

Towards Large Eddy Simulation of Boundary Layer Flows at High Reynolds Number: Statistical Modeling of the Inner Layer

by

Johan Larsson

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Mechanical Engineering

Waterloo, Ontario, Canada, 2006

© Johan Larsson 2006

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Most fluid flows of practical interest involve and are affected by turbulence. One of the most promising computational methods for the prediction of turbulent flows is the so-called large eddy simulation (LES) methodology. Experience over the past decades have shown the capability of LES to provide accurate predictions for several types of flow at a reasonable computational cost. It has also become clear, however, that the LES methodology fails when applied to boundary layer flows at high Reynolds numbers. Since many engineering applications fall in exactly that category, this failure is often considered the most severe bottleneck of LES.

The present thesis is an attempt to move towards a solution of this problem. Inspired by the idealized picture of a turbulent boundary layer, a statistical model is used for the approximately universal turbulence in the inner boundary layer, whereas the more flow dependent outer boundary layer is solved by LES. Ideally, this results in a computational method that provides accurate predictions of rather general turbulent flows, while maintaining a tractable computational cost. In practice, the results are a vast improvement compared to LES without any inner layer modeling, but a transition layer appears where the state of the turbulence changes from being modeled statistically to resolved by LES. This so-called ‘artificial buffer layer’ results in the skin friction being consistently underpredicted by 10-15%.

The physics and dynamics of this artificial buffer layer are investigated and characterized, and it is argued that there exist several similarities with true buffer layer turbulence. Additional forcing of the momentum equations is used as a means to trigger resolved turbulence motions more quickly, and it is demonstrated that the results are better: the artificial buffer layer is smaller, the skin friction is accurately predicted, and the dynamics in the inner layer have more correct length scales. The results with the additional forcing are very sensitive to the forcing amplitude, and a simple control algorithm for this parameter is proposed and tested with favourable results.

Acknowledgements

I gratefully acknowledge my co-supervisors, Drs Fue-Sang Lien and Eugene Yee, for their contributions to this work. Special thanks to Dr Lien for encouraging me to publish the work on the multigrid method, and to Dr Yee for being the most accomplished reviewer/editor I have ever met.

I would like to thank Dr Gord Stublely for taking the time to act as a mentor in many ways, including about teaching and academic life in general. I would also like to thank my fellow graduate students in the group, and especially Andrea Scott and Andrew Keats. You have both played big parts in making my time at Waterloo fun and rewarding.

Furthermore, I wish to thank Dr Lars Davidson at Chalmers University of Technology for interesting discussions about the near wall problem of LES, and for inviting me to Chalmers for a short stay in 2004.

During my time at Waterloo I have been supported financially by the Natural Science and Engineering Research Council of Canada and the Ontario Graduate Scholarship Program (both in the form of scholarships), the University of Waterloo (in the form of scholarships and awards), and the Department of National Defence (in the form of a research assistantship).

Finally, my deepest thanks to my wife Lara, for your love, support, and ability to expand my horizons. Can't wait for our next adventure together!

Nomenclature

Roman Symbols

a	matrix coefficient for the Poisson equation
A	forcing amplitude
B	log-law additive constant, 5.2
J	$\sqrt{-1}$
k	wavenumber, $2\pi/\lambda$
L	integral length scale
N_i	number of cells in direction i
p	pressure
Re	Reynolds number
S_{ij}	rate-of-strain tensor
u_i	velocity (resolved velocity from chapter 3)
u_τ	friction velocity, $\sqrt{\nu \partial \bar{u} / \partial y _w}$
y_{int}	location of LES/RANS interface
y_f	location of maximum forcing

Greek Symbols

γ	normalized mean velocity gradient, $y^+ \partial \bar{u}^+ / \partial y^+$
δ	channel half width
Δ	filter width or grid size
ΔU^+	size of the velocity shift
ε	dissipation of kinetic energy
ϵ	error
η_k	Kolmogorov length scale
κ	von Karman constant, 0.41
λ	wavelength
μ	smoothing factor
ν	viscosity
$\rho_{ij}(l_x, l_z)$	two-point correlation
τ_{ij}	residual stress tensor
ϕ	generic function
ω_i	vorticity

Subscripts

i, j, k	tensor indices
I, J, K	grid indices in the x , y , and z directions
w	quantity at a solid wall

Other Symbols

$\bar{\cdot}$	average
\cdot'	fluctuation around the average
$\langle \cdot \rangle$	plane-average over the wall parallel directions
$\hat{\cdot}$	test filter or Fourier transform
$\tilde{\cdot}$	filter or discrete approximation
\cdot^+	quantity normalized by the viscous scales ν and u_τ
\sim	'similar to', or 'scales with'

Abbreviations

AR	cell aspect ratio
CFL	Courant-Friedrichs-Lewy number
DES	detached eddy simulation
DNS	direct numerical simulation
FFT	fast Fourier transform
LES	large eddy simulation
MPI	message passing interface
RANS	Reynolds averaged Navier-Stokes

Contents

Abstract	iii
Acknowledgements	iv
Nomenclature	v
1 Introduction	1
1.1 Turbulence and Modeling Approaches	2
1.2 Boundary Layer Turbulence and LES	5
1.3 Motivation and Objectives	7
1.4 Fully Developed Plane Channel Flow	8
1.5 Outline	9
2 Mathematical Formulation	10
2.1 Large Eddy Simulation	10
2.1.1 Energy Transfer	12
2.1.2 Closure Model	12
2.1.3 Evolution Equations for LES and RANS	15
2.2 Shear Stress Balance in Plane Channel Flow	15
3 Numerical Method I: Discretization	18
3.1 Spatial Discretization	18
3.2 Temporal Discretization	20
3.2.1 Stability	21
3.2.2 Fractional Step Treatment	23
3.3 Boundary and Initial Conditions	24
3.4 FFT Solver for the Poisson Equation	25
3.5 Validation	26
4 Numerical Method II: Multigrid Solver	28
4.1 Introduction and Review of Multigrid Algorithms	28
4.1.1 Objectives	31
4.2 Multigrid Algorithm	31
4.2.1 Conditional Semicoarsening	32
4.2.2 Structured Topology	34

4.2.3	An Odd Number of Cells	35
4.2.4	Intergrid and Coarse Grid Operators	37
4.3	Numerical Experiments	38
4.3.1	Choice of Parameters	38
4.3.2	Cost Scaling	40
4.3.3	Parallel Efficiency	42
4.3.4	Complex Grids	44
4.4	Summary and Conclusions	45
5	Wall Resolved LES	47
5.1	Introduction and Review of Plane Channel Flow	47
5.2	Methodology	50
5.2.1	Initial Condition and Transition	51
5.3	Single-Point Statistics	53
5.4	Multi-Point Statistics	58
5.5	Turbulence Structure	61
5.6	Summary	63
6	Hybrid LES/RANS	64
6.1	Introduction and Review of Approximate Near Wall Modeling	64
6.1.1	Wall Stress Models that use Algebraic Relations	65
6.1.2	Wall Stress Models that use Auxiliary Grids	66
6.1.3	Models that Adapt the Residual Stress Model	67
6.1.4	Wall Stress Models based on Control Theory	68
6.1.5	Objectives	69
6.2	Methodology	69
6.2.1	Resolution Requirements	71
6.3	Results Without Forcing	74
6.3.1	The Velocity Shift	74
6.4	Results With Forcing	77
6.4.1	Forcing Database	78
6.4.2	Single-Point Statistics	80
6.4.3	Multi-Point Statistics	83
6.4.4	Turbulence Structure	86
6.4.5	Effective Forcing Length Scales	87
6.4.6	Amplitude Sensitivity	89
6.4.7	A Control Algorithm for the Amplitude	91
6.5	Summary and Conclusions	97
7	Summary	99
7.1	Future Directions	100
	Bibliography	102

List of Tables

4.1	Complex grid cases. Details of the near wall treatments are given in chapters 5 (Channel C) and 6 (remaining cases).	44
5.1	Wall resolved LES cases	51
6.1	Amplitude control algorithm cases. The left column indicates which parameter is varied, while the remaining parameters are fixed. Note that some cases are repeated for clarity.	94

List of Figures

1.1	Model spectra of homogeneous and isotropic turbulence. The spectra are normalized, premultiplied, and plotted versus the normalized large scale wavelength. The vertical lines and the arrows show the 99th percentiles of the energy and dissipation, respectively. Thus, 99% of each quantity is contained in the region marked by each arrow. —: Energy spectrum function $\hat{E}/(u'^2L)$; --: Dissipation spectrum function $\hat{\epsilon}/(\epsilon\eta_k)$ at two Reynolds numbers (higher to the left).	3
1.2	Regions in turbulent channel flow at high Reynolds number, broadly following Wosnik <i>et al.</i> (2000).	5
1.3	Geometry for fully developed plane channel flow.	8
3.1	Grid and variable arrangement. The pressure, eddy viscosity, and the velocity components that satisfy the momentum equations are stored at the cell centers (filled circles). The convecting velocities (uppercase U) are stored at the cell faces. The I and J refer to the cell and face indices: for example, $(I + 1/2, J)$ is the face between cells (I, J) and $(I + 1, J)$	19
3.2	Difference between the numerical and analytical solutions for decaying vortices. \circ : error defined as $\epsilon = \ u - u_{\text{analytical}}\ $; —: second order slope; --: third order slope; -·-: fourth order slope.	27
4.1	Example of the conditional semicoarsening, and the effect of requiring a structured topology.	34
4.2	Sequence of grids created by $\mu_{\text{lim}} = 0.85$. Labels above the grids indicate planes being left uncoarsened in the x -direction by each algorithm (a due to coefficient anisotropy, o due to an odd number of cells).	36
4.3	Number of cells at each grid level for a unit cube with 96^3 cells and $\text{AR}_{\text{max}} = 100$ on the finest grid. \circ : $\mu_{\text{lim}} = 0.72$; \square : $\mu_{\text{lim}} = 0.78$; \diamond : $\mu_{\text{lim}} = 0.87$; \triangle : $\mu_{\text{lim}} = 0.96$; $+$: $\mu_{\text{lim}} = 1.00$; --: slope for $(2 : 1)^3$ coarsening ratio.	37
4.4	Convergence history of the residual norm $\ r\ $. $+$: $\text{AR} = 1$, $\omega = 1.15$; \circ : $\text{AR}_{\text{max}} = 10$, $\omega = 1.15$, $\mu_{\text{lim}} = 0.71$; \square : $\text{AR}_{\text{max}} = 10$, $\omega = 1.15$, $\mu_{\text{lim}} = 0.87$; \diamond : $\text{AR}_{\text{max}} = 10$, $\omega = 1.5$, $\mu_{\text{lim}} = 1$	39
4.5	Computational cost and convergence factor for grid with $\text{AR}_{\text{max}} = 10$	39
4.6	Contours of \mathcal{T} scaled by the lowest value for each grid.	41

4.7	Computational cost versus the number of grid points. \circ : AR = 1; +: AR _{max} = 100; -: equation (4.12), AR = 1; --: equation (4.12), AR _{max} = 100.	42
4.8	Parallel efficiency η for different grids. Open markers for AR = 1, filled markers for AR _{max} = 100.	43
4.9	Time-averaged streamlines around the wall mounted cube. The flow is from left to right.	44
4.10	Computational cost versus the number of cells. Open markers use $\mu_{\text{lim}} = 0.92$, $\omega = 1.45$. Filled markers use the best combination for each case. \square : channel A; \diamond : channel B; \circ : channel C; \triangle : wall mounted cube; -: equation (4.12), AR = 1; --: equation (4.12), AR _{max} = 100.	45
5.1	Time histories after initializing to a laminar profile with random white noise added on grid fine- y . -: with model; --: without model.	52
5.2	Time history after initializing to a laminar profile with random white noise added on grid fine- y with a model. Quantities are averaged over all grid cells without proper volume weighting. -: C from the dynamic procedure divided by the equivalent value from the standard Smagorinsky model with van Driest damping (using $C = 0.13^2$); --: ν_{les}/ν	52
5.3	Mean velocity. Open markers with model, filled markers without model. \circ : base grid; \square : fine- xz grid; +: DNS at $Re_\tau \approx 550$ by del Alamo <i>et al.</i> (2004); --: $\bar{u}^+ = y^+$ and $\bar{u}^+ = \ln(y^+)/0.41 + 5.2$, respectively.	53
5.4	Resolved normal stresses. Open markers with model, filled markers without model. From top to bottom $\overline{u'u'}$, $\overline{w'w'}$, $\overline{v'v'}$ (offset for clarity). \circ : base grid; \square : fine- xz grid; +: DNS at $Re_\tau \approx 550$ by del Alamo <i>et al.</i> (2004).	54
5.5	Average eddy viscosity and dynamic Smagorinsky coefficient. \circ : base grid; \square : fine- xz grid; \diamond : fine- y grid.	55
5.6	Magnitude of the vorticity fluctuations. Open markers with model, filled markers without model. \circ : base grid; \square : fine- xz grid; +: DNS at $Re_\tau \approx 550$ by del Alamo <i>et al.</i> (2004).	56
5.7	Resolved ($-\overline{u'v'}^+$, upper curves) and modeled ($\overline{\tau}_{12}^+$, lower curves) shear stresses. Open markers with model, filled markers without model. \circ : base grid; \square : fine- xz grid; +: DNS at $Re_\tau \approx 550$ by del Alamo <i>et al.</i> (2004).	57
5.8	Contours of the premultiplied one dimensional spectra $k\hat{E}(k)$. 15 contours between 0 and the common maximum, with the axes the same columnwise. The left column are the streamwise spectra, the right column are the spanwise. Top row: base grid with model; second row: fine- xz grid with model; third row: base grid without model; bottom row: fine- xz grid without model.	59
5.9	Snapshots of the streamwise velocity fluctuation u' at $y^+ \approx 12$. Contours one standard deviation (i.e., u_{rms} , taken from the base case) above and below zero, where darker regions imply higher velocity.	62

5.10	Snapshot of the streamwise vorticity ω_x^+ at $x^+ \approx 2360$ for base grid with model. Contours spaced 0.1 units apart, with solid and dashed lines representing positive and negative vorticity, respectively.	63
6.1	Sketch of grids for the most common approximate near wall treatments. In wall stress models, the evolution equations are solved on the left grid. The wall shear stress is either estimated from some algebraic relations or by solving the thin boundary layer equations on an auxiliary grid (the middle grid). In both cases the flow of information is indicated by the arrows. For hybrid LES/RANS, the evolution equations are solved on the right grid, and the turbulence model switches to a RANS model near the wall.	65
6.2	Blended eddy-viscosity for LES/RANS simulations. \times : $\bar{\nu}_{les}$ from the Smagorinsky model; $+$: $\bar{\nu}_{rans}$ from the mixing length model; \circ : blended $\bar{\nu}_t$ from (6.5); $-\cdot-$: location of the interface $y_{int}^+ \approx 120$. The width parameter $\Delta_{int}^+ \approx 240$	71
6.3	Influence of the streamwise grid resolution. \circ : $N_x = 24$; \diamond : $N_x = 32$; \square : $N_x = 48$; \times : $N_x = 64$; ∇ : $N_x = 96$; \triangle : $N_x = 128$	72
6.4	Influence of the spanwise grid resolution. \circ : $N_z = 16$; \diamond : $N_z = 24$; \square : $N_z = 32$; \times : $N_z = 48$; ∇ : $N_z = 64$	73
6.5	Mean velocity and resolved normal stresses. \circ : LES/RANS without forcing; \diamond : LES/RANS with forcing; $-\cdot-$: locations of the interface ($y_{int}/\delta = 0.06$) and that of maximum forcing ($y_f/\delta = 0.094$), respectively.	75
6.6	A closer look at the velocity shift. \circ : LES/RANS without forcing; \diamond : LES/RANS with forcing; $-\cdot-$: boundaries of zones A-D discussed in text.	76
6.7	Correlation coefficient and turbulence production. \circ : LES/RANS without forcing; \diamond : LES/RANS with forcing; $+$: DNS at $Re_\tau \approx 950$ by del Alamo <i>et al.</i> (2004); $-\cdot-$: locations of the interface ($y_{int}^+ \approx 120$) and that of maximum forcing ($y_f^+ \approx 190$), respectively.	81
6.8	Magnitude of the vorticity fluctuations. \circ : LES/RANS without forcing; \diamond : LES/RANS with forcing; $-\cdot-$: locations of the interface ($y_{int}^+ \approx 120$) and that of maximum forcing ($y_f^+ \approx 190$), respectively.	83
6.9	Contours of the premultiplied one dimensional spectra $k\hat{E}(k)$. 20 contours between 0 and each individual maximum, with the axes the same columnwise. The left column are the streamwise spectra, the right column are the spanwise. Top row: wall resolved LES at $Re_\tau \approx 500$ (base grid of chapter 5); second row: LES/RANS without forcing; bottom row: LES/RANS with forcing; $-\cdot-$: locations of the interface ($y_{int}/\delta = 0.06$) and that of maximum forcing ($y_f/\delta = 0.094$), respectively.	84
6.10	Contours of the two-point correlation of the streamwise velocity $\rho_{11}(l_x, l_z)$. Three contours at (from left to right) 0.9, 0.6, 0.3 of the maximum at each y location. $-$: LES/RANS without forcing; $-\cdot-$: LES/RANS with forcing; $+$: resolved LES at $Re_\tau \approx 500$ (base grid of chapter 5); $-\cdot-$: locations of the interface ($y_{int}/\delta = 0.06$) and that of maximum forcing ($y_f/\delta = 0.094$), respectively.	85

6.11	Snapshots of the streamwise velocity fluctuation u' . Contours one standard deviation (i.e., u_{rms} , taken from the case without forcing) above and below zero, where darker regions imply higher velocity. The left column are without forcing, the right with forcing. Top row: $y^+ \approx 230$; second row: $y^+ \approx 120$; bottom row: $y^+ \approx 12$	87
6.12	Size of the velocity shift for forcing fields with different length scale contents, as defined by the spectral cut-off wavelength λ^c . \square : only modes with $\lambda < \lambda^c$; \times : only modes with $\lambda > \lambda^c$; $---$: full forcing field (i.e., with all resolvable modes); $- \cdot -$: without forcing.	88
6.13	Mean velocity and shear stress for different amplitudes. \circ : $A = 0$, i.e. without forcing; \diamond : $A = A_{\text{opt}} = 0.335$; \square : $A = 0.668$; $- \cdot -$: locations of the interface ($y_{\text{int}}/\delta = 0.06$) and that of maximum forcing ($y_f/\delta = 0.094$), respectively.	89
6.14	Size of the velocity shift as a function of the forcing amplitude.	90
6.15	Time histories for the control algorithm with different parameters, offset for clarity, where the dashed lines mark zero for each curve. $y_{\text{ref}}^+ = y_1^+ = 62$, $y_h^+ = 360$. From bottom to top: $A = A_{\text{opt}} = 0.335$, i.e. fixed; $(K_P, K_I) = (0, 0.3)$; $(K_P, K_I) = (1, 1)$; $(K_P, K_I) = (0, 3)$; $(K_P, K_I) = (30, 1)$	95
6.16	Mean velocity profiles with a logarithmic slope subtracted, i.e. $\alpha(y^+)$ as defined by (6.26). \circ : $A = 0$, i.e. without forcing; \diamond : $A = A_{\text{opt}} = 0.335$; \square : $A = 0.563$; $- \cdot -$: locations of the interface ($y_{\text{int}}^+ \approx 120$) and that of maximum forcing ($y_f^+ \approx 190$), respectively.	96

Chapter 1

Introduction

Most fluid flows of practical interest involve and are affected by turbulence. This includes the flow of air over the wing of an airplane, the cooling of the engine in an automobile, or the process of mixing different fluids together, to name just a few. From an engineering point of view, the effect of turbulence is different in these cases. It increases, respectively, the drag of the airplane, the rate of heat transfer for the cooling system, and the rate of mixing in the final example. The engineer, then, wants to reduce the effect or presence of turbulence in some situations, while increasing it in others. The underlying point, however, is that knowledge of the properties of the turbulence is necessary for successful engineering in all of the situations mentioned.

Such engineering knowledge of turbulence generally comes from three different areas: turbulence theory, physical experiments, and numerical simulations. These tools have fruitfully interacted over many decades, and it is safe to say that all three are needed for proper engineering design of systems involving turbulence. Each tool has its own advantages. The use of theory is both the cheapest and most reliable when applicable, and, more importantly, it offers insights into the important physics of a flow. Physical experiments allow for the evaluation of complex systems, such as flows involving acoustic resonance or chemical reactions. Finally, numerical simulations typically provide very detailed information, since the full flow field can be accessed for analysis. There are, of course, disadvantages with each as well. For example, most theory on turbulence is only strictly applicable to very idealized situations, while both experiments and numerical simulations are approximations to the flow being considered. One such area of approximation is related to the boundary conditions. To study the flow around an airplane, for example, both wind tunnel experiments and numerical simulations introduce approximate free-stream boundary conditions: solid walls at a finite distance and some mathematical approximation, respectively.

The main advances in the field of turbulence were arguably accomplished through theoretical and experimental studies in the first half of the 20th century. Examples that come to mind¹ include the concept of the boundary layer by Prandtl; the linear stability theories by, among others, Rayleigh, Orr, Sommerfeld, and Taylor; the beginnings of a

¹Goldstein (1969) provides an interesting recollection of the advances in fluid mechanics during this time.

theory for turbulence by Taylor and Kolmogorov; and the logarithmic law for the mean velocity profile in a turbulent boundary layer by Prandtl, von Karman, and Millikan.

With the emergence of computers it became possible to compute solutions for turbulent flows numerically. This technique emerged in the 1960s, and has been growing rapidly ever since. Different solution strategies use different levels of approximations, leading to vastly different uses of computational resources. One strategy is to solve for only the statistics of the flow. With this approach, the computational requirements are relatively small, and hence very complex flows can be computed. This approach has been used in industrial applications for around two decades, and accounts for the vast majority of the use of simulations in industry today (c.f. Larsson, 2001; Hanjalic, 2005). A second strategy is to compute the evolution, rather than merely the statistics, of the flow, which provides information about the dynamics of the turbulence. Due to the much higher computational requirements, this strategy has found very limited use in industry, but has been essential in fundamental studies of turbulence (c.f. Moin & Mahesh, 1998).

With the exponential increase in computer power over the last decades, we are soon entering an era when the second strategy for turbulence simulations is becoming feasible in engineering applications. There are several advantages to making use of this approach. In theory, the turbulence modeling introduces much less empiricism, and it makes possible the study of inherently unsteady phenomena such as fluid-structure interaction, the generation of sound, and the control of turbulence. In a very general sense, the objective of this thesis is to take yet another (small) step towards such simulations.

1.1 Turbulence and Modeling Approaches

The difficulties involved in computational modeling of turbulence, and the differences between different modeling strategies, are most easily appreciated when considering an idealized form of turbulence (isotropic and homogeneous) in spectral space. Leaving details aside, the energy spectrum function \hat{E} and the viscous dissipation spectrum function $\hat{\varepsilon}$ describe the contribution from eddies of size (wavelength) λ to the kinetic energy and the rate of dissipation thereof, respectively.

Model spectra for both spectrum functions are shown in figure 1.1. The multi-scale character of turbulence implies that different dynamics or quantities scale with different parameters. The large and energetic eddies scale with the velocity u' (which is of the same order as the square root of the kinetic energy) and the integral length scale L (which is similar to a characteristic dimension of the flow). Since most of the energy is contained in these large scales, plots of \hat{E} at different Reynolds numbers Re collapse (approximately) when normalized by u' and L and plotted versus λ/L . The small and dissipative eddies scale with the dissipation rate ε and the Kolmogorov length scale $\eta_k \equiv (\nu^3/\varepsilon)^{1/4}$, where ν is the viscosity. Most of the dissipation of kinetic energy occurs at the small scales, and hence $\hat{\varepsilon}$ at different Re collapse when normalized by these quantities and plotted versus λ/η_k .

To bring out the effect of increasing Reynolds number, the spectrum functions in




Figure 1.1: Model spectra of homogeneous and isotropic turbulence. The spectra are normalized, premultiplied, and plotted versus the normalized large scale wavelength. The vertical lines and the arrows show the 99th percentiles of the energy and dissipation, respectively. Thus, 99% of each quantity is contained in the region marked by each arrow. —: Energy spectrum function $\hat{E}/(u'^2L)$; ---: Dissipation spectrum function $\hat{\varepsilon}/(\varepsilon\eta_k)$ at two Reynolds numbers (higher to the left).

figure 1.1 are both plotted versus the large scale wavelength λ/L . This implies that only the energy spectra collapse (or remain essentially invariant for different Re), while the dissipation spectra are ‘translated’ to smaller wavelengths as the Reynolds number increases. Note that the spectra are plotted in their premultiplied forms: leaving details aside, the areas under the graphs equal the kinetic energy and the dissipation rate, respectively². Included in the figure are the 99th percentile wavelengths for each quantity.

Kolmogorov’s hypothesis (c.f. Pope, 2000) states that only the energetic scales are strongly affected by the initial and boundary conditions (i.e., dependent on the particular flow), whereas the scales with negligible energy are assumed to be approximately universal (i.e., similar in every flow). From figure 1.1, the scales with $\lambda/L \lesssim 0.002$ have negligible energy. Thus, according to Kolmogorov’s hypothesis, those scales are approximately universal.

For the higher Reynolds number in the figure, the dissipation is negligible for $\lambda/L \gtrsim 0.0004$. Thus, the flow dependent scales are inviscid and independent of Re , and the dissipative scales are universal. The scales in the inertial subrange ($0.0004 \lesssim \lambda/L \lesssim 0.002$) are, accordingly, both inviscid and universal.

For the lower Reynolds number, however, the situation is different. The energetic and dissipative subranges partly overlap, and there is no longer an inertial subrange separating them. Therefore, the flow dependent scales are not inviscid, and, hence, their dynamics are dependent on Re . Similarly, the dissipative scales are not universal.

²Spectra and their premultiplied forms are discussed briefly in section 5.4.

With this (very) brief introductory view of turbulence, the differences between the most common modeling strategies become more clear. The most straightforward is direct numerical simulation (DNS), in which all scales of motion are resolved. With no explicit modeling involved, only errors due to the numerical scheme and the imposed initial and boundary conditions are present. Since these errors can be controlled to some degree, a properly performed DNS represents the turbulence very accurately. The computational cost, however, is heavily dependent on the Reynolds number, since the ratio between the energetic and dissipative length scales increases with Re (as can be seen in figure 1.1). The ratio of these length scales can be shown to be $L/\eta_k \sim Re^{3/4}$ (c.f. Tennekes & Lumley, 1972), and hence the number of grid points in three dimensions is $\sim Re^{9/4}$. If the numerical time step is assumed to be proportional to the cell size, then the total computational cost is $\sim Re^3$. Note that different numerical methods may yield different constants of proportionality in this relation, but the main point is that the cost increases rapidly with Re . Also, at high Re most of the computational effort is devoted to the motions in the inertial and dissipative subranges, i.e. the approximately universal scales.

Due to this high and Reynolds number dependent computational cost, the only turbulent flows computed by DNS to date have been canonical (‘building block’) flows at low Re . For example, the very recent DNS of a fully developed plane channel flow by Hoyas & Jimenez (2005) was performed at $Re_b \approx 40000$ (based on the bulk velocity U_b and the channel half width δ). This simulation used approximately 10^{10} grid points, and required 2048 processors for 4 months. To put this Reynolds number in perspective, the theory by Wosnik *et al.* (2000) suggests that $Re_b \gtrsim 70000$ is required for a fully developed plane channel flow to begin being representative of truly high Reynolds numbers.

A second turbulence modeling strategy is based on the Reynolds averaged Navier-Stokes (RANS) equations. Here, the evolution equations are averaged, and the effects of all scales of turbulence motions must be modeled. The cost becomes independent of Re (for the idealized turbulence of figure 1.1 – for certain flows and RANS models the cost could be weakly dependent on Re), and much smaller than for DNS. Since the effects of the flow dependent scales are modeled, however, it is impossible to construct universally applicable RANS models. This has not prevented RANS from becoming the by far most common modeling approach used in industry to date, since experience has shown certain models to give accurate results on certain flows, especially so when they have been ‘tuned’ for the flow in question.

The third modeling strategy, and the topic of this thesis, is large eddy simulation (LES). The underlying concept of LES is a direct application of figure 1.1: the energetic, flow dependent scales are computed directly (without modeling), while the universal scales are modeled. This implies that LES, in theory, can provide accurate results using a universal model. Also, and as importantly, the computational cost becomes independent of the Reynolds number, since the directly computed (or resolved) scales are inviscid and, therefore, independent of Re . Studies over the last few decades have largely shown LES to live up to these promises for free shear flows, such as jets, mixing layers, and wakes. For boundary layer flows, however, the results have been much less encouraging.




Figure 1.2: Regions in turbulent channel flow at high Reynolds number, broadly following Wosnik *et al.* (2000).

1.2 Boundary Layer Turbulence and LES

A sketch of a boundary layer is shown in figure 1.2, where δ is the boundary layer thickness or the channel half width. In the near wall region, the relevant length and velocity scales for the turbulence are the distance from the wall y and the friction velocity u_τ , respectively (u_τ will be defined in section 2.2). With the viscous length scale ν/u_τ one can form $y^+ \equiv yu_\tau/\nu$, which is essentially a local Reynolds number (c.f. Pope, 2000). The effects of viscosity are small for $y^+ \gtrsim 30$, which defines the essentially inviscid ‘outer layer’. The effects of the outer flow are small for $y/\delta \lesssim 0.1$ (c.f. Wosnik *et al.*, 2000), which defines the ‘inner layer’³.

If the inner and outer layers overlap (as they do in figure 1.2), then the flow dependent region ($y/\delta \gtrsim 0.1$) is inviscid and the motions in the inner layer are essentially universal (following a similar line of reasoning as for the idealized turbulence in the previous section). Also, the size of the energetic eddies scales with the wall distance y , while the dissipative eddies scale with the Kolmogorov scale η_k . The Kolmogorov scale grows very slowly with the wall distance (as $\sim y^{1/4}$, c.f. Pope, 2000), and, thus, the ratio of these length scales depends strongly on y (as $\sim y^{3/4}$). If the model spectrum of figure 1.1 is assumed to be representative of the near wall motions (which it is not), and the length scale separation is $\sim y^{3/4}$, then the dissipative and energetic scales of motion will approach each other in size as the wall is approached. Near the wall, then, there is no longer a separation in length scales – in a sense, the energetic, dynamically important scales are no longer ‘large’, and the interpretation of LES is no longer unambiguous.

The perhaps most common approach in LES of boundary layers is to require the direct resolution of all energetic scales. This is called ‘wall resolved LES’ in this thesis.

³Note that the limits of the inner and outer layers used here are approximate and somewhat arbitrary – other values have appeared in the literature.

Due to the lack of separation between the energetic and dissipative length scales, this implies that some viscous scales must be resolved as well, and hence the computational cost is no longer independent of Re . To estimate the cost, the approach by Chapman (1979) is followed here for the case of fully developed plane channel flow.

In the inner boundary layer the solution scales with the viscous length scale ν/u_τ , which means that the necessary grid resolution in viscous units Δ^+ is constant, where Δ is the cell size. For a computational domain size that is constant in outer units (i.e., scaled by δ), this leads to $\sim \delta/\Delta = \delta^+/\Delta^+ \sim \delta^+$ cells in the directions parallel to the wall. In the direction normal to the wall, a constant number of cells is needed up to some value of y^+ in the overlap layer. Thus, the required number of cells in the inner boundary layer is

$$N_{\text{inner}} \sim \delta^{+2} \sim Re_b^{1.75}, \quad (1.1)$$

where the empirical relation (5.2) by Dean (1978) has been used. Including the temporal dimension leads to a computational cost that scales as $\sim Re_b^{2.3}$, again by simply assuming that the time step scales with the cell size, and thus that the number of time steps required is $\sim N_{\text{inner}}^{1/3}$. This implies that LES that resolves the inner boundary layer is almost as expensive as DNS, and hence some authors use the term ‘quasi-DNS’ for wall resolved LES.

In the inviscid outer boundary layer the solution scales with the outer length scale δ , and hence the number of cells required is independent of the Reynolds number. Note that this result is unique to channel flows – in a spatially developing boundary layer, Chapman (1979) found that N_{outer} depends weakly on Re . The main point, however, is that the computational cost of LES is essentially independent of Re for the outer layer.

The observation that the computational cost of the outer boundary layer is essentially independent of Re suggests a second approach in LES of boundary layers: to only compute directly the motions in the outer layer, while resorting to modeling in the inner. Conceptually, this approach nicely follows the general ideas of LES. The inviscid, large scale motions are resolved, whereas the viscous, small scale motions are modeled. In the literature, this is often referred to as LES with ‘near wall modeling’ or ‘approximate near wall treatment’.

As a side note, the term ‘large eddy simulation’ in a sense suggests that the near wall region, with its small scales of motion, should be modeled. Simulations of type quasi-DNS, which resolve not only the large scales but also the dynamically important small scales near the wall, should perhaps be referred to as ‘dynamically important eddy simulations’ instead. In this thesis, however, the more common terms are used.

Assuming that the dynamics in the inner layer are more or less universal and that only the outer layer is strongly flow dependent, LES with approximate near wall treatment retains the favourable characteristics of LES applied to free shear flows. Since the flow dependent motions are resolved directly, the method is theoretically capable of providing accurate results with simple, universal models (for the small scale and near wall turbulence).

1.3 Motivation and Objectives

This promise of high accuracy at a constant (albeit high) cost has prompted an increasing amount of research over the past decade into LES with approximate near wall treatment. For example, Piomelli (1999) stated in a relatively recent review that ‘Perhaps the most urgent challenge that needs to be met, in order to apply LES to technologically relevant flows, is the modeling of the wall layer’. Similarly, Jimenez & Moser (2000) stated that ‘This problem of wall-boundary conditions... continues therefore to be the main roadblock to the practical application of LES’.

A review of the literature on approximate near wall treatments in LES will be presented in chapter 6, and will show that there are several outstanding issues. One example is the appearance of an ‘artificial buffer layer’ around the modeling interface in hybrid LES/RANS (one particular type of approximate near wall treatment). This has been seen in many studies using different numerical methods, turbulence models, etc, and hence this artificial buffer layer is a robust artefact of the near wall modeling. One effect of this is a consistent underprediction of the skin friction by 10-15% (Nikitin *et al.*, 2000; Piomelli *et al.*, 2003) in attached boundary layers. To see the engineering importance of such errors, consider the fact that the drag of a modern airplane or ship is largely due to the skin friction (Kim, 2003), and hence errors in its prediction have direct impacts on subsequent design and development decisions. Consider also the fact that flow separation induced by an adverse pressure gradient (which, for example, occurs around a wing at high angle of attack, or in the flow around modern automobiles) is highly dependent on the skin friction of the upstream boundary layer, and hence such errors will affect the location and strength of separated flow zones directly.

The main objective of this thesis is to contribute to the understanding and to improve on the modeling of the inner boundary layer in LES of boundary layer flows. More specifically, the present work has resulted in the following contributions:

1. An in-depth study into the physics and dynamics of the artificial buffer layer in hybrid LES/RANS, which has contributed to an increased understanding of these phenomena.
2. A thorough evaluation, on a fundamental level, of the effects of forcing in reducing the size and impact of the artificial buffer layer in LES/RANS, which will guide future developments of generally applicable forcing models.
3. The development of a simple yet effective algorithm for determination of the most important parameter in forcing models (the amplitude) based on control theory.

Furthermore, while not directly related to approximate near wall treatments in LES, the present work has also resulted in a contribution towards the efficient solution of the evolution equations in LES through:

4. The development of a conditional semicoarsening multigrid algorithm for the pressure Poisson equation that is highly parallel, applicable in complex geometries, and that has a sound theoretical basis.




Figure 1.3: Geometry for fully developed plane channel flow.

1.4 Fully Developed Plane Channel Flow

The test case chosen for the work in this thesis is the fully developed plane channel flow, i.e. the flow between two infinitely large plates. A sketch of the flow geometry is shown in figure 1.3, where x , y , and z are the coordinates in the streamwise, wall normal, and spanwise directions, respectively. The velocity components in these directions are, respectively, u , v , and w . The channel half width is δ , and the flow is driven by a pressure gradient in the streamwise direction.

The plane channel flow is chosen for several reasons. First, it is the perhaps simplest and most idealized boundary layer flow imaginable. The flow is statistically homogeneous in the streamwise, spanwise, and temporal directions. This decreases the computational cost and, equally importantly, allows for some important analytical relations that aid in the interpretation of the results. Second, the flow is perhaps the most thoroughly studied canonical flow out there, and there is a wealth of knowledge in the literature. Third, despite the geometrical simplicity, channel flow is a very sensitive test case for LES.

The suitability of channel flow both as a test case for LES and as a canonical flow for studies of boundary layer dynamics is evidenced by the rich literature on it. The earliest attempts at LES of wall bounded flows by Deardorff (1970) and Schumann (1975) used channel flow as a test case, as did the first wall resolved LES by Moin & Kim (1982). As computers evolved, several DNS of channel flows have appeared, at ever increasing Re . The number of numerical studies of channel flow is (almost) equaled by the number of experimental ones, with ever increasing detail in the measurements obtained. A review is presented in chapter 5.

