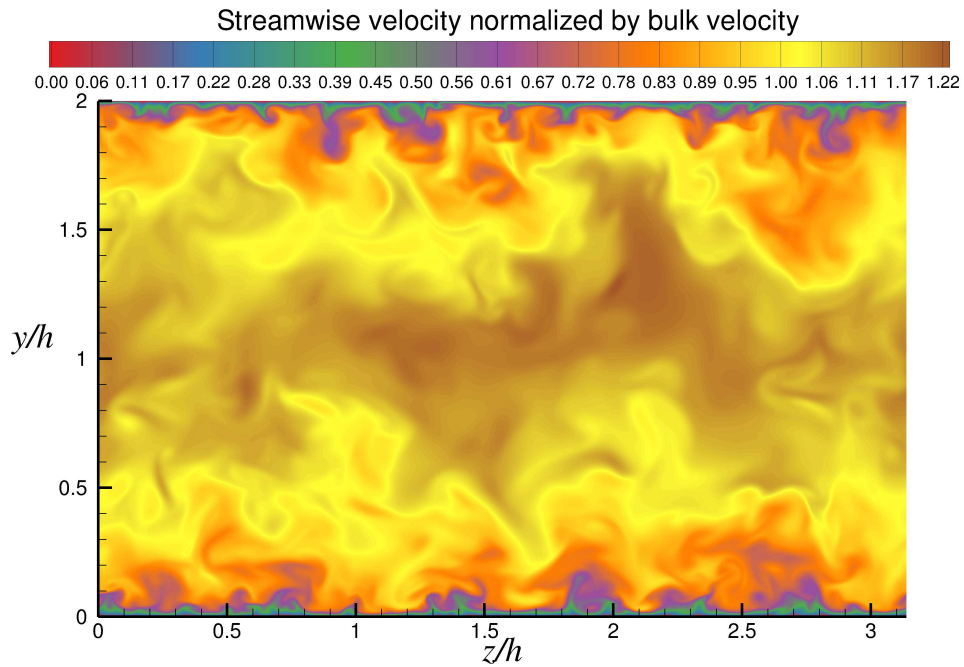
\*\*See <https://www.microway.com/hpc-tech-tips/comparing-nvlink-vs-pci-e-nvidia-tesla-p100-gpus-openpower-servers/> for a comparison between NVLink versus PCI-Express performance.

the computation, 18 NVIDIA Tesla V100 GPUs are used on the Summit system. Each GPU solves a grid block of  $256^3$  points. The second-order time integration scheme [7] is used with a CFL number of 0.4. Explicit filtering is applied at every 5 time steps with a filtering parameter of  $\sigma = 0.09$ . Note that the finer overall grid resolution used here, relative to that in the speed bump case, and the simpler configuration keeps the simulation stable with less frequent filtering. To ensure full convergence of the time-averaged results, the flow statistics are averaged over  $1009h/u_{\text{bulk}}$ .

Figure 31 depicts the instantaneous snapshots of the turbulent channel flow. As seen here, the flow structures appear very smooth and provide evidence that the minimal filtering applied in the present case is able to provide a solution free of any numerical wiggles. Figure 32 compares our mean streamwise velocity and Reynolds stress component profiles with the data from Moser et al. [17] as well as Vreman and Kuerten [18]. Both of these groups used spectral methods to solve the incompressible flow equations for the turbulent channel flow at the same  $Re_\tau$ . As seen in the comparisons, our results show an excellent agreement with their results.




(a) Transverse velocity magnitude contours




(b) Streamwise velocity contours

**Fig. 31 Instantaneous snapshots of turbulent channel flow.**



(a) Mean velocity



(b) Reynolds stress components

Fig. 32 Turbulent channel flow mean streamwise velocity and Reynolds stress component profiles in wall units.



## Acknowledgments

This work was sponsored by the NASA Transformational Tools and Technologies Project of the Transformative Aeronautics Concepts Program under the Aeronautics Research Mission Directorate. The calculations were made possible by the computing resources provided by the NASA High-End Computing Program through the NASA Advanced Supercomputing Division at Ames Research Center. This work also used resources of the Oak Ridge Leadership Computing Facility (OLCF) at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. The access to computational resources was provided under the OLCF Director's Discretion Projects Program. The authors acknowledge valuable discussions with Dr. Philippe Spalart. The RANS solutions were provided by Drs. Michael Strelets and Prahladh Iyer.

