

UNIVERSIDAD POLITÉCNICA DE MADRID
Escuela Técnica Superior de Ingenieros Aeronáuticos



ENTRAINMENT EFFECTS IN
TURBULENT BOUNDARY LAYERS

TESIS DOCTORAL

Guillem Borrell Nogueras

Ingeniero Aeronáutico

Madrid, septiembre de 2015



POLITÉCNICA

Tribunal nombrado por el Sr. Rector Magfco. de la Universidad Politécnica de Madrid, el día.....de.....de 20....

Presidente: _____

Vocal: _____

Vocal: _____

Vocal: _____

Secretario: _____

Suplente: _____

Suplente: _____

Realizado el acto de defensa y lectura de la Tesis el día.....de.....de 20 ...
en la E.T.S.I. /Facultad.....

Calificación

EL PRESIDENTE

LOS VOCALES

EL SECRETARIO

UNIVERSIDAD POLITÉCNICA DE MADRID
Escuela Técnica Superior de Ingenieros Aeronáuticos

Entrainment Effects in Turbulent Boundary Layers

Tesis Doctoral

Guillem Borrell Nogueras

Ingeniero Aeronáutico

Madrid, Agosto de 2015

Abstract

This thesis studies the characteristics of the outer region of zero-pressure-gradient turbulent boundary layers at moderate Reynolds numbers. Two relatively established theories are put to test. The wall similarity theory states that with the presence of roughness, turbulent motion is mostly affected by the additional drag caused by the roughness, and that other secondary effects are restricted to a region very close to the wall. The consensus is that this theory is valid, but only as a first approximation. At the edge of the boundary layer there is a thin layer caused by the interaction between the turbulent eddies and the irrotational fluid of the free stream, called turbulent/non-turbulent interface. The bulk of results about this layer suggest the presence of some localized shear, with properties that make it distinguishable from the rest of the turbulent flow. The properties of the interface are likely to change if the rate of spread of the turbulent boundary layer is amplified, an effect that is usually achieved by increasing the drag. Roughness and entrainment are therefore linked, and the local features of the turbulent/non-turbulent interface may explain the reason why rough-wall boundary layers deviate from the wall similarity theory precisely far from the wall.

To study boundary layers at a higher Reynolds number, a new high-resolution code for the direct numerical simulation of a zero pressure gradient turbulent boundary layers over a flat plate has been developed. This code is able to simulate a wide range of Reynolds numbers from $Re_\tau = 100$ to 2000 while showing a linear weak scaling up to around two million threads in the BG/Q architecture. Special attention has been paid to the generation of proper inflow boundary conditions. The results are in good agreement with existing numerical and experimental data sets.

The turbulent/non-turbulent interface of a boundary layer is analyzed by thresholding the vorticity magnitude field. The value of the threshold is considered a parameter in the analysis of the surfaces obtained from isocontours of the vorticity magnitude. Two different regimes for the surface can be distinguished depending on the threshold, with a gradual topological transition across which its geometrical properties change significantly. The width of the transition scales well with δ_{99}^+ when $u_\tau^2/\nu\delta_{99}^{+1/2}$ is used as a unit of vorticity. The properties of the flow relative to the position of the vorticity magnitude isocontour are analyzed within the same range of thresholds, using the ball distance field, which can be obtained regardless of the size of the domain and complexity of the

interface. The properties of the flow at a given distance to the interface also depend on the threshold, but they are similar regardless of the Reynolds number. The interaction between the turbulent and the non-turbulent flow occurs in a thin layer with a thickness that scales with the Kolmogorov length. Deeper into the turbulent side, the properties are undistinguishable from the rest of the turbulent flow.

A zero-pressure-gradient turbulent boundary layer with a volumetric near-wall forcing has been simulated. The forcing has been designed to increase the wall friction without introducing any obvious geometrical effect. The actual simulation is split in two domains, a smaller one in charge of the generation of correct inflow boundary conditions, and a second and larger one where the forcing is applied. One-point and two-point statistics do not collapse beyond the logarithmic layer, but those differences can be explained as a consequence of the geometrical complexity of the intermittent region, and by the fact that the scaling with the wall-normal coordinate is no longer dominant. The geometrical effects can be avoided using the turbulent/non-turbulent interface as a reference frame, and the minimum distance respect to it. The conditional analysis of the vorticity field with the alternative reference frame recovers the scaling with δ_{99} and ν/u_τ already present in the logarithmic layer, the only two length-scales allowed if Townsend's wall similarity hypothesis is valid.