1.5 Outline

In the view of the author, a thesis serves mainly two purposes: it should present the work done in a self-contained and complete fashion, and it should highlight the contributions to the scientific body of knowledge. Since these two purposes are not necessarily compatible, a somewhat ‘asymmetric’ treatment is given here. The novel parts of the work are dealt with in an in-depth and rigorous manner, while other areas are treated more briefly. For example, the numerical code is rather standard within the field of LES, but one component (the multigrid solver for the Poisson equation) is novel – the amount of space devoted to these topics reflect their novelty (or lack thereof) rather than their relative importance for the present study (or, for that matter, the time spent developing them).

Chapter 2 will present the mathematical foundations for LES, including the turbulence model used here. The major (and conventional) part of the numerical method will be presented in chapter 3, along with a validation test case. The multigrid solver for the Poisson equation for pressure will be presented separately in chapter 4, along with some validation and performance tests.

The remaining chapters present results from numerical simulations. In chapter 5 results from wall resolved LES are presented. The purpose of these is partly to validate the numerical method, but primarily to develop a reference database for later use. The main contributions of this thesis are presented in chapter 6, in which the hybrid LES/RANS methodology is developed and applied to high Reynolds number channel flows. The results illustrate some outstanding problems with the method, and the use of additional forcing as a means to remove these problems is investigated. Finally, in chapter 7 the findings will be summarized, along with recommendations for future work.

A thorough literature survey is an essential component of any thesis. Due to the relatively broad spectrum of this thesis, surveys of the relevant literature will be presented within each chapter.

Chapter 2

Mathematical Formulation

This thesis deals exclusively with the incompressible flow of Newtonian fluids, and in Cartesian geometries. Within these restrictions, the Navier-Stokes equations governing the conservation of mass and momentum can be expressed in Cartesian tensor notation as (c.f. Pope, 2000)

$$\frac{\partial u_j}{\partial x_j} = 0 \quad (2.1)$$

and

$$\frac{\partial u_i}{\partial t} + \frac{\partial u_i u_j}{\partial x_j} = -\frac{\partial p}{\partial x_i} + \nu \frac{\partial^2 u_i}{\partial x_j^2}, \quad (2.2)$$

where p is the pressure divided by the constant density. These equations, and their solutions, are formally defined in an infinite dimensional (function) space. For numerical simulations, the equations are represented discretely in a finite dimensional phase space. If solved directly, i.e. in a DNS, the full range of length scales from the integral scale L down to the Kolmogorov scale η_k must be resolved. This implies that the dimension of the discrete system, or equivalently the number of degrees of freedom¹, is a function of the ratio L/η_k , and hence the computational cost is both large and increases drastically with the Reynolds number as discussed in the previous chapter.

2.1 Large Eddy Simulation

The main concept of LES is to reduce the computational complexity (i.e., the number of degrees of freedom, or the dimension of the system) by resolving only the flow dependent, large scales. This is typically accomplished by applying a low-pass filter to the Navier-Stokes equations that removes the smaller, supposedly universal, scales. One common class of filters is defined by the use of convolutions as

$$\tilde{\phi}(x, y, z, t) = \int G_{\tilde{\Delta}}(x - x', y - y', z - z', t - t') \phi(x', y', z', t') dx' dy' dz' dt', \quad (2.3)$$

¹The degrees of freedom here could, for example, be the values of the four variables u_i and p at every computational grid point and time step. In this case, the number of degrees of freedom is $4N_{xyz}N_t$, where N_{xyz} is the total number of grid points and N_t is the total number of time steps. As a rough estimate, the computational cost can be assumed to be proportional to the number of degrees of freedom.

where $G_{\tilde{\Delta}}(x, y, z, t)$ is a filter kernel with filter width $\tilde{\Delta}$. Note that the filtering operation defined by (2.3) commutes with differentiation since $G_{\tilde{\Delta}}$ is only a function of $x - x'$ etc (c.f. Geurts, 2004). This is achieved for the channel flows studied here by only filtering in the wall parallel directions, and hence the filter kernel is of the form

$$G_{\tilde{\Delta}}(x, y, z, t) = \delta_D(y)\delta_D(t)G_{\tilde{\Delta},xz}(x, z) , \quad (2.4)$$

where δ_D refers to the Dirac delta function (c.f. Debnath & Mikusinski, 1999).

Before continuing with the development of the LES equations, one should reflect on this reduction of the number of degrees of freedom through filtering. Such reductions are commonplace in physics: it is, for example, similar to the application of the continuum hypothesis² when deriving the Navier-Stokes equations in the first place (c.f. Tennekes & Lumley, 1972). The second observation is that the most severe reduction in the number of degrees of freedom can be achieved by averaging (instead of filtering). In this sense, the RANS and LES equations are derived in exactly the same fashion, the only difference being the precise choice of filtering operation (or kernel). For example, the kernel corresponding to averaging over the three statistically homogeneous directions is

$$G(x, y, z, t) = \frac{\delta_D(y)}{L_x L_z T} , \quad (2.5)$$

where L_x and L_z are the streamwise and spanwise domain sizes, and T is the total time of the simulation. For now, the derivation will continue using ‘LES terminology’.

Application of the filter (2.3) to the Navier-Stokes equations (2.1) and (2.2) yields

$$\frac{\partial \tilde{u}_j}{\partial x_j} = 0 \quad (2.6)$$

and, when re-arranging in traditional fashion,

$$\frac{\partial \tilde{u}_i}{\partial t} + \frac{\partial \tilde{u}_i \tilde{u}_j}{\partial x_j} = -\frac{\partial \tilde{p}}{\partial x_i} + \nu \frac{\partial^2 \tilde{u}_i}{\partial x_j^2} - \frac{\partial \tau_{ij}}{\partial x_j} , \quad (2.7)$$

where

$$\tau_{ij} = \widetilde{u_i u_j} - \tilde{u}_i \tilde{u}_j \quad (2.8)$$

is the residual (subgrid) stress tensor. This residual stress term arises due to the lack of commutation between multiplication and filtering, and can not be expressed in terms of the filtered (or resolved) velocity field \tilde{u}_i – and, hence, it must be modeled. Note also that the derivation made no use of the exact form of the filter kernel $G_{\tilde{\Delta}}$, only the property that it commutes with differentiation. The residual stress tensor τ_{ij} , however, depends on the filter chosen.

²The continuum hypothesis is essentially a filtering over a small volume. This filtering reduces the complexity by removing information about individual molecules and only accounting for their effect on average (through a molecular viscosity).

2.1.1 Energy Transfer

While equations (2.7) clearly show that the residual stress τ_{ij} influences the momentum equations, it also affects the transfer of energy between the resolved and unresolved scales. Multiplying (2.7) by \tilde{u}_i , and defining the kinetic energy of the resolved scales as $\tilde{q}^2 \equiv \tilde{u}_i \tilde{u}_i / 2$, yields (after minor manipulations) the equation for the kinetic energy of the resolved field

$$\frac{\partial \tilde{q}^2}{\partial t} + \frac{\partial \tilde{q}^2 \tilde{u}_j}{\partial x_j} = \frac{\partial}{\partial x_j} \left(-\tilde{p} \tilde{u}_j + \nu \frac{\partial \tilde{q}^2}{\partial x_j} - \tau_{ij} \tilde{u}_i \right) - \nu \frac{\partial \tilde{u}_i}{\partial x_j} \frac{\partial \tilde{u}_i}{\partial x_j} + \tau_{ij} \frac{\partial \tilde{u}_i}{\partial x_j}. \quad (2.9)$$

The first term on the right hand side is in divergence form, and thus it conserves \tilde{q}^2 on a global level³. The second term on the right hand side is the viscous dissipation of resolved scale energy. This term is typically very small in LES since the low-pass filtering removes the motions in the dissipative range (see figure 1.1 and the discussion in the previous chapter). The last term contains interactions between the residual stress and the resolved velocity field, and thus describes the transfer of energy between the two. In contrast to the viscous dissipation, the term can be both negative (transfer of energy from the resolved scales, or ‘forward scatter’) and positive (transfer of energy to the resolved scales, or ‘backscatter’) On average it is negative (c.f. Sagaut, 2002), but experience shows that backscatter can occur locally in space and time. Despite the fact that it does not always decrease the resolved scale energy, the term is commonly referred to as the ‘subgrid scale dissipation’ ε_{sgs} . Using the fact that τ_{ij} is a symmetric tensor, it can be written as

$$\varepsilon_{\text{sgs}} \equiv -\tau_{ij} \frac{\partial \tilde{u}_i}{\partial x_j} = -\tau_{ij} \tilde{S}_{ij}, \quad \tilde{S}_{ij} \equiv \frac{1}{2} \left(\frac{\partial \tilde{u}_i}{\partial x_j} + \frac{\partial \tilde{u}_j}{\partial x_i} \right), \quad (2.10)$$

where \tilde{S}_{ij} is the resolved rate-of-strain tensor. To conclude, it is now clear that the residual stress tensor τ_{ij} affects the transfers of both momentum and resolved scale kinetic energy. Models for τ_{ij} should ideally predict both effects accurately, but experience shows that this rarely is the case.

2.1.2 Closure Model

While there are many possibilities for the modeling of τ_{ij} , the historically most common is to assume a constitutive relation of form

$$\tau_{ij} - \frac{\tau_{kk}}{3} \delta_{ij} = -2\nu_{\text{les}} \tilde{S}_{ij}, \quad (2.11)$$

where ν_{les} is the eddy viscosity. The resolved rate-of-strain tensor \tilde{S}_{ij} has zero trace, and hence only the deviatoric (or anisotropic) part of the residual stress tensor can be modeled. The isotropic part $(\tau_{kk}/3)\delta_{ij}$ is absorbed into a modified pressure in (2.7).

³To see this, integrate over the whole domain and make use of the divergence theorem. The only contribution will then come from the boundaries. For periodic boundaries (in the wall parallel directions for the channel flow) the contributions cancel out, while at the solid boundaries $\tilde{u}_i = 0$. The net effect is then neither creation nor destruction of \tilde{q}^2 , but merely transport within the domain.

The constitutive relation (2.11) assumes proportionality between τ_{ij} and \tilde{S}_{ij} , i.e. that their principal axes are aligned. This implies a perfect correlation between the stress and the strain, whereas *a priori* studies of DNS data show the true correlation to be rather low (see for example the review by Meneveau & Katz, 2000). Similarly, the constitutive relation implies $\varepsilon_{\text{sgs}} = 2\nu_{\text{les}}\tilde{S}_{ij}\tilde{S}_{ij} \geq 0$, and hence the possibility of backscatter is excluded (for $\nu_{\text{les}} \geq 0$).

The perhaps most common model for the eddy viscosity ν_{les} in (2.11) is the Smagorinsky model (c.f. Geurts, 2004)

$$\nu_{\text{les}} = C\tilde{\Delta}^2|\tilde{S}|, \quad |\tilde{S}| \equiv \sqrt{2\tilde{S}_{ij}\tilde{S}_{ij}}. \quad (2.12)$$

For isotropic turbulence, the value of the model coefficient C (which is actually the square of the original one) can be determined to be $C \approx 0.029$ by requiring the energy transfer ε_{sgs} to equal the viscous dissipation ε (Pope, 2000). This value has been found to be too large (i.e., it produces too much dissipation from the resolved scales) in other flows, such as in the presence of mean shear, near solid boundaries, or in transitional flows (Piomelli, 1999). To see why this is the case, recall that the residual stress term τ_{ij} is the difference between two low-pass filtered quantities. It is then easy to show that the term contains mainly small scale information. The model for τ_{ij} given by (2.11) and (2.12), on the other hand, depends on the filtered rate-of-strain tensor \tilde{S}_{ij} . In a plane channel flow, for example, \tilde{S}_{ij} contains both the mean shear $\partial\bar{u}/\partial y$ and the fluctuating rates-of-strain – and, hence, the modeled τ_{ij} contains more than just small scale information. One *ad hoc* modification is to reduce C in such situations. The simulations in chapter 6, for example, use $C = 0.13^2 = 0.0169$ in the core of the channel, which was found to give more accurate resulting flow fields.

One of the most promising developments of the last decade within the field of LES is the concept of estimating C directly from the resolved velocity field in the so-called ‘dynamic procedure’ by Germano *et al.* (1991). Recall that filtering the Navier-Stokes equations once at filter width $\tilde{\Delta}$ yields an evolution equation for \tilde{u}_i with the residual stress tensor τ_{ij} . Similarly, filtering a second time at accumulated filter width $\hat{\Delta}$ yields an equation for \hat{u}_i with the residual stress tensor

$$T_{ij} = \widehat{\tilde{u}_i\tilde{u}_j} - \hat{\tilde{u}}_i\hat{\tilde{u}}_j. \quad (2.13)$$

The relation between the two stress tensors (commonly called the ‘Germano identity’) can be written as (Germano *et al.*, 1991)

$$\mathcal{L}_{ij} \equiv \widehat{\tilde{u}_i\tilde{u}_j} - \hat{\tilde{u}}_i\hat{\tilde{u}}_j = T_{ij} - \hat{\tau}_{ij}. \quad (2.14)$$

The key point in the dynamic procedure is the recognition that \mathcal{L}_{ij} is expressed entirely using the resolved field \tilde{u}_i , and hence it can be computed directly in LES (by application of a filter). If it is assumed that the residual stress can be modeled by the same underlying model at both filtering levels with the same value of the model coefficient C , then the Germano identity (2.14) can be used to compute C . This assumption, in essence, assumes that there is some form of similarity across different length scales

(c.f. Meneveau & Katz, 2000), which is true if both filter widths are within the inertial subrange. In principle, the procedure is applicable to many models for τ_{ij} . Applying it here for the Smagorinsky model yields (where the first line is a simple restatement of (2.11) and (2.12))

$$\tau_{ij} - \frac{\tau_{kk}}{3}\delta_{ij} = -2C\tilde{\Delta}^2|\tilde{S}|\tilde{S}_{ij} , \quad (2.15)$$

$$T_{ij} - \frac{T_{kk}}{3}\delta_{ij} = -2C\hat{\Delta}^2|\hat{S}|\hat{S}_{ij} . \quad (2.16)$$

Substituting this into the Germano identity (2.14) yields (for the deviatoric part)

$$\mathcal{L}_{ij}^d \equiv \mathcal{L}_{ij} - \frac{\mathcal{L}_{kk}}{3}\delta_{ij} = 2C\widehat{\tilde{\Delta}^2|\tilde{S}|\tilde{S}_{ij}} - 2C\hat{\Delta}^2|\hat{S}|\hat{S}_{ij} . \quad (2.17)$$

In the channel flows considered here, C is taken as constant in the wall parallel directions, i.e. $C = C(y, t)$. With filtering only in the wall parallel directions, and since the filter width $\tilde{\Delta}$ is constant in those directions as well, $C\tilde{\Delta}^2$ can be taken outside of the test filtering which yields

$$\mathcal{L}_{ij}^d - CM_{ij} = 0 , \quad M_{ij} = 2\tilde{\Delta}^2\widehat{|\tilde{S}|\tilde{S}_{ij}} - 2\hat{\Delta}^2|\hat{S}|\hat{S}_{ij} . \quad (2.18)$$

Since the tensors are symmetric with zero trace, (2.18) consists of five independent equations which can not all be satisfied simultaneously. Instead, C can be determined by minimizing the mean square error (Lilly, 1992)

$$\langle \epsilon^2 \rangle = \langle (\mathcal{L}_{ij}^d - CM_{ij})^2 \rangle = \langle \mathcal{L}_{ij}^d \mathcal{L}_{ij}^d \rangle - 2C\langle \mathcal{L}_{ij}^d M_{ij} \rangle + C^2\langle M_{ij} M_{ij} \rangle , \quad (2.19)$$

where $\langle \cdot \rangle$ is the average over the wall parallel directions. This is a parabola in C , and hence the minimum can be found by setting $\partial\langle \epsilon^2 \rangle/\partial C = 0$, which yields

$$C = \frac{\langle \mathcal{L}_{ij}^d M_{ij} \rangle}{\langle M_{ij} M_{ij} \rangle} . \quad (2.20)$$

The filter $\hat{\cdot}$ is taken as a top-hat filter with width $\hat{\Delta} = 2\tilde{\Delta}$ in the present thesis.

The dynamic procedure combined with the Smagorinsky model has been applied to many flows, and shown accurate results in most cases. Since the dynamic calculation of C does not affect the linear constitutive relation, which erroneously assumes perfect correlation between τ_{ij} and \tilde{S}_{ij} , one might wonder why it works as well as it does. Jimenez & Moser (2000) argued that the dynamic procedure finds the value of C that ensures an accurate prediction of ε_{sgs} . The modeled stresses, however, were found to be only about 20% of their true values by Baggett *et al.* (1997). Both these studies concluded that LES with a dynamic Smagorinsky model will give accurate results whenever the modeled stresses are negligible compared to the resolved ones. Baggett *et al.* (1997) argued that this is the case for basic shear flows if the filter width $\tilde{\Delta}$ is one order of magnitude smaller than the integral length scale L .

It is possible to accurately model both the dissipation and the stresses simultaneously if a more complex constitutive relation is used. One example of this is the mixed scale-similarity model by Bardina *et al.* (1980), which has been shown to produce better results in several studies (c.f. Meneveau & Katz, 2000).

Such models are not considered in this thesis, for two reasons. First, the wall resolved (quasi-DNS) simulations in chapter 5 use fine grids since their purpose is to provide accurate data for comparisons. With these fine grids, the modeled stresses are rather small, and hence the dynamic Smagorinsky model is expected to perform well. Second, the main focus of the simulations in chapter 6 is the approximate near wall treatment, and hence additional complexities are kept at a minimum – in fact, not even the dynamic procedure will be used there.

2.1.3 Evolution Equations for LES and RANS

Throughout this thesis, the constitutive relation (2.11) is used to parametrize the residual stress tensor τ_{ij} . The filtered evolution equations can then be written as

$$\frac{\partial u_j}{\partial x_j} = 0 \quad (2.21)$$

and

$$\frac{\partial u_i}{\partial t} + \frac{\partial u_i u_j}{\partial x_j} = -\frac{\partial p}{\partial x_i} + \nu \frac{\partial^2 u_i}{\partial x_j^2} + \frac{\partial 2\nu_t S_{ij}}{\partial x_j}, \quad (2.22)$$

where the $\tilde{\cdot}$ denoting a filtered component has been dropped and the isotropic part of the residual stress $(\tau_{kk}/3)\delta_{ij}$ has been absorbed into the pressure.

The derivation of the evolution equations for RANS exactly follow that for LES, except for a replacement of the filter operation $\tilde{\cdot}$ by an averaging operation $\bar{\cdot}$. The exact definition of this average is not needed for the derivations, but suffice it to say that it does not have to be a temporal average (which is a rather common misconception).

Since (2.22) applies for both LES and RANS, the eddy viscosity has been written in the more generic form ν_t . For the remainder of the thesis, quantities are filtered unless otherwise noted (i.e., the $\tilde{\cdot}$ will be dropped from here on). Resolved quantities are decomposed as $\phi = \bar{\phi} + \phi'$, where $\bar{\phi}$ is the average and ϕ' is the fluctuation around this average.

2.2 Shear Stress Balance in Plane Channel Flow

The fully developed plane channel flow is in many ways an ideal test case for simulations. It has three homogeneous directions, i.e. directions along which statistics do not change, and there are no issues regarding inflow conditions which can be problematic in LES. There are also (uncommon in the field of turbulence) some analytical results for this flow that are useful for interpretations. One of the most useful is the shear stress balance that will be derived here.

The only assumption necessary is that of statistical homogeneity in x , z , and t . This implies that

$$\frac{\partial \bar{\cdot}}{\partial x} = \frac{\partial \bar{\cdot}}{\partial z} = \frac{\partial \bar{\cdot}}{\partial t} = 0 \quad (2.23)$$

for all quantities except the streamwise mean pressure gradient. This mean pressure gradient is necessary to drive the flow, and hence the pressure can be decomposed (without approximation) as

$$p(x, y, z, t) = P(x) + p_{\text{per}}(x, y, z, t), \quad \frac{\partial \bar{p}}{\partial x} = \frac{dP}{dx}, \quad \frac{\partial \bar{p}_{\text{per}}}{\partial x} = 0. \quad (2.24)$$

The continuity equation for the mean flow becomes

$$\frac{\partial \bar{u}_j}{\partial x_j} = \frac{\partial \bar{v}}{\partial y} = 0 \quad (2.25)$$

when using (2.23). Integrating this from the lower wall, and making use of the boundary condition $v = 0$, yields $\bar{v} = 0$ everywhere.

Averaging the streamwise momentum equation (2.7) using the conditions developed so far yields

$$-\frac{\overline{u'v'}}{\partial y} - \frac{dP}{dx} + \nu \frac{\partial^2 \bar{u}}{\partial y^2} - \frac{\partial \overline{\tau_{12}}}{\partial y} = 0. \quad (2.26)$$

Integrating this from the lower wall to y , and making use of the boundary conditions $u = v = 0$, yields

$$-\overline{u'v'} - \frac{dP}{dx}y + \nu \frac{\partial \bar{u}}{\partial y} - \nu \frac{\partial \bar{u}}{\partial y} \Big|_w - \overline{\tau_{12}} + \overline{\tau_{12}}|_w = 0. \quad (2.27)$$

The average wall shear stress is $\nu \partial \bar{u} / \partial y|_w \equiv u_\tau^2$, which defines the friction velocity u_τ . The average residual stress tensor (2.8) is zero at the wall due to the boundary conditions, and hence consistent models for τ_{ij} should satisfy this as well. One then gets

$$-\overline{u'v'} - \frac{dP}{dx}y + \nu \frac{\partial \bar{u}}{\partial y} - u_\tau^2 - \overline{\tau_{12}} = 0. \quad (2.28)$$

With δ denoting the channel half width, setting $y = 2\delta$ and making use of the boundary conditions yields

$$-\frac{dP}{dx} = \frac{u_\tau^2}{\delta}, \quad (2.29)$$

which is simply a balance between the driving force (the pressure gradient) and the retarding force (the wall friction). Substituting this into (2.28) yields

$$-\overline{u'v'} + \nu \frac{\partial \bar{u}}{\partial y} - \overline{\tau_{12}} = u_\tau^2 \left(1 - \frac{y}{\delta}\right), \quad (2.30)$$

which shows the necessary balance between the different shear stresses in fully developed channel flow. Note that this is an exact result (it only depends on the assumption of statistically homogeneous flow), and that it is valid for all turbulence models for τ_{ij} that satisfy the proper boundary condition for the stress tensor.

With the eddy viscosity hypothesis (2.11) the shear stress balance becomes

$$-\overline{u'v'} + \nu \frac{\partial \bar{u}}{\partial y} + \nu_t \left(\frac{\partial u}{\partial y} + \frac{\partial v}{\partial x} \right) = u_\tau^2 \left(1 - \frac{y}{\delta}\right). \quad (2.31)$$

If the fluctuations of ν_t and the strain rate are assumed to be uncorrelated, then

$$-\overline{u'v'} + (\nu + \overline{\nu}_t) \frac{\partial \overline{u}}{\partial y} \approx u_\tau^2 \left(1 - \frac{y}{\delta}\right) \quad (2.32)$$

or when re-arranged

$$\frac{\partial \overline{u}}{\partial y} \approx \frac{u_\tau^2(1 - y/\delta) - (-\overline{u'v'})}{\nu + \overline{\nu}_t} . \quad (2.33)$$

This relation shows an important fact about channel flows: that the mean velocity gradient is determined by the resolved shear stress and the eddy viscosity. The shear stress balances (2.30) and (2.33) will be used to interpret the LES results in chapters 5 and 6.

Chapter 3

Numerical Method I: Discretization

The numerical method used in this study is rather standard in the context of LES. It is described here not for the purpose of any contribution towards improved LES, but rather for the purpose of completeness. Note that the $\tilde{\cdot}$ denoting a resolved component has been dropped for clarity.

The present numerical method is similar to the method developed by Zang *et al.* (1994) in that a colocated grid arrangement is used. Such a grid arrangement can result in ‘pressure checkerboarding’ unless special care is taken – here the approach by Zang *et al.* (1994) is followed to ensure strong coupling between the pressure and velocity fields. The spatially discretized system is integrated in time using a fractional step method to ensure that the continuity constraint is satisfied. The time integration is mixed implicit/explicit and similar to the one developed by Spalart *et al.* (1991).

3.1 Spatial Discretization

The grid and variable arrangement is shown in figure 3.1. Integrating the continuity equation (2.21) over a computational cell D and using the divergence theorem yields (c.f. Ferziger & Peric, 1997)

$$\int_D \frac{\partial u_j}{\partial x_j} dV = \int_{\partial D} u_j n_j dA \approx \sum_{f_{xyz}} U_j n_j A_f = 0, \quad (3.1)$$

where ∂D is the surface of the cell, A_f is the area of the cell face, and f_{xyz} indicates that the sum is to be taken over faces with normals in all three coordinate directions. An uppercase U_j is used for the convecting velocity, i.e. the velocity that satisfies the continuity equation (3.1). As shown in figure 3.1, this velocity is stored at the cell faces. The approximation in the final step above is second order accurate if U_j is considered the velocity at the midpoint of the face.




Figure 3.1: Grid and variable arrangement. The pressure, eddy viscosity, and the velocity components that satisfy the momentum equations are stored at the cell centers (filled circles). The convecting velocities (uppercase U) are stored at the cell faces. The I and J refer to the cell and face indices: for example, $(I + 1/2, J)$ is the face between cells (I, J) and $(I + 1, J)$.

The momentum equation (2.22) similarly yields

$$\int_D \frac{\partial u_i}{\partial t} dV = \int_{\partial D} \left\{ -u_i u_j n_j - p n_i + \nu \frac{\partial u_i}{\partial x_j} n_j + \nu_t \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) n_j \right\} dA \quad (3.2)$$

which can be approximated as

$$V_c \frac{\partial u_i}{\partial t} \approx C(U_j n_j, u_i) + P_i(p) + D_{\text{imp}}(u_i) + D_{\text{exp}}(u_i), \quad (3.3)$$

where V_c is the cell volume and

$$C(U_j n_j, u_i) = \sum_{f_{xz}} -U_j n_j \check{u}_i A_f + \sum_{f_y} -U_j n_j \tilde{u}_i A_f, \quad (3.4)$$

$$P_i(p) = \sum_{f_{xyz}} -\tilde{p} n_i A_f, \quad (3.5)$$

$$D_{\text{imp}}(u_i) = \sum_{f_y} (\nu + \tilde{\nu}_t) \frac{\widetilde{\partial u_i}}{\partial x_j} n_j A_f, \quad (3.6)$$

$$D_{\text{exp}}(u_i) = \sum_{f_{xz}} (\nu + \tilde{\nu}_t) \frac{\widetilde{\partial u_i}}{\partial x_j} n_j A_f + \sum_{f_{xyz}} \tilde{\nu}_t \frac{\widetilde{\partial u_j}}{\partial x_i} n_j A_f. \quad (3.7)$$

Here the diffusion terms have been split into parts that will be treated implicitly and explicitly, respectively. All the quantities above except U_j are defined and stored at the cell centers, and hence their values at the cell faces must be approximated. The quantities denoted by $\tilde{\cdot}$ are discretized using second order accurate approximations (i.e., linear interpolation or central difference), while the momentum velocity \check{u}_i is approximated using a higher order approximation in the wall parallel directions.

Using the notation in figure 3.1, these approximations can be written in the following forms (in two dimensions for clarity) at the face $(I + 1/2, J)$ (and similarly for other faces) using the generic function ϕ :

$$\left. \frac{\partial \phi}{\partial x} \right|_{(I+1/2, J)} = \frac{\phi_{(I+1, J)} - \phi_{(I, J)}}{x_{(I+1)} - x_{(I)}} , \quad (3.8)$$

$$\tilde{\phi}_{(I+1/2, J)} = \phi_{(I, J)} + \left. \frac{\partial \phi}{\partial x} \right|_{(I+1/2, J)} (x_{(I+1/2)} - x_{(I)}) , \quad (3.9)$$

$$\left. \frac{\partial \phi}{\partial y} \right|_{(I+1/2, J)} = \frac{\tilde{\phi}_{(I+1/2, J+1)} - \tilde{\phi}_{(I+1/2, J-1)}}{y_{(J+1)} - y_{(J-1)}} , \quad (3.10)$$

$$\begin{aligned} \check{\phi}_{(I+1/2, J)} = & c_1 (\phi_{(I, J)} + \phi_{(I+1, J)}) + c_2 (\phi_{(I-1, J)} + \phi_{(I+2, J)}) \\ & + c_3 (\phi_{(I-2, J)} + \phi_{(I+3, J)}) . \end{aligned} \quad (3.11)$$

Note that brackets and capital letters have been used to indicate grid indices (as opposed to tensor indices), and that the grid coordinates only depend on a single index due to the Cartesian grid geometry. The first three expressions use second order accurate approximations, whereas the last is the fourth order accurate dispersion relation preserving scheme by Tam & Webb (1993). The coefficients $(c_1, c_2, c_3) = (0.63637324, -0.16289319, 0.02651995)$ of this scheme were chosen not only for a high order of accuracy (as is commonly done), but also for a minimal dispersion error. Note that all discrete approximations are centered, which ensures that the numerical diffusion is minimized (c.f. Geurts, 2004).

3.2 Temporal Discretization

The semi-discrete momentum equation (3.3) is integrated in time using the method by Spalart *et al.* (1991), where the wall normal diffusion is treated in a Crank-Nicolson like fashion and the remaining terms are treated explicitly by a three stage Runge-Kutta scheme. This yields, for substep (or stage) m ,

$$\begin{aligned} V_c \frac{u_i^m - u_i^{m-1}}{\Delta t} = & \alpha_m D_{\text{imp}}(u_i^{m-1}) + \beta_m D_{\text{imp}}(u_i^m) + (\alpha_m + \beta_m) P_i(p^m) \\ & + \gamma_m \left\{ C(U_j^{m-1} n_j, u_i^{m-1}) + D_{\text{exp}}(u_i^{m-1}) \right\} \\ & + \rho_m \left\{ C(U_j^{m-2} n_j, u_i^{m-2}) + D_{\text{exp}}(u_i^{m-2}) \right\} , \quad m = 1, 2, 3 . \end{aligned} \quad (3.12)$$

Note that the pressure is more of a mathematical than physical variable in the incompressible Navier-Stokes equations (see for example Ferziger & Peric, 1997, for a discussion on the role of pressure), and hence it is treated here using a formally first order accurate approximation. The index m denotes the substep, where $m = 0$ and $m = 3$ refer to the solutions at times n and $n + 1$, respectively. The coefficients were determined by Spalart *et al.* (1991) as

$$\begin{array}{lll} \alpha_1 = \frac{29}{96} & \alpha_2 = \frac{-3}{40} & \alpha_3 = \frac{1}{6} \\ \beta_1 = \frac{37}{160} & \beta_2 = \frac{5}{24} & \beta_3 = \frac{1}{6} \\ \gamma_1 = \frac{8}{15} & \gamma_2 = \frac{5}{12} & \gamma_3 = \frac{3}{4} \\ \rho_1 = 0 & \rho_2 = \frac{-17}{60} & \rho_3 = \frac{-5}{12} \end{array}$$

which makes the scheme (nonlinearly) third order accurate for the explicit terms and second order accurate for the implicit ones.

3.2.1 Stability

The implicit treatment of the wall normal diffusion and the pressure makes the temporal integration unconditionally stable for those terms, i.e. it is stable for arbitrarily large time steps Δt . The explicit treatment of the remaining terms, however, imposes stability constraints on the time step. While a full stability analysis of (3.12), including the non-uniformity of the grid, is possible, a simpler analysis in one dimension based on the concept of a modified wavenumber (c.f. Moin, 2001) is presented here. While simplified, this analysis conveys the essential components of the issue of numerical stability. It will be presented here in a somewhat brief form, with details of the algebra suppressed for brevity.

Consider the spatially discrete momentum equation (3.3) on a uniform grid in one dimension (say x) and its exact counterpart. Neglecting the implicitly treated terms $D_{\text{imp}}(u)$ and $P(p)$, and linearizing by assuming U and ν_t to be constant, yields the exact equation

$$\frac{\partial u}{\partial t} = -U \frac{\partial u}{\partial x} + \nu_{\text{tot}} \frac{\partial^2 u}{\partial x^2}, \quad (3.13)$$

and the semi-discrete

$$\frac{\partial u}{\partial t} = \frac{C_{\text{lin}}^{1D}(U, u)}{V_c} + \frac{D_{\text{exp,lin}}^{1D}(u)}{V_c}, \quad (3.14)$$

where $\nu_{\text{tot}} = \nu + \nu_t$. The interpolation formula (3.11) yields

$$\begin{aligned} \frac{C_{\text{lin}}^{1D}(U, u)}{V_c} &= -\frac{UA_f}{V_c} (\check{u}_{(I+1/2)} - \check{u}_{(I-1/2)}) \\ &= -\frac{U}{\Delta x} \{ (c_1 - c_2) (u_{(I+1)} - u_{(I-1)}) + (c_2 - c_3) (u_{(I+2)} - u_{(I-2)}) \\ &\quad + c_3 (u_{(I+3)} - u_{(I-3)}) \}. \end{aligned} \quad (3.15)$$

The second term in $D_{\text{exp}}(u)$ vanishes by the continuity equation for constant ν_t , and hence in a similar way

$$\frac{D_{\text{exp,lin}}^{1D}(u)}{V_c} = \frac{\nu_{\text{tot}}}{\Delta x^2} (u_{(I+1)} - 2u_{(I)} + u_{(I-1)}) . \quad (3.16)$$

Now taking $u = \exp(jkx)$, where $j \equiv \sqrt{-1}$ and k is a wavenumber (i.e., taking u as a Fourier mode), yields for the exact equation

$$\frac{\partial u}{\partial t} = -jkUu - \nu_{\text{tot}}k^2u , \quad (3.17)$$

and for the semi-discrete one

$$\frac{\partial u_{(I)}}{\partial t} = -jk_{\text{c,mod}}Uu_{(I)} - \nu_{\text{tot}}k_{\text{d,mod}}^2u_{(I)} , \quad (3.18)$$

where the modified wavenumbers are

$$k_{\text{c,mod}} = 2 \frac{(c_1 - c_2) \sin(k\Delta x) + (c_2 - c_3) \sin(2k\Delta x) + c_3 \sin(3k\Delta x)}{\Delta x} , \quad (3.19)$$

$$k_{\text{d,mod}}^2 = 2 \frac{1 - \cos(k\Delta x)}{\Delta x^2} . \quad (3.20)$$

The key observation here is that the discrete and the exact forms are similar, except that the modified wavenumbers are used in the discrete form. These modified wavenumbers are functions of the relative grid resolution $k\Delta x$ and of the particular numerical approximation used. The maximum values represent the largest possible ‘amplification’ due to spatial differentiation (actually, the spectral radii of the discretization matrices), and are $\max\{k_{\text{c,mod}}\} \approx 1.73/\Delta x$ for the Tam & Webb (1993) scheme and $\max\{k_{\text{d,mod}}^2\} = 4/\Delta x^2$.

The stability of Runge-Kutta schemes is commonly studied on the model equation

$$\frac{\partial u}{\partial t} = \psi u , \quad (3.21)$$

where ψ is a complex amplification factor. The three-stage scheme used here is stable for $|\text{Im}(\psi\Delta t)| \leq \sqrt{3}$ for purely imaginary ψ and for $-2.5 \leq \text{Re}(\psi\Delta t) \leq 0$ for purely real ψ (c.f. Moin, 2001). Comparing these limits to each term in (3.18) then yields the stability constraints

$$|\text{Im}(-jk_{\text{c,mod}}U\Delta t)| \leq \sqrt{3} \Rightarrow \max\{k_{\text{c,mod}}\}U\Delta t \approx 1.73 \frac{U\Delta t}{\Delta x} \leq \sqrt{3} \quad (3.22)$$

and

$$-2.5 \leq \text{Re}(-\nu_{\text{tot}}k_{\text{d,mod}}^2\Delta t) \leq 0 \Rightarrow \max\{k_{\text{d,mod}}^2\}\nu_{\text{tot}}\Delta t = 4 \frac{\nu_{\text{tot}}\Delta t}{\Delta x^2} \leq 2.5 . \quad (3.23)$$

Note that the convective spectral radius multiplied by the time step traditionally is termed the Courant-Friedrichs-Lewy number, i.e. $\text{CFL} \equiv \max\{k_{\text{c,mod}}\}U\Delta t$.

3.2.2 Fractional Step Treatment

The lack of a temporal derivative in the continuity equation implies that this equation can be seen as a constraint (rather than an evolution equation), which complicates the temporal integration. Put another way, despite the fact that there are 4 equations and 4 unknowns (u_i and p), closer inspection shows that there really only are 2 independent evolution variables. If two velocity components are known, the third can be obtained from the continuity equation, while the pressure can be obtained from a Poisson equation (that will be derived later). To deal with this problem, a fractional step method is used to advance the solution to the next time level in a predictor/corrector fashion. This treatment causes the formal order of accuracy to drop to second order in time (Le & Moin, 1991).