## References

- [1] Spalart, P. R., Belyaev, K. V., Garbaruk, A. V., Shur, M. L., Strelets, M. K., and Travin, A. K., "Large-Eddy and Direct Numerical Simulations of the Bachalo-Johnson Flow with Shock-Induced Separation," *Flow, Turbulence and Combustion*, Vol. 99, No. 3–4, 2017, pp. 865–885.
- [2] Slotnick, J. P., "Integrated CFD Validation Experiments for Prediction of Turbulent Separated Flows for Subsonic Transport Aircraft," NATO Science and Technology Organization, Meeting Proceedings RDP, STO-MP-AVT-307, October 2019.
- [3] Bogey, C., and Bailly, C., "A Family of Low Dispersive and Low Dissipative Explicit Schemes for Flow and Noise Computations," *Journal of Computational Physics*, Vol. 194, No. 1, 2004, pp. 194–214.
- [4] Berland, J., Bogey, C., Marsden, O., and Bailly, C., "High-Order, Low Dispersive and Low Dissipative Explicit Schemes for Multiple-Scale and Boundary Problems," *Journal of Computational Physics*, Vol. 224, No. 2, 2007, pp. 637–662.
- [5] Bogey, C., and Bailly, C., "A Shock-Capturing Methodology Based on Adaptive Spatial Filtering for High-Order Non-Linear Computations," *Journal of Computational Physics*, Vol. 228, No. 5, 2009, pp. 1447–1465.
- [6] Gottlieb, S., Shu, C.-W., and Tadmor, E., "Strong Stability-Preserving High-Order Time Discretization Methods," *SIAM Review*, Vol. 43, No. 1, 2001, pp. 89–112.
- [7] Verstappen, R. W. C. P., and Veldman, A. E. P., "Direct Numerical Simulation of Turbulence at Lower Costs," *Journal of Engineering Mathematics*, Vol. 32, No. 2–3, 1997, pp. 143–159.
- [8] Uzun, A., and Malik, M. R., "Large-Eddy Simulation of Flow over a Wall-Mounted Hump with Separation and Reattachment," *AIAA Journal*, Vol. 56, No. 2, 2018, pp. 715–730.
- [9] Uzun, A., and Malik, M. R., "Wall-Resolved Large-Eddy Simulations of Transonic Shock-Induced Flow Separation," *AIAA Journal*, Vol. 57, No. 5, 2019, pp. 1955–1972.
- [10] Spalart, P. R., and Watmuff, J. H., "Experimental and Numerical Study of a Turbulent Boundary Layer with Pressure Gradients," *Journal of Fluid Mechanics*, Vol. 249, 1993, pp. 337–371.
- [11] Badri Narayanan, M. A., and Ramjee, V., "On the Criteria for Reverse Transition in a Two-Dimensional Boundary Layer Flow," *Journal of Fluid Mechanics*, Vol. 35, No. 2, 1969, pp. 225–241.
- [12] Narasimha, R., and Sreenivasan, K. R., "Relaminarization in Highly Accelerated Turbulent Boundary Layers," *Journal of Fluid Mechanics*, Vol. 61, No. 3, 1973, pp. 417–447.
- [13] Spalart, P. R., "Numerical Study of Sink-Flow Boundary Layers," *Journal of Fluid Mechanics*, Vol. 172, 1986, pp. 307–328.
- [14] Patel, V. C., and Head, M. R., "Reversion of Turbulent to Laminar Flow," *Journal of Fluid Mechanics*, Vol. 34, No. 2, 1968, pp. 371–392.
- [15] Warnack, D., and Fernholz, H. H., "The Effects of a Favourable Pressure Gradient and of the Reynolds Number on an Incompressible Axisymmetric Turbulent Boundary Layer. Part 2. The Boundary Layer with Relaminarization," *Journal of Fluid Mechanics*, Vol. 359, 1998, pp. 357–381.
- [16] Muck, K. C., Hoffmann, P. H., and Bradshaw, P., "The Effect of Convex Surface Curvature on Turbulent Boundary Layers," *Journal of Fluid Mechanics*, Vol. 161, 1985, pp. 347–369.
- [17] Moser, R. D., Kim, J., and Mansour, N. N., "Direct Numerical Simulation of Turbulent Channel Flow up to  $Re_\tau = 590$ ," *Physics of Fluids*, Vol. 11, No. 4, 1999, pp. 943–945.
- [18] Vreman, A. W., and Kuerten, J. G. M., "Statistics of Spatial Derivatives of Velocity and Pressure in Turbulent Channel Flow," *Physics of Fluids*, Vol. 26, No. 8, 2014, pp. 085103–1/29.