Resumen

Esta tesis estudia el comportamiento de la región exterior de una capa límite turbulenta sin gradientes de presiones. Se ponen a prueba dos teorías relativamente bien establecidas. La teoría de semejanza para la pared supone que en el caso de haber una pared rugosa, el fluido sólo percibe el cambio en la fricción superficial que causa, y otros efectos secundarios quedarán confinados a una zona pegada a la pared. El consenso actual es que dicha teoría es aproximadamente cierta. En el extremo exterior de la capa límite existe una región producida por la interacción entre las estructuras turbulentas y el flujo irrotacional de la corriente libre llamada interfaz turbulenta/no turbulenta. La mayoría de los resultados al respecto sugieren la presencia de fuerzas de cortadura ligeramente más intensa, lo que la hace distinta al resto del flujo turbulento. Las propiedades de esa región probablemente cambien si la velocidad de crecimiento de la capa límite aumenta, algo que puede conseguirse aumentando la fricción en la pared. La rugosidad y la ingestión de masa están entonces relacionadas, y el comportamiento local de la interfaz turbulenta/no turbulenta puede explicar el motivo por el que las capas límite sobre paredes rugosas no se comportan como en el caso de tener paredes lisas precisamente en la zona exterior.

Para estudiar las capas límite a números de Reynolds lo suficientemente elevados, se ha desarrollado un nuevo código de alta resolución para la simulación numérica directa de capas límite turbulentas sin gradiente de presión. Dicho código es capaz de simular capas límite en un intervalo de números de Reynolds entre $Re_\tau = 100 - 2000$ manteniendo una buena escalabilidad hasta los dos millones de hilos en superordenadores de tipo Blue Gene/Q. Se ha guardado especial atención a la generación de condiciones de contorno a la entrada correctas. Los resultados obtenidos están en concordancia con los resultados previos, tanto en el caso de simulaciones como de experimentos.

La interfaz turbulenta/no turbulenta de una capa límite se ha analizado usando un valor umbral del módulo de la vorticidad. Dicho umbral se considera un parámetro para analizar cada superficie obtenida de un contorno del módulo de la vorticidad. Se han encontrado dos regímenes distintos en función del umbral escogido con propiedades opuestas, separados por una transición topológica gradual. Las características geométricas de la zona escalan con δ_{99}^+ cuando $u_\tau^2/\nu\delta_{99}^{+1/2}$ es la unidad de vorticidad. Las propiedades del fluido relativas a la posición del contorno de vorticidad han sido analizados para una serie de umbrales utilizando el campo de distancias esféricas, que puede obtenerse con independencia

de la complejidad de la superficie de referencia. Las propiedades del fluido a una distancia dada del interfaz también dependen del umbral de vorticidad, pero tienen características parecidas con independencia del número de Reynolds. La interacción entre la turbulencia y el flujo no turbulento se restringe a una zona muy fina con un espesor del orden de la escala de Kolmogorov local. Hacia el interior del flujo turbulento las propiedades son indistinguibles del resto de la capa límite.

Se ha simulado una capa límite sin gradiente de presiones con una fuerza volumétrica cerca de la pared. La el forzado ha sido diseñado para aumentar la fricción en la pared sin introducir ningún efecto geométrico obvio. La simulación consta de dos dominios, un primer dominio más pequeño y a baja resolución que se encarga de generar condiciones de contorno correctas, y un segundo dominio mayor y a alta resolución donde se aplica el forzado. El estudio de los perfiles y los coeficientes de autocorrelación sugieren que los dos casos, el liso y el forzado, no colapsan más allá de la capa logarítmica por la complejidad geométrica de la zona intermitente, y por el hecho que la distancia a la pared no es una longitud característica. Los efectos causados por la geometría de la zona intermitente pueden evitarse utilizando el interfaz como referencia, y la distancia esférica para el análisis de sus propiedades. Las propiedades condicionadas del flujo escalan con δ_{99} y ν/u_τ , las dos únicas escalas contenidas en el modelo de semejanza de pared de Townsend, consistente con estos resultados.