To simplify notation, first define the right hand side of (3.12) as

$$\text{RHS (3.12)} \equiv \beta_m D_{\text{imp}}(u_i^m) + (\alpha_m + \beta_m) P_i(p^m) + R_i . \quad (3.24)$$

A predicted velocity field u_i^* that does not satisfy continuity is first computed by solving

$$V_c \frac{u_i^* - u_i^{m-1}}{\Delta t} = \beta_m D_{\text{imp}}(u_i^*) + (\alpha_m + \beta_m) P_i(p^{m-1}) + R_i . \quad (3.25)$$

The effect of the pressure gradient is removed by

$$V_c \frac{\tilde{u}_i - u_i^*}{\Delta t} = -(\alpha_m + \beta_m) P_i(p^{m-1}) . \quad (3.26)$$

Note that these two predictor steps often are combined in fractional step methods. They are separated here simply for convenience, since the issue of boundary conditions for u_i^* is simplified in this way. The velocity field is finally corrected as

$$V_c \frac{u_i^m - \tilde{u}_i}{\Delta t} = (\alpha_m + \beta_m) P_i(p^m) , \quad (3.27)$$

where the pressure p^m must be determined such that the final velocity field satisfies the continuity equation. Such an equation can be derived (Zang *et al.*, 1994) by first writing (3.27) in the equivalent finite difference form for the convecting velocity at a face (since a coupling between the pressure and the convecting velocity is sought)

$$\frac{U_i^m - \check{U}_i}{\Delta t} = -(\alpha_m + \beta_m) \frac{\partial \widetilde{p^m}}{\partial x_i} . \quad (3.28)$$

Applying the divergence operator $\text{div}(\phi_i) = \sum_f \phi_i n_i A_f$ yields

$$\sum_{f_{xyz}} \frac{U_i^m n_i A_f - \check{U}_i n_i A_f}{\Delta t} = -(\alpha_m + \beta_m) \sum_{f_{xyz}} \frac{\partial \widetilde{p^m}}{\partial x_i} n_i A_f . \quad (3.29)$$

Requiring the final velocity field to satisfy continuity, i.e. applying (3.1), yields the Poisson equation for the pressure

$$-\sum_{f_{xyz}} \frac{\partial \widetilde{p^m}}{\partial x_i} n_i A_f = -\sum_{f_{xyz}} \frac{\check{U}_j n_j A_f}{(\alpha_m + \beta_m) \Delta t} . \quad (3.30)$$

The algorithm is then to compute \check{u}_i , interpolate $\check{U}_j n_j$ from \check{u}_i using the appropriate interpolation schemes as described above, compute the source term of the Poisson equation and solve for the pressure, and then finally correct the velocity fields using (3.27) for u_i^m and (3.28) for U_i^m . The eddy viscosity is computed at the end of each time step, and held constant for all three substeps.

The main computational cost of this algorithm is in solving the Poisson equation (3.30). The predictor step (3.25) treats only the wall normal diffusion term implicitly, which generates a tridiagonal system that can be solved very quickly by a direct solver such as the Thomas algorithm for tridiagonal matrices (Ferziger & Peric, 1997). The Poisson equation, however, is elliptic in character, and hence it is more computationally expensive to solve. In the present code two solvers are used for this equation: a direct solver based on fast Fourier transforms (FFT) which is very efficient for plane channel flows, and a multigrid solver based on conditional semicoarsening that is less efficient (compared to the FFT) but applicable to arbitrary geometries. The latter solver has also been parallelized.

3.3 Boundary and Initial Conditions

The plane channel flows considered in this thesis make use of only two boundary conditions. In the wall parallel directions, the unknowns are assumed to be periodic as

$$\phi(x) = \phi(x + L_x), \quad \forall x, \quad (3.31)$$

for some imposed periodicity length (i.e., domain size) L_x (and similarly in the z -direction).

At the solid walls, the dynamic boundary conditions are $u|_w = w|_w = 0$, while the kinematic boundary condition is $v|_w = 0$. Inserting the dynamic conditions into the continuity equation yields $\partial v / \partial y|_w = 0$ as well.

No boundary condition is needed at the wall for the pressure equation, since (3.28) already defines the pressure gradient at the wall. With zero transpiration velocity, (3.28) yields

$$\left. \frac{\partial \widetilde{p}^m}{\partial x_i} n_i \right|_w = \left. \frac{\check{U}_i n_i}{(\alpha_m + \beta_m) \Delta t} \right|_w. \quad (3.32)$$

Inserting this into the Poisson equation (3.30) for a cell adjacent to the solid wall yields

$$- \sum_{f_{\text{inner}}} \frac{\partial \widetilde{p}^m}{\partial x_i} n_i A_f = - \sum_{f_{\text{inner}}} \frac{\check{U}_j n_j A_f}{(\alpha_m + \beta_m) \Delta t}, \quad (3.33)$$

where f_{inner} denotes inner faces, i.e. that the sum is not to be taken over the boundary faces. After solving for p^m at the cell centers, the pressure at the wall can be extrapolated using (3.32).

Fully developed channel flow is, by definition, statistically independent of the initial condition. Thus, the only purpose of the initial condition (i.e., u_i at $t = 0$) is to generate fully developed turbulence as quickly as possible. In the present study, the initial

conditions are either a laminar flow with random noise added or an instantaneous field from a different simulation, interpolated onto the grid. Some discussion of the initial condition and the subsequent transition to turbulence will be given in section 5.2.1.

3.4 FFT Solver for the Poisson Equation

The discrete Poisson equation (3.30) with the central difference approximation (3.8) can be written in the form (in three dimensions)

$$\begin{aligned} a_{(I,J,K)}p_{(I,J,K)} + a_{(I,J,K)}^{I+}p_{(I+1,J,K)} + a_{(I,J,K)}^{I-}p_{(I-1,J,K)} + a_{(I,J,K)}^{J+}p_{(I,J+1,K)} \\ + a_{(I,J,K)}^{J-}p_{(I,J-1,K)} + a_{(I,J,K)}^{K+}p_{(I,J,K+1)} + a_{(I,J,K)}^{K-}p_{(I,J,K-1)} = q_{(I,J,K)} , \end{aligned} \quad (3.34)$$

where the $\{a_{(I,J,K)}\}$ are the matrix coefficients. These matrix coefficients are completely determined by the geometry of the grid: they are constant for all time steps, and they are uniform in the streamwise and spanwise directions since the grids are uniform in those directions (due to the statistical homogeneity). With this uniformity in the wall parallel directions, the matrix coefficients can be written in the abbreviated forms

$$\begin{aligned} a_{(J)} &\equiv a_{(I,J,K)} , \\ a_{(J)}^I &\equiv a_{(I,J,K)}^{I+} = a_{(I,J,K)}^{I-} , \\ a_{(J)}^K &\equiv a_{(I,J,K)}^{K+} = a_{(I,J,K)}^{K-} , \\ a_{(J)}^{J+} &\equiv a_{(I,J,K)}^{J+} , \\ a_{(J)}^{J-} &\equiv a_{(I,J,K)}^{J-} , \end{aligned} \quad (3.35)$$

and the discrete Poisson equation (3.34) can then be written as

$$\begin{aligned} a_{(J)}p_{(I,J,K)} + a_{(J)}^I (p_{(I+1,J,K)} + p_{(I-1,J,K)}) + a_{(J)}^{J+}p_{(I,J+1,K)} \\ + a_{(J)}^{J-}p_{(I,J-1,K)} + a_{(J)}^K (p_{(I,J,K+1)} + p_{(I,J,K-1)}) = q_{(I,J,K)} . \end{aligned} \quad (3.36)$$

The uniformity of the coefficients can be exploited for efficient solution of (3.36). First, the pressure and the source term are expanded in two dimensional discrete Fourier series as

$$p_{(I,J,K)} = \frac{1}{N_x N_z} \sum_{n,m} \hat{p}_{(n,J,m)} \exp \left\{ j2\pi \left(\frac{(I-1)n}{N_x} + \frac{(K-1)m}{N_z} \right) \right\} , \quad (3.37)$$

where again $j \equiv \sqrt{-1}$ and similarly for $q_{(I,J,K)}$. Inserting this into (3.36) yields

$$\begin{aligned} \frac{1}{N_x N_z} \sum_{n,m} \left[\exp \left\{ j2\pi \left(\frac{(I-1)n}{N_x} + \frac{(K-1)m}{N_z} \right) \right\} \cdot \right. \\ \left. \left\{ b_{(n,J,m)} \hat{p}_{(n,J,m)} + a_{(J)}^{J+} \hat{p}_{(n,J+1,m)} + a_{(J)}^{J-} \hat{p}_{(n,J-1,m)} - \hat{q}_{(n,J,m)} \right\} \right] = 0 , \quad \forall I, J, K , \end{aligned} \quad (3.38)$$

where

$$b_{(n,J,m)} \equiv a_{(J)} + a_{(J)}^I 2 \cos\left(\frac{2\pi n}{N_x}\right) + a_{(J)}^K 2 \cos\left(\frac{2\pi m}{N_z}\right) \quad (3.39)$$

has been defined for simplified notation. Orthogonality of the individual Fourier modes implies that (c.f. Moin, 2001)

$$b_{(n,J,m)} \hat{p}_{(n,J,m)} + a_{(J)}^{J+} \hat{p}_{(n,J+1,m)} + a_{(J)}^{J-} \hat{p}_{(n,J-1,m)} = \hat{q}_{n,J,m} , \quad \forall n, J, m . \quad (3.40)$$

The key observation here is that the transformed Poisson equation (3.40) is a decoupled system of $N_x N_z$ tridiagonal systems, each of which can be solved quickly by the Thomas algorithm. After solution of (3.40), the pressure field is obtained by an inverse Fourier transform.

One should note that $\sum_J \hat{q}_{(0,J,0)} = 0$ must be satisfied for a solution to exist (this is satisfied automatically in the present method), and that the non-uniqueness of the solution is reflected in (3.40) by $b_{(0,N_y,0)} = 0$ after the forward sweep of the Thomas algorithm – this implies that $\hat{p}_{(0,N_y,0)}$ can be set to an arbitrary value.

3.5 Validation

It is important to validate a numerical method on simple benchmark test cases. The code has been developed continuously during the course of the work, and various test cases have been computed after major developments. Shown here is a flow with an analytical solution that illustrates the different orders of accuracy for the convective and viscous terms.

The decaying vortices described by

$$\begin{aligned} u &= -\cos(kx) \sin(ky) \exp\left(-\frac{2k^2 t}{Re}\right) , \\ v &= \sin(kx) \cos(ky) \exp\left(-\frac{2k^2 t}{Re}\right) , \\ p &= -\frac{1}{4} (\cos(2kx) + \cos(2ky)) \exp\left(-\frac{4k^2 t}{Re}\right) \end{aligned} \quad (3.41)$$

is an exact solution to the two dimensional Navier-Stokes equations on a periodic domain with dimensions $2\pi/k$ (Zang *et al.*, 1994). The Reynolds number is defined here as $Re \equiv u_{\max} 2\pi/(k\nu)$, where $u_{\max} = 1$ is the maximum initial velocity.

Two cases are considered here: one viscously dominant with $Re = 10$ and one convectively dominant with $Re = 1000$, both with $k = 1$. Grids with $N_x = N_y \in \{12, 16, 24, 32, 48, 64, 96, 128\}$ (a Bulirsch sequence) cells in each direction are used, and the time step is reduced proportionately to keep the CFL number constant. The equations are integrated to times $T \approx 0.077$ and 1.54 (at which the velocity fields have decayed by about 10% and 25%), respectively, and the error is quantified by the L_2 norm of the difference between the numerical and analytical solutions and shown in figure 3.2.

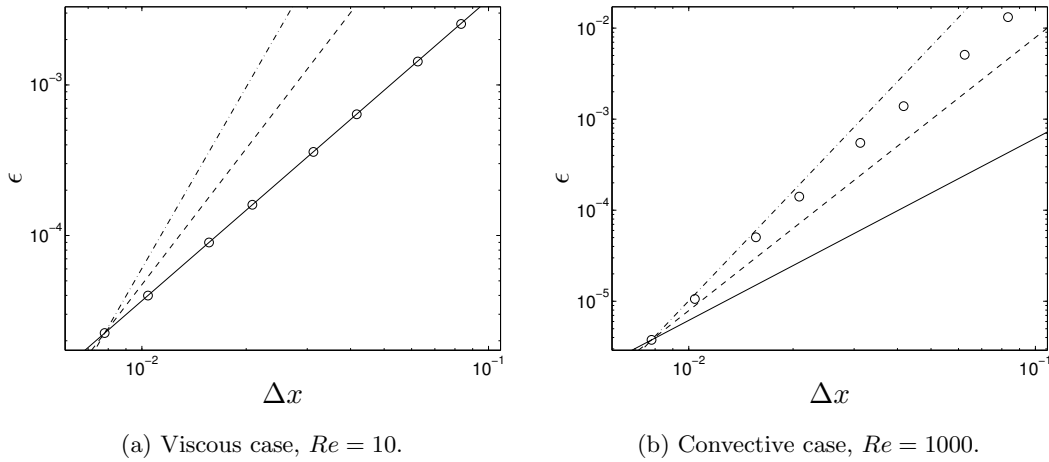


Figure 3.2: Difference between the numerical and analytical solutions for decaying vortices. \circ : error defined as $\epsilon = \|u - u_{\text{analytical}}\|$; $-$: second order slope; $- -$: third order slope; $- \cdot -$: fourth order slope.

The code is clearly second order accurate for the viscous case, which simply reflects the fact that the viscous term is treated by a low order approximation. The high Re case results in an order of accuracy between three and four (about 3.5 from a curve fit). While the convective interpolation scheme is formally of fourth order accuracy, the remaining components of the code (such as the approximation of surface integrals by the midpoint value times the area) are of second order accuracy, and thus the fourth order accuracy is not reached in practice.

The rationale for using a higher order approximation for one term only can of course be questioned. In LES, however, the resolved motions are essentially inviscid, and thus the convective term is dominant for most of these motions. It would be interesting to study the influence of errors from each term individually, but that has not been considered in this work.

Chapter 4

Numerical Method II: Multigrid Solver

The test cases presented in this thesis are all plane channel flows, for which the Poisson equation can be solved very quickly using fast Fourier transforms as discussed in section 3.4. The code, however, was designed for use on more complex flows as well, and hence an alternative solver for the Poisson equation was developed as well. This multigrid solver uses a novel conditional semicoarsening algorithm to construct the coarse grids, and is presented here in some detail. It has previously been published as Larsson *et al.* (2005a) and Larsson *et al.* (2005c).

4.1 Introduction and Review of Multigrid Algorithms

Multigrid techniques are commonly used to accelerate the convergence when solving discrete elliptical problems iteratively. To get a sense for the underlying concept of multigrid algorithms, consider the discrete Poisson equation (3.34), and the following conceptual iterative algorithm to obtain the solution p from the initial guess p^0 :

1. Apply a *smoother* to the equation. This smoother need not reduce the error $\epsilon^0 = p - p^0$ in magnitude, but must yield an error that is smooth.
2. Since the error is now smooth, it can be represented on (or *restricted* to) a coarser grid.
3. Solve for the error ϵ^0 on the coarse grid. Due to the lower number of cells on this grid, the solution can be obtained more quickly.
4. Having found the error ϵ^0 on the coarse grid, interpolate (or *prolongate*) it back onto the original fine grid and update the pressure as $p^1 = p^0 + \epsilon^0$.
5. Smooth the pressure field again to eliminate any high-frequency error that may have resulted from the coarse grid correction process.

This conceptual algorithm involves only two grids, the fine and the coarse. The solution on the coarse grid (step 3), however, is most efficiently obtained by recursive application

of the full algorithm outlined here, and hence a sequence of ever coarser grids is used in multigrid algorithms.

The conceptual algorithm shows the essential idea behind multigrid: that a smooth error can be approximated on a coarser grid, on which the solution can be obtained much more quickly (c.f. Trottenberg *et al.*, 2001). It also implies that multigrid algorithms fail whenever the error is not smooth enough, since in that case the high-frequency components of the error will be aliased on the coarse grid, leading to poor convergence or even divergence of the algorithm. Note that the smoothness of the error is intimately related to the particular coarse grid used. If, for example, the coarse grid is only coarser in one direction, then the error need only be smooth in that direction. In fact, the high-frequency components of the error are typically defined as those that can not be represented on the related coarse grid (Trottenberg *et al.*, 2001).

For discrete problems that are isotropic¹, even simple smoothers (e.g., pointwise Gauss-Seidel) efficiently yield a smooth error in all directions, and hence a grid that is coarser in all directions (full coarsening) can be used. In three dimensions, this yields a coarse grid with 8 times fewer grid points if standard 2:1 coarsening is used². For anisotropic problems, however, the situation becomes more complex. Pointwise smoothers then typically only yield a smooth error in some direction(s). A grid that is coarser in all directions can then not represent the error adequately, and hence the whole underlying concept of multigrid fails. Anisotropic problems arise either due to anisotropy of the underlying physics, or perhaps more commonly due to the use of stretched grids, which makes the discrete problem anisotropic. The latter is the case for the Poisson equation considered here. Note that the coefficient anisotropy in this case is directly related to the cell aspect ratio $AR \equiv \max\{\Delta x, \Delta y, \Delta z\} / \min\{\Delta x, \Delta y, \Delta z\}$.

There are various approaches to ensure good multigrid convergence on anisotropic problems. All these approaches increase the computational cost compared to the optimal algorithm for isotropic problems, which simply is a reflection of the fact that anisotropic problems are more difficult to solve. The most common approaches can be grouped according to which component(s) of the conceptual multigrid algorithm they modify (from the isotropic case):

1. Adaptation of the smoother, while maintaining full coarsening. The increased cost is, naturally, associated with the more elaborate smoother necessary. Typically, block-implicit smoothers (e.g., plane-smoothers) are used, which update a collection of points simultaneously.
2. Adaptation of the coarse grid, while maintaining a pointwise smoother. This is generally called *semicoarsening*, since the coarse grid can only be coarser in some direction(s). The increased cost is associated with the denser coarse grids.
3. Combinations of the two previous approaches.

¹The discrete Poisson equation (3.34) is isotropic if $a_{(I,J,K)}^{I\pm} = a_{(I,J,K)}^{J\pm} = a_{(I,J,K)}^{K\pm}$, and anisotropic if at least one of the matrix coefficients is different from the others.

²2:1 coarsening simply implies that each cell on the coarse grid is related to two cells on the fine grid, in each coordinate direction.

To ensure a smooth error in all directions using a plane-smoother, the planes must be swept in all directions (i.e., in the xy -, xz -, and yz -planes) sequentially. Such alternating plane-smoothers in combination with full coarsening were used by, for example, McBryan *et al.* (1991) and Emvin (1997). The latter study is particularly relevant here, since it focused on a problem very similar to the current one: the solution of the Poisson equation on stretched, structured grids in three dimensional Cartesian geometries. On uniform and isotropic grids, Emvin (1997) found the alternating plane-smoother to be about 3 times less efficient³ than a pointwise Gauss-Seidel smoother (which is close to optimal for that case). On stretched grids, however, the alternating plane-smoother essentially maintained its convergence rate, whereas the pointwise Gauss-Seidel smoother failed. The results by Emvin (1997) will be used for comparisons here, due to the similarity of the problems studied.

Schaffer (1998) used a plane-smoother in two (the x and y) directions, and coarsening only in the third (z) direction. The plane-smoother, in turn, was a two dimensional multigrid algorithm with a line-smoother in the x direction coupled with coarsening only in the y direction. The resulting algorithm was rather robust in terms of yielding fast convergence for different anisotropies. Brown *et al.* (2000) studied the parallel performance of Schaffer (1998)'s algorithm using domain decomposition. When the domain was decomposed in all three directions, the parallel efficiency was very low due to the fact that both the plane- and the line-smoothers crossed domain boundaries. Decomposition only in the y and z directions (where only the plane-smoother crossed domain boundaries) improved the situation somewhat. Only when the domain was decomposed in the z direction only, i.e. such that neither smoother crossed any domain boundaries, was the parallel efficiency reasonably high. This shows one of the drawbacks associated with the use of block-implicit smoothers, since they are, by nature, less amenable to parallelization.

This leads to the final approach to make multigrid perform well for anisotropic problems: to maintain the pointwise smoother, and to adapt the construction of the coarse grid. The key component of such semicoarsening methods is, naturally, the coarse grid construction. Loosely speaking, if the grid is coarsened by 'too much', the lack of error smoothness will prevent good convergence. Conversely, if the grid is coarsened by 'too little', the coarse grids will be unnecessarily dense, and the computational cost will be large.

Horton & Vandewalle (1993) developed a highly parallel algorithm by combining a pointwise smoother with adaptive semicoarsening. They considered uniform but anisotropic grids (i.e., Δx and Δy constant but not equal), and hence the coarse grids automatically had a structured topology due to the uniformity. The adaptive part of the semicoarsening consisted of comparing the matrix coefficients to a predefined limit value, which was taken from *a priori* Fourier analysis.

Elias *et al.* (1997) considered the coarse grid construction within the context of unstructured grids, and hence allowed the coarse grids to have an unstructured topology. They argued that each matrix coefficient can be interpreted as the inverse of a physical

³Convergence efficiency is defined here in terms of the time required to reduce the residual by one order of magnitude.

time scale that determines the rate of propagation of information during the smoothing step. They further argued that the error becomes smooth in directions associated with the smallest time scale (largest coefficient) for each cell, and hence that the size of the coefficients can be used to construct the coarse grid. They then constructed the coarse grid based on a comparison between the matrix coefficients and a predefined limit value.

4.1.1 Objectives

The present work is aimed at developing a highly parallel multigrid algorithm for the Poisson equation on arbitrarily stretched Cartesian structured grids. As shown by Brown *et al.* (2000), while plane-smoothers can efficiently smooth the error in all directions, they are not inherently parallel. For this reason, a pointwise smoother is used here.

The semicoarsening strategy used here is similar to the ones by Horton & Vandewalle (1993) and Elias *et al.* (1997) in that it is based on the size of the matrix coefficients. Here, however, a stronger mathematical coupling between the matrix coefficients and the error smoothness is sought. Within the necessary approximations, this strong coupling implies that the smoothness of the error can be controlled, and hence that the convergence of the algorithm can be ensured.

Further differences between the present work and the work by Elias *et al.* (1997) and Horton & Vandewalle (1993) include the requirement of structured coarse grids and the fact that arbitrarily stretched grids are considered, respectively.

4.2 Multigrid Algorithm

To achieve a highly parallel algorithm, the red-black point Gauss-Seidel (RBPGS) smoother (Trottenberg *et al.*, 2001) is used in the present work. The implementation pays careful attention to the parallel efficiency by minimizing the communication overhead. Each smoothing sweep consists of the following steps, where red points are those where the sum of the indices $I + J + K$ is odd, and black points are those where the sum is even:

1. Smooth the red grid points adjacent to the processor boundaries.
2. Initiate communication of those points.
3. Smooth the red grid points away from the processor boundaries (i.e., in the interior of each processor's domain).
4. Finalize communication of the red grid points.
5. Repeat for the black grid points.

All communication is handled by the message passing interface (MPI) and is of the 'non-blocking' kind (see for example the MPI manual, 2003).

The performance of a smoother is best quantified in terms of the smoothing factor μ (Trottenberg *et al.*, 2001). This is essentially defined as the worst reduction of the high

frequency components of the error, where the high frequency components are those that can not be represented on the coarse grid. With full coarsening, the RBPGS smoother is very effective (i.e., yields low μ) for isotropic problems. As discussed above, however, the smoothing behaviour deteriorates severely on anisotropic problems.

4.2.1 Conditional Semicoarsening

The goal of the conditional semicoarsening can be stated as to, given a fine grid, construct a coarse grid that:

1. is such that the error on the fine grid is sufficiently smooth for all cells and directions that are coarsened, i.e. that the smoothing factor

$$\mu_{(I,J,K)} \leq \mu_{\text{lim}} , \quad \forall I, J, K , \quad (4.1)$$

where μ_{lim} is some maximum allowable value;

2. has as few cells as possible;
3. has a structured topology.

In order to proceed, a relation between the matrix coefficients and the smoothing factor is needed. Such a relation was found by Yavneh (1995) using Fourier analysis on infinite rectangular grids, i.e. grids where the cell sizes are constant but not equal in the different directions. Here Yavneh (1995)'s main result is first presented in its original form for infinite rectangular grids, and then the necessary adaptations for use in the present conditional semicoarsening algorithm will be given.

If D_c is the set of directions that are coarsened⁴, then Yavneh (1995) derived

$$\mu = \max \left\{ (1 - \check{a}^{\min})^2 \vartheta(\check{a}^{\min}, \xi, \#D_c) , \zeta(\xi) \right\} , \quad (4.2)$$

where

$$\check{a}^{\min} \equiv \min_{D \in D_c} \{ \check{a}^D \} , \quad (4.3)$$

$$\vartheta(\check{a}^{\min}, \xi, \#D_c) = \begin{cases} [1 + \check{a}^{\min} / \{2(1 - \check{a}^{\min})\}]^{1/\xi} , & \text{if } \#D_c = 1 \\ 1 & , \text{if } \#D_c > 1 \end{cases} , \quad (4.4)$$

and $\zeta(\xi)$ is a lower bound that is only dominant when a large number of smoothing sweeps are performed. Furthermore, ξ is the total number of smoothing sweeps, $\#D_c$ is the number of elements in D_c , and \check{a}^D is the matrix coefficient in direction D , normalized such that $\sum_D \check{a}^D = 1$.

In order to use formula (4.2) in the present application, some approximations are necessary. First, if Yavneh's analysis is viewed as a *local* Fourier analysis, then the results can be applied to each cell individually in non-uniform grids. The matrix coefficients for opposing cell faces (e.g., $a_{(I,J,K)}^{I+}$ and $a_{(I,J,K)}^{I-}$) are then not exactly equal, but will be

⁴For example, if only the x - and z -directions are coarsened, then $D_c = \{I, K\}$

approximately so on smoothly stretched grids. The normalized matrix coefficients can then be defined as

$$\check{a}_{(I,J,K)}^D = \frac{a_{(I,J,K)}^D}{\sum_D a_{(I,J,K)}^D} = \frac{a_{(I,J,K)}^D}{a_{(I,J,K)}^I + a_{(I,J,K)}^J + a_{(I,J,K)}^K}, \quad D = I, J, K, \quad (4.5)$$

with the (arbitrarily chosen) definition

$$a_{(I,J,K)}^D \equiv \sqrt{\frac{a_{(I,J,K)}^{D+} + a_{(I,J,K)}^{D-}}{2}}, \quad D = I, J, K. \quad (4.6)$$

Second, if the lower bound $\zeta(\xi)$ is neglected, Yavneh's formula (4.2) becomes

$$\mu \approx (1 - \check{a}^{\min})^2 \vartheta(\check{a}^{\min}, \xi, \#D_c). \quad (4.7)$$

This involves ϑ , which depends on $\#D_c$. Since the purpose of the conditional semicoarsening is to determine D_c , ϑ is unknown. However, as shall be shown here, the approximation $\vartheta \approx 1$ is appropriate for the present use.

If $\#D_c > 1$, then $\vartheta = 1$, and the approximation is exact. If $\#D_c = 1$, then $\check{a}^{\min} = \check{a}^{D_c}$ by (4.3). For high values of \check{a}^{D_c} , (4.7) then yields very low predicted smoothing factors regardless of whether ϑ is included or not, and thus the error introduced by using $\vartheta \approx 1$ does not alter the condition on the smoothing factor (4.1). For low values of \check{a}^{D_c} , on the other hand, $\vartheta(\check{a}^{D_c}, \xi, 1) \approx 1$ anyway, and thus the approximation is valid for this case as well. For the values of μ_{lim} and ξ used later in this chapter, for example, $\vartheta \lesssim 1.01$ for those values of \check{a}^{D_c} that generate smoothing factors close to μ_{lim} . Therefore, the approximation $\vartheta \approx 1$ is valid in the present context of conditional semicoarsening.

Neglecting ϑ in (4.7) then yields (for individual cells due to the local view of the Fourier analysis)

$$\mu_{(I,J,K)} \approx \left(1 - \check{a}_{(I,J,K)}^{\min}\right)^2 = \max_{D \in D_c} \left\{ \left(1 - \check{a}_{(I,J,K)}^D\right)^2 \right\}, \quad (4.8)$$

where definition (4.3) has been used.

Using relations (4.1) and (4.8), the first two goals of the conditional semicoarsening will be satisfied with the following algorithm:

Algorithm 1: conditional semicoarsening set-up

```

for all cells  $(I, J, K)$  do
  set  $D_c = \{I, J, K\}$  (initialize to full coarsening)
  compute  $\check{a}_{(I,J,K)}^D$ ,  $D \in \{I, J, K\}$ , by (4.6) and (4.5)
  for directions  $D = I, J, K$  do
    if  $(1 - \check{a}_{(I,J,K)}^D)^2 > \mu_{\text{lim}}$  then
      remove  $D$  from  $D_c$  for cell  $(I, J, K)$ 
    end if
  end do
end do

```

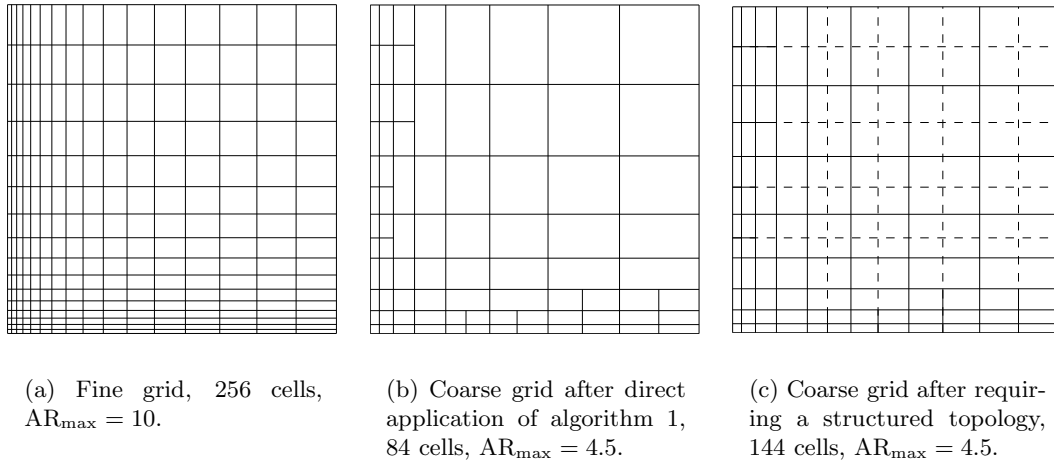


Figure 4.1: Example of the conditional semicoarsening, and the effect of requiring a structured topology.

This algorithm will generate D_c for every cell such that, if the grid is coarsened accordingly, the first two goals are met. Figure 4.1(a) shows a sample fine grid, and figure 4.1(b) shows the resulting coarse grid after direct application of the algorithm above. The cells on the coarse grid are less anisotropic, and the number of cells decreases by a factor of 3. The grid, however, has an unstructured topology, which the current solver can not handle.

At this stage one can compare the present method with the method by Elias *et al.* (1997). Their ‘agglomeration’ method to create the coarse grids is essentially based on physical arguments, while the present method is entirely mathematical. Despite this, both methods create rather similar grids (not shown here) given similar coefficient anisotropies. While not surprising, this shows the interplay between mathematics and physics – or, rather, how the mathematics reflect the underlying physics.

4.2.2 Structured Topology

The requirement of having a structured coarse grid further limits the coarsening. From figure 4.1, it is clear that the structure of the grid is maintained if whole planes of cells (as opposed to individual cells) are considered. For the sample grid of the figure, this implies that the cells adjacent to the dashed grid lines should be left uncoarsened despite satisfying the condition on the smoothing factor.

Before giving the explicit algorithm, first define P_l^D as the plane of cells where the value of index D is l . Also, for the moment, the number of cells in each direction N_D is assumed to be even. In the current work, only 2:1 coarsening is considered, and the coarsening is restricted to merging only cells in the planes P_{l-1}^D and P_l^D , $l = 2, 4, \dots, N_D$. The algorithm for constructing coarse grids that are structured is then:

Algorithm 2: construction of a structured coarse grid

```

for directions  $D = I, J, K$  do
  for  $l = 2, 4, \dots, N_D$  do
    if  $D \in D_c$  for all cells in  $P_{l-1}^D$  and  $P_l^D$  then
      coarsen these cells in direction  $D$ 
    else
      leave these cells uncoarsened in  $D$ 
    end if
  end do
end do

```

The resulting structured coarse grid is shown in figure 4.1(c). The maximum cell aspect ratio is the same as for the unstructured coarse grid, since the same condition on $\mu_{(I,J,K)}$ is enforced. Due to the extra limitations on the coarsening, the coarsening ratio is 1.8 (as opposed to 3 for the unstructured grid). Hence, there is a trade-off between using a computationally cheap structured solver with unnecessarily dense coarse grids, or a more expensive unstructured solver with more ‘optimal’ coarse grids. Only the former option is considered here.

4.2.3 An Odd Number of Cells

Whenever the number of cells in some direction is odd, there are essentially two options available: either to merge three cells into one somewhere, or to leave one plane of cells uncoarsened. In the present work the latter approach is chosen. To avoid leaving more planes than necessary uncoarsened, the following algorithm is used to decide which plane to leave due to an odd number of cells:

Algorithm 3: dealing with odd N_D

```

for directions  $D = I, J, K$  do
  if  $N_D$  is odd then
    for  $l = 1, 3, 5, \dots, N_D$  do
       $M_l = \max_{P_l^D} \left\{ (1 - \check{a}_{(I,J,K)}^D)^2 \right\}$ 
    end do
    leave plane  $P_m^D$  with  $m = \arg \max_l \{M_l\}$  uncoarsened
  end if
end do

```

The complete semicoarsening algorithm starts by dealing with the case of an odd number of cells. From then on, the other algorithms are applied with those planes ignored, i.e. as if N_D was even.

Figure 4.2 shows a sample sequence of coarse grids generated by the conditional semicoarsening algorithm. Note the planes left uncoarsened for different reasons, and

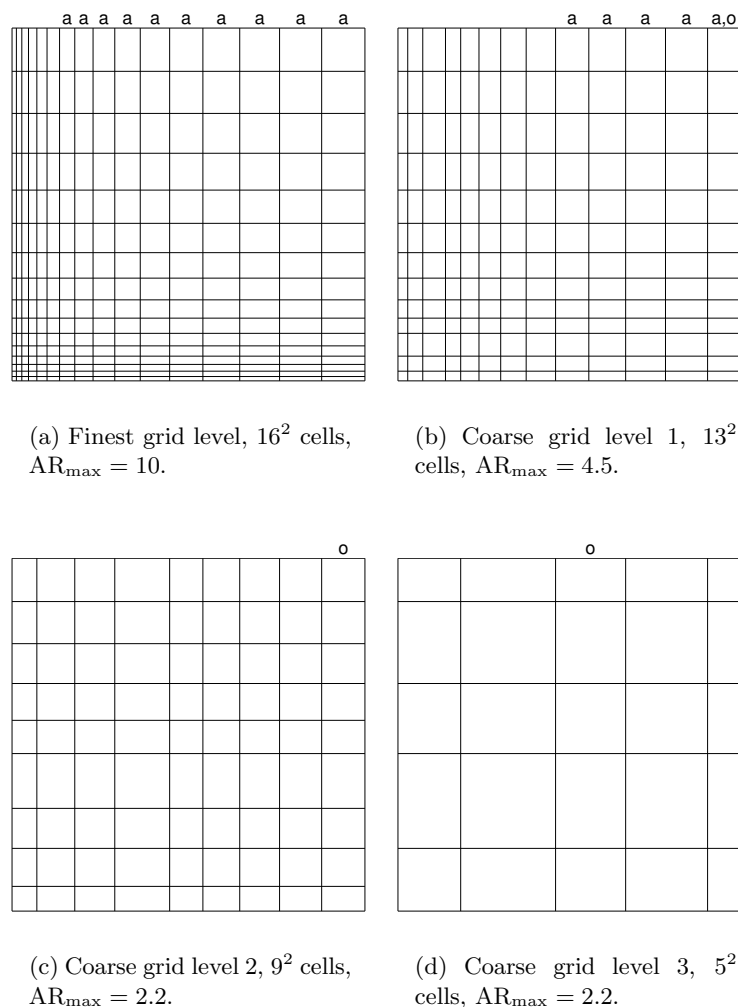


Figure 4.2: Sequence of grids created by $\mu_{\text{lim}} = 0.85$. Labels above the grids indicate planes being left uncoarsened in the x -direction by each algorithm (**a** due to coefficient anisotropy, **o** due to an odd number of cells).

how the coarse grids become more and more isotropic. The grid (and, hence, the coefficient) anisotropy is measured by the maximum cell aspect ratio AR_{\max} .

The only parameter of the algorithm is the value of μ_{lim} . The optimum value of this parameter will be determined by numerical experiments on a three dimensional unit cube in section 4.3.1. Note that $\mu_{\text{lim}} = 1$ corresponds to full coarsening.

The number of cells in the coarse grids is, of course, dependent on μ_{lim} . This is illustrated in figure 4.3, where the number of cells N on each grid level is shown for some values of μ_{lim} . When the parameter is less than 1, the coarsening ratio for the first coarse grids in the sequence is quite low as condition (4.1) severely limits the coarsening process. At some point in the process, the grids are close enough to being isotropic

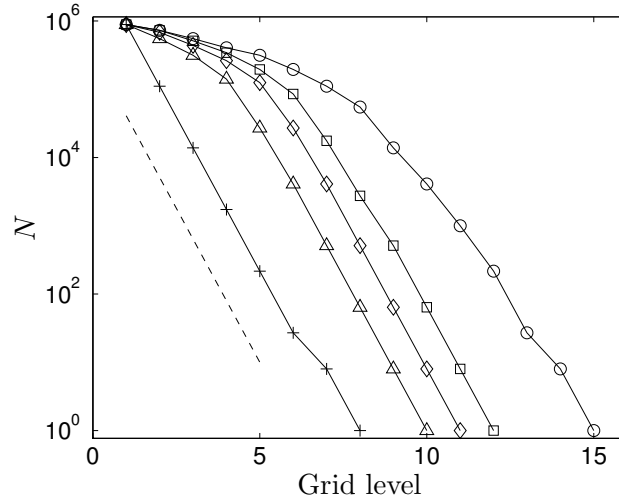


Figure 4.3: Number of cells at each grid level for a unit cube with 96^3 cells and $AR_{\max} = 100$ on the finest grid. \circ : $\mu_{\text{lim}} = 0.72$; \square : $\mu_{\text{lim}} = 0.78$; \diamond : $\mu_{\text{lim}} = 0.87$; \triangle : $\mu_{\text{lim}} = 0.96$; $+$: $\mu_{\text{lim}} = 1.00$; $--$: slope for $(2 : 1)^3$ coarsening ratio.

that the coarsening proceeds almost without limitation, and hence the coarsening ratio approaches the standard $(2 : 1)^3$.

The work involved in finding D_c everywhere is proportional to the number of cells N , and the cost of constructing the coarse grid is similar to that using full coarsening. When used for the pressure equation in fractional step type algorithms, the matrix coefficients do not change during the simulation. Due to this, the coarse grids can be constructed before the simulation, and hence the cost of the set-up phase is completely negligible.

4.2.4 Intergrid and Coarse Grid Operators

For problems with smoothly varying coefficients, standard tri-linear prolongation and coarse grid operators defined by direct discretization on the coarse grids are commonly used (Trottenberg *et al.*, 2001). For problems with coefficient discontinuities, however, operator-dependent prolongation and Galerkin coarse grid operators have been shown to improve the convergence rate (de Zeeuw, 1990). In the present work, coefficient discontinuities occur whenever there is a large difference in cell size between adjacent cells. For purposes of numerical accuracy, it is standard practice in DNS and LES to use grids with smooth stretching, and hence the discontinuities are expected to be small on the finest grid. Due to the fact that the semicoarsening algorithm typically coarsens small cells and leaves larger ones, the difference in cell size between adjacent cells is normally kept within a factor of 2. For these reasons, standard tri-linear prolongation and coarse grid operators defined by direct discretization on the coarser grids are used in this work. These choices also ensure that the coarse grid stencils consist of only 7 points, whereas a Galerkin operator would consist of 27. The parallelization of the smoother is better with a 7-point stencil, since the two-colour RBPGS smoother completely decouples the

grid points. The restriction operator is residual summation, which makes physical sense in the context of a finite volume method (Trottenberg *et al.*, 2001). Only V-cycles are considered in this work.

4.3 Numerical Experiments

Numerical experiments are performed first to determine suitable values of μ_{lim} and the number of smoothing sweeps (ξ_1, ξ_2) (before and after the coarse grid correction, respectively). After that, the performance of the algorithm on some model problems will be evaluated.

Yavneh (1996) found that the smoothing behaviour of the RBPGS smoother on anisotropic problems with full coarsening can be improved by the use of an over-relaxation factor ω . Therefore, ω is included as a parameter in the numerical experiments.

4.3.1 Choice of Parameters

To find suitable values of the parameters, the discretized Poisson equation (3.34) is solved on a unit cube domain (i.e., $x, y, z \in [0, 1]$) for grids with various degrees of stretching. The source term is set to zero and the initial field is randomized. The grid coordinates are given by

$$x^l = \frac{1}{2} \left\{ 1 + \frac{\tanh\{\Gamma(2l/N_x - 1)\}}{\tanh(\Gamma)} \right\}, \quad l = 0, \dots, N_x, \quad (4.9)$$

where x^l is the coordinate of the l th grid line, Γ is a stretching parameter, and the remaining directions are stretched identically. All grids consist of $N = 64^3$ cells. Initial tests showed that the performance is relatively insensitive to the number of smoothing sweeps, but that $\xi \equiv \xi_1 + \xi_2 = 4$ or 5 gives good results for the different grids. In the following, $(\xi_1, \xi_2) = (2, 3)$ is used exclusively.

An example of the convergence history is shown in figure 4.4, where $\|r\|$ is the L_2 -norm of the residual. On the stretched grid, the rate of convergence using full coarsening ($\mu_{\text{lim}} = 1$) is very poor – in fact, the over-relaxation parameter ω needs to be increased to ensure convergence for this case. With the conditional semicoarsening, the rate of convergence is improved considerably.

To measure the performance, some definitions are needed. The convergence factor φ is defined as (Trottenberg *et al.*, 2001)

$$\varphi = \left(\frac{\|r^n\|}{\|r^m\|} \right)^{1/(n-m)}, \quad (4.10)$$

where $\|r^m\|$ is the residual norm after iteration m . From the definition, φ is the average reduction in residual norm between iterations (MG-cycles) m and n . The asymptotic rate of convergence is reached after about 4 iterations (see figure 4.4), and hence $m = 4$ is used throughout this paper, while n is taken as the number of iterations to convergence (typically about 10).




Figure 4.4: Convergence history of the residual norm $\|r\|$. +: AR = 1, $\omega = 1.15$; o: AR_{max} = 10, $\omega = 1.15$, $\mu_{lim} = 0.71$; □: AR_{max} = 10, $\omega = 1.15$, $\mu_{lim} = 0.87$; ◇: AR_{max} = 10, $\omega = 1.5$, $\mu_{lim} = 1$.

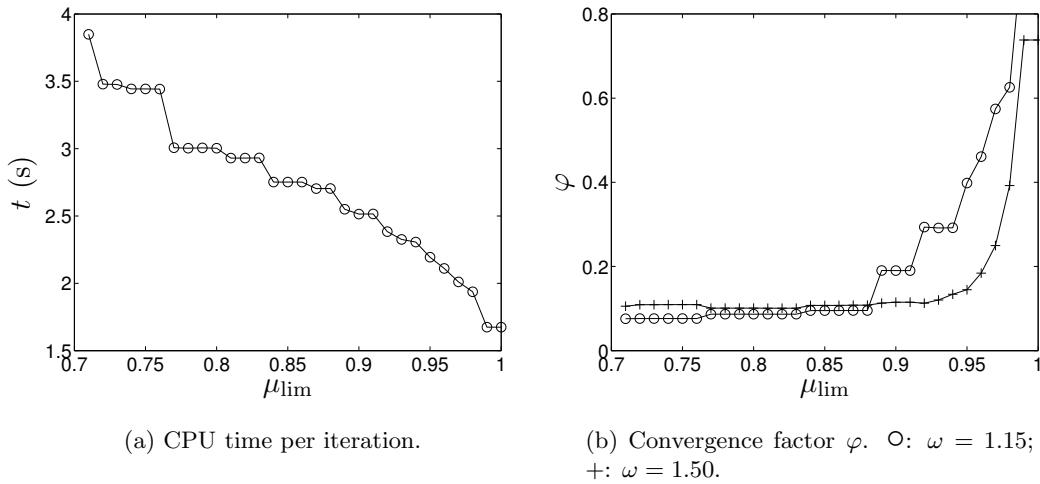


Figure 4.5: Computational cost and convergence factor for grid with AR_{max} = 10.

The cost of each iteration depends on the number of cells in the coarse grids, and hence on μ_{lim} . There are two possible measures of this cost: either work units (WU) or direct timing of the code. Work units attempt to measure the cost of an iteration in terms of smoothing sweeps on the finest grid, while often neglecting residual evaluations and intergrid transfers. With the cheap smoother used here, these operations will be of about the same complexity, and hence direct timing of the code is used here for simplicity.

The CPU time per iteration t is plotted as a function of μ_{lim} in figure 4.5(a) for a representative grid. The cost t decreases with increasing μ_{lim} , since higher values of

μ_{lim} are less restrictive on the coarser grids. The algorithm then generates grids with fewer cells, which decreases the cost per iteration. The convergence factor φ , which is plotted in figure 4.5(b), increases with μ_{lim} as expected. This is simply due to the fact that the coarser grids have been coarsened in directions where the smoother can not adequately smooth the error. Note that as ω is increased, the convergence improves for high μ_{lim} but becomes worse for low μ_{lim} . The more aggressive coarsening with a higher μ_{lim} generates grids that are, on average, more suitable for a higher over-relaxation parameter in conformance with the results of Yavneh (1996). Since the convergence factor is clearly a function of both μ_{lim} and ω , the optimal value of each parameter depends on the value of the other.

To get a good measure of the convergence efficiency, the CPU time per iteration is normalized as

$$\mathcal{T} = \frac{t}{-\log_{10} \varphi}, \quad (4.11)$$

which means that \mathcal{T} is the time it takes to reduce the residual by one order of magnitude. The normalized cost \mathcal{T} for three grids with different maximum aspect ratios is shown in figures 4.6(a) - 4.6(c). The optimum combination of μ_{lim} and ω is different for the different grids, but some trends can be seen. First, higher values of μ_{lim} require higher values of ω , which is an effect of the coarse grids being closer to those generated by full coarsening. For such grids, Yavneh (1996) suggested over-relaxation parameters of 1.75 and above. Second, the optimum value of μ_{lim} increases with the maximum aspect ratio of the grid. This is due to the fact that grids with higher stretching have larger numbers of coarse grids, and hence higher computational costs. This increased cost becomes more pronounced for lower values of μ_{lim} , which results in an optimum value that is higher.

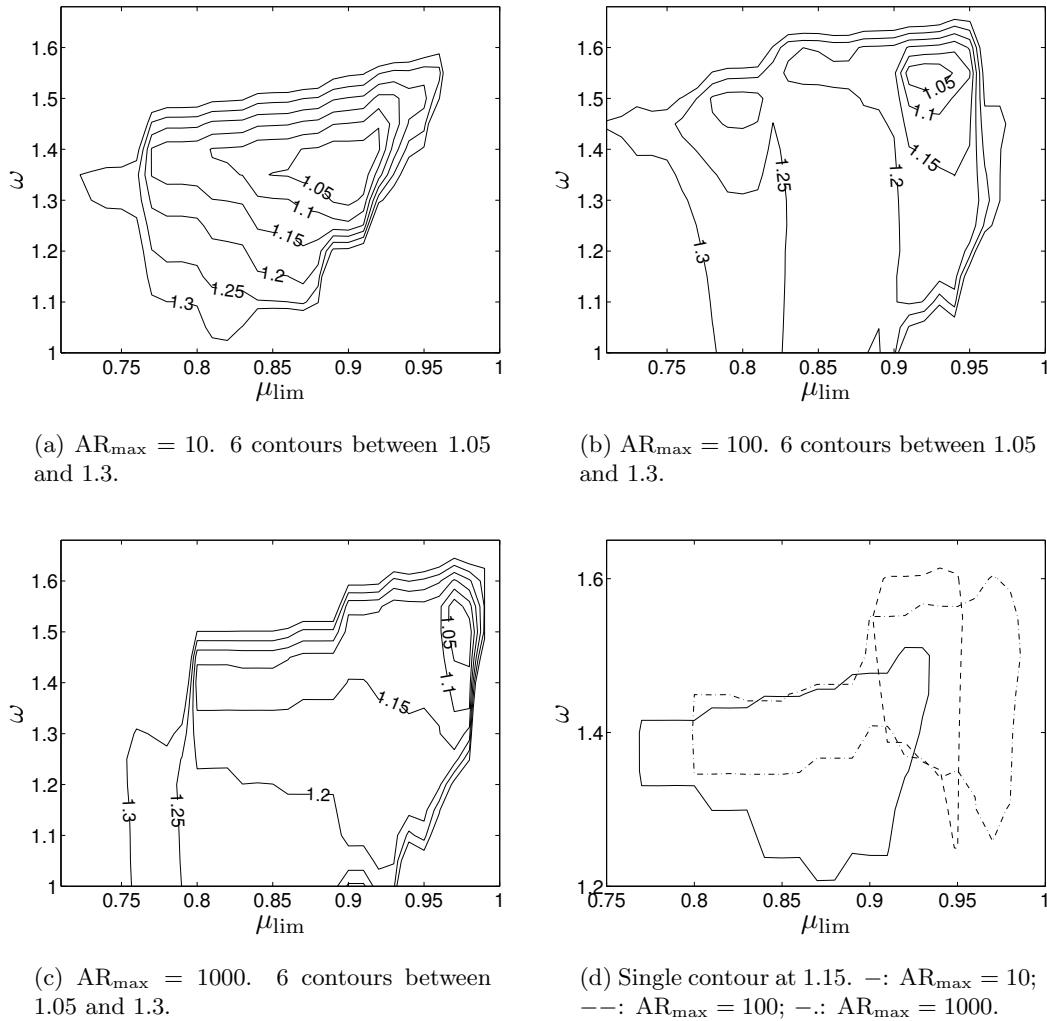
A single contour of \mathcal{T} at a value 15% higher than the optimum is plotted for each grid in figure 4.6(d). From this, $\mu_{\text{lim}} = 0.92$ and $\omega = 1.45$ are taken as compromise values, yielding convergence within 15% of the optimum for each grid.

4.3.2 Cost Scaling

Optimal multigrid methods have computational costs that scale linearly with the number of unknowns N . To assess the cost scaling of the present method, the discretized Poisson equation (3.34) is solved on grids similar to those used in the previous section, with $N_D \in \{16, 32, 48, 64, 80, 96\}$. Two different degrees of stretching are used: $\text{AR} = 1$ (isotropic) and $\text{AR}_{\text{max}} = 100$. The compromise values of $\mu_{\text{lim}} = 0.92$ and $\omega = 1.45$ are used for the stretched grids, while the isotropic grids use the theoretically optimal $\omega = 1.15$.

The normalized computational cost is shown in figure 4.7, where $N = N_D^3$. The cost for the uniform grid scales linearly with N , which is characteristic of optimal multigrid algorithms. The cost for the stretched grid does not scale linearly, but rather like $\mathcal{T} \sim N^{1.18}$. Using curve fits, the normalized CPU time on a single CPU is

$$\mathcal{T} \approx \begin{cases} 4.7 \cdot 10^{-6} N & \text{(s), AR} = 1 \\ 1.8 \cdot 10^{-6} N^{1.18} & \text{(s), AR}_{\text{max}} = 100 \end{cases} . \quad (4.12)$$

Figure 4.6: Contours of \mathcal{T} scaled by the lowest value for each grid.

Hence, the present conditional semicoarsening algorithm does not show optimal multigrid efficiency. For the grid sizes considered here, the cost on the highly stretched grid is about 2 to 5 times that of the isotropic problem.

These results can be compared to the method by Emvin (1997), who used an alternating plane-smoother with full coarsening. The computational cost scaled linearly with N for both a uniform and a moderately stretched ($AR_{\max} = 10$) grid. The cost, however, was not independent of the problem anisotropy. The stretched grids consistently required about 1.5 times more work than the uniform grids (using the alternating plane-smoother for both cases), and about 4.5 times more work than the uniform grid with a pointwise smoother (similar to the $AR = 1$ case here). Emvin (1997) unfortunately did not present results for grids with higher degrees of stretching, but one must assume that the cost would continue increasing further with increasing AR_{\max} .




Figure 4.7: Computational cost versus the number of grid points. \circ : $\text{AR} = 1$; $+$: $\text{AR}_{\max} = 100$; $-$: equation (4.12), $\text{AR} = 1$; $--$: equation (4.12), $\text{AR}_{\max} = 100$.

Thus, taking Emvin (1997)’s method as representative of algorithms that adapt the smoother on anisotropic problems, the present conditional semicoarsening method shows a worse (and non-optimal) cost *scaling* with N , but an actual cost *level* that is comparable for the grid sizes considered here. As will be seen later, the true benefit of the present conditional semicoarsening algorithm is its high parallel efficiency.

4.3.3 Parallel Efficiency

The code was parallelized using MPI and run on up to 8 CPUs on a cluster of AMD Athlon XP machines. The communication was handled by an ethernet switch. With $\mathcal{T}(N, p)$ denoting the normalized CPU time for N cells on p processors, the parallel efficiency can be defined as

$$\eta(N, p) = \frac{\mathcal{T}(N, 1)}{p\mathcal{T}(N, p)}. \quad (4.13)$$

The results are shown in figure 4.8. Figure 4.8(a) shows η versus the number of processors p in traditional fashion. As expected, the efficiency improves with increasing grid size, since the relative time spent on inter-processor communication then decreases. The efficiency is slightly higher for the stretched grids, since there are more grid levels with large numbers of cells compared to the uniform grids. The increase of η for the 32^3 stretched grids with $p = 8$ is probably an effect of cache performance. In this work, no cache optimization has been performed, and any such increases in performance are completely fortuitous.

A slightly different view is shown in figure 4.8(b), where the results have been plotted versus the number of cells per processor N/p . Apart from some irregularities, again likely due to cache performance, the results collapse rather well. The parallel efficiency is above 0.9 for $N/p \gtrsim 4000$, and stays above 0.7 even for $N/p = 512$ – a quite high value for such a coarse grid.

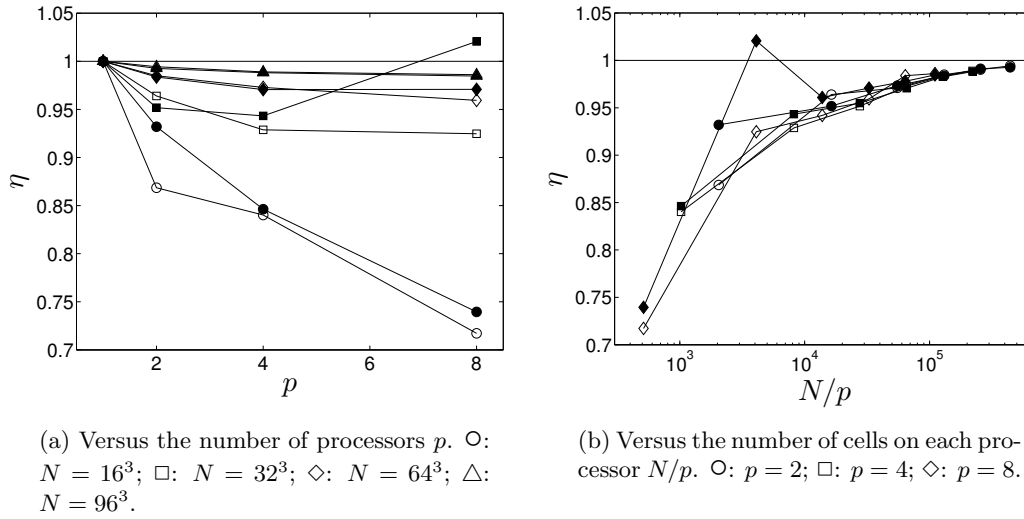


Figure 4.8: Parallel efficiency η for different grids. Open markers for $\text{AR} = 1$, filled markers for $\text{AR}_{\max} = 100$.

For comparison, Brown *et al.* (2000) studied the parallel performance of Schaffer (1998)'s method which combines a plane-smoother, a line-smoother, and semicoarsening, all in fixed directions. They solved a Poisson equation on uniform but anisotropic grids using a 7-point grid stencil, which is similar to the present case. However, Brown *et al.* (2000) defined the parallel efficiency in a slightly different way, as $\mathcal{T}(N, 1)/(\mathcal{T}(pN, p))$, which makes a direct comparison difficult. Only if $\mathcal{T} \sim N$, and thus $\mathcal{T}(pN, p) = p\mathcal{T}(N, p)$, is the parallel efficiency defined by Brown *et al.* (2000) equivalent to (4.13). This condition holds neither here nor in Brown *et al.* (2000). With these caveats, a very rough comparison can still be made. For $N/p \approx 8000$, Brown *et al.* (2000) reported parallel efficiencies of 0.7-0.8 when the domain decomposition was done such that the plane- and line-smoothers never crossed domain boundaries; 0.5-0.6 when only the line-smoother never crossed domain boundaries; and 0.2-0.3 for the general domain decomposition. Compare this to the present results, where $\eta \approx 0.93$ for $N/p \approx 8000$ with no restrictions on the domain decomposition. This shows the inherent parallelism of the present method, which is mainly due to the parallel properties of the pointwise smoother.

To really test the parallel efficiency, computations on a larger number of CPUs would be necessary. Brown *et al.* (2000), for example, go up to $8 \cdot 10^6$ cells on up to 64 processors. That being said, the present results, especially figure 4.8(b), certainly suggest that the present method might be efficient for very large parallel computations.

Note also that the parallel algorithm is equivalent to the sequential one, in that the smoother exchanges information at the block boundaries during every smoothing sweep. In other words, the solution is exactly the same after each iteration independent of the number of processors used.

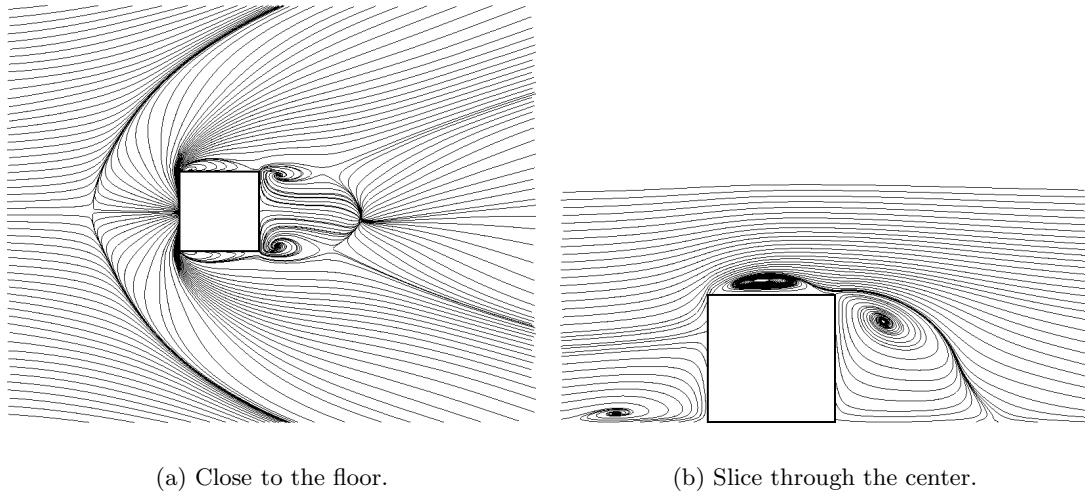


Figure 4.9: Time-averaged streamlines around the wall mounted cube. The flow is from left to right.

Case	Near wall treatment	N	AR_{\max}
Channel A	Schumann (1975)	49,152	8
Channel B	hybrid LES/RANS	98,304	393
Channel C	wall resolved LES	196,608	32
Wall mounted cube	instantaneous log-law	436,696	35

Table 4.1: Complex grid cases. Details of the near wall treatments are given in chapters 5 (Channel C) and 6 (remaining cases).

4.3.4 Complex Grids

One benefit of the present algorithm is that it is trivially applicable to more complex geometries (that use structured grids). To demonstrate this, the flow around a cube mounted on one wall of a rectangular channel is solved by LES. This flow has been studied experimentally and numerically by many authors (c.f. Rodi *et al.*, 1997). The grid consists of 17 blocks and is stretched in all directions. While only the multigrid performance on the Poisson equation for the pressure is of interest here, some flow results are shown in figure 4.9 to give an impression of the geometrical complexity.

Three different plane channel flows are also considered, all with different near wall treatments which has quite dramatic effects on the grids. Details of all 4 cases are listed in table 4.1, and the computational cost of solving the Poisson equation is shown in figure 4.10. Two results are plotted for each case: one using the compromise values of $\mu_{\text{lim}} = 0.92$ and $\omega = 1.45$, and one using the combination found to give the best performance for each case. As can be seen, the compromise values yield convergence that is about 30% worse compared to the optimum for each case. On the one hand, this shows that some tuning of these parameters is necessary to obtain the fastest




Figure 4.10: Computational cost versus the number of cells. Open markers use $\mu_{\text{lim}} = 0.92$, $\omega = 1.45$. Filled markers use the best combination for each case. \square : channel A; \diamond : channel B; \circ : channel C; \triangle : wall mounted cube; $-$: equation (4.12), $\text{AR} = 1$; $--$: equation (4.12), $\text{AR}_{\text{max}} = 100$.

convergence possible. On the other hand, it also shows that the compromise values yield reasonable results even for widely different cases. This is taken as a sign that the algorithm is relatively, but not perfectly, robust as far as different grid anisotropies are concerned.

Also plotted in the figure are the curve fits (4.12) to the results for the three dimensional unit cube. It is seen that the performance for the complex grid cases agrees quite well with the unit cube results. Channels A and C have maximum cell aspect ratios between 1 and 100 and are located between the two curve fits, while channel B, as expected due to its higher grid stretching, has a higher computational cost.

4.4 Summary and Conclusions

A conditional semicoarsening multigrid algorithm that maintains a reasonably high convergence rate even on highly stretched grids has been proposed and investigated. The algorithm displays high parallel efficiencies, mainly due to the inherent parallelism of the RBPGS smoother. The smoothing behaviour predicted by local Fourier analysis is used to determine how to coarsen the grid, and while the grids are not optimal in terms of how they reflect the local anisotropy, they do maintain a structured topology on all grid levels. The fact that the conditional semicoarsening is based on a local Fourier analysis results in a strong coupling between the anisotropy of the matrix coefficients and the smoothness of the error. Several numerical experiments were performed to determine suitable values for the parameters involved. The combination $\mu_{\text{lim}} = 0.92$ and $\omega = 1.45$ was found to yield convergence within a computational cost of about 30% of the optimum for all cases studied. The optimum combination was found to be case-dependent, and some tuning on a case-by-case basis is required to find this

combination. This being said, the fact that performance within 30% of the optimum can be achieved by one combination for the many cases studied here suggests that the algorithm is relatively robust with respect to grid anisotropies.

The computational cost does not scale linearly with the number of grid cells, but rather like $\sim N^{1.18}$ for grids with a maximum cell aspect ratio of 100. While the scaling is not optimal, the actual cost level is competitive with other multigrid algorithms for anisotropic problems, especially in parallel situations. Parallel efficiencies of above 90% were found for grids with as few as 4096 cells per processor.

Only structured coarse grids were considered here, but the algorithm is in some ways more naturally suited to allowing unstructured coarse grids. Since there would be less limitations on the coarsening process for such grids, it is conceivable that this would result in a more favourable cost scaling.

Chapter 5

Wall Resolved LES

While the main topic of this thesis is the modeling of the inner boundary layer in LES, wall resolved computations are performed for a few reasons. First, these simulations provide further validation of the numerical method, in that the results can be compared to existing DNS results. Second, these results can later be used for comparisons when assessing the approximate near wall treatments in chapter 6. Finally, they provide a flow database from which physically accurate forcing fields (to be used in chapter 6) can be extracted.

5.1 Introduction and Review of Plane Channel Flow

Plane channel flow is defined as the flow between two infinitely large plates separated by a distance of 2δ , and is, of course, an idealization rather than a real flow. It could correspond to a few flows that exist in nature, such as the flow between two concentric cylinders for high D/δ (where D is the average diameter of the cylinders) or the flow in rectangular ducts for high L_z/δ (where L_z is the spanwise size of the duct, as shown in figure 1.3). Dean (1978) surveyed a large collection of experiments on rectangular duct flows, and concluded that the flow statistics become independent of L_z for $L_z/\delta \gtrsim 14$. Hence, the idealized channel flow considered here is an appropriate model of the flow in such high aspect ratio ducts.

Plane channel flow can be characterized by (at least) three different Reynolds numbers:

$$Re_b \equiv \frac{U_b \delta}{\nu}, \quad Re_c \equiv \frac{U_c \delta}{\nu}, \quad Re_\tau \equiv \frac{u_\tau \delta}{\nu} \equiv \delta^+, \quad (5.1)$$

where U_b is the average bulk velocity, U_c is the average centerline velocity, and u_τ is the friction velocity. Note that the length scale used here is the channel half width δ , whereas the duct height 2δ is used in parts of the literature. The different Reynolds numbers in (5.1) are dependent parameters, in that only (any) one of them can be specified in an experiment or a computation, whereas the remaining two are determined by the resulting flow. The survey by Dean (1978) found that the empirical relations

$$Re_c \approx 1.27 Re_b^{0.988}, \quad Re_\tau \approx 0.175 Re_b^{0.875} \quad (5.2)$$

accurately described the experimental data. Transition to turbulence was found to occur in the range $600 \lesssim Re_b \lesssim 1400$.

At sufficiently high Reynolds number, plane channel flow can be divided into several sublayers, all sketched in figure 1.2. This division is inspired partly by experimental observations, but primarily by theoretical analysis of the average streamwise momentum equation. The average streamwise momentum equation becomes inviscid for $y^+ \gtrsim 30$ (the outer layer), whereas the effect of the driving pressure gradient is small for $y/\delta \lesssim 0.1$ (the inner layer). Matching the solutions to these outer and inner layer equations yields the famous log-law (c.f. Tennekes & Lumley, 1972)

$$\bar{u}^+ = \frac{1}{\kappa} \ln(y^+) + B, \quad (5.3)$$

which is valid in the overlap layer $30 \lesssim y^+ \lesssim 0.1\delta^+$. Dean (1978) suggested that

$$\kappa = 0.41, \quad B = 5.2 \quad (5.4)$$

provided the best fit to the experimental data in his survey.

Wei & Willmarth (1989) performed experiments in the range $190 \lesssim Re_\tau \lesssim 1900$ to investigate the effects of the Reynolds number very near the wall, and found that while the mean velocity profiles essentially collapsed when normalized using inner variables (i.e., u_τ and ν), the Reynolds stresses did not. More specifically, the wall normal stress component $\overline{v'v'^+}$ displayed the least agreement between the different Reynolds numbers.

Such Reynolds number effects have been found by many experimental and computational studies, and is an active area of research. A theory that attempts to explain these effects was developed recently by Wosnik *et al.* (2000), who argued that the equations for single- and multi-point statistics¹ become inviscid at different distances from the wall. The single-point equations become inviscid for $y^+ \gtrsim 30$, which is the definition of the lower limit of the outer layer in figure 1.2. Wosnik *et al.* (2000), however, argued that the multi-point equations become inviscid farther out, for $y^+ \gtrsim 300$. As a direct consequence, flow statistics are potentially dependent on the Reynolds number for $y^+ \lesssim 300$, which is consistent with the findings of Wei & Willmarth (1989) and others. Furthermore, a true inertial sublayer (in which data should collapse at different Reynolds numbers) can, then, only be found for $y^+ \gtrsim 300$. From figure 1.2, such an inertial sublayer can only be expected if $0.1\delta^+ > 300$, i.e. for $Re_\tau > 3000$.

Starting with the wall resolved LES by Moin & Kim (1982), there have been several published DNS and wall resolved LES studies of plane channel flow. The rather limited increase in Reynolds number over the years shows the high computational cost of resolving the inner layer. The DNS by Kim *et al.* (1987) was performed at $Re_\tau \approx 180$, Moser *et al.* (1999) performed DNS up to $Re_\tau \approx 590$, del Alamo *et al.* (2004) went up to $Re_\tau \approx 950$, and the most recent DNS by Hoyas & Jimenez (2005) pushed the limit to $Re_\tau \approx 2000$. The rate of increase for LES has been faster, partly due to the formally

¹Single-point statistics are formed using information at a single point, such as \bar{u} , $\overline{u'v'}$ etc. Multi-point statistics use information at multiple points, such as the two-point correlation $u'(x)u'(x+l_x)$. Since spectra are defined as the Fourier transform of two-point correlations, spectral information can be thought of as multi-point statistics.

slightly lower resolution requirements and partly due to the fact that the purpose of DNS studies is to provide extremely accurate results for all quantities, whereas the purpose of LES studies often only requires accurate prediction of some lower order statistics. The landmark LES by Moin & Kim (1982) was performed at $Re_\tau \approx 640$, which was an impressive achievement at the time. Later Piomelli (1993) performed wall resolved LES up to $Re_\tau \approx 2000$. The resolution, however, was quite coarse, and hence one might wonder whether the rather accurate results were partly fortuitous. Kravchenko *et al.* (1996) made use of coarser grids in the core of the channel in a very interesting numerical method, and as a consequence were able to reach $Re_\tau \approx 4000$.

Some important strengths of computational studies were highlighted by these studies. While it is very difficult to measure accurately at small scales in experiments, there is no such difficulty in numerical simulations. Hence, these studies allowed for the examination of the behaviour very near the wall, and for the direct calculation of the dissipation ε . It is also difficult to measure pressure fluctuations in experiments (except for at the wall), and hence these computational studies were the first to directly compute all the terms in the exact equation for the turbulence kinetic energy. A final strength is the amount of data made available by a simulation. It is, for example, difficult to obtain multi-point statistics from experiments. Such multi-point statistics are easily computed in a numerical study, and have, among many other things, been used to develop low dimensional models of near wall turbulence (e.g. Podvin, 2001) and so-called ‘optimal’ residual stress models for LES by Völker *et al.* (2002).

One of the most important contributions from DNS and LES studies has been the increased understanding of the turbulence structure in the near wall region, as pointed out by Moin & Mahesh (1998). The dominant features in this region are quasi-streamwise vortices and high- and low-speed streaks (c.f. Robinson, 1991), which interact and regenerate continuously. While the sizes of these structures vary, the most energetic ones are of size $\lambda_z^+ \approx 100$ and $\lambda_x^+ \approx 800$ in the DNS by del Alamo & Jimenez (2003).

Due to the statistical homogeneity of channel flows, the most common numerical approach is to use periodic boundary conditions in the wall parallel directions. This assumed periodicity affects the solution unless the domain is large enough to contain the largest structures. Jimenez & Moin (1991) found that the smallest domains that sustained turbulence were of sizes $L_x^+ \approx 250 - 350$ and $L_z^+ \approx 100$, where the streamwise value displayed more sensitivity to the Reynolds number. These domain sizes correspond to the approximate size of the near wall structures, especially in the spanwise direction, and contain exactly one streak and a pair of quasi-streamwise vortices. For this reason, Jimenez & Moin (1991) called it the ‘minimal flow unit’ of plane channel flow. While this provides a definite lower limit for the necessary domain size to sustain turbulence, accurate results demand much larger domains. For $Re_\tau \gtrsim 360$, most studies have used $L_x/\delta = 2\pi$ and $L_z/\delta = \pi$ (e.g. Moser *et al.*, 1999; Gullbrand & Chow, 2003), but recent work by del Alamo & Jimenez (2003) showed that even larger domains are needed to capture the largest structures present².

²The common use of multiples of π stems from the fact that most of the original DNS studies used spectral codes.

5.2 Methodology

The numerical method described in chapter 3 is used, and the Smagorinsky model (equations (2.11) and (2.12)) with the dynamic procedure (equation (2.20) and the ones leading up to it) is used to model the residual stress tensor. The filter width $\tilde{\Delta}$ is taken as the cubic root of the cell volume. The quality of the statistics is measured partly by their convergence in time, and partly by whether they satisfy the appropriate symmetries across the channel. Averaging is performed over time and the wall parallel directions.

The grids are uniform in the wall parallel directions, and stretched in the wall normal direction. The grid coordinates in this direction are given by

$$y^l = 1 + \frac{\tanh\{\Gamma(2l/N_y - 1)\}}{\tanh(\Gamma)}, \quad l = 0, \dots, N_y, \quad (5.5)$$

where y^l is the y -coordinate of the l th grid line, and Γ is a stretching parameter.

Estimation of errors is always important, be it in experimental or numerical investigations. The errors here are of essentially three kinds: errors due to insufficient resolution in space and time (either numerical errors or a lack of capturing the physically relevant structures), errors due to the finite domain size, and finally errors in the turbulence modeling.

The domain sizes used here are rather conventional, and not as large as those used by del Alamo & Jimenez (2003). The reason is the current focus on the approximate near wall modeling, for which the largest scales should be less important. Since many LES studies have reported accurate results with similar or even smaller domain sizes (e.g. Piomelli, 1993; Kravchenko *et al.*, 1996; Gullbrand & Chow, 2003), this error is assumed to be small in the present calculations.

To assess the remaining errors, simulations on different grids and without the turbulence model are performed. Two different grids that are finer than the base case are considered: one that is refined in the wall parallel directions only, and one that is refined in the wall normal direction only. Due to the high computational cost, a smaller domain in the spanwise direction is used for the first of these cases. Removal of the turbulence model is a (rather drastic) way to test its influence on the results. While it does not bring out the errors due to the turbulence modeling, it at least provides some measure of the importance of the modeling.