Acknowledgements

I'd like to thank everyone who made this possible.

Any scientific work stands on the shoulders of giants, and I have been lucky enough to have one giant teaching me how Science is meant to be done. From him I learned that *why* is a mighty and useful word.

I am deeply thankful to my parents, Antoni and Maria Isabel, who encouraged me to reach my goals without shortcuts. I'd like to thank my sister Maria Isabel, who gave me some sanity in the hardest times, and Beatriz, for standing by my side and cheering me when I needed it the most. There is no greater achievement in life than making the ones you love proud, your grandmother in particular.

I was once told that when you are the smartest person in the room, you are in the wrong room. That's why I am so happy to have shared space and time with Adrián, Juan, Óscar, Mark, Juan Carlos, Sergio, Pablo, Ayse, Florian, Jorge, Miguel P., Miguel H., Alberto, Fer, Julio, Yvan, Kike, Jose, Siwei, Atsushi, Jason... I also have to thank the staff at BSC and JSC, because I know that keeping supercomputers alive is not an easy job.

This work of was funded by CICYT under grant TRA2009-11498, and by the European Research Council under grant ERC-2010.AdG-20100224. Figure 3.2 was obtained with the help of the Barcelona Supercomputing Center. The computational resources of the Argonne Leadership Computing Facility at Argonne National Laboratory were supported by the U.S. Department of energy under contract No DE-AC02-06CH11357, and the computational resources of the Juelich Forschungszentrum were provided by the PRACE initiative.

Als meus pares. Que quan treia un deu, volien un onze.

‘Anything becomes interesting if you look at it long enough.’

Gustave Flaubert

‘If you look at a thing long enough, it loses all of its meaning.’

Andy Warhol

Contents

Abstract	i
Resumen	iii
Acknowledgements	v
Contents	xi
List of Tables	xiii
List of Figures	xv
Glossary	xvii
1 Introduction	1
1.1 The turbulent motion of fluids	1
1.2 Turbulent wall bounded flows.	3
1.3 Turbulent boundary layers.	7
1.4 Turbulent entrainment of irrotational fluid.	8
1.5 Boundary layers over rough walls.	10
1.6 Turbulence as a fractal.	11
1.7 Direct Numerical Simulation of turbulent flows.	12
1.8 Goals	14
1.9 Structure of this thesis.	15
2 DNS of turbulent boundary layers at high Re	17
2.1 Introduction	17

2.2	The numerical code	20
2.3	Scalability and portability	26
2.4	Parallel I/O	28
2.5	Validation	28
2.6	Conclusions and future work	30
3	Properties of the T/NT interface	33
3.1	Introduction	33
3.2	Description of the data.	37
3.3	The turbulent/non-turbulent interface detection criterion.	39
3.4	The geometry of the T/NT interface.	46
3.5	Conditional analysis of the vorticity field.	53
3.6	Other velocity gradients	71
3.7	Conclusions.	76
4	The effect of increased entrainment	79
4.1	Introduction	79
4.2	The numerical experiment	83
4.3	Average properties.	87
4.4	Velocity field structure	97
4.5	The effects of the additional entrainment in the intermittent zone	104
4.6	Conclusions	110
5	Conclusions and future work	113
	Bibliography	119

List of Tables

2.1	Computational setup for the auxiliary BL_1 and main BL_2 boundary layers in Sillero <i>et al.</i> (2013)	22
2.2	Time spent in communication during global transposes	25
2.3	OpenMP scalability test	26
2.4	Data collected from the profiled test cases	32
3.1	Summary of important parameters of the simulation.	38
3.2	Properties of $Q(y_I)$ for the different values of $\langle y_I \rangle$ found in the literature. . . .	43
4.1	Summary of the important parameters of the four simulations used more often in this study.	84