The differences in the results will be presented throughout this chapter. In summary, the finer wall normal resolution does not change any flow statistics, while the finer wall parallel resolution does change the results slightly. This change, however, is considered small enough for the base grid to give reasonably accurate results, especially in the light of the purpose of these simulations – had the purpose been a detailed investigation of the near wall turbulence, for example, finer grids should certainly be used.

Choi & Moin (1994) showed that a time step of size $\Delta t^+ \equiv \Delta t u_\tau^2 / \nu \approx 0.4$ resulted in negligible error in their DNS of the minimal flow unit using a numerical method similar to the one used here. A similar test on the minimal flow unit, though not shown here, results in the same conclusion for the present method. Since this is the time step used here, no further tests are deemed necessary.

Case	(N_x, N_y, N_z)	$(L_x, L_z)/\delta$	Δx^+	Δy_{\min}^+	Δy_{\max}^+	Δz^+	Re_τ
base	(72, 80, 64)	$(2\pi, 3)$	44	0.8	33	24	506
base, without model			47	0.9	35	25	541
fine- xz	(128, 80, 64)	$(2\pi, 2)$	25	0.8	33	16	512
fine- xz , without model			27	0.9	35	17	543
fine- y	(72, 124, 64)	$(2\pi, 3)$	44	0.3	24	24	506
fine- y , without model			47	0.3	25	25	543

Table 5.1: Wall resolved LES cases

The important parameters for the different cases are summarized in table 5.1. The Reynolds number based on the bulk velocity is $Re_b = 9000$, which results in $Re_\tau \approx 500$. The mean pressure gradient is adjusted at each time step to maintain a constant U_b .

5.2.1 Initial Condition and Transition

Fully developed channel flows are, by definition, independent of the initial conditions. There is one caveat, though, in that the initial conditions must be such that the flow transitions to turbulence. Linear stability theory (c.f. Drazin & Reid, 1981) predicts that plane channel flow is unstable to infinitesimal disturbances³ for $Re_b > 3848$, and hence any initial condition should eventually lead to transition and subsequent turbulence due to the round-off errors in any computer. This process, however, can be very slow in numerical simulations, since the unstable modes have to grow from the round-off amplitude of 10^{-16} (with double precision), and the growth rate is very low. For example, the most unstable mode with $\lambda_x/\delta = 1$ at $Re_b = 6667$ grows by two orders of magnitude in $tU_b/\delta \approx 800$ (Drazin & Reid, 1981), a typical total time of the present simulations.

In practice, this transition can be accelerated by explicitly introducing large fluctuations in the initial laminar field. If these fluctuations are eigenmodes of the Orr-Sommerfeld equation (c.f. Drazin & Reid, 1981), then the resulting transition will be physically representative of natural transition. Since the transition phase is of no interest here, random white noise with 10% intensity is added instead. While less physically appealing, it leads to transition after only about $tU_b/\delta \approx 15$.

While the transition is unphysical, it nonetheless shows some interesting phenomena. Figure 5.1 shows the evolution of two quantities from the given initial condition for one of the grids. The wall shear stress (figure 5.1(a), normalized as the friction Reynolds number) is initially low, but increases rapidly at $tU_b/\delta \approx 10$ when the flow undergoes transition. This increase happens later for the case with a residual stress model, which is simply due to the added dissipation of the model. Also, the whole transition process happens more slowly when the model is used.

The wall normal velocity fluctuation near the center of the channel is shown in figure 5.1(b). It decays initially, since all but a few modes (of the several hundred thousand) of the initial disturbances are stable. Towards the end of the increase in

³In the literature this limit is commonly given as $Re_c > 5772$. For laminar flow $U_c = 3U_b/2$, which implies $Re_b > 3848$.

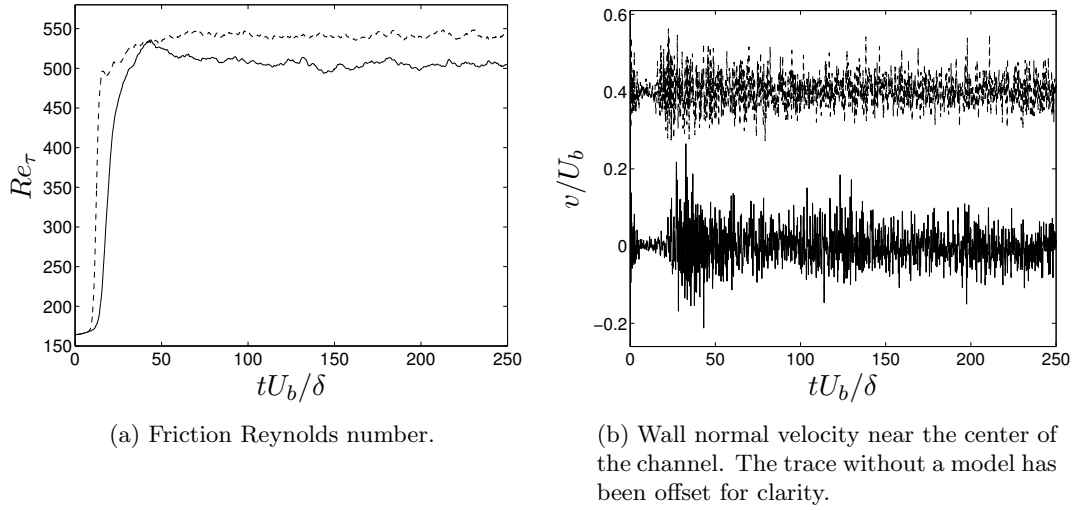


Figure 5.1: Time histories after initializing to a laminar profile with random white noise added on grid fine- y . —: with model; ---: without model.




Figure 5.2: Time history after initializing to a laminar profile with random white noise added on grid fine- y with a model. Quantities are averaged over all grid cells without proper volume weighting. —: C from the dynamic procedure divided by the equivalent value from the standard Smagorinsky model with van Driest damping (using $C = 0.13^2$); ---: ν_{les}/ν .

the wall shear stress, however, the velocity fluctuation near the center of the channel increases again to a fully turbulent state. The fact that this increase happens later in the middle of the channel simply shows that the instability is located closer to the wall.

To see how the dynamic procedure responds to this evolution, the Smagorinsky coefficient C and the eddy viscosity ν_{les} are shown in figure 5.2. These values have been




Figure 5.3: Mean velocity. Open markers with model, filled markers without model. \circ : base grid; \square : fine- xz grid; $+$: DNS at $Re_\tau \approx 550$ by del Alamo *et al.* (2004); $--$: $\bar{u}^+ = y^+$ and $\bar{u}^+ = \ln(y^+)/0.41 + 5.2$, respectively.

averaged over all cells, without any volume weighting, and thus these quantities have no physical significance. After an initial (very brief) increase, both quantities decay as the stable modes of the initial disturbance decay. During the transition phase, however, the quantities increase in response to the presence of resolved turbulence. This shows that the dynamic procedure is sensitive to the state of the flow, which is indeed the reason for its success. The eddy viscosity reaches a fully developed state after about 70 time units. This probably explains why the case with a model shows a slower development of the wall shear stress in figure 5.1(a), since the flow has to continue adjusting to the eddy viscosity up to $tU_b/\delta \approx 70$.

5.3 Single-Point Statistics

The mean velocity is shown in figure 5.3 for some representative cases. The remaining cases are essentially indistinguishable from the ones shown: the fine- y profile coincides with the base one, while all the cases without a model coincide. Both the base and the fine- xz cases agree well with the DNS by del Alamo *et al.* (2004), except for the fact that the wake is too small for both. This absence of a wake was found by Cabot *et al.* (1999) as well, using grids with similar resolutions as here. They concluded that it was due to insufficient resolution in the core of the channel. More interestingly, \bar{u}^+ is about 0.25 lower for the fine- xz grid (compared to the base grid) for $y^+ \gtrsim 15$. This indicates that the buffer layer is slightly underresolved.

Removing the model causes a rather large underprediction of the velocity gradient in the buffer layer, and a lower velocity gradient up to $y^+ \approx 100$. Clearly the grids are too coarse to properly resolve the buffer layer turbulence. Gullbrand & Chow (2003) obtained qualitatively similar results when removing the model. Interestingly, a wake

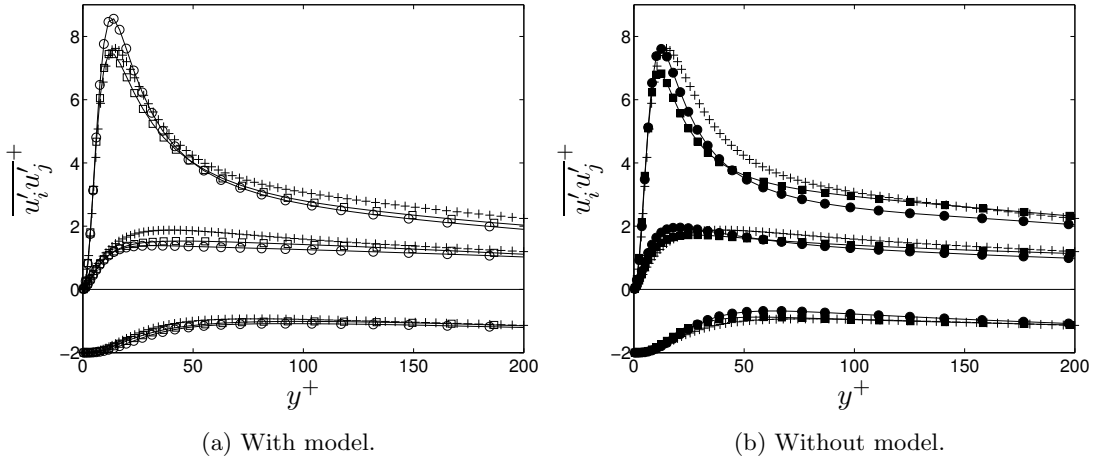


Figure 5.4: Resolved normal stresses. Open markers with model, filled markers without model. From top to bottom $\overline{u'u'}$, $\overline{w'w'}$, $\overline{v'v'}$ (offset for clarity). \circ : base grid; \square : fine- xz grid; $+$: DNS at $Re_\tau \approx 550$ by del Alamo *et al.* (2004).

is recovered when the model is removed, although the reason for this is not clear. One possibility is that the lower effective viscosity essentially translates into a slightly higher effective resolution.

The resolved normal stresses are shown in figure 5.4, where, again, the results for the fine- y grid (not shown) coincide with the base grid. The streamwise stress is overpredicted in the buffer layer for the base grid, but is closer to the DNS result for the fine- xz grid. Both the spanwise and wall normal components are slightly underpredicted in the present results. This implies a too high level of anisotropy in the buffer layer, especially for the base case. This is a very common result in LES, and is due to the grid resolution in the wall parallel directions (c.f. Kravchenko *et al.*, 1996).

The agreement with DNS in the buffer layer becomes worse when the model is removed. The anisotropy is now too low, with underpredicted $\overline{u'u'}$ and overpredicted $\overline{v'v'}$ and $\overline{w'w'}$. The fact that the peak $\overline{u'u'}$ occurs closer to the wall also suggests that the near wall turbulence in the buffer layer (quasi-streamwise vortices and streaks) has moved closer to the wall. Note that the resolved kinetic energy (i.e., half the sum of the three normal stress components) actually decreases when the model is removed, despite the dissipative character of the model. This seemingly contradictory result is due to a simultaneous change in the production of resolved kinetic energy $-\overline{u'v'}\partial\bar{u}/\partial y$, which is lower without a model due to the lower mean velocity gradient.

The eddy viscosity and the dynamically computed Smagorinsky coefficient C are shown in figure 5.5 for all three grids. The average eddy viscosity is lower than ν for all cases, which shows both that the grid resolution is relatively fine and that the Reynolds number is rather low. The eddy viscosity is identical in the base and fine- y cases, but the mean value of C differs between the two. This illustrates a key concept of the dynamic procedure, in that it finds the value of C that provides the appropriate dissipation. The fact that the statistics of the two cases coincide indicates that the two flow fields are




Figure 5.5: Average eddy viscosity and dynamic Smagorinsky coefficient. \circ : base grid; \square : fine- xz grid; \diamond : fine- y grid.

similar (and that the fields are well resolved in the wall normal direction). The eddy viscosity should, therefore, be similar in the two cases, which indeed is the case. The difference in the wall normal grid size causes differences in the value of the filter width $\tilde{\Delta}$ for the two grids, and hence the dynamic procedure finds different values of C .

The eddy viscosity is lower for the fine- xz grid, since the finer grid resolution in the important directions requires a lower ν_{les} . One also notes that the dynamic procedure correctly reduces C in the very near wall region, since $C = 0$ is needed at the wall to satisfy the boundary condition for τ_{ij} .

The structure of the near wall turbulence causes the vorticity to be a natural variable to examine. The quasi-streamwise vortices in the buffer layer are slightly inclined with respect to the wall (c.f. Robinson, 1991), but are close enough to the streamwise direction that the streamwise vorticity $\omega_x \equiv \partial w/\partial y - \partial v/\partial z$ is a good measure of their strength. Similarly, the high- and low-speed streaks near the wall are fairly elongated and primarily associated with fluctuations of u . Thus, the wall normal vorticity $\omega_y \equiv \partial u/\partial z - \partial w/\partial x \approx \partial u/\partial z$ in this region, which implies that ω_y is a good measure of the strength of the streaks.

The magnitudes of the streamwise and wall normal vorticity fluctuations are shown in figure 5.6. The DNS by del Alamo *et al.* (2004) shows a peak $\omega_{y,\text{rms}}$ at $y^+ \approx 14$, which indicates that this is where the streaks are the strongest. This peak coincides with the location of the peak $\overline{u'u'}$ in figure 5.4. The present results show substantially lower $\omega_{y,\text{rms}}$, although the peak value of $\overline{u'u'}$ agrees well with the DNS. This suggests that the streaks are wider in the present LES, and the fact that the fine- xz grid shows higher levels of $\omega_{y,\text{rms}}$ supports this. Interestingly, the turbulence model has little effect on the wall normal vorticity.

The effect of finer grid resolution is to increase $\omega_{y,\text{rms}}^+$ by about 7% at $y^+ \approx 12$, both with and without the model. From figures 5.4(a) and 5.4(b) the equivalent decreases of




Figure 5.6: Magnitude of the vorticity fluctuations. Open markers with model, filled markers without model. \circ : base grid; \square : fine- xz grid; $+$: DNS at $Re_\tau \approx 550$ by del Alamo *et al.* (2004).

u_{rms}^+ is about 6% for both cases. Assuming

$$\omega_{y,\text{rms}}^+ \sim \frac{u_{\text{rms}}^+}{\lambda_z^+} \quad (5.6)$$

for some average streak spacing λ_z^+ suggests that λ_z^+ decreases by 10-15% when the grid is refined, and the fact that these changes occur both with and without the model indicates that such a decrease of the streak spacing is one main effect of the finer grid resolution.

The streamwise vorticity magnitude $\omega_{x,\text{rms}}$ has a local maximum at $y^+ \approx 18$ and a local minimum at $y^+ \approx 5$ in the DNS results. Kim *et al.* (1987) argue that these locations correspond to the average center and edge of the quasi-streamwise vortices, respectively. The present results accurately predict the location of the local minimum, but erroneously predict the local maximum at $y^+ \approx 10 - 12$. This suggests that the streamwise vortices in the present simulations both are located too close to the wall and are too small in diameter. While the location of the local maximum does not agree with the DNS, it does agree with the location of peak $\omega_{y,\text{rms}}$ for all cases. This indicates that the whole near wall turbulence cycle is too close to the wall in the present simulations.

Interestingly, whereas the turbulence model had little effect on the wall normal vorticity, it is instead the grid resolution that appears to have little effect on the streamwise vorticity. That being said, the fine- xz case with a model shows higher levels of $\omega_{x,\text{rms}}$ for $y^+ \gtrsim 5$ in comparison to the base case with a model.

Removal of the turbulence model increases the streamwise vorticity by about 40% throughout the buffer layer, and causes an overprediction compared to the DNS for $y^+ \lesssim 12$. This is consistent with simultaneous increases of v_{rms}^+ (by about 30%) and w_{rms}^+ (by about 25%). Since the streamwise vortices are a main component of the wall




Figure 5.7: Resolved ($-\overline{u'v'}^+$, upper curves) and modeled ($\overline{\tau}_{12}^+$, lower curves) shear stresses. Open markers with model, filled markers without model. \circ : base grid; \square : fine- xz grid; $+$: DNS at $Re_\tau \approx 550$ by del Alamo *et al.* (2004).

normal velocity fluctuations, a large increase of ω_x likely leads to a higher resolved shear stress.

Figure 5.7 shows the resolved and modeled shear stresses in the stress balance relation for channel flows (2.30). The resolved stress clearly dominates, in fact it is larger than the modeled at all grid points. Although the viscous stress is not shown, it is the dominant component for $y^+ \lesssim 10$. It is also larger than the average modeled stress across the channel, due to the fact that the eddy viscosity is smaller than the molecular one for these grids (see figure 5.5(a)). The modeled stress is about 10% of the resolved stress at $y^+ \approx 10 - 15$, and this ratio then decreases down to about 1% in the core of the channel.

This low level indicates that the model has a very limited effect on the transfer of momentum. The poor results without a model, however, show that the overall effect of the model is rather large. Thus, the most likely conclusion is that the main effect of the model is to provide the necessary dissipation of resolved kinetic energy. Equivalently, the grids used here are fine enough to directly resolve the dominant stress-generating structures, but not nearly fine enough to accurately predict the dissipation.

The effect of refining the grid is, as expected, simultaneous increases and decreases of the resolved and modeled shear stresses, respectively. The case without a model overpredicts the shear stress compared to the DNS, which is consistent with the overpredicted streamwise vorticity discussed above. The re-arranged shear stress balance (2.33) then shows how this overly large resolved shear stress translates into a too low mean velocity gradient, as evidenced in figure 5.3.

Finally, it is clear that different quantities show different sensitivity to the grid resolution – the present grids give reasonably accurate mean velocity profiles and resolved stresses, but underpredict the vorticity fluctuations rather drastically. Since the vorticity involves velocity derivatives, which amplify the smaller scales, this underprediction

is not unexpected even for wall resolved LES.

5.4 Multi-Point Statistics

While single-point statistics bring out many important features of the resolved flow fields, they do not convey information about the length scales at which the processes take place. Such information can be brought out by considering multi-point statistics instead.

Two-point correlations can be defined in a very general sense, but are restricted to points separated only in the wall parallel directions in this thesis. With l_x and l_z denoting the separation distance in the streamwise and spanwise directions, respectively, the two-point correlation between two resolved velocity components is (c.f. Tennekes & Lumley, 1972)

$$r_{ij}(l_x, l_z) \equiv \overline{u'_i(x, y, z, t)u'_j(x + l_x, y, z + l_z, t)}, \quad (5.7)$$

where r_{ij} is a function only of the separation due to the statistical homogeneity. Note that the resolved Reynolds stresses are $\overline{u'_i u'_j} = r_{ij}(0, 0)$, and that the resolved kinetic energy is $r_{ii}(0, 0)/2$. Also, since the resolved shear stress is negative, $r_{12}(0, 0) < 0$. In this thesis, referring to a ‘high’ value of $r_{12}(0, 0)$ implies a large negative value.

The two-point correlation coefficient is defined as

$$\rho_{ij}(l_x, l_z) \equiv \frac{r_{ij}(l_x, l_z)}{u'_{i,\text{rms}}u'_{j,\text{rms}}}. \quad (5.8)$$

For isotropic and homogeneous turbulence, the energy spectrum function $\hat{E}(k)$ is defined as the three dimensional Fourier transform of half the trace of the two-point correlation, with the directional information in the wavenumber vector averaged out (c.f. Tennekes & Lumley, 1972). More simply put, $\hat{E}(k)$ is the kinetic energy between wavenumbers k and $k + dk$, and the kinetic energy $\overline{u'_i u'_i}/2 = \int_0^\infty \hat{E}(k)dk$. The averaging of directional information implies that k then represents the magnitude of the wavenumber, i.e. $k \geq 0$.

For the channel flows considered here, the Fourier transform is only well defined in the wall parallel directions. One can then consider one dimensional spectra $\hat{E}_j(k_j)$, where the Fourier transform is taken in direction j . Due to the periodic boundary conditions, this Fourier transform is actually a Fourier series with wavenumbers $k_j^{(n)} = 2\pi n/L_j$, but such technicalities are ignored here for simplicity.

For plotting purposes, the so-called premultiplied form $k\hat{E}(k)$ can be considered. Note that

$$\frac{\overline{u'_i u'_i}}{2} = \int_0^\infty \hat{E}(k)dk = \int_{-\infty}^\infty k\hat{E}(k)d(\ln k) = \int_{-\infty}^\infty k\hat{E}(k)d(\ln \lambda) \quad (5.9)$$

with $\lambda = 2\pi/k$. When $k\hat{E}(k)$ is plotted in a linear scale versus $\ln \lambda$, the area under graph represents the kinetic energy contribution from different wavelength intervals.

The one dimensional energy spectra are plotted in their premultiplied forms in figure 5.8. Considering first the cases with a model, the overall peak in the spectra occurs

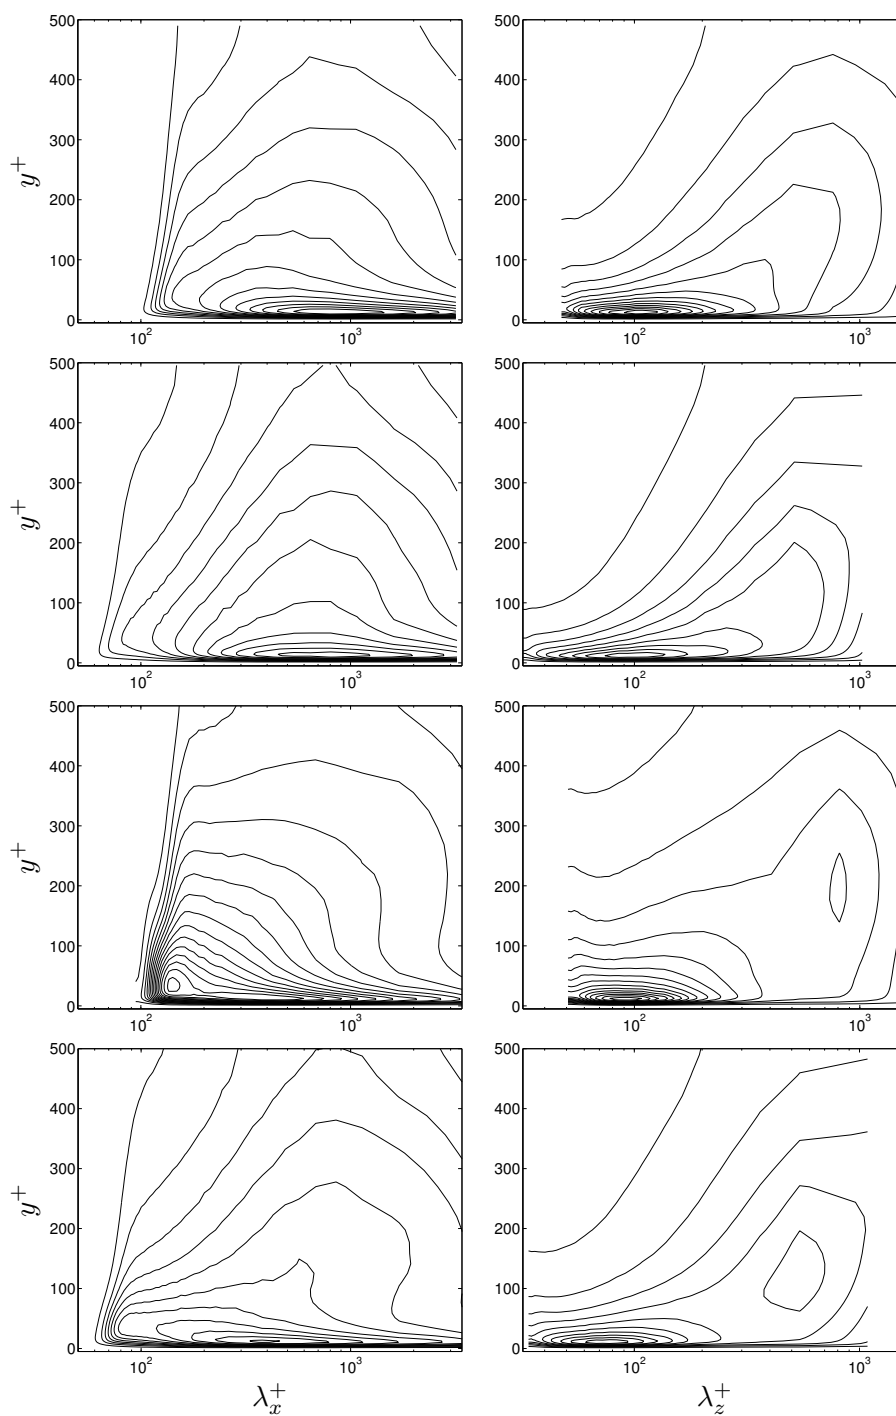


Figure 5.8: Contours of the premultiplied one dimensional spectra $k\hat{E}(k)$. 15 contours between 0 and the common maximum, with the axes the same columnwise. The left column are the streamwise spectra, the right column are the spanwise. Top row: base grid with model; second row: fine- xz grid with model; third row: base grid without model; bottom row: fine- xz grid without model.

at $y^+ \approx 12 - 14$ for wavelengths $\lambda_x^+ \approx 800$ and $\lambda_z^+ \approx 110$ for the base grid, and at only slightly smaller wavelengths ($\lambda_z^+ \approx 100$) for the fine- xz grid. This 10% decrease of λ_z^+ corresponds well to the estimate based on the wall normal vorticity magnitude in the previous section. The spectra look qualitatively similar for the two grids, which is encouraging. Both cases, however, show signs of the domain sizes being insufficient in both directions, which is consistent with the findings of del Alamo & Jimenez (2003). The fact that most of the single-point statistics are still rather well predicted, despite the limited domain sizes, shows that the most important dynamical phenomena are contained within relatively smaller wavelengths.

It is interesting to note the behaviour in the small scale limit. In the spanwise direction, there is considerable energy at the numerical cut-off in the buffer layer for both grids. In the streamwise direction, however, the energy is very low at the numerical cut-off for both grids. One might then conclude that the grid resolution is sufficiently fine in the streamwise direction. On the other hand, the fact that the streamwise spectra are different at small wavelengths shows that there are resolution effects at play. Whether it is the finer spanwise resolution that generates smaller streamwise motions (which would suggest that the streamwise resolution is sufficiently fine) or whether the streamwise resolution affects the flow such that the energy decays too quickly as $\lambda_x \rightarrow 2\Delta x$ is impossible to tell. Either way, around $y^+ \approx 50$ the dominant streamwise wavelength is somewhat lower for the base grid compared to the fine- xz grid. Apart from that, the dominant streamwise wavelength stays more or less constant across the channel, at what corresponds to $\lambda_x/\delta \approx 1.6$. Here one can note that del Alamo & Jimenez (2003) found peaks at larger wavelengths of $\lambda_x/\delta \approx 3$ in their DNS on a much larger domain ($L_x/\delta = 8\pi$, $L_z/\delta = 4\pi$). The fact that the smaller domains used here result in lower wavelengths at the peaks is consistent with expectations.

The dominant spanwise wavelength grows to $\lambda_z^+ \approx 700$ towards the center of the channel for the base grid. This corresponds to $\lambda_z/\delta \approx 1.4$, which is close to half the domain width. In contrast to the streamwise direction, the peak spanwise wavelength agrees well with the DNS by del Alamo & Jimenez (2003).

Consider next the cases without a residual stress model, the removal of which results in rather different spectra. The peaks occur slightly closer to the wall, at $y^+ \approx 12$. While a minor difference, it is consistent with the peaks at slightly lower y^+ for the single-point statistics. The peak spanwise wavelength in the buffer layer is $\lambda_z^+ \approx 90$ for the base grid and $\lambda_z^+ \approx 80$ for the fine- xz grid (again, the decrease is consistent with the estimate based on the wall normal vorticity magnitude).

The peak streamwise wavelength at $y^+ \approx 12$ is $\lambda_x^+ \approx 400$ for both grids, considerably shorter than for the cases with a model. The vertical and spanwise velocity fluctuations have their peaks at even smaller wavelengths $\lambda_x^+ \approx 200$ (not shown). The fact that the v' and w' components are associated with smaller streamwise structures is consistent with DNS findings by e.g. Jimenez *et al.* (2004), but the wavelengths are too small in the present results. This shows that the present grids are too coarse for DNS, and that the added dissipation from the residual stress model improves the structure of the buffer layer turbulence by increasing the length scales. Furthermore, the base grid has a peak in the streamwise spectrum at $y^+ \approx 35$ for $\lambda_x^+ \approx 150$. This is likely due to a pile-up of energy near the numerical cut-off. Interestingly, the fine- xz grid shows no signs of such

a pile-up, which suggests that there is some critical resolution for which this occurs. When the model is included, the added dissipation prevents such pile-up of energy.

Towards the center of the channel, the dominant wavelengths approach those for the cases with a model. This is not unexpected, since the resolution here is much finer relative to the integral scales. For both grids there is a spurious peak for rather large λ_z^+ around $y^+ \approx 200$ (base grid) and 100 (fine- xz grid). The reason for this is not known, but one may note that the streamwise spectra develop large scale energy at those values of y^+ , respectively. It is, then, possible that the spurious peaks in the spanwise spectra are related to this formation of larger structures.


5.5 Turbulence Structure

Snapshots of the streamwise velocity fluctuation u' in the buffer layer are shown in figure 5.9 for three cases. The elongated structure of the streaks is clearly seen, as is the rather random nature of their relative positions. The low speed streaks are longer (in the streamwise direction) and narrower (in the spanwise) than the high speed streaks for the cases with a model, which agrees with the review by Robinson (1991). There are no qualitative differences between the results on the two grids (with a model), but there appears to be somewhat more small scale structures on the finer grid, as one would expect. Removing the model creates considerably more small scale activity due to the lower dissipation.


For all cases there are structures of similar size as the grid resolution, but this is particularly true for the case without a model. This implies that the turbulence creates the smallest scales possible, and hence that these simulations are, by nature, poorly resolved in this region. Only by refining the grid towards typical DNS resolutions ($\Delta x^+ \approx 8$, $\Delta z^+ \approx 5$ for spectral methods, which corresponds to finite-volume cell sizes of about half of that), where the molecular viscosity is strong enough to dissipate the energy, can this issue be removed.

Contours of the streamwise vorticity are shown in figure 5.10. There is a strong vortex at $z^+ \approx 1390$, $y^+ \approx 12$ (with negative vorticity). The solid lines in figure 5.9(a) show the approximate location of this vortex. Notice the high speed streaks immediately to the left (above in figure 5.9(a)) and right (below in the figure) of this vortex. The streak at $z^+ \approx 1450$ is lifted away from the wall by the vortex, whereas on the opposite side a streak is formed by the motions towards the wall (at $z^+ \approx 1350$).


This shows the interplay between the streaks and the streamwise vortices, and how numerical simulations aid in the elucidation of these mechanisms. A wall resolved LES clearly retains these dynamics, albeit not with the accuracy of DNS. The fact that these near wall motions, that clearly scale on viscous (and hence Reynolds number dependent) length scales, must be captured by wall resolved LES illustrates the high computational cost of such simulations. The objective of approximate near wall treatments is to model these dynamics instead, thereby allowing for much coarser grids in the near wall region.



(a) Base grid with model. The solid lines at $x^+ \approx 2360$ and $z^+ \approx 1390$ mark the locations of the slice and the strong vortex in figure 5.10, respectively.



(b) Base grid without model.



(c) Fine- xz grid with model.

Figure 5.9: Snapshots of the streamwise velocity fluctuation u' at $y^+ \approx 12$. Contours one standard deviation (i.e., u_{rms} , taken from the base case) above and below zero, where darker regions imply higher velocity.




Figure 5.10: Snapshot of the streamwise vorticity ω_x^+ at $x^+ \approx 2360$ for base grid with model. Contours spaced 0.1 units apart, with solid and dashed lines representing positive and negative vorticity, respectively.

5.6 Summary

The wall resolved LES results in accurate statistics and qualitatively correct instantaneous flow structures. The effect of the residual stress model is primarily to provide dissipation, since the modeled stresses are much smaller than the resolved ones throughout the channel – in fact, they are even smaller than the viscous ones, at least on average. Removal of the residual stress model changes primarily the flow in the region very near the wall, say for $y^+ \lesssim 30$, which then translates into an underpredicted mean velocity profile throughout the channel. The fact that the mean velocity gradient is decreased implies that the removal of the model increases the resolved stress more than it decreases the modeled one. This is consistent with the fact that the modeled stresses are very small, and that the main effect of the model is to provide the appropriate dissipation of resolved kinetic energy.

Chapter 6

Hybrid LES/RANS

The simulations of the previous chapter illustrate the sensitivity of the results to the resolution in the buffer layer when resolving the inner boundary layer in LES. Since the necessary grid resolution scales on viscous length scales, the computational cost becomes highly dependent on the Reynolds number – as discussed in chapter 1, the number of grid cells necessary is $\sim Re_\tau^2 \sim Re_b^{1.75}$.

One way to remove this scaling of the computational cost is to model, rather than resolve, the inner boundary layer. Resolution of the outer boundary layer requires a fixed number of cells independent of the Reynolds number, and hence the computational cost scaling when modeling the inner layer depends almost exclusively on the resolution requirements of the modeling strategy there. The strategies covered here all have costs either independent of Re or at most $\sim \log Re_\tau$.

When only resolving the outer layer motions, the grid size is based on the outer length scale δ . Since the scales in the inner layer decrease dramatically as the wall is approached, the ratio of the grid size to the inner (viscous) length scale increases as well, and hence the filtering operation in LES becomes increasingly like an ensemble average¹. Therefore, it is appropriate to view the near wall turbulence in a statistical sense. With some exceptions (see section 6.1.4, or the reviews by Cabot & Moin (1999) and Piomelli & Balaras (2002)), most approximate near wall treatments in LES proposed to date are based on this statistical view of the turbulence in the inner boundary layer.

Parts of the work reported here have been published as Larsson *et al.* (2005b), will appear as Larsson *et al.* (2006a), and have been submitted for publication as Larsson *et al.* (2006b).

6.1 Introduction and Review of Approximate Near Wall Modeling

There are many possible modeling strategies that follow the general concept outlined above, i.e. that use a statistical model in the inner layer. The three most common such approaches will be covered in the following sections, and a sketch of the grids used in

¹More practically, one could similarly say that each grid cell contains many turbulence structures near the wall.




Figure 6.1: Sketch of grids for the most common approximate near wall treatments. In wall stress models, the evolution equations are solved on the left grid. The wall shear stress is either estimated from some algebraic relations or by solving the thin boundary layer equations on an auxiliary grid (the middle grid). In both cases the flow of information is indicated by the arrows. For hybrid LES/RANS, the evolution equations are solved on the right grid, and the turbulence model switches to a RANS model near the wall.

these approaches is shown in figure 6.1. Other near wall modeling strategies that do not rely on this statistical view have also been proposed. One such method is reviewed briefly in section 6.1.4.

6.1.1 Wall Stress Models that use Algebraic Relations

One group of models uses algebraic equations to relate the wall shear stress to the instantaneous velocity somewhere in the overlap layer, and hence there is no need to solve the evolution equations below this point. These methods are analogous to the ‘wall functions’ commonly used in RANS, but some modifications are needed for their application in LES. Since the first grid point is placed in the overlap layer, at a constant value of y/δ , the computational cost of these methods is independent of the Reynolds number.

The log-law given by (5.3) is valid in the overlap layer $30 \lesssim y^+ \lesssim 0.1\delta^+$ for average quantities. The simplest algebraic wall stress model is to assume that each grid cell adjacent to the wall is coarse enough to effectively represent an ensemble average, and hence to apply (5.3) directly to each cell in an instantaneous fashion. This procedure is common in atmospheric flows (Piomelli & Balaras, 2002), and was used by Krajnovic & Davidson (2003) to compute the flow around a bluff body.

The pioneering LES by Schumann (1975) introduced a novel boundary treatment that is more physically appealing. Instead of relying on the assumption that each grid cell effectively represents an ensemble average of the turbulence, it relies on the weaker assumption that the wall shear stress and the velocity in the cell adjacent to the wall are in phase as

$$\frac{u_i}{\langle u_i \rangle} = \frac{\tau_{w,i}}{\langle \tau_{w,i} \rangle}, \quad (6.1)$$

where $\tau_{w,i} = \nu \partial u_i / \partial y|_w$ and $\langle \cdot \rangle$ is some type of averaging operation. Taking this to be an average over the wall parallel directions, both u_i and $\langle u_i \rangle$ are available from the instantaneous velocity field. Solving the log-law (5.3) for $\langle \tau_{w,i} \rangle$ then allows for the calculation of the instantaneous wall shear stress by (6.1) (Grötzbach, 1987).

Piomelli *et al.* (1989) suggested improvements to this boundary condition, but by and large the original formulation by Schumann (1975) still constitutes a rather accurate boundary condition for the flows it was designed for, i.e. attached boundary layers in approximate equilibrium. Worse performance for separated and non-equilibrium flows has been reported by Balaras *et al.* (1996) and Wang & Moin (2002).

Wu & Squires (1998) applied relation (6.1) in a slightly different way in their LES of the boundary layer over a swept bump. Instead of computing the average wall shear stress from the log-law, they instead took $\langle \tau_{w,i} \rangle$ from an *a priori* RANS calculation. This approach is then no longer predictive, but it does remove the assumption of the log-law being valid on average.

6.1.2 Wall Stress Models that use Auxiliary Grids

A second group of models solves the LES evolution equations only in the outer boundary layer, and then solves a simplified set of equations (typically the thin boundary layer equations) using a RANS turbulence model on an auxiliary grid all the way down to the wall. These auxiliary equations relate the instantaneous velocity in the overlap layer from the LES (which is given as a boundary condition) to the wall shear stress, which is then provided as a boundary condition back to the LES. While the approach is similar to the algebraic wall stress models described above in that it establishes a relation between an outer layer velocity and the wall shear stress, more physics are included due to the more elaborate near wall treatment (Balaras *et al.*, 1996; Wang & Moin, 2002).

This approach was originally proposed by Balaras *et al.* (1996), who applied it to plane channel flow, rotating channel flow, and the flow in a square duct. The results for the plane channel flow were similar to results obtained by the Schumann (1975) boundary condition, but they found improved results for the more complex flows studied. For these flows the mean velocity profile does not obey the standard log-law, and hence the use of a near wall treatment that does not enforce such a log-law leads to improved results. Wang & Moin (2002) applied a similar method to the flow over a hydrofoil, and reported improved results compared to algebraic models, especially near the small separation zone.

Due to the RANS treatment, the auxiliary grid can be arbitrarily coarse in the wall parallel directions, but must resolve the wall normal direction properly. If the first grid point (in the LES grid) is placed at a constant y/δ (in the overlap layer),

then the auxiliary grid must extend up to the same y/δ . The first cell in the auxiliary grid must have a constant size in viscous units, typically $\Delta y_{\text{aux},1}^+ \approx 1$. With geometric grid stretching, it is then straightforward to show that the number of cells in the wall normal direction is $\sim \log Re_\tau$ (Nikitin *et al.*, 2000), and thus the computational cost is not independent of the Reynolds number. The logarithmic scaling, however, is still an enormous improvement compared to wall resolved LES. Note also that there is no need to solve a Poisson equation (the most expensive part of LES algorithms) for the thin boundary layer equations, and thus the cost increase is rather low.

While this approach has performed well for a reasonably wide range of flows, the use of the thin boundary layer approximation prevents these methods from being universally applicable. It is conceivable that future developments will solve this problem, and hence this approach is a promising one deserving future attention.

6.1.3 Models that Adapt the Residual Stress Model

In a third group of models, the full evolution equations are solved all the way down to wall, with adaptation of the turbulence model in the inner layer. This group of models in some ways represent the most straightforward application of the sketched boundary layer in figure 1.2, in that the model is changed as the wall is approached. Following the fact that the near wall turbulence can be viewed in a statistical sense, RANS models (usually based on the eddy viscosity hypothesis (2.11)) are commonly used in this region. There are many specific models in this group, which typically go by names such as ‘hybrid LES/RANS’ (Davidson & Dahlström, 2004; Temmerman *et al.*, 2005) or ‘detached eddy simulation (DES)’ (Nikitin *et al.*, 2000). The computational cost of this approach is higher than the ones discussed previously due to the need to solve the full evolution equations (including the Poisson equation) all the way down to the wall. With suitable grid stretching, however, the cost only increases logarithmically with the Reynolds number (as discussed in the previous section). In theory, LES/RANS could provide accurate results for a wide range of flows, the range of applicability being determined by the specific turbulence models used. In practice, however, most studies have shown that while the results are vastly improved compared to simulations without any near wall modeling (on similarly coarse grids), the skin friction is consistently underpredicted by around 10-15% in attached boundary layers (Nikitin *et al.*, 2000; Piomelli *et al.*, 2003). The reason for this underprediction is that an artificial buffer layer develops between the RANS and LES regions in the domain, where the turbulence ‘transitions’ from being mainly modeled to mainly resolved. This transition is accompanied by a shift in the mean velocity profile to a higher value, which results in a too low skin friction. While a detailed discussion will be given later, the velocity shift can be seen in figure 6.5(a) (the curve denoted by circles, around $y^+ \approx 200$).

Baggett (1998) performed LES on coarse grids with and without near wall RANS modeling, and found unphysical streak-like structures in the near wall region. These ‘superstreaks’ had spanwise dimensions of two grid points, overly large streamwise dimensions, and were associated with highly anisotropic resolved motions (i.e., too high $\overline{u'u'}$ and too low $\overline{v'v'}$ and $\overline{w'w'}$). He further argued that these superstreaks cause a de-correlation of the streamwise and wall normal velocity fluctuations, which results in

a too low resolved shear stress and, consequently, a shift in the velocity profile. Hamba (2003) obtained the velocity shift using both LES/RANS as discussed here and with the modeling regions swapped (i.e., LES near the wall and RANS in the core of the channel). Several studies (e.g. Piomelli *et al.*, 2003; Davidson & Dahlström, 2004; Temmerman *et al.*, 2005) have experimented with the location of the modeling interface, all resulting in minimal change of the size of the velocity shift. Similarly, experiments with different interface conditions (e.g. Davidson & Dahlström, 2004; Tucker & Davidson, 2004) and different turbulence models (e.g. Nikitin *et al.*, 2000; Tucker & Davidson, 2004; Temmerman *et al.*, 2005) have failed to remove the velocity shift. Furthermore, grid refinements in section 6.2.1 will show that the solution reaches a grid independent state that still has a velocity shift. The main conclusion of these studies, then, is that the velocity shift is a rather robust feature of LES/RANS, with little dependence on the details of the specific models, interface locations or conditions, or grid resolutions used.

In light of this robustness, there have been recent attempts to remove the velocity shift by additional forcing of the momentum equations. The role of this forcing can be seen in different ways: as a way to provide better boundary conditions to the LES region, as a way to account for the backscatter of energy from the unresolved motions, or simply as a way to help trigger resolved motions with the correct physical structure. Piomelli *et al.* (2003) viewed the forcing as a way to account for backscatter, and argued that it, therefore, should have length and time scales on the order of the grid size and time step, respectively. Hence, their forcing field was essentially white noise, and they reported a lower velocity shift and (by visual inspection) that the superstreaks were broken up. Batten *et al.* (2004) and Davidson & Billson (2006) proposed stochastic models that generate forcing fields with predefined spectra and anisotropy, and reported quicker transition to resolved turbulence (Batten *et al.*, 2004) and a lower velocity shift (Davidson & Billson, 2006), respectively. Davidson & Dahlström (2004) took the forcing field from a DNS database in order to get a field with physical structures. While they reported a lower velocity shift, their particular choice of forcing term acted as a source of mean momentum, thereby destroying the conservation of momentum inherent in the Navier-Stokes equations.

6.1.4 Wall Stress Models based on Control Theory

The approximate near wall treatments discussed up to this point all view the inner layer in a statistical sense, and attempt to model the physics of this layer in different ways. A radically different approach was taken by Nicoud *et al.* (2001), who instead posed the problem in a completely mathematical fashion. In essence, they viewed it as a control problem, with the resolved velocity field as the state variables and the instantaneous wall shear stresses as the control variables. They then defined a cost function as the mean square difference between the plane-averaged velocities $\langle u \rangle$ and $\langle w \rangle$ and a reference state given by the log-law (5.3) (and $w = 0$). They then used control theory to determine the wall shear stresses such that this cost function was minimized.

This method uses grids that are coarse in all directions, and thus there is a certain similarity to the algebraic wall stress models discussed in section 6.1.1. The key dif-

ference, however, is the lack of any physical reasoning in the method by Nicoud *et al.* (2001) – it is a completely mathematical approach, that requires no knowledge about the physical processes.

One benefit of the approach based on control theory is the fact that numerical errors are accounted for implicitly, since the cost function measures the error in the discrete velocity field. The biggest drawback is that the method is not predictive, in that it drives the flow towards the specified reference solution. A second drawback is the high computational cost of solving the control problem – Nicoud *et al.* (2001) reported a 20 times higher cost compared to a ‘non-controlled’ LES. Later, Templeton *et al.* (2006) successfully reduced the cost of the approach to only 2-3 times higher than a non-controlled LES.

These drawbacks aside, methods based on control theory appear promising, and should be considered in future work.

6.1.5 Objectives

The hybrid LES/RANS approach is chosen in the present study, primarily due to the theoretically large range of applicability. The main focus here is on the artificial buffer layer that develops around the modeling interface, both in terms of insight into its physics and dynamics, and in terms of finding ways to remove it.

There remains some open questions about the artificial buffer layer and the super-streaks in LES/RANS. For example, while it is clear that a too low level of resolved shear stress causes the velocity shift, the precise cause of the low shear stress is not as clear. The relationship between the low shear stress and the superstreaks also needs clarification. Piomelli *et al.* (2003), for example, suggested that the connection is weaker than conjectured by Baggett (1998) (as discussed above). Another question is to what extent the dynamics of the artificial buffer layer are similar to those of the ‘true’ buffer layer (i.e., $5 \lesssim y^+ \lesssim 30$, see figure 1.2). The first objective of the present study is to gain more insight into the artificial buffer layer by attempting to answer these questions.

A second objective is to study the effects of forcing more closely. With one exception (to the author’s knowledge, only Davidson & Dahlström, 2004), the studies discussed above all examine the effects of forcing using stochastic models. This makes the interpretation difficult, since the concept of forcing can not be evaluated separately from the particular forcing model used. By using forcing fields with more of the correct physics, and fewer modeling assumptions, the uncertainties in the interpretation are minimized in the present study. Hence, the present use of forcing is not intended as a practically useful model, but rather as a tool for basic insight that can guide future developments of generally applicable forcing models.

6.2 Methodology

Adoption of the eddy viscosity hypothesis (2.11) means that the only distinction between the RANS- and LES-regions is in the model for ν_t . Many previous investigations (e.g. Nikitin *et al.*, 2000; Hamba, 2003; Davidson & Dahlström, 2004; Temmerman *et al.*, 2005) into the LES/RANS concept have found qualitatively similar results using

a reasonably wide array of eddy viscosity turbulence models, indicating that details of the models themselves are of secondary importance. For this reason, simple algebraic models are used in this work.

In the outer layer, the Smagorinsky model given by (2.12) is used. To reduce the complexity of the problem, the dynamic procedure is not used, and thus the Smagorinsky coefficient C is fixed. Initial tests showed that a value of $C = 0.13^2$ gives accurate results in the outer layer.

In the inner layer, a simple mixing length model is used. Such a mixing length model can be derived from the fact that the relevant velocity and length scales in the (inviscid) overlap layer are u_τ and y , respectively (see for example Wilcox, 2000, who presents an interesting discussion about the eddy viscosity hypothesis and the mixing length). Here a slightly different, and less rigorous, approach is used.

For a pure RANS solution, the resolved shear stress $-\overline{u'v'} = 0$. The re-arranged shear stress balance (2.33) then yields

$$\frac{\partial \bar{u}}{\partial y} \approx \frac{u_\tau^2}{\nu + \nu_{\text{rans}}} , \quad y \ll \delta . \quad (6.2)$$

Differentiation of the log-law (5.3) yields $\partial \bar{u} / \partial y = u_\tau / \kappa y$. Inserting this into (6.2), i.e. requiring the RANS model to return a logarithmic velocity profile for $y^+ \gtrsim 30$, yields

$$\frac{u_\tau}{\kappa y} \approx \frac{u_\tau^2}{\nu_{\text{rans}}} \Rightarrow \nu_{\text{rans}} \approx \kappa u_\tau y , \quad 30 \lesssim y^+ \ll \delta^+ . \quad (6.3)$$

The reason for choosing this particular derivation of the mixing length model can now be seen, since it brings out the connection to the log-law – both the functional form of ν_{rans} , and that the adjustable constant κ is the inverse logarithmic slope.

It is easy to see that this expression for ν_{rans} yields too high values in the viscous region $y^+ \lesssim 30$. In the viscous sublayer ($y^+ \lesssim 5$), for example, the mean velocity gradient should be $\partial \bar{u} / \partial y \approx u_\tau^2 / \nu$ from the definition of the wall shear stress. Inserting this into relation (6.2) shows that ν_{rans} must vanish near the wall in order to yield the correct mean velocity profile. This problem can be solved by including a damping function for the eddy viscosity². The most common such damping function is the one by van Driest (c.f. Wilcox, 2000), which yields

$$\nu_{\text{rans}} = \kappa u_\tau y D(y^+) , \quad D(y^+) = \{1 - \exp(-y^+ / A^+)\}^2 , \quad A^+ = 19 . \quad (6.4)$$

The final issue is the merging or interfacing of the LES and RANS eddy viscosities (and, hence, of τ_{ij} by the different models). Various proposals for this interfacing have appeared in the literature, some which enforce continuity of ν_t (e.g. Nikitin *et al.*, 2000; Temmerman *et al.*, 2005) and some which do not (e.g. Davidson & Peng, 2003). Since the results are similar in these studies, and since the exact interface conditions appear to have a marginal effect on the appearance of an artificial buffer layer, the eddy viscosities are simply blended here as (a similar blending was used by, e.g., Baggett, 1998)

$$\nu_t = f(y) \nu_{\text{les}} + (1 - f(y)) \nu_{\text{rans}} \quad (6.5)$$

²It is actually the mixing length κy that is damped, but this is a point of minor importance here.




Figure 6.2: Blended eddy-viscosity for LES/RANS simulations. \times : $\bar{\nu}_{les}$ from the Smagorinsky model; $+$: $\bar{\nu}_{rans}$ from the mixing length model; \circ : blended $\bar{\nu}_t$ from (6.5); $-\cdot-$: location of the interface $y_{int}^+ \approx 120$. The width parameter $\Delta_{int}^+ \approx 240$.

using the (arbitrarily chosen) blending function

$$f(y) = \frac{1}{2} \left\{ 1 + \tanh \left(4.595 \frac{y - y_{int}}{\Delta_{int}} \right) \right\}. \quad (6.6)$$

With this blending, the eddy viscosity smoothly transitions from a (high) RANS value near the wall to a (lower) LES value away from the wall. The midpoint of the transition, termed the ‘interface’, is located at y_{int} , and the total width of the transition is roughly Δ_{int} .

A sample distribution of the eddy viscosity is shown in figure 6.2. The Smagorinsky ν_{les} is more or less constant throughout the channel, whereas ν_{rans} grows linearly outside of where the van Driest damping function is active. At the wall, $\nu_t = 0$ as required by the boundary condition on τ_{ij} (and as required for the correct, linear, velocity profile in the viscous sublayer). The blending parameters used in the figure result in a smooth transition of ν_t . While this smoothness is not necessary on physical grounds, it does result in a mean velocity profile that is continuously differentiable everywhere.

The numerical method used in chapter 5 is used here as well, and the grids use the same stretching function (5.5). The domain sizes are rather conventional, typically $L_x/\delta \approx 6$ and $L_z/\delta \approx 2$, following the arguments in chapter 5.

6.2.1 Resolution Requirements

While the influence of the grid resolution on the results of wall resolved LES has been studied rather extensively in the literature, less is known about the resolution requirements of hybrid LES/RANS. First of all, the wall normal resolution is similar in both LES and RANS, and hence the grids used here have similar stretching and grid sizes as those used for the wall resolved LES in chapter 5. In the wall parallel directions, the situation is different: in the RANS region there is, in theory, no lower limit on the required

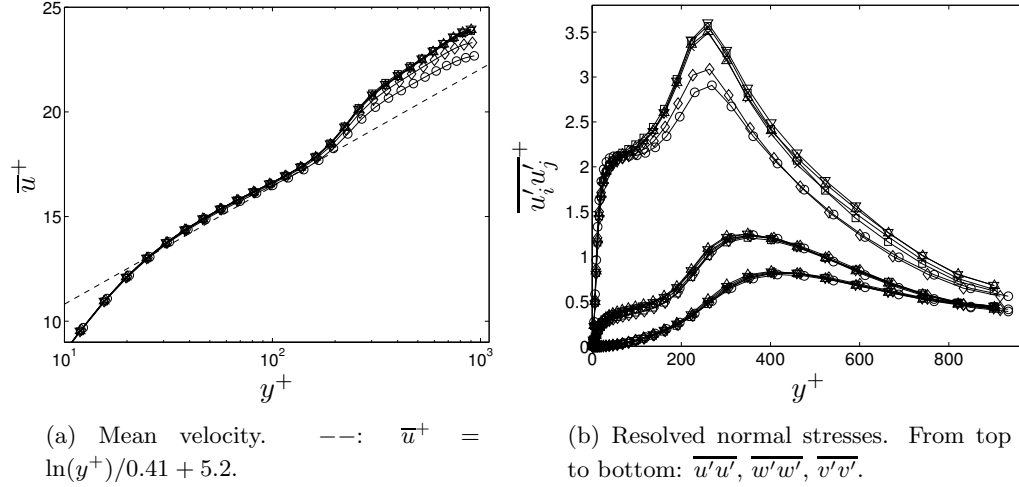


Figure 6.3: Influence of the streamwise grid resolution. \circ : $N_x = 24$; \diamond : $N_x = 32$; \square : $N_x = 48$; \times : $N_x = 64$; ∇ : $N_x = 96$; \triangle : $N_x = 128$.

resolution, whereas the resolution (in outer units) in the LES region needs to be similar to that used in the wall resolved LES (one should note that the resolution of the LES region scales with δ since this region is inviscid in LES/RANS). There are, however, some complicating factors. First, the integral length scale grows approximately linearly with y (c.f. Pope, 2000), and hence the least resolved part of the LES region will be immediately above the modeling interface. This situation does not correspond directly to anything in a wall resolved LES, and the grid requirements arrived at in chapter 5 need not necessarily apply here. Second, while the inner layer in a wall resolved LES generates small scale motions that to some extent propagate outwards, the RANS treatment of the inner layer in LES/RANS limits such small scale activity. More accurately, while there may be small scales (that are triggered by the outer layer LES), their sizes and dynamics need not be dictated by physics.

For these reasons the resolution requirements in the wall parallel directions are investigated here. The Reynolds number based on the bulk velocity is $Re_b = 19800$, which yields $Re_\tau \approx 1000$. The domain size is $L_x/\delta = 6$ and $L_z/\delta = 2$, and the blending parameters are set to $y_{\text{int}}/\delta = 0.15$ and $\Delta_{\text{int}}/\delta = 0.30$, which yields an interface at $y_{\text{int}}^+ \approx 150$.

Streamwise Resolution

The influence of the streamwise resolution on the mean velocity profile and the resolved normal stresses is shown in figure 6.3. The spanwise resolution is fixed at $\Delta_z/\delta \approx 0.031$ ($N_z = 64$), which is similar to the finest grid used for the wall resolved LES in the previous chapter. The number of cells in the streamwise direction is taken from the Bulirsch sequence $N_x \in \{24, 32, 48, 64, 96, 128\}$, which yields $\Delta x/\delta \approx 0.250 \rightarrow 0.047$, where the finest resolution is, again, similar to the finest grid used in the previous

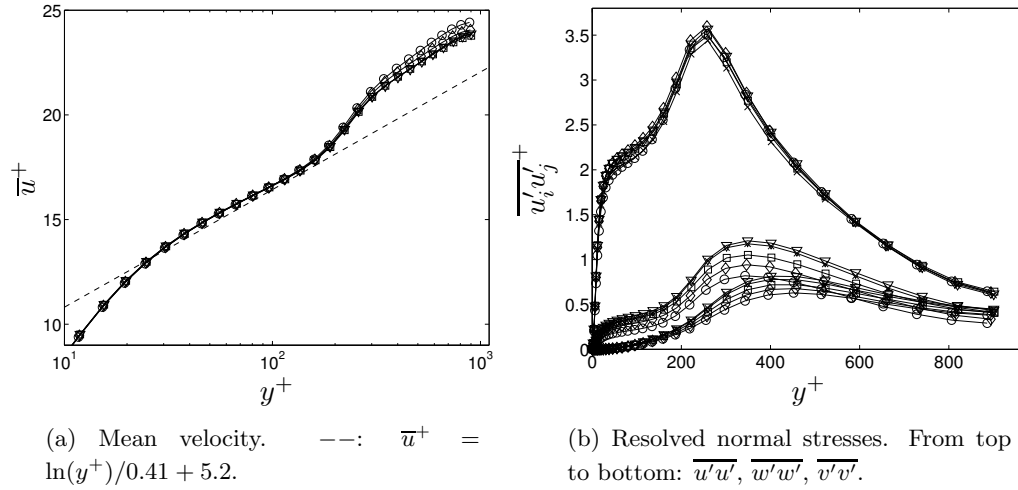


Figure 6.4: Influence of the spanwise grid resolution. \circ : $N_z = 16$; \diamond : $N_z = 24$; \square : $N_z = 32$; \times : $N_z = 48$; ∇ : $N_z = 64$.

chapter (in outer units).

To within the statistical convergence error, the solutions become (statistically) grid independent for $N_x \geq 48$, corresponding to $\Delta x/\delta \lesssim 0.125$, for both the first- and second-order statistics.

Spanwise Resolution

The influence of the spanwise grid resolution is shown in figure 6.4, where $N_x = 48$ is used based on the previous finding. Again, the number of cells is taken from a Bulirsch sequence $N_z \in \{16, 24, 32, 48, 64\}$, which yields $\Delta z/\delta \approx 0.125 \rightarrow 0.031$.

The results approach the grid independent solution in a different manner compared to when the streamwise resolution was studied. The mean velocity profile now approaches the grid independent state from above (albeit hard to see in the figure), rather than from below as in figure 6.3. Also, while varying Δx mainly caused changes in the streamwise resolved normal stress $\overline{u'u'}$, varying Δz mainly affects the spanwise and wall normal components. From this can be inferred that the spanwise resolution affects the quasi-streamwise vortices, while the streamwise resolution primarily affects the streaks. The only puzzling part of this is the fact that $\overline{u'u'}$ increases with decreasing Δx – one would perhaps expect a larger Δx to generate longer, and therefore stronger, streaks instead.

The mean velocity profile becomes grid independent for $N_z \geq 32$ to within the statistical convergence error, corresponding to $\Delta z/\delta \lesssim 0.063$. The resolved normal stresses appear to be slightly more sensitive, and approach grid independence for $N_z \geq 48$ ($\Delta z/\delta \lesssim 0.042$) instead.

Temporal Resolution

The influence of the size of the time step is tested by successive refinement of Δt . Due to the explicit treatment of the momentum equations, the time step must satisfy the stability constraints. Lowering the time step further results in only minor differences (not shown), with independent statistics for $\max\{\text{CFL}\} \lesssim 1.0$.

6.3 Results Without Forcing

The simulations used to study the LES/RANS concept in more detail are performed at a higher Reynolds number of $Re_b \equiv U_b \delta / \nu = 43500$, which yields a Reynolds number based on the friction velocity $Re_\tau \equiv \delta^+ \approx 2000$. This is done to ensure that the modeling interface can be placed in the overlap layer, since effects of the core flow should be minimized in order to simplify the interpretation. The blending parameters are set to $y_{\text{int}}/\delta = 0.06$ and $\Delta_{\text{int}}/\delta = 0.12$, which yields an interface located at $y_{\text{int}}^+ \approx 120$.

The domain size used here is approximately the same as that used above, $L_x/\delta \approx 6.333$ and $L_z/\delta \approx 2.253$. The reason for choosing these rather odd values has to do with the introduction of the forcing field in section 6.4, and a full explanation is postponed until then.

The number of cells is $(48, 80, 36)$ in the streamwise, wall normal, and spanwise directions, respectively, which yields cell sizes of $(\Delta x, \Delta y, \Delta z)/\delta \approx (0.13, 5.9 \cdot 10^{-4} \rightarrow 0.080, 0.063)$. This is in rough agreement with the grid independent resolution found above. In viscous units, the cell sizes are $(\Delta x^+, \Delta y^+, \Delta z^+) \approx (260, 1.2 \rightarrow 160, 130)$. This is not exceptionally coarse by any means, but certainly too coarse for wall resolved LES. Thus, this represents a good test case for approximate near wall modeling.

The mean velocity profile is shown in figure 6.5(a). For comparison, the profile from a simulation on the same grid without any near wall modeling (i.e., with $f = 1$ in (6.5)) using the dynamic procedure is shown as well. The skin friction is severely underpredicted in the latter case (by about 40%), showing that some form of additional near wall modeling is necessary for LES on coarse grids. The LES/RANS profile is considerably better, but the velocity shifts to a higher value immediately outside of the interface. This velocity shift causes about 10% underprediction of the skin friction coefficient.

6.3.1 The Velocity Shift

To bring out the variations of the mean velocity gradient more clearly, one can consider the quantity

$$\gamma = y^+ \frac{\partial \bar{u}^+}{\partial y^+} . \quad (6.7)$$

If the velocity profile follows the log-law $\bar{u}^+ = \ln(y^+)/\kappa + B$, then $\gamma = 1/\kappa$. This quantity is plotted in figure 6.6(a).

Consider first the DNS result by del Alamo *et al.* (2004). There is a peak in γ around $y^+ \approx 10$ in the buffer layer, followed by a decrease up to $y^+ \approx 70$. From there, γ again increases in the wake region. There is no truly logarithmic region (which would have γ

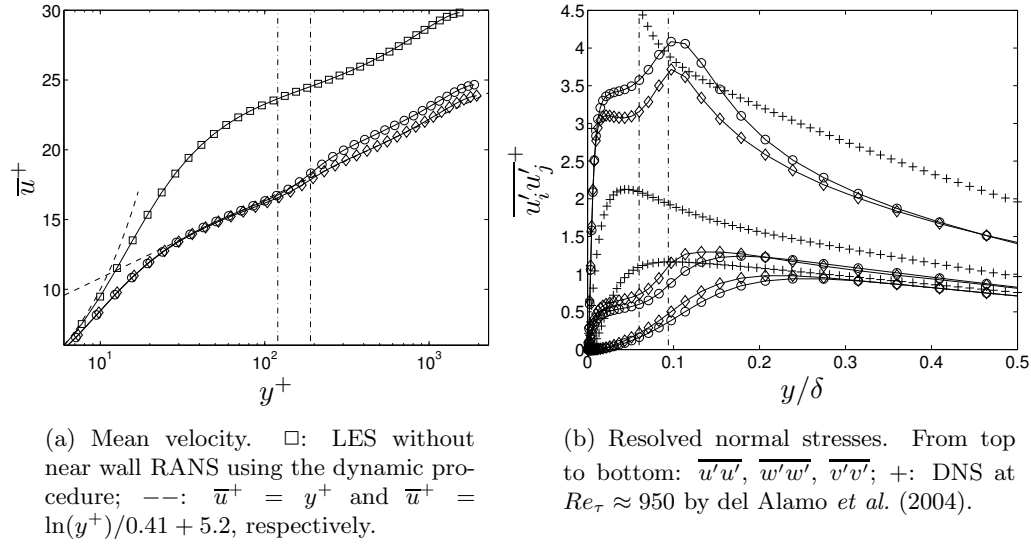


Figure 6.5: Mean velocity and resolved normal stresses. \circ : LES/RANS without forcing; \diamond : LES/RANS with forcing; $---$: locations of the interface ($y_{\text{int}}/\delta = 0.06$) and that of maximum forcing ($y_f/\delta = 0.094$), respectively.

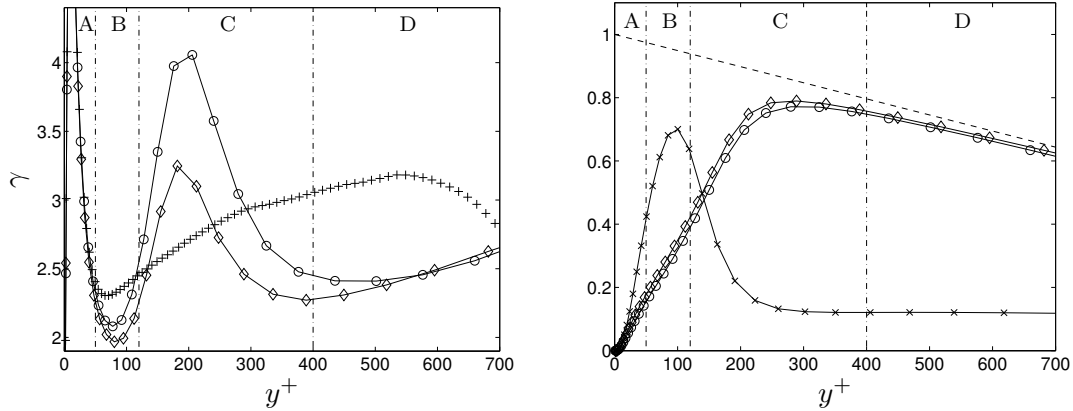
= constant), and hence it is hard to distinguish where the buffer layer ends and where the wake in the core of the channel begins. Here, the local minimum of γ is loosely considered to mark the beginning of the wake region.

The LES/RANS result has a different shape of the γ profile, in that there is a second peak outside of the interface around $y^+ \approx 190$. Outside of this peak γ first decreases and then increases again in the wake region. Qualitatively, there appears to be four distinct regions in the solution: an accurate RANS region ($y^+ \lesssim 50$, labeled A in figure 6.6); a region where the velocity gradient is too small ($50 \lesssim y^+ \lesssim 120$, labeled B); a region where the velocity gradient is too large ($120 \lesssim y^+ \lesssim 400$, labeled C); and a LES wake region ($y^+ \gtrsim 400$, labeled D). Here the boundary between regions C and D comes from the same loose definition given above.

This behaviour of the velocity gradient is most easily understood by considering the balance of the resolved and modeled shear stresses. When considering only the region away from the wall (i.e., the inviscid part), the re-arranged shear stress balance (2.33) yields

$$\frac{\partial \bar{u}}{\partial y} \approx \frac{u_\tau^2(1 - y/\delta) - (-\overline{u'v'})}{\bar{\nu}_t}. \quad (6.8)$$

The regions identifying the velocity gradient behaviour discussed above are marked in figure 6.6(b), which shows the resolved shear stress $-\overline{u'v'}$ and the eddy viscosity $\bar{\nu}_t$. Immediately outside of the interface, the resolved shear stress is heavily damped by the presence of the RANS region, which together with a small $\bar{\nu}_t$ results in a too high velocity gradient as a consequence of relation (6.8). The resolved shear stress reaches a typical LES level at $y^+ \approx 400$, which is consistent with the approximate beginning of



(a) γ defined by (6.7). +: DNS at $Re_\tau \approx 950$ by del Alamo *et al.* (2004).

(b) Resolved shear stress $-\overline{u'v'}^+$ (symbols as in main caption) and scaled eddy viscosity $0.025 \cdot \overline{\nu}_t^+$ (x, from the case without forcing). --: $1 - y^+/\delta^+$. Note that the eddy viscosity for the case with forcing (not shown) is similar to the one shown.

Figure 6.6: A closer look at the velocity shift. o: LES/RANS without forcing; d: LES/RANS with forcing; - - -: boundaries of zones A-D discussed in text.

the LES wake region.

The opposite is true immediately below the interface, in that $\overline{\nu}_t$ is large but $-\overline{u'v'}$ is non-negligible. This results in a lower velocity gradient, but the net effect here on the velocity profile is smaller.

The resolved normal stresses are shown in figure 6.5(b). In the core of the channel the levels of $\overline{u'u'}$ and $\overline{w'w'}$ are below the corresponding DNS results of del Alamo *et al.* (2004), whereas the wall normal component $\overline{v'v'}$ has a similar level as in the DNS. This is qualitatively correct, since only the largest scales are included in the LES/RANS. The lower levels could also be due to the relatively small domain of the LES/RANS, which limits the existence of the very large scales, or it could be due to the difference in Reynolds number. The streamwise and spanwise components increase immediately outside of the interface, and reach maximums at $y/\delta \approx 0.1$ and 0.2 , respectively. They are both somewhat damped in the RANS region, but still reach about one third of their corresponding DNS levels. Due to the kinematic boundary condition, the wall normal component has its maximum farther out at $y/\delta \approx 0.25$. Note that the order in which the normal components reach their maximums is the same as outside the true buffer layer in the DNS.

Overall, figures 6.5 and 6.6 indicate that the velocity field reaches some form of ‘LES equilibrium’ for $y^+ \gtrsim 400$ ($y/\delta \gtrsim 0.2$), where the solution bears the hallmarks of an accurate LES field (correct $\partial\overline{u}/\partial y$, all resolved stress components close to their respective DNS levels, etc). The region $50 \lesssim y^+ \lesssim 400$, on the other hand, constitutes a form of artificial buffer layer, where the turbulence changes from being mainly modeled to being mainly resolved. Again, there is a similarity between the true and artificial

buffer layers in that they are associated with a change between different types of shear stresses: to resolved from viscous and modeled, respectively.

At this point one can compare the present results to those of Temmerman *et al.* (2005). They use different turbulence models and a different way of interfacing these models, but their $\bar{\nu}_t$ profile is similar to the present one in figures 6.2 and 6.6(b). It is then not surprising that both studies report maximum mean velocity gradients at the same locations: outside of the peak in $\bar{\nu}_t$, where the eddy viscosity first reaches the lower value given by the LES model.

In light of relation (6.8), one might wonder whether it is possible to adjust ν_t such that the resulting mean velocity profile becomes correct. Mason & Thomson (1992) attempted just that. They first ran a simulation with a model for ν_t that resulted in a velocity shift similar to the one found here. They then used relation (6.8) to compute the ‘ideal’ eddy viscosity profile that would give the correct mean velocity gradient everywhere, and used this ‘ideal’ eddy viscosity in a second simulation. The results, however, were disappointing: the artificial buffer layer simply moved outwards, and the size of the velocity shift remained essentially the same.

The reason for this is the direct effect on the dissipation of the resolved motions by the eddy viscosity model (as discussed in chapters 2 and 5). An increased ν_t further reduces the resolved shear stress $-\overline{u'v'}$, which simply moves the artificial buffer layer farther away from the wall. Even if it were possible to find an ‘ideal’ ν_t profile, it would probably be hard to do so in an *a priori* manner.

It then seems likely that a successful solution to this problem must allow for the simultaneous prediction of appropriate levels of both the dissipation and the modeled shear stress. With these phenomena being distinctly different from each other, and unlikely to be accurately predicted by any one model, one might conjecture that successful approaches need at least two ‘degrees of freedom’ (in the sense of different model terms) in the model. One possibility would be to move away from the eddy viscosity hypothesis, either by using a non-linear constitutive relation (c.f. Wilcox, 2000) or a model with ‘scale-similarity’ type terms (Bardina *et al.*, 1980). A second possibility is to increase the resolved shear stress through additional forcing in the momentum equations.

6.4 Results With Forcing

The goal when using additional forcing is to increase the resolved shear stress in such a way that the net effects of the artificial buffer layer are minimized. With the merging of two different turbulence models, some form of transitional region will probably always persist, so a reasonable goal is to aim for accurate results in the LES region some distance away from the interface, and for accurate global quantities like the skin friction coefficient. In the current results, this implies that the shift in the mean velocity profile should be removed.

As discussed in the introduction to this chapter, the objective of the present work is to investigate the effects of forcing on a fairly fundamental level. In order to introduce as few modeling assumptions as possible, the forcing fields are taken from a reference simulation, rather than being generated by some stochastic model as has been the case

in several previous studies (Piomelli *et al.*, 2003; Batten *et al.*, 2004; Davidson & Billson, 2006). This ensures that the forcing fields have the correct length and time scales, phase relationships, spectra, and so on.

6.4.1 Forcing Database

Ideally, the forcing field should be taken from a DNS using the same domain and Reynolds number as the LES/RANS simulation, since this would create the most accurate forcing field possible. Such a DNS would, however, be prohibitively expensive, and hence the wall resolved LES of chapter 5 (using the base grid) is used instead. To justify this, there are a few points that need to be considered.

First, for the LES/RANS, γ has a peak at $y^+ \approx 190$ (figure 6.6(a)). Together with relation (6.8), this suggests that the forcing should be added mainly around this location. Second, the coarseness of the LES/RANS grid in the streamwise and spanwise directions means that only the largest scales of the forcing field taken from the reference simulation can be resolved. These large scales can be captured accurately with a wall resolved LES, especially at the locations relatively far from the wall ($y^+ \sim 190$) that are of interest. The forcing fields can then be taken from a wall resolved LES, rather than a more expensive DNS.

The difference in Reynolds number between the forcing (from the wall resolved LES) and the LES/RANS means that the fields are not perfectly compatible. Apart from the Reynolds number dependence of most statistics, one also notes that the region $y^+ \sim 190$ is located in the overlap layer in the LES/RANS, whereas it is in the wake region in the wall resolved LES. The forcing fields are, nevertheless, considered to be more accurate than what could have been achieved by any stochastic model.