List of Figures

1.1	Sketch of the coordinates based on the position of the wall	4
1.2	Evidence of the existence of the logarithmic layer	6
1.3	Regimes of the evolution of a boundary layer.	7
1.4	Example of intermittent signal obtained by a hot wire	8
1.5	Enstrophy contour $\omega^+ = 0.4$ of a boundary layer at $Re_\tau \simeq 2000$	12
1.6	Evolution of the performance of supercomputers	13
2.1	Schema of the computational domain and boundary conditions	20
2.2	Elemental domains of the domain decomposition.	22
2.3	Node mappings for BG/P	24
2.4	Scalability tests.	27
2.5	Validation tests.	29
3.1	Sketch of the simulation and the boundary conditions.	38
3.2	Vorticity magnitude isosurfaces	40
3.3	Quantitative analysis of the intermittent region	41
3.4	Position of the T/NT interface	44
3.5	Premultiplied joint PDF of vorticity and \tilde{k} , for a boundary layer and a temporal jet respectively.	46
3.6	Sketch of the three basic geometrical features in the vorticity isosurface: (a) handles, (b) pockets and (c) bubbles.	47
3.7	(a) Set $\Omega_{\omega>}$ of voxels where $\omega > \omega_0$. (b) Set $\Omega_{\omega<}$ of voxels where $\omega < \omega_0$. (c) The set Ω_i of voxels that contain the interface, obtained from (3.12).	48
3.8	Fractal dimension of the interface	50
3.9	Genus of the interface	52

3.10	Sketches of the distance definitions	54
3.11	Sketches of the properties of the vertical distance	55
3.12	(a) Set Ω_i of voxels that contain the interface. (b) Set Ω_p of points used to approximate the T/NT interface.	55
3.13	Cross-stream sections of the distance fields	56
3.14	Analysis of the surface $\Delta = 0$	57
3.15	Premultiplied joint probability density function of vorticity and distance, $\omega F_{\omega,\Delta}$	59
3.16	Analysis of the overhangs present in the T/NT interface	60
3.17	Conditional vorticity profiles	62
3.18	Effect of the threshold on the conditional vorticity profiles	64
3.19	Sketch of the geometrical meaning of equation (3.22)	65
3.20	Analysis of the vorticity within pockets.	67
3.21	Premultiplied joint probability density function of the vorticity and ball distance in the turbulent side of the interface ($\omega F_{\omega,\Delta_b}$)	69
3.22	Premultiplied joint PDFs: $S\Gamma_{y,S}$, $\omega\Gamma_{y,\omega}$	71
3.23	Fractal dimension and genus of the strain interface.	72
3.24	Conditional mean profiles for the vorticity and strain interfaces.	74
3.25	Properties of strain and stretching within the T/NT interface layer.	75
4.1	Sketch of the twin box configuration	84
4.2	Spatial evolution of the Clauser wake parameter G	88
4.3	Spatial evolution of the wall-normal term of the turbulent kinetic energy velocity budget	90
4.4	Profiles of turbulent kinetic energy production and dissipation	91
4.5	Comparison of the wall-normal profiles for the cases given in table 4.1	93
4.6	Sections of C_{uu}	98
4.7	Sections of C_{vv}	100
4.8	Sections of C_{ww}	101
4.9	Joint PDF of vorticity magnitude ω^* and distance to the wall y/δ_{99} , $\Gamma_{\omega,y}$	104
4.10	Geometrical properties of the T/NT interface of BL_f and BL_s	106
4.11	Premultiplied joint PDF of vorticity and minimum distance $\omega F_{\omega,\Delta_b}$	108

Glossary

- A The area that limits a region in space of volume V . 2
- $C_{\phi\phi}$ Autocorrelation coefficient. See equation (4.10). 96
- G Clauser wake parameter. See equation (4.9). 87
- L_i Relevant parameters that define a roughness pattern other than its average height. 79
- Re_λ Reynolds number based on the Taylor microscale. 38
- Re_τ Friction Reynolds number. 5
- Re_θ Reynolds number based on the momentum thickness θ . 18
- U_∞ Magnitude of the velocity of the free stream. 4
- V An arbitrary volume. 2
- $\Gamma_{\omega,y}$ Joint probability density function of vorticity magnitude and wall-normal distance. 40, 42
- ω Vorticity. 9, 39
- $\langle \cdot \rangle$ Ensemble average operator. 4
- δ^* Displacement thickness. 8
- δ_{99} Boundary layer thickness. 7, 37
- ϵ Dissipation of kinetic energy per unit mass and time. 2
- η Kolmogorov length. 37
- γ Intermittency parameter. 42
- \hat{x} Eddy turnover distance. See equation (4.8). 87
- κ Kármán constant. 5, 42
- λ Taylor microscale. 37