The turbulence intensities in the wall resolved LES are about 5-8% at $y^+ \approx 190$, which means that Taylor's 'frozen turbulence' hypothesis is approximately valid in this region (c.f. Pope, 2000). Using this to limit the size of the database, the velocity field is stored at one streamwise location only, as $u_{i,\text{db}}^+(y^+, z^+, t^+)$ in viscous units. The true forcing (per unit mass) in an Eulerian frame is simply the acceleration $\partial u_i / \partial t$, and hence the database is differentiated in time. To avoid aliasing when adding the forcing to the much coarser LES/RANS grid, the database is filtered in the temporal and spanwise directions. This yields

$$a_{i,\text{f}}^+ = G_t * G_z * \frac{\partial u_{i,\text{db}}^+}{\partial t^+}, \quad (6.9)$$

where G_t and G_z are filter kernels. The filtering and temporal differentiation are done using fast Fourier transforms (FFT), and the filters are of the spectral cut-off type. Finally, the forcing field is scaled and windowed in the wall normal direction, and Taylor's hypothesis is used to get the streamwise variation via the convection velocity U_{conv}^+ . This yields

$$a_i^+(x^+, y^+, z^+, t^+) = A W(y^+) a_{i,\text{f}}^+ \left(y^+, z^+, t^+ - \frac{x^+}{U_{\text{conv}}^+} \right), \quad (6.10)$$

where A is an amplitude parameter and $W(y^+)$ is a wall normal window defined by

$$W(y^+) = 2 \frac{(y^+/y_f^+)^n}{1 + (y^+/y_f^+)^{n-m}}, \quad (n, m) = (8, -5.5). \quad (6.11)$$

The parameter y_f determines the location around which the forcing is centered. Some previous studies (Piomelli *et al.*, 2003; Davidson & Dahlström, 2004; Davidson & Billson, 2006) centered the forcing around the interface. Here, it is assumed that the maximum increase of the resolved shear stress will occur at the location of maximum forcing. This assumption, together with relation (6.8) and the fact that γ peaks at $y^+ \approx 190$ (figure 6.6(a)), instead implies that $y_f/\delta = 0.094$ ($y_f^+ \approx 190$) is a more proper choice.

The resulting velocity fields are relatively insensitive to the width of the window (i.e., to the values of the parameters n and m in (6.11)), which is consistent with the findings of Piomelli *et al.* (2003) who reported no significant differences between an exponential window and one similar to (6.11). The results are, however, much more sensitive to the value of the amplitude parameter A . Large sensitivity to the amplitude was also reported by both Piomelli *et al.* (2003) and Davidson & Dahlström (2004). In the following, the ‘optimal’ amplitude A_{opt} was determined by trial-and-error, and the results shown are for that case. A discussion about the sensitivity is postponed until a later section, along with a proposed method to compute the value of A_{opt} using a simple control algorithm. The forcing is finally added as a source term to the momentum equations (2.22) in the dimensional form $a_i = a_i^+ u_\tau^3 / \nu$.

There are a few subtleties associated with the introduction of the forcing into the LES/RANS. First, in order to preserve the spatial structure of the forcing field, the convection velocity U_{conv}^+ in (6.10) must remain constant in space and time³. Secondly, the forcing field is not necessarily periodic on the LES/RANS domain, which may lead to a redistribution of energy among the modes (Fourier components). This is a very minor point in general, but since one purpose of the present study is to look at the effects of individual modes of forcing, the redistribution could possibly lead to erroneous conclusions. For this reason, simple steps are taken to ensure periodicity of the forcing field in this thesis.

Ensuring Periodicity of the Forcing Field

The forcing field a_i^+ is periodic in the spanwise direction and assumed to be periodic in the temporal one on the domain of the database, i.e. on $L_{z,\text{db}}^+ \approx 1519$ and $T_{\text{db}}^+ \approx 6537$. The spanwise wavelength of the modes is $\lambda_{z,l}^+ \equiv L_{z,\text{db}}^+/l$, where l is the mode number (integer). For a mode to be periodic on the LES/RANS domain size L_z , one must have

$$\alpha_l \lambda_{z,l}^+ = \frac{L_z u_\tau}{\nu} \Rightarrow \alpha_l \frac{L_{z,\text{db}}^+}{l} = \frac{L_z}{\delta} \delta^+ \quad (6.12)$$

³To see this, consider the two-point correlation of a_i for two values of y , i.e. $\overline{a_i(y_1, t) a_i(y_2, t)} \propto \overline{a_{i,f}(y_1, t - x/U_{\text{conv}}(y_1)) a_{i,f}(y_2, t - x/U_{\text{conv}}(y_2))}$ using simplified notation. If U_{conv} is a function of y , then the two points (y_1, t) and (y_2, t) will correspond to points in the database that could be very far apart in time. This would lead to essentially zero correlation, even if $y_1 \approx y_2$.

for some integer α_l . Note that the value of δ^+ is part of the solution and hence unknown at this stage. For this reason, a (constant) target value of $\delta_{\text{target}}^+ \approx 2023$ is used⁴ to non-dimensionalize the LES/RANS coordinates when introducing the forcing, i.e. in equations (6.10) and (6.11). Being used only for this purpose, the exact value of δ_{target}^+ affects the results in a very minor way. Picking the spanwise domain size as

$$\frac{L_z}{\delta} = 3 \frac{L_{z,\text{db}}^+}{\delta_{\text{target}}^+} \approx 2.253 \quad (6.13)$$

yields $\alpha_l = 3l$ by (6.12), and hence periodicity of all modes is ensured.

The wavelength in the streamwise direction implied by Taylor's hypothesis is $\lambda_{x,l}^+ \equiv U_{\text{conv}}^+ T_{\text{db}}^+ / l$ for the l th mode, and hence there must exist some integer β_l such that

$$\beta_l \lambda_{x,l}^+ = \frac{L_x u_\tau}{\nu} \Rightarrow \beta_l \frac{U_{\text{conv}}^+ T_{\text{db}}^+}{l} = \frac{L_x}{\delta} \delta_{\text{target}}^+ \quad (6.14)$$

for mode l to be periodic on the LES/RANS domain. Picking⁵ $U_{\text{conv}}^+ \approx 17.63$, which roughly corresponds to the mean velocity around y_f^+ , and

$$\frac{L_x}{\delta} = \frac{1}{9} \frac{U_{\text{conv}}^+ T_{\text{db}}^+}{\delta_{\text{target}}^+} \approx 6.333 \quad (6.15)$$

yields $\beta_l = l/9$ by (6.14), and hence only modes $l = 0, \pm 9, \pm 18, \pm 27 \dots$ are periodic.

To avoid aliasing of energy when introducing the forcing on the coarse LES/RANS grid, only modes that satisfy the Nyquist criterion $\lambda \geq 2\Delta$ should be used. The grid for the LES/RANS with forcing is the same as the one described for the case without forcing, with 48 and 36 cells in the streamwise and spanwise directions, respectively. Relations (6.13) and (6.15) then imply that only spanwise modes $l = 0, \pm 1, \dots, \pm 6$ and temporal modes $l = 0, 9, \dots, \pm 207, \pm 216$ will be used. This defines the spectral filter kernels G_z and G_t used to filter the database.

6.4.2 Single-Point Statistics

The mean velocity profile and its gradient (as measured by γ) are shown in figures 6.5(a) and 6.6(a), respectively. The forcing essentially removes the velocity shift, in that the mean velocity profile closely follows the log-law for $y^+ \gtrsim 300$, but the artificial buffer layer clearly still exists. Figure 6.6(a) shows how the forcing reduces γ primarily in the region where the forcing is added, with the largest reduction taking place around the location of maximum forcing. This confirms the reasoning behind picking the value of y_f based on the local maximum of γ , rather than based on the location of the interface y_{int} (as was done by some previous studies (Piomelli *et al.*, 2003; Davidson & Dahlström, 2004; Davidson & Billson, 2006) – note, however, that these studies used turbulence models that generate slightly different ν_t profiles, and hence the location of maximum

⁴The exact value is $\delta_{\text{target}}^+ = u_{\tau,\text{target}} \delta / \nu = (u_{\tau,\text{target}} / U_b) Re_b$, with $u_{\tau,\text{target}} / U_b = 0.0465$.

⁵The exact value is $U_{\text{conv}}^+ = (U_{\text{conv}} / U_b) (U_b / u_{\tau,\text{target}})$, with $U_{\text{conv}} / U_b = 0.82$.

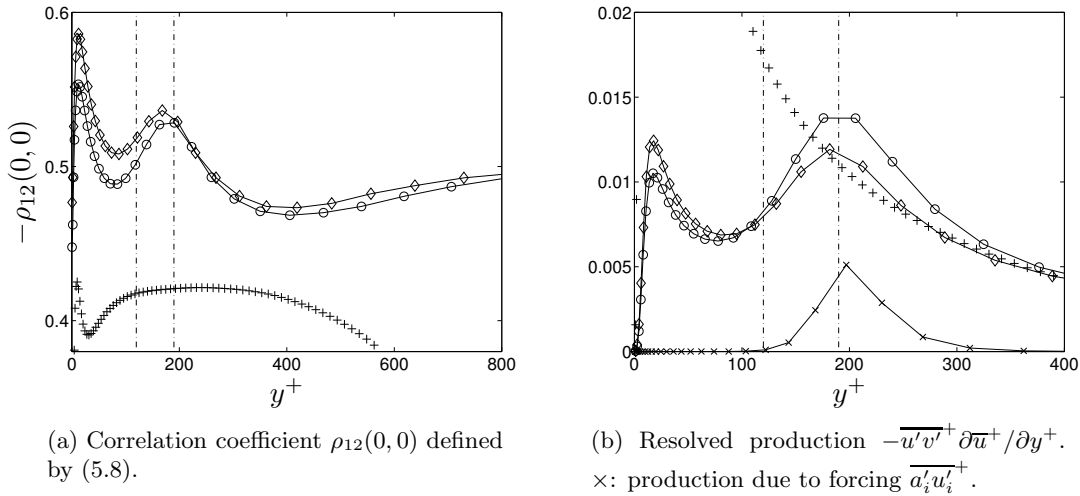


Figure 6.7: Correlation coefficient and turbulence production. \circ : LES/RANS without forcing; \diamond : LES/RANS with forcing; $+$: DNS at $Re_\tau \approx 950$ by del Alamo *et al.* (2004); $-\cdot-$: locations of the interface ($y_{\text{int}}^+ \approx 120$) and that of maximum forcing ($y_f^+ \approx 190$), respectively.

γ might be closer to the interface in their results). The mean velocity gradient finally reaches its LES equilibrium value around $y^+ \approx 300$.

The resolved normal stresses are shown in figure 6.5(b), and the forcing is seen to make the stresses less anisotropic. Since this anisotropy is related to the superstreaks (Baggett, 1998), this implies that their strength has been reduced. The maximum change of the streamwise and spanwise components occurs around the peak in forcing at y_f . The wall normal component behaves slightly differently due to the kinematic boundary condition, with the maximum increase occurring farther out. Although the normal stresses are inactive (in the sense of not affecting the mean flow) in fully developed channel flow, figure 6.5(b) still suggests that the forcing decreases the size of the artificial buffer layer, in that $\overline{v'v'}$ and $\overline{w'w'}$ reach their respective peaks closer to the interface.

The single-point correlation coefficient $\rho_{12}(0,0)$ given by (5.8) is shown in figure 6.7(a). Before considering the effect of the forcing, one notes that the (negative) correlation coefficient is much larger than in the DNS results throughout the channel. In the LES equilibrium region, this is likely due to the lack of small scale, approximately isotropic, motions in the LES/RANS, which decreases the denominator of (5.8) while hardly affecting the numerator. The interpretation is more difficult near the wall, where the resolved shear stress is strongly affected, but at the very least the high correlation suggests that the superstreaks have a structure capable of carrying substantial shear stress. Also, the fact that $\rho_{12}(0,0)$ has a peak in the middle of the artificial buffer layer (i.e., where γ has a peak) offers further indication that it is not the structure of the superstreaks that causes the velocity shift.

Consider also the effect of the forcing using relation (5.8). The (negative) correlation coefficient increases by only 1% at $y^+ \approx y_f^+$ with the use of forcing (figure 6.7(a)). At that location, u_{rms} decreases by about 5% while v_{rms} increases by about 10% (figure 6.5(b)). The 6% increase of $-\overline{u'v'}$ (figure 6.6(b)) is, then, primarily due to the increased magnitude of the wall normal velocity component.

These observations are interesting since they contradict Baggett (1998)'s suggestion that a primary cause of the velocity shift is a too low correlation coefficient. The present results instead show that the correlation first of all is rather high in the artificial buffer layer, and secondly that it is the lack of resolved wall normal motions that is the cause of the velocity shift, not a low correlation between u' and v' .

There is also a similarity between the true and artificial buffer layers with respect to the relation between γ and $-\rho_{12}(0,0)$. The correlation coefficient from the DNS has a maximum at $y^+ \approx 10$ and a minimum at $y^+ \approx 30$. These locations roughly coincide with the peak and leveling off of γ , respectively. The LES/RANS has two maximum/minimum combinations in the $-\rho_{12}(0,0)$ profile, which coincide with the equivalent points in the true and artificial buffer layers, respectively. This indicates that the superstreaks have some of the 'correct' dynamics.

The production of resolved kinetic energy is shown in figure 6.7(b) for the region around the interface. It should first be noted that the total (i.e., resolved plus modeled) production in a channel flow is accurate whenever the mean velocity gradient is accurately predicted, which is simply due to the fact that the total shear stress is determined by the shear stress balance (2.30). Figure 6.6(a) then shows that the total production is slightly underpredicted for $50 \lesssim y^+ \lesssim 120$, and more substantially overpredicted for $120 \lesssim y^+ \lesssim 400$. Figure 6.7(b) instead shows the production of *resolved* kinetic energy only, which is more relevant when investigating the transition to resolved turbulence around the interface. The production due to the forcing $\overline{a'_i u'_i}$ is slightly less than half of the resolved production at the forcing peak, and less than one third of the total production (not shown). In contrast, the white noise forcing in Piomelli *et al.* (2003) was found to result in a production about five times larger than the total production at the forcing peak. This high level could be due to the lack of physical structure in their forcing fields, or (perhaps more likely) due to the fact that their forcing was more concentrated around small scales (that dissipate quickly with less effect on the resolved shear stress – and hence a higher input of energy was needed to remove the velocity shift).

The forcing also changes the resolved production indirectly, by decreasing it around the forcing peak and increasing it in the RANS region. The decrease at $y^+ \approx 190$ is due to the lower mean velocity gradient, while the increase in the RANS region is due to a higher resolved shear stress. The fact that the resolved production has a small peak in the true buffer layer at $y^+ \approx 18$ again suggests that the superstreaks have some of the correct dynamics.

Also notable is the fact that the resolved kinetic energy $\overline{u'_i u'_i}/2$ (not shown, but can be inferred from figure 6.5(b)) actually decreases with the use of forcing over essentially the whole channel. Since the production, including that due to the forcing, increases, this shows that the dissipation of resolved kinetic energy must also increase with the use of forcing. Below it will be shown that the forcing decreases the length scales of the




Figure 6.8: Magnitude of the vorticity fluctuations. \circ : LES/RANS without forcing; \diamond : LES/RANS with forcing; $-\cdot-$: locations of the interface ($y_{\text{int}}^+ \approx 120$) and that of maximum forcing ($y_f^+ \approx 190$), respectively.

resolved motions throughout the artificial buffer layer, which is consistent with a higher dissipation.

The magnitudes of the vorticity fluctuations are shown in figure 6.8. The levels are much lower than for the wall resolved LES in chapter 5 (figure 5.6): between 75-90% lower depending on the location. Both the streamwise and the wall normal vorticities are lower by similar amounts. This shows that the quasi-streamwise vortices are heavily damped in the RANS region, and that the streaks are overly wide (since $\overline{u'u'}$ is high, see the discussion in section 5.3).

The streamwise vorticity peaks outside the interface at $y^+ \approx 250$. The forcing increases $\omega_{x,rms}^+$ by up to 15%, and causes the peak to shift towards the interface. The wall normal vorticity profile is very similar to the $\overline{u'u'}$ profile (figure 6.5(b)), which simply reflects the dominance of streak-like motions in the RANS region and the artificial buffer layer, with $\omega_{y,rms}^+ \sim u_{\text{rms}}^+/\lambda_z^+$. The forcing decreases $\omega_{y,rms}^+$ by 2-5%, primarily around the location of maximum forcing. At the interface, $\omega_{y,rms}^+$ decreases by 5% with forcing, while u_{rms}^+ decreases by 8%. From the estimate $\omega_{y,rms}^+ \sim u_{\text{rms}}^+/\lambda_z^+$ one can then infer that the forcing decreases the average streak spacing only marginally.

6.4.3 Multi-Point Statistics

To see the spectral content of the resolved turbulence, and how this changes with wall distance, the premultiplied one dimensional spectra are shown in figure 6.9. For comparison, the spectra from the wall resolved LES of chapter 5 are shown as well (this time in outer units).

The spectra for the LES/RANS are quite different compared to the wall resolved LES. Consider first the case without forcing. Near the wall, for $y/\delta \lesssim 0.05$ ($y^+ \lesssim 100$), there is a small peak in energy for streamwise length scales $\lambda_x/\delta \approx 3$ (i.e, half the




Figure 6.9: Contours of the premultiplied one dimensional spectra $k\hat{E}(k)$. 20 contours between 0 and each individual maximum, with the axes the same columnwise. The left column are the streamwise spectra, the right column are the spanwise. Top row: wall resolved LES at $Re_\tau \approx 500$ (base grid of chapter 5); second row: LES/RANS without forcing; bottom row: LES/RANS with forcing; $-\cdot-$: locations of the interface ($y_{\text{int}}/\delta = 0.06$) and that of maximum forcing ($y_f/\delta = 0.094$), respectively.

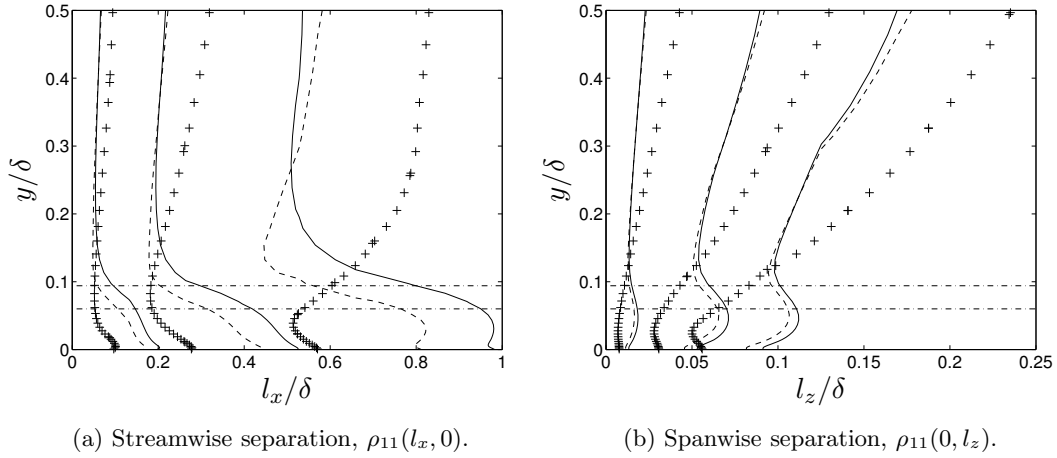


Figure 6.10: Contours of the two-point correlation of the streamwise velocity $\rho_{11}(l_x, l_z)$. Three contours at (from left to right) 0.9, 0.6, 0.3 of the maximum at each y location. —: LES/RANS without forcing; - - : LES/RANS with forcing; +: resolved LES at $Re_\tau \approx 500$ (base grid of chapter 5); - · - : locations of the interface ($y_{\text{int}}/\delta = 0.06$) and that of maximum forcing ($y_f/\delta = 0.094$), respectively.

domain size). The spanwise spectrum at this height is, though hard to see in the figure, rather flat, and hence the spanwise length scale is on the order of the grid spacing. These motions are unphysical and constitute the superstreaks described by Baggett (1998). The presence of these superstreaks, with scales determined by the numerics rather than by physics, is potentially problematic in LES/RANS, since they are convected and diffused into the LES region.

Outside of the interface there is a peak for spanwise length scales $\lambda_z/\delta \approx 0.6$ at $y/\delta \approx 0.1$ and streamwise length scales $\lambda_x/\delta \approx 0.8$ at $y/\delta \approx 0.15$. The streamwise wavelength corresponds to about 6 cell sizes. Since the dispersion error of the convection scheme is large for scales with fewer than three to four cells per wavelength, the streamwise peak occurs near the numerical cut-off. It is interesting to note that it is the streamwise direction which seems most affected by the resolution, rather than the spanwise as is common in LES. In the LES equilibrium region ($y/\delta \gtrsim 0.2$) the dominant length scales are similar to those in the wall resolved LES.

The forcing changes the spectra only for $y/\delta \lesssim 0.25$, and primarily in the streamwise direction. The peaks here occur much closer to the interface compared to the case without forcing, which suggests that the artificial buffer layer has decreased in size. Note that the multiple peaks around $y \approx y_f$ correspond to the discrete modes in the forcing field. There appears to be more small scale energy below the location of the forcing peak, indicating that the superstreaks have decreased in size. To see this more clearly, contours of the two-point correlation coefficient of the streamwise velocity are shown in figure 6.10. The wall resolved LES shows that the length scale starts increasing towards the wall in the upper part of the buffer layer, at $y/\delta \approx 0.05$ ($y^+ \approx 25$). In the LES/RANS, this increase occurs at the corresponding point in the artificial buffer

layer instead, at $y/\delta \approx 0.15$ ($y^+ \approx 300$). The forcing moves this point closer to the interface, and also decreases the length scale all the way to the wall. This is true for both directions, but the difference is largest for the streamwise one – this is hardly surprising, since the spanwise length scale is already on the order of the grid size. Overall, the forcing certainly decreases the vertical size of the artificial buffer layer and the length of the superstreaks, albeit not to the point of the latter being physically correct.

Figures 6.9 and 6.10 together show that there is a length scale mismatch between the LES and RANS regions. Since the near wall grid in LES/RANS is necessarily coarse (otherwise a wall resolved LES could be performed), the motions in the RANS region can never have the physically correct length scales. Therefore, some transition region (i.e., artificial buffer layer) will likely always be present in these cases. The forcing reduces the size of this transition, and the size of the mismatch in length scales. The fact that the spectra are unchanged in the core of the channel is encouraging, since the forcing should not affect the accurate LES region of the flow.

6.4.4 Turbulence Structure

Snapshots of the streamwise velocity in three planes parallel to the wall are shown in figure 6.11, both with and without forcing. Near the wall the structures are very elongated, with lengths of up to $\lambda_x^+ \sim 6000$. The low speed streaks are longer and slightly narrower compared to the high speed streaks, which is in qualitative agreement with true near wall turbulence. This suggests that the motions in the artificial buffer layer have at least some of the correct physics, albeit at unphysical length scales. The snapshots at $y^+ \approx 12$ and $y^+ \approx 120$ look similar, especially without forcing, which indicates that the streaks are rather ‘tall’ structures, i.e. that they are well correlated in the wall normal direction. At $y^+ \approx 260$ ($y/\delta \approx 0.13$), more small scale structures are present, and the correlation to the two lower planes is much lower. This is also the wall distance at which the streamwise spectra in figure 6.9 shows a transition to more small scale energy.

The forcing appears to create smaller resolved scales, but without removing the excessively long streaks near the wall. The similarity between the two lower planes also appears to be lower, indicating that the forcing reduces the height of the near wall structures somewhat.

Having seen that the turbulence in the RANS region and in the artificial buffer layer share some qualitative features with the true buffer layer, one might hypothesize that the near wall region corresponds to the turbulence at an effectively lower Reynolds number. The eddy viscosity (figure 6.2) is $\nu_t/\nu \lesssim 25$ for $y^+ \lesssim 120$, and hence the effective Reynolds number in this region is lower by a similar factor. This is in very rough agreement with the fact that the streak length $\lambda_x^+ \sim 6000$ is about 10 times larger than the typical lengths found for the wall resolved LES in chapter 5.




Figure 6.11: Snapshots of the streamwise velocity fluctuation u' . Contours one standard deviation (i.e., u_{rms} , taken from the case without forcing) above and below zero, where darker regions imply higher velocity. The left column are without forcing, the right with forcing. Top row: $y^+ \approx 230$; second row: $y^+ \approx 120$; bottom row: $y^+ \approx 12$.

6.4.5 Effective Forcing Length Scales

In order to develop a general model for the forcing, a basic requirement is to know at which length scales the forcing should be applied. Some models proposed in the literature are based on the assumption that the forcing represents backscatter (Piomelli *et al.*, 2003), and hence should have length scales on the order of the grid size. Such small scales would also potentially be effective at breaking up the superstreaks. On the other hand, one could also argue that the forcing should be applied to the largest scales which carry most of the energy and Reynolds stresses. These larger scales might cause a quicker formation of LES equilibrium motions.

In an effort to shed some light on which length scales should be included in a general forcing model, different filter kernels G_z and G_t are used to produce forcing fields with different length scale contents. In essence, since spectral cut-off filters are used, this amounts to the selective removal of modes from the database. The effectiveness of the forcing is measured by the size of the velocity shift, defined here as the difference between the mean velocity and the corresponding value of the log-law at $y^+ = 500$, i.e.

$$\Delta U^+ \equiv \bar{u}^+|_{y^+=500} - \left(\frac{\ln(500)}{0.41} + 5.2 \right). \quad (6.16)$$




Figure 6.12: Size of the velocity shift for forcing fields with different length scale contents, as defined by the spectral cut-off wavelength λ^c . \square : only modes with $\lambda < \lambda^c$; \times : only modes with $\lambda > \lambda^c$; $---$: full forcing field (i.e., with all resolvable modes); $---$: without forcing.

This position is outside of the artificial buffer layer while still not too far into the wake region in the core of the channel.

Figure 6.12(a) shows the size of the velocity shift for cases where spanwise modes have been removed from the forcing database. The amplitude parameter A is held constant, which means that the overall energy of the forcing decreases as modes are removed. One should note that ΔU^+ is quite sensitive to the convergence of the statistics, as evidenced by the wiggles in the figure. This error is estimated as ± 0.03 .

The results clearly show that spanwise forcing modes with $\lambda_z/\delta \lesssim 0.2$ are ineffective at reducing the size of the velocity shift. This wavelength corresponds to about three cell sizes, which suggests that the ineffectiveness of these modes is due to the inability of the numerical scheme to accurately represent motions at these length scales. More interesting is the fact that the larger length scales are relatively effective in the sense discussed here. Forcing only at wavelengths $\lambda_z/\delta \gtrsim 0.5$ still reduces the velocity shift by more than 50%. Even the first mode, with $\lambda_z = \infty$, has an effect on the solution: when only this mode is included, the velocity shift is decreased by 17% (to $\Delta U^+ \approx 0.73$).

Streamwise modes are removed in exactly the same fashion, and the results are shown in figure 6.12(b). Note that the ‘stepwise’ pattern of $\Delta U^+(\lambda_x^c)$ is an artefact of the use of Taylor’s hypothesis in combination with the use of a short domain when creating the forcing database, causing some modes to be less energetic than others. The temporal differentiation also causes the largest modes to have little energy. With these caveats, the results are similar to those for the spanwise direction. Modes with three or fewer cells per wavelength ($\lambda_x/\delta \lesssim 0.4$) are ineffective at reducing ΔU^+ , while modes with $\lambda_x \sim \delta$ are rather effective.

Within the limitations of the current forcing fields, one can only say that modes with wavelengths in the range $3\Delta \lesssim \lambda \lesssim \delta$ are effective at reducing the size of the

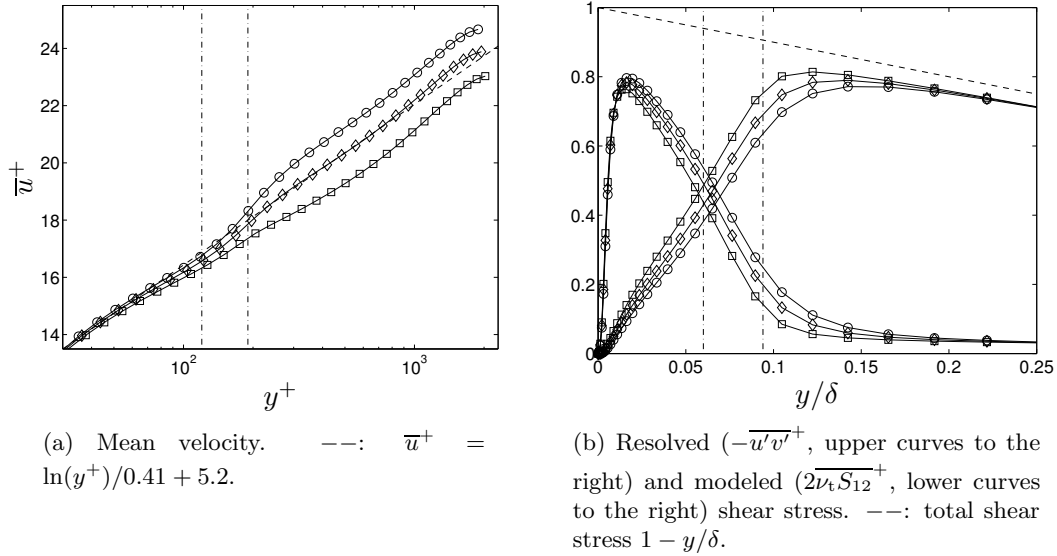


Figure 6.13: Mean velocity and shear stress for different amplitudes. \circ : $A = 0$, i.e. without forcing; \diamond : $A = A_{\text{opt}} = 0.335$; \square : $A = 0.668$; $-\cdot-$: locations of the interface ($y_{\text{int}}/\delta = 0.06$) and that of maximum forcing ($y_f/\delta = 0.094$), respectively.

velocity shift. It remains a distinct possibility that the upper limit on the forcing length scale should be much larger, but to test this would require forcing fields with more large scale content. One should also note that the findings here say little about which forcing length scales should be used from a *physical* point of view, and only which length scales effectively decrease the velocity shift with the forcing amplitudes used here.

6.4.6 Amplitude Sensitivity

The results with forcing are quite sensitive to the value of the amplitude parameter A . To illustrate this sensitivity, the mean velocity profiles using three different values of A are shown in figure 6.13(a). A too high amplitude generates too much resolved shear stress, resulting in an underpredicted mean velocity profile and gradient. A too low amplitude, on the other hand, does not quite remove the velocity shift. This is seen more clearly in figure 6.14, which shows the size of the velocity shift ΔU^+ for different amplitudes. Interestingly, the response is an approximately linear function for forcing amplitudes around the optimal value $A_{\text{opt}} = 0.335$ (defined by $\Delta U^+ = 0$). To quantify the sensitivity, say that one needs an amplitude A such that $|\Delta U^+(A)| \leq 0.1 \Delta U^+(A = 0) \approx 0.088$. Figure 6.14 then implies that A must be within $\pm 8\%$ of A_{opt} . This presents a difficulty for the development of generally applicable forcing models, which should give accurate results with the amplitude determined *a priori*.

Up to this point, this difficulty has been avoided by use of trial-and-error to find A_{opt} . In other words, A was adjusted until $\Delta U^+ \approx 0$, and all results discussed previously in this chapter used that amplitude. Other studies, e.g. Piomelli *et al.* (2003), also used




Figure 6.14: Size of the velocity shift as a function of the forcing amplitude.

trial-and-error in a similar fashion.

Some proposed model expressions (i.e., *a priori* relations) for A have appeared in the literature. Davidson & Dahlström (2004) used an amplitude that was essentially a ratio between the modeled kinetic energy (of their two-equation model) at the interface and the energy of the forcing field, and reported a high degree of sensitivity to this model expression. Batten *et al.* (2004) also used an amplitude that was proportional to the modeled turbulence, in this case the modeled stress tensor directly.

These studies assume proportionality between A and the modeled kinetic energy. This implicitly takes the view that the forcing in some sense represents the transfer of energy from the modeled scales (i.e., backscatter), since two flows with similar resolved velocity fields but different levels of modeled kinetic energy would result in different amplitudes.

If the forcing is viewed not in the context of backscatter, but rather as a means to trigger resolved motions, then one might instead expect A_{opt} to be dependent on the resolved (rather than modeled) scales of motion. In this view, the proposed model expressions by Davidson & Dahlström (2004) and Batten *et al.* (2004) would be flawed, as follows from the simple hypothetical example given above.

Mason & Thomson (1992) used inertial range arguments to find an expression for the amplitude that depended on the dissipation rate ε and the characteristic length scale of the unresolved motions. With the assumption of equilibrium for the unresolved scales, and more importantly the existence of an inertial range, $\varepsilon \approx \varepsilon_{\text{sgs}}$. The rate of energy transfer from the resolved scales ε_{sgs} depends on the resolved scales with the turbulence models used here, and, thus, as does Mason & Thomson (1992)'s amplitude. One should note, however, that Mason & Thomson (1992) used an algebraic wall stress model, for which the resolved motions are much different from the ones near the interface in LES/RANS. As shown in section 6.4.3, the length (and time) scales in the interface region of LES/RANS are unphysically large, whereas the algebraic wall model used by Mason & Thomson (1992) generated much smaller scales in the near wall region. Based

on this argument, Keating & Piomelli (2006) concluded that the amplitude expression and the underlying physical arguments used by Mason & Thomson (1992) do not apply within LES/RANS.

6.4.7 A Control Algorithm for the Amplitude

The sensitivity of the results to the forcing amplitude suggests that even physically sound models for the amplitude would require tuning on a case-by-case basis, since even small changes in the flow, grid, or turbulence models could change the relation between the optimal amplitude A_{opt} and whatever quantities would be included in a model.

In light of this, one might instead turn to a control algorithm for A , as was done by Keating & Piomelli (2006). Before considering the details of the method used here, and how it differs from the method by Keating & Piomelli (2006), some general thoughts about this approach should be considered. Note that this control algorithm is very different from the wall model by Nicoud *et al.* (2001) described in section 6.1.4. Here, the inner layer is modeled statistically, and the control algorithm is only used to find the appropriate value of A .

A control algorithm for the amplitude would have the form $A = A(\epsilon)$ where ϵ is some error that needs to be specified. With a proper choice of the control function, this error can be driven towards zero. As discussed earlier in this chapter, the main problem with the LES/RANS concept is the shift in the velocity profile outside of the interface. The most natural definition of the error should then be sensitive to this velocity shift. The size of the velocity shift, however, requires prior knowledge of the proper mean velocity profile (given by the log-law for channel flows), suggesting that such an error definition would remove the predictive capability of the method.

Such a statement is too pessimistic. In all situations where one would consider using LES/RANS, it is implicitly assumed (see the discussion in section 1.2) that the RANS model accurately describes the mean flow for $y/\delta \lesssim 0.1$ – after all, were this not the case, the whole underlying concept of using a ‘universal’ RANS model near the wall would fail. If the RANS equations are solved from the wall up to $y/\delta \approx 0.1$ in an auxiliary solver, then the resulting velocity profile $\bar{u}_{\text{ref}}(y)$ should, by the implicit assumption underlying LES/RANS, constitute an accurate reference solution relative to which the error can be defined. Note that the computational cost associated with such an auxiliary RANS solver would be very small compared to the full LES/RANS solver.

For the cases considered here, this idea can be implemented in a simpler fashion. The rather non-rigorous derivation of the mixing length RANS model for ν_{rans} in section 6.2 showed the link between the RANS model and the slope of the log-law $1/\kappa$. Also, close inspection of figure 6.13(a) reveals that the mean velocity is accurate to within 0.1 units for $y^+ \lesssim 70$. One can then write the approximate reference solution (without the need for auxiliary numerical solution) as

$$\bar{u}_{\text{ref}}(y^+) = \bar{u}(y_{\text{ref}}^+) + \frac{u_\tau}{\kappa} \ln \left(\frac{y^+}{y_{\text{ref}}^+} \right), \quad y_{\text{ref}}^+ \leq y^+ \lesssim 0.1\delta^+, \quad 30 \lesssim y_{\text{ref}}^+ \lesssim 70. \quad (6.17)$$

The average quantities required to compute the reference solution (both $\bar{u}(y_{\text{ref}})$ and for $u_\tau \equiv \sqrt{\nu \partial \bar{u} / \partial y|_w}$) are unavailable during the simulation. Replacing these by the corresponding plane-averages $\langle \cdot \rangle$ yields the re-defined reference solution

$$\langle u_{\text{ref}} \rangle(y) = \langle u \rangle(y_{\text{ref}}) + \frac{u_{\tau,p}}{\kappa} \ln \left(\frac{y}{y_{\text{ref}}} \right), \quad (6.18)$$

where $u_{\tau,p} \equiv \sqrt{\nu \partial \langle u \rangle / \partial y|_w}$. The difference between (6.17) and (6.18) can be shown to be a factor $\bar{u}_{\tau,p} / u_\tau - 1$ multiplying the logarithmic term when u_τ is used to non-dimensionalize the y -coordinate. For the present cases, this error is $\sim 10^{-4}$.