- ∇ Gradient vector operator. 2
- $\nabla \wedge$ Rotor vector operator. 9
- ∇^2 Laplace operator. 2
- ν Kinematic viscosity. 2, 37
- ω Vorticity magnitude. 11, 39
- ω_0 Vorticity value used to threshold a vorticity magnitude field. 40
- ρ Density. 2
- θ Momentum thickness. 8, 82
- \mathbf{u} Velocity vector. 2, 37
- \mathbf{x} Position coordinate vector. 2
- c_f Skin drag friction coefficient. 8
- k Average height of the roughness elements in a rough-wall boundary layer. 79
- p Pressure. 2
- u Component of \mathbf{u} along the x coordinate. 4, 37
- u_τ Friction velocity. 4, 37
- v Component of \mathbf{u} along the y coordinate. 4, 37
- w Component of \mathbf{u} along the w coordinate. 4, 37
- x Coordinate in the direction of the free stream. 3, 37
- y Wall-normal coordinate. 3, 37
- z Coordinate normal to the plane defined by x , and y . 4, 37

Acronyms

ALCF Advanced Leadership Computing Facility. 18

BG/P Blue Gene/P supercomputer. 23

BG/Q Blue Gene/Q supercomputer. 26

BL Boundary Layer. 21, 37

DNS Direct Numerical Simulation. 12

DoF Degrees of freedom. 22

GF Giga FLOPS. 13

JSC Jülich Supercomputing Center. 18

MPI Message Passing Interface. 17

PIV Particle Image Velocimetry. 34

PTV Particle Tracking Velocimetry. 34

T/NT turbulent/non-turbulent interface. 79

Chapter 1

Introduction

Turbulence is a fundamental phenomenon in fluids. It is the most common and resilient mechanism for which a fluid dissipates its kinetic energy, and turns it into heat. Turbulence has been into intense research since its key role was discovered, and it is now a fundamental part of engineering. Despite turbulence modeling is ubiquitous in many practical applications of Computational Fluid Dynamics, many details of the nature of turbulent motion remain uncovered.

This thesis deals with some aspects of a particular instance of turbulent flow, the zero-pressure-gradient turbulent boundary layer, that have not been yet understood. Despite it is one of the many realizations of turbulence, it is of great interest in aeronautics, turbomachinery and meteorology.

1.1 The turbulent motion of fluids

Turbulence has been a relatively obvious phenomenon to the observation of nature. There has always been a word for the disordered motion of wind and water in the chaos and turmoil of adverse weather, called *turbare* by Gonzalo de Berceo (1197-1264). The first experiments explicitly targeted to test turbulence were due to Leonardo da Vinci (1452-1519), although they were aimed to a pictorial purpose. He famously sketched the whirls of a current of water, describing it as tangled hair.

The scientific study of Turbulence in fluid started with Hagen (1854), and Darcy (1857), with the goal of understanding the pressure drop of fluids moving in closed ducts. The first precise phenomenological description is due to their research. They found that the

drag force suffered by the fluid had two components, one proportional to the velocity of the fluid, and a second one proportional to the square of the velocity. This was then a paradox, because the observation implied that the drag force does not vanish even when the viscosity is nominally zero.

Reynolds (1895) contributed the fundamental parameter for the appearance of disordered motion in fluids, named after him. He also pioneered the mathematical treatment of the problem. He found that the small features of the motion were not measurable at that time, and he proposed to decompose the Navier-Stokes equation as a temporal average (the Reynolds average), and a fluctuating term. This formulation became the foundation of turbulence modeling. The first approximate description of the nature of turbulent motion is due to Richardson (1922), who proposed a hierarchy of structures, or eddies, of different sizes. The largest ones contain a significant amount of kinetic energy, but the gradients within are not strong enough to produce any dissipation. Kinetic energy dissipation therefore not direct. The big eddies transfer the kinetic energy to smaller eddies, and those to even smaller eddies until the gradients are strong enough for the kinetic energy to be dissipated by viscosity. This description is not necessarily accurate, but it provides the essential concepts needed to solve the previously mentioned paradox. An eddy can be defined as a coherent