The error in the control algorithm can now be defined as

$$\begin{aligned} \epsilon(t) &\equiv \frac{1}{U_b(y_h - y_l)} \int_{y_l}^{y_h} \{ \langle u \rangle(y) - \langle u_{\text{ref}} \rangle(y) \} dy \\ &= \frac{1}{U_b(y_h - y_l)} \int_{y_l}^{y_h} \left\langle u(y) - u(y_{\text{ref}}) - \frac{u_{\tau,p}}{\kappa} \ln \left(\frac{y}{y_{\text{ref}}} \right) \right\rangle dy. \end{aligned} \quad (6.19)$$

This is, of course, only one of many possible definitions of the error. Keating & Piomelli (2006) used a different error, one not involving the mean velocity. First, they used the detached eddy simulation (DES) method (Spalart *et al.*, 1997; Nikitin *et al.*, 2000). They then based their error on the observation that the shear stress changes from being mainly modeled to mainly resolved in the artificial buffer layer (also shown earlier in this chapter). They further argued that this change should occur between the interface location (called y_{int} here) and the location of maximum eddy viscosity y_ν , and thus defined the error as

$$\epsilon_{\text{keating}} = \int_{y_{\text{int}}}^{y_\nu} \left(-\langle u'v' \rangle - \langle \nu_t \rangle \frac{\partial \langle u \rangle}{\partial y} \right) dy, \quad (6.20)$$

where the quantities were averaged over the wall parallel directions (and a short backward temporal average). While being physically appealing (in that it measures the change of the shear stresses), and leading to good results in Keating & Piomelli (2006), (6.20) appears to be rather specific to the turbulence model and interface condition in DES. Note, for example, that the models and blending used in this thesis yield $y_\nu < y_{\text{int}}$, for which (6.20) does not make sense. Perhaps more importantly, the argument that the change in shear stresses should occur between the interface location and the position of maximum eddy viscosity appears (to the present author) rather *ad hoc*. Figure 6.13(b) shows that the shear stress profiles are essentially linear around the interface. The definition of $\epsilon_{\text{keating}}$, together with the fact that the control algorithm forces $\epsilon_{\text{keating}} \rightarrow 0$, then implies that the error defined by (6.20) causes the modeled and resolved shear stresses to intersect (i.e., be equal) at $y \approx (y_{\text{int}} + y_\nu)/2$. Hence, the integration limits basically control this intersection point, and thus one would expect the results to be rather dependent on the exact choice of these limits.

The error (6.19) proposed here, on the other hand, is directly related to the most severe deficiency in LES/RANS: the velocity shift. A thorough comparison between these different error definitions is not attempted here, but would be of interest in future work.

Control Strategy

The control strategy should provide the amplitude A as a function of the error ϵ . The most efficient control strategies utilize knowledge of the process itself, and some such knowledge is available here, for example the relation given by figure 6.14. This relation, however, is unique to the particular flow, grid, and models used, and hence a more general control strategy is used. One basic and general such strategy is a proportional-integral-derivative (PID) controller (c.f. Åström & Wittenmark, 1990). Derivative information is by nature very noisy, especially so for the turbulence considered here, and hence only a proportional-integral (PI) controller is used. Such a controller can be written as

$$A(t) = K_P \epsilon(t) + \underbrace{\frac{K_I U_b}{\delta} \int_0^t \epsilon(s) ds}_{A_I(t)}, \quad K_P, K_I \geq 0, \quad (6.21)$$

where the constants need to be positive for a positive error to generate a positive amplitude. The temporal integral is most easily evaluated by noting that

$$\frac{dA_I}{dt} = \frac{K_I U_b}{\delta} \epsilon(t) \quad (6.22)$$

which can be discretized by an implicit Euler method as

$$\begin{aligned} A_I^n &= A_I^{n-1} + \Delta t \frac{K_I U_b}{\delta} \epsilon^n, \\ A^n &= K_P \epsilon^n + A_I^n. \end{aligned} \quad (6.23)$$

Note that it is the integral part that removes the error in the long-time limit, since for a stationary solution $dA_I/dt = (K_I U_b/\delta)\epsilon = 0$. The proportional part is less active in the long-time limit, and mainly plays a part by speeding up the transient behaviour. As a side note, Keating & Piomelli (2006) used an integral controller (i.e., (6.23) with $K_P = 0$).

It remains to determine suitable values for the parameters involved (the reference location y_{ref} and the integration limits y_l and y_h in the error definition, and the control parameters K_P and K_I), and, equally importantly, to determine the sensitivity of the results to these parameters. After all, one major reason for considering a control algorithm is the high sensitivity to A – if the results are as sensitive to any of the error or control parameters, little has been gained.

Table 6.1 lists results from cases where each parameter is varied systematically. For the first two sets of tests, the error parameters $y_{\text{ref}}^+ = y_l^+ = 62$ and $y_h^+ = 360$ are used. The lower value is consistent with the range given above, and the upper limit is at the

Parameter	K_I	K_P	y_{ref}^+	y_1^+	y_h^+	\bar{A}	A'_{rms}/\bar{A}	$\epsilon_{\text{rms}} \cdot 10^3$	ΔU^+
K_I	0.3	0	62	62	360	0.322	0.038	2.5	0.054
	1					0.327	0.059	2.8	0.068
	3					0.318	0.201	4.2	0.063
	10					0.241	1.350	15	0.014
K_P	1	0	62	62	360	0.327	0.059	2.8	0.068
		0.3				0.326	0.075	3.2	0.073
		1				0.327	0.064	2.6	0.069
		3				0.327	0.070	2.7	0.064
		10				0.327	0.093	2.4	0.061
		30				0.318	0.193	2.0	0.055
		100				0.288	0.642	2.0	0.066
y_{ref}^+	1	1	36	250	600	0.335	0.096	5.6	-0.021
			52			0.351	0.085	5.1	-0.017
			62			0.368	0.112	6.1	-0.084
			74			0.395	0.077	4.6	-0.168
y_1^+	1	1	52	130	600	0.336	0.089	4.0	0.029
				180		0.343	0.088	4.7	-0.006
				250		0.351	0.085	5.1	-0.017
				390		0.326	0.117	6.3	0.003
y_h^+	1	1	52	250	290	0.359	0.072	4.0	-0.076
					390	0.357	0.077	4.6	-0.054
					600	0.351	0.085	5.1	-0.017
					780	0.349	0.093	5.3	-0.017

Table 6.1: Amplitude control algorithm cases. The left column indicates which parameter is varied, while the remaining parameters are fixed. Note that some cases are repeated for clarity.

end of the artificial buffer layer. Initial estimates for the control parameters can be obtained by considering the adjustment time scale for ϵ when a fixed amplitude A_{opt} is used. For that case, the error adjusts to zero in about $T_\epsilon U_b/\delta \approx 10$ (figure 6.15(a)). Requiring the initial amplitude $A(0) = K_P \epsilon(0) \sim A_{\text{opt}}$ and the approximate growth $A_I(T_\epsilon) \approx (K_I U_b/\delta) T_\epsilon \epsilon(0) \sim A_{\text{opt}}$ then yields

$$K_P \sim \frac{A_{\text{opt}}}{\epsilon(0)} \approx 30, \quad K_I \sim \frac{A_{\text{opt}} \delta}{T_\epsilon U_b \epsilon(0)} \approx 3. \quad (6.24)$$

Control Parameters

Varying the control parameters K_P and K_I over rather large ranges yields flow statistics that are virtually indistinguishable. This is reflected in table 6.1 by the almost unchanged velocity shift ΔU^+ (recall that ΔU^+ is sensitive to the convergence of the statistics), but is also true for the higher order statistics. Note that ΔU^+ is not exactly zero, and that the average amplitudes are lower than A_{opt} . This is simply a reflection of the fact that the definitions of ϵ and ΔU^+ are different. Some sample time histories




Figure 6.15: Time histories for the control algorithm with different parameters, offset for clarity, where the dashed lines mark zero for each curve. $y_{\text{ref}}^+ = y_1^+ = 62$, $y_h^+ = 360$. From bottom to top: $A = A_{\text{opt}} = 0.335$, i.e. fixed; $(K_P, K_I) = (0, 0.3)$; $(K_P, K_I) = (1, 1)$; $(K_P, K_I) = (0, 3)$; $(K_P, K_I) = (30, 1)$.

of the error and the amplitude are shown in figure 6.15. The error fluctuates around zero for all cases, including when the amplitude is fixed. The fluctuations become more pronounced for larger values of K_I , but slightly more suppressed for larger values of K_P . The amplitude oscillations, on the other hand, increase when either parameter is increased. As expected, the initial transient is longer for lower values of K_I . Since the flow statistics are essentially identical regardless of the parameter values, the combination $(K_P, K_I) = (1, 1)$ is chosen here. This yields a short initial transient and reasonable fluctuation levels for ϵ and A , as seen in table 6.1. Overall, the control algorithm is not sensitive to the values of the control parameters.

Error Parameters

The sensitivity of the results to the error parameters is investigated in the last three sets of cases in table 6.1. The influence of the error parameters is perhaps best understood by considering the approximate (since $u_\tau \approx \overline{u_{\tau,p}}$) average error

$$\bar{\epsilon} \approx \frac{1}{U_b^+(y_h^+ - y_1^+)} \int_{y_1^+}^{y_h^+} \left\{ \overline{u^+}(y^+) - \overline{u^+}(y_{\text{ref}}^+) - \frac{1}{\kappa} \ln \left(\frac{y^+}{y_{\text{ref}}^+} \right) \right\} dy^+ \quad (6.25)$$

or equivalently

$$\bar{\epsilon} \approx \frac{1}{U_b^+(y_h^+ - y_1^+)} \int_{y_1^+}^{y_h^+} \{ \alpha(y^+) - \alpha(y_{\text{ref}}^+) \} dy^+, \quad \alpha(y^+) \equiv \overline{u^+}(y^+) - \frac{1}{\kappa} \ln(y^+), \quad (6.26)$$




Figure 6.16: Mean velocity profiles with a logarithmic slope subtracted, i.e. $\alpha(y^+)$ as defined by (6.26). \circ : $A = 0$, i.e. without forcing; \diamond : $A = A_{\text{opt}} = 0.335$; \square : $A = 0.563$; $-\cdot-$: locations of the interface ($y_{\text{int}}^+ \approx 120$) and that of maximum forcing ($y_{\text{f}}^+ \approx 190$), respectively.

where $\alpha(y^+)$ is the mean velocity profile with a logarithmic slope subtracted. This quantity is shown in figure 6.16, and can be used to infer the sensitivity to the error parameters. From relation (6.26), the average error is approximately the area between a curve in the figure and a constant line at its value at y_{ref}^+ , between the integration limits. Since $\alpha(y^+)$ is essentially constant for $250 \lesssim y^+ \lesssim 800$, one can infer that the results are rather insensitive to the integration limits as long as they are chosen in this range. This is confirmed in table 6.1, with only minor changes in ΔU^+ (and the remaining statistics, though not shown). Note that $|\Delta U^+|$ stays within the range of acceptable values given above for all values of y_1^+ and y_h^+ tested.

A similar situation applies to y_{ref}^+ , for which $|\Delta U^+|$ is larger than the value given above only for $y_{\text{ref}}^+ \approx 74$. Due to the curvature of $\alpha(y^+)$ in this region, $\alpha(y^+ = 74)$ is about 0.15 units lower than $\alpha(y^+ = 36)$ and $\alpha(y^+ = 52)$. As a direct consequence of the error interpretation given above, this translates into a velocity shift that is lower by about 0.15 units. Thus, the results are slightly more sensitive to y_{ref}^+ (compared to the other parameters). That being said, there is little reason for not choosing $y_{\text{ref}}^+ \approx 50$, i.e. immediately outside of the buffer layer.

Overall, the control algorithm for the forcing amplitude quickly finds the appropriate value of A , and is rather robust with respect to the parameters involved. Despite varying the parameter values over large ranges (e.g., 300% for y_1^+), the resulting statistics as measured by ΔU^+ are within the arbitrary limit 0.088 units. Contrast this to the sensitivity to A directly, where the amplitude must be within 8% of the optimal value to produce similar accuracy.

6.5 Summary and Conclusions

This chapter has focused on hybrid LES/RANS and on the artificial buffer layer that develops around the interface between the two modeling regions. The solution behaviour in the artificial buffer layer is studied, and it is argued (but not shown) that it will be hard to eliminate the resulting velocity shift within the context of eddy viscosity models: since the velocity shift depends directly on the balance between the resolved and modeled shear stresses, and since these quantities are both affected by the eddy viscosity (through the modeled dissipation and stress, respectively), it is unlikely that any one model for ν_t can be calibrated to provide both quantities accurately. While not shown here, it is consistent with the fact that the velocity shift has appeared in a range of studies using different turbulence models and interface conditions.

In the second part of the study, additional forcing of the momentum equations is considered. This forcing introduces a second degree of freedom in the modeling by increasing the resolved shear stress without affecting the modeled eddy viscosity. The forcing field is taken from a wall resolved LES, and is hence more physically sound than the stochastic forcing used in previous studies. The results with forcing are better in essentially all aspects: the velocity shift is removed, the artificial buffer layer is smaller in (vertical) size, and the strength and size of the near wall superstreaks are reduced. In contrast to some previous studies, the forcing is centered around the location of maximum scaled mean velocity gradient (as opposed to around the interface), with the underlying assumption that the locations of peak forcing, maximum increase of the resolved shear stress, and maximum decrease of the mean velocity gradient will all coincide. The results show this hypothesis to be essentially correct.

The success of the forcing suggests that this might be a way forward towards generally applicable LES/RANS models for high Reynolds number flows. The main problem with this concept appears, at least in the present study, to be the sensitivity of the results to the forcing amplitude. An *a priori* model expression for the amplitude is a necessity for predictive models, and the rather high sensitivity to this parameter suggests that universal models might be hard to develop. This problem is solved here by a proposed control algorithm for the amplitude. This is an inherently mathematical procedure, and the only physics involved is in the definition of the reference velocity profile and the error. The control algorithm quickly finds the appropriate amplitude, and it is remarkably robust with respect to the values of the parameters involved. It is argued that the approach is consistent with the use of LES/RANS in the first place, and that no extra empiricism is introduced.

There are several similarities between the true and artificial buffer layers: the mean velocity gradient scaled by the wall distance (i.e., γ) has peaks in the middle of both, the resolved streamwise stress component has a peak in both, the spanwise stress component increases rapidly through both, and the correlation coefficient between the streamwise and wall normal velocity fluctuations has a peak in both. In addition, the streamwise length scale of the u' component also increases towards the wall in both the true and artificial buffer layers. Since both buffer layers are associated with a change to resolved shear stress (from viscous and modeled, respectively), it is perhaps not surprising that there is a range of similarities between them. This, nevertheless, indicates that the

artificial buffer layer dynamics might be better understood in the context of the true buffer layer. Consider, for example, the fact that Jimenez *et al.* (2004) found evidence of very long streaks in the buffer layer of their ‘autonomous’ near wall turbulence DNS. These streaks had low wall normal velocity, but contributed to the shear stress due to their organized structure (i.e., the high correlation between u' and v'). This is very similar to the superstreaks in LES/RANS, and suggests that there is a common explanation for their existence.

The present results also contradict the suggestion by Baggett (1998) that the superstreaks are inefficient at carrying resolved shear stress due to a low correlation between u' and v' . Here, the (negative) correlation coefficient is high in the artificial buffer layer, in fact it even has a small peak there. Also, when adding the forcing, the correlation coefficient changes only by a small amount while the wall normal velocity fluctuation changes considerably more. This indicates that it is not the correlation, but rather the low level of v' that is the main reason for the low level of resolved shear stress. However, more study is needed into the artificial buffer layer dynamics before this issue can be settled.

On a side note, the importance of the wall normal component in generating a higher resolved shear stress opens up the possibility of using a non-zero wall transpiration velocity as a means to accomplish this.

Selective removal of forcing modes suggests that scales smaller than three to four cell sizes are ineffective at the removal of the velocity shift, while even the largest scales are rather effective. This suggests that general forcing models should provide these larger scales in order to be effective, and not just scales around the filter width. It also lends support to the idea that the concept of forcing should be viewed as a means to quickly trigger large, energetic structures, rather than in the context of backscatter.

Chapter 7

Summary

This thesis presents progress towards the application of large eddy simulation (LES) to flows of engineering interest. Due to the non-linear nature of the Navier-Stokes equations, predictions of turbulent flows are among the most difficult problems in engineering despite a century of vigorous research. The LES methodology offers hope, since it, in theory, solves for all flow dependent motions directly, while only modeling the more universal ones. Since the flow dependent motions are inviscid in nature, the computational cost is high but independent of the Reynolds number. The main bottleneck of LES is its application to attached boundary layers. In such flows, additional modeling of the inner boundary layer is necessary to retain the favourable characteristics of the method – resolution of all flow dependent motions at a cost independent of the Reynolds number. In the present thesis, a statistical turbulence model is used near the wall, and interfaced to the LES model at a height y_{int} above the wall. While the results are much improved compared to cases using no additional near wall modeling, an artificial buffer layer is formed around the modeling interface. The resolved motions are rather unphysical below this interface, and adjust throughout the artificial buffer layer.

To increase the resolved shear stress in the artificial buffer layer, additional forcing is introduced in the momentum equations. This has been attempted in previous studies using stochastic models for the forcing, but is done here using forcing fields from a reference LES. Thus, the forcing fields used here are more physically accurate than those used in other studies. This ensures that the present study of forcing can focus on the concept itself, rather than merely being an evaluation of any particular stochastic model. One could argue that the present results, using physically accurate forcing fields, constitute a form of ‘upper boundary’ for what may be achieved using the LES/RANS with forcing concept. The results with forcing are better in essentially all aspects: the velocity shift is removed, the artificial buffer layer is smaller in size, and the superstreaks in the near wall region are weaker and shorter. In contrast to previous studies, the forcing is introduced around the location of maximum mean velocity gradient here, based on the necessary balance of the stresses in the streamwise momentum equation. The results prove this reasoning to be correct.

The study reveals several similarities between the artificial and ‘true’ buffer layers: the mean velocity gradient (scaled by the wall distance) has peaks in the middle of both; the resolved streamwise stress component has peaks in both; the spanwise stress compo-

ment increases rapidly through both; the correlation coefficient between the streamwise and wall normal velocity components has peaks in both; and the streamwise length scale increases towards the wall in both. Since both buffer layers are associated with a change in shear stress, this is perhaps not surprising. It does, however, suggest that recent findings on ‘true’ near wall turbulence might be applicable to the artificial buffer layer as well.

The conjecture about the low resolved shear stress being caused by ‘de-correlated’ motions in the superstreaks by Baggett (1998) does not seem correct in the present study. Instead, the most likely explanation for the low resolved shear stress in the artificial buffer layer is the overly damped wall normal velocity fluctuations. Overall, the study of the artificial buffer layer contributes to the understanding of it, and provides a useful background for future attempts to reduce its impact in hybrid LES/RANS.

The results with forcing are very sensitive to the forcing amplitude, and it is argued that it will be difficult to find accurate *a priori* models for the amplitude due to this inherent sensitivity. As a way around this problem, a control algorithm that quickly finds the appropriate amplitude is proposed and tested. An error that measures the size of the velocity shift is specified, and it is argued that this error is consistent with LES/RANS and that it does not increase the empiricism of the method. With this error, the control algorithm generates accurate results for a wide range of parameter values, which shows the robustness of the method.

In a separate contribution, a conditional semicoarsening multigrid method for the efficient solution of the Poisson equation on structured grids is proposed and tested. The method does not yield optimal multigrid cost scaling, but is highly parallel, simple to implement, and applicable to complex geometries. While implemented here for structured grid topologies only, the method actually lends itself more naturally to unstructured grids. It should also, with proper modifications, be applicable to problems with strong discontinuities in the matrix coefficients.

The conditional semicoarsening method establishes a rather strong mathematical connection between the behaviour of the smoother and the proper coarse grids. While the underlying idea of coarsening the grid only in the directions of greatest coefficient strength has been proposed before based on other arguments, the mathematical connection derived here places the concept on firmer ground.

7.1 Future Directions

As per usual in research, every study raises at least as many questions as it may answer. In the author’s view, some of the most promising questions or ideas for future work are:

- Application of existing near wall modeling techniques in LES, including the one proposed here, to more complex flows. While the plane channel flow is one of the most challenging flows for LES, the time to evaluate more complex flows has come. For example, can the techniques handle impinging and separating flows properly? Are the error and the control algorithm proposed here valid for more complex flows, or do they need adjustment?

- A thorough comparison of different near wall modeling concepts, and an assessment of their strengths and weaknesses. For example, what are the inherent differences between the wall stress models (with RANS equations on an auxiliary grid) and the hybrid LES/RANS concept? Is the problem of the artificial buffer layer specific to LES/RANS, or does it appear with wall stress models as well?

Similarly, for the conditional semicoarsening multigrid method proposed here, some future questions are:

- Application of the conditional semicoarsening concept to unstructured grids. Is the computational cost scaling closer to optimal for unstructured grids, where the coarse grids are no longer unnecessarily dense?
- Application of the method to problems with coefficient discontinuities.
- The conditional semicoarsening presented here is mathematical in nature, but yields similar coarse grids as the agglomeration method based on physical arguments by Elias *et al.* (1997). The similarity suggests that there is a connection between these different arguments, and it would be instructive to examine this connection.
- The mathematical connection between the smoothing behaviour and the coarse grid suggests that it might be possible to prove the convergence of the conditional semicoarsening method for some class of problems.

Bibliography

- ÅSTRÖM, K. J. & WITTENMARK, B. 1990 *Computer-Controlled Systems*. Englewood Cliffs, New Jersey: Prentice Hall.
- BAGGETT, J. S. 1998 On the feasibility of merging LES with RANS for the near-wall region of attached turbulent flows. In *Annual Research Briefs*, pp. 267–277. Center for Turbulence Research.
- BAGGETT, J. S., JIMENEZ, J. & KRAVCHENKO, A. G. 1997 Resolution requirements in large-eddy simulations of shear flows. In *Annual Research Briefs*, pp. 51–66. Center for Turbulence Research.
- BALARAS, E., BENOCCI, C. & PIOMELLI, U. 1996 Two-layer approximate boundary conditions for large-eddy simulations. *AIAA Journal* **34** (6), 1111–1119.
- BARDINA, J., FERZIGER, J. H. & REYNOLDS, W. C. 1980 Improved subgrid scale models for large eddy simulation. *AIAA Paper 80-1357* .
- BATTEN, P., GOLDBERG, U. & CHAKRAVARTHY, S. 2004 Interfacing statistical turbulence closures with large-eddy simulation. *AIAA Journal* **42** (3), 485–492.
- BROWN, P. N., FALGOUT, R. D. & JONES, J. E. 2000 Semicoarsening multigrid on distributed memory machines. *SIAM Journal on Scientific Computing* **21** (5), 1823–1834.
- CABOT, W., JIMENEZ, J. & BAGGETT, J. S. 1999 On wakes and near-wall behavior in coarse large-eddy simulation of channel flow with wall models and second-order finite-difference methods. In *Annual Research Briefs*, pp. 343–354. Center for Turbulence Research.
- CABOT, W. & MOIN, P. 1999 Approximate wall boundary conditions in the large-eddy simulation of high Reynolds number flow. *Flow, Turbulence and Combustion* **63**, 269–291.
- CHAPMAN, D. R. 1979 Computational aerodynamics development and outlook. *AIAA Journal* **17** (12), 1293–1313.
- CHOI, H. & MOIN, P. 1994 Effects of the computational time step on numerical solutions of turbulent flow. *Journal of Computational Physics* **113**, 1–4.

- DAVIDSON, L. & BILLSON, M. 2006 Hybrid LES-RANS using synthesized turbulent fluctuations for forcing in the interface region. *International Journal of Heat and Fluid Flow* **in press**.
- DAVIDSON, L. & DAHLSTRÖM, S. 2004 Hybrid LES-RANS: An approach to make LES applicable at high Reynolds number. In *Proceedings of CHT-04*. Norway, April.
- DAVIDSON, L. & PENG, S. H. 2003 Hybrid LES-RANS modelling: a one-equation SGS model combined with a $k - \omega$ model for predicting recirculating flows. *International Journal for Numerical Methods in Fluids* **43**, 1003–1018.
- DE ZEEUW, P. M. 1990 Matrix-dependent prolongations and restrictions in a blackbox multigrid solver. *Journal of Computational and Applied Mathematics* **33** (1), 1–27.
- DEAN, R. B. 1978 Reynolds number dependence of skin friction and other bulk flow variables in two-dimensional rectangular duct flow. *Journal of Fluids Engineering* **100**, 215–223.
- DEARDORFF, J. W. 1970 A numerical study of three-dimensional turbulent channel flow at large Reynolds numbers. *Journal of Fluid Mechanics* **41**, 453–480.
- DEBNATH, L. & MIKUSINSKI, P. 1999 *Introduction to Hilbert Spaces with Applications*. Academic Press.
- DEL ALAMO, J. C. & JIMENEZ, J. 2003 Spectra of the very large anisotropic scales in turbulent channels. *Physics of Fluids* **15** (6), L41–L43.
- DEL ALAMO, J. C., JIMENEZ, J., ZANDONADE, P. & MOSER, R. D. 2004 Scaling of the energy spectra of turbulent channels. *Journal of Fluid Mechanics* **500**, 135–144.
- DRAZIN, P. G. & REID, W. H. 1981 *Hydrodynamic Stability*. Cambridge University Press.
- ELIAS, S. R., STUBLEY, G. D. & RAITHBY, G. D. 1997 An adaptive agglomeration method for additive correction multigrid. *International Journal for Numerical Methods in Engineering* **40** (5), 887–903.
- EMVIN, P. 1997 The full multigrid method applied to turbulent flow in ventilated enclosures using structured and unstructured grids. PhD thesis. Chalmers University of Technology.
- FERZIGER, J. H. & PERIC, M. 1997 *Computational Methods for Fluid Dynamics*. Springer Verlag.
- GERMANO, M., PIOMELLI, U., MOIN, P. & CABOT, W. H. 1991 A dynamic subgrid-scale eddy viscosity model. *Physics of Fluids* **3** (7), 1760–1765.
- GEURTS, B. J. 2004 *Elements of direct and large-eddy simulation*. Edwards.
- GOLDSTEIN, S. 1969 Fluid mechanics in the first half of this century. *Annual Review of Fluid Mechanics* **1**, 1–28.

- GRÖTZBACH, G. 1987 Direct numerical and large eddy simulation of turbulent channel flows. In *Encyclopedia of Fluid Mechanics* (ed. N. P. Cheremisinoff). West Orange, NJ.
- GULLBRAND, J. & CHOW, F. K. 2003 The effect of numerical errors and turbulence models in large-eddy simulations of channel flow, with and without explicit filtering. *Journal of Fluid Mechanics* **495**, 323–341.
- HAMBA, F. 2003 A hybrid RANS/LES simulation of turbulent channel flow. *Theoretical and Computational Fluid Dynamics* **16**, 387–403.
- HANJALIC, K. 2005 Will RANS survive LES? A view of perspectives. *Journal of Fluids Engineering* **127**, 831–839.
- HORTON, G. & VANDEWALLE, S. 1993 A space-time multigrid method for parabolic P.D.E.s. Technical report IMMD 3, 6/93. Universität Erlangen-Nürnberg.
- HOYAS, S. & JIMENEZ, J. 2005 Scaling of the velocity fluctuations in turbulent channels up to $Re_\tau = 2003$. In *Annual Research Briefs*, pp. 351–356. Center for Turbulence Research.
- JIMENEZ, J., DEL ALAMO, J. C. & FLORES, O. 2004 The large-scale dynamics of near-wall turbulence. *Journal of Fluid Mechanics* **505**, 179–199.
- JIMENEZ, J. & MOIN, P. 1991 The minimal flow unit in near-wall turbulence. *Journal of Fluid Mechanics* **225**, 213–240.
- JIMENEZ, J. & MOSER, R. D. 2000 Large-eddy simulations: Where are we and what can we expect? *AIAA Journal* **38** (4), 605–612.
- KEATING, A. & PIOMELLI, U. 2006 A dynamic stochastic forcing method as a wall-layer model for large-eddy simulation. *Journal of Turbulence* **7**, N12.
- KIM, J. 2003 Control of turbulent boundary layers. *Physics of Fluids* **15** (5), 1093–1105.
- KIM, J., MOIN, P. & MOSER, R. 1987 Turbulent statistics in fully developed channel flow at low Reynolds number. *Journal of Fluid Mechanics* **177**, 133–166.
- KRAJNOVIC, S. & DAVIDSON, L. 2003 Numerical study of the flow around a bus-shaped body. *Journal of Fluids Engineering* **125**, 500–509.
- KRAVCHENKO, A. G., MOIN, P. & MOSER, R. 1996 Zonal embedded grids for numerical simulations of wall-bounded turbulent flows. *Journal of Computational Physics* **127**, 412–423.
- LARSSON, J. 2001 Fluid dynamic development at Volvo Car Corporation. In *Time Compression Technologies Conference*. Manchester, UK.
- LARSSON, J., LIEN, F. S. & YEE, E. 2005a Conditional semicoarsening multigrid algorithm for the Poisson equation on anisotropic grids. *Journal of Computational Physics* **208**, 368–383.

- LARSSON, J., LIEN, F. S. & YEE, E. 2005*b* Large eddy simulation of high Reynolds number channel flow on coarse grids. In *13th Annual Conference of the Computational Fluid Dynamics Society of Canada*, pp. 61–68. St. Johns, Newfoundland, Canada.
- LARSSON, J., LIEN, F. S. & YEE, E. 2005*c* Parallel multigrid algorithm with conditional semicoarsening for fractional step type solvers. In *13th Annual Conference of the Computational Fluid Dynamics Society of Canada*, pp. 298–305. St. Johns, Newfoundland, Canada.
- LARSSON, J., LIEN, F. S. & YEE, E. 2006*a* Additional forcing at the interface in hybrid LES/RANS. In *14th Annual Conference of the Computational Fluid Dynamics Society of Canada*. Kingston, Ontario, Canada.
- LARSSON, J., LIEN, F. S. & YEE, E. 2006*b* Large eddy simulation with near wall modeling: the artificial buffer layer and the effects of forcing. *Submitted to Phys. Fluids* .
- LE, H. & MOIN, P. 1991 An improvement of fractional step methods for the incompressible Navier-Stokes equations. *Journal of Computational Physics* **92**, 369–379.
- LILLY, D. K. 1992 A proposed modification of the Germano subgrid-scale closure method. *Physics of Fluids* **4** (3), 633–635.
- MASON, P. J. & THOMSON, D. J. 1992 Stochastic backscatter in large-eddy simulations of boundary layers. *Journal of Fluid Mechanics* **242**, 51–78.
- MCBRYAN, O. A., FREDERICKSON, P. O., LINDEN, J., SCHÜLLER, A., SOLCHENBACH, K., STÜBEN, K., THOLE, C.-A. & TROTTENBERG, U. 1991 Multigrid methods on parallel computers – a survey of recent developments. *Impact of Computing in Science and Engineering* **3**, 1–75.
- MENEVEAU, C. & KATZ, J. 2000 Scale-invariance and turbulence models for large-eddy simulation. *Annual Review of Fluid Mechanics* **32**, 1–32.
- MOIN, P. 2001 *Fundamentals of Engineering Numerical Analysis*. Cambridge University Press.
- MOIN, P. & KIM, J. 1982 Numerical investigation of turbulent channel flow. *Journal of Fluid Mechanics* **118**, 341–377.
- MOIN, P. & MAHESH, K. 1998 Direct numerical simulation: A tool in turbulence research. *Annual Review of Fluid Mechanics* **30**, 539–578.
- MOSER, R. D., KIM, J. & MANSOUR, N. N. 1999 Direct numerical simulation of turbulent channel flow up to $Re_\tau = 590$. *Physics of Fluids* **11** (4), 943–945.
- MPI MANUAL 2003 *MPI-2: Extensions to the Message-Passing Interface*. Message Passing Interface Forum, <http://www.mpi-forum.org>.

- NICOUD, F., BAGGETT, J. S., MOIN, P. & CABOT, W. 2001 Large eddy simulation wall-modeling based on suboptimal control theory and linear stochastic estimation. *Physics of Fluids* **13** (10), 2968–2984.
- NIKITIN, N. V., NICOUD, F., WASISTHO, B., SQUIRES, K. D. & SPALART, P. R. 2000 An approach to wall modeling in large-eddy simulations. *Physics of Fluids* **12** (7), 1629–1632.
- PIOMELLI, U. 1993 High Reynolds number calculations using the dynamic subgrid-scale stress model. *Physics of Fluids* **5** (6), 1484–1490.
- PIOMELLI, U. 1999 Large-eddy simulation: Achievements and challenges. *Progress in Aerospace Sciences* **35** (4), 335–362.
- PIOMELLI, U. & BALARAS, E. 2002 Wall-layer models for large-eddy simulations. *Annual Review of Fluid Mechanics* **34**, 349–374.
- PIOMELLI, U., BALARAS, E., PASINATO, H., SQUIRES, K. D. & SPALART, P. R. 2003 The inner-outer layer interface in large-eddy simulations with wall-layer models. *International Journal of Heat and Fluid Flow* **24**, 538–550.
- PIOMELLI, U., FERZIGER, J., MOIN, P. & KIM, J. 1989 New approximate boundary conditions for large eddy simulations of wall-bounded flows. *Phys. Fluids A* **1** (6), 1061–1068.
- PODVIN, B. 2001 On the adequacy of the ten-dimensional model for the wall layer. *Physics of Fluids* **13** (1), 210–224.
- POPE, S. B. 2000 *Turbulent Flows*. Cambridge University Press.
- ROBINSON, S. K. 1991 Coherent motions in the turbulent boundary layer. *Annual Review of Fluid Mechanics* **23**, 601–639.
- RODI, W., FERZIGER, J. H., BREUER, M. & POURQUIE, M. 1997 Status of large eddy simulation: Results of a workshop. *Journal of Fluids Engineering* **119**, 248–261.
- SAGAUT, P. 2002 *Large Eddy Simulation for Incompressible Flows*. Springer Verlag.
- SCHAFFER, S. 1998 A semicoarsening multigrid method for elliptic partial differential equations with highly discontinuous and anisotropic coefficients. *SIAM Journal on Scientific Computing* **20** (1), 228–242.
- SCHUMANN, U. 1975 Subgrid scale model for finite difference simulations of turbulent flows in plane channels and annuli. *Journal of Computational Physics* **18**, 376–404.
- SPALART, P. R., JOU, W.-H., STRELETS, M. & ALLMARAS, S. R. 1997 Comments on the feasibility of LES for wings, and on a hybrid RANS/LES approach. In *Advances in DNS/LES* (ed. C. Liu & Z. Liu). Greyden, Columbus, OH.

- SPALART, P. R., MOSER, R. D. & ROGERS, M. M. 1991 Spectral methods for the Navier-Stokes equations with one infinite and two periodic directions. *Journal of Computational Physics* **96**, 297–324.
- TAM, C. K. W. & WEBB, J. C. 1993 Dispersion-relation-preserving finite difference schemes for computational acoustics. *Journal of Computational Physics* **107** (2), 262–281.
- TEMMERMAN, L., HADZIABDIC, M., LESCHZINER, M. & HANJALIC, K. 2005 A hybrid two-layer URANS-LES approach for large eddy simulation at high Reynolds numbers. *International Journal of Heat and Fluid Flow* **26**, 173–190.
- TEMPLETON, J. A., WANG, M. & MOIN, P. 2006 An efficient wall model for large-eddy simulation based on optimal control theory. *Physics of Fluids* **18**, 025101.
- TENNEKES, H. & LUMLEY, J. L. 1972 *A First Course in Turbulence*. MIT Press.
- TROTTEMBERG, U., OOSTERLEE, C. & SCHULLER, A. 2001 *Multigrid*. Academic Press.
- TUCKER, P. G. & DAVIDSON, L. 2004 Zonal $k - l$ based large eddy simulations. *Computers & Fluids* **33**, 267–287.
- VÖLKER, S., MOSER, R. D. & VENUGOPAL, P. 2002 Optimal large eddy simulation of turbulent channel flow based on direct numerical simulation statistical data. *Physics of Fluids* **14** (10), 3675–3691.
- WANG, M. & MOIN, P. 2002 Dynamic wall modeling for large-eddy simulation of complex turbulent flows. *Physics of Fluids* **14** (7), 2043–2051.
- WEI, T. & WILLMARTH, W. W. 1989 Reynolds-number effects on the structure of a turbulent channel flow. *Journal of Fluid Mechanics* **204**, 57–95.
- WILCOX, D. C. 2000 *Turbulence Modeling for CFD*. DCW Industries.
- WOSNIK, M., CASTILLO, L. & GEORGE, W. K. 2000 A theory for turbulent pipe and channel flows. *Journal of Fluid Mechanics* **421**, 115–145.
- WU, X. & SQUIRES, K. D. 1998 Prediction of the three-dimensional turbulent boundary layer over a swept bump. *AIAA Journal* **36** (4), 505–514.
- YAVNEH, I. 1995 Multigrid smoothing factors for red-black Gauss-Seidel relaxation applied to a class of elliptic operators. *SIAM Journal on Numerical Analysis* **32** (4), 1126–1138.
- YAVNEH, I. 1996 On red-black SOR smoothing in multigrid. *SIAM Journal on Scientific Computing* **17** (1), 180–192.
- ZANG, Y., STREET, R. L. & KOSEFF, J. R. 1994 A non-staggered grid, fractional step method for time-dependent incompressible Navier-Stokes equations in curvilinear coordinates. *Journal of Computational Physics* **114**, 18–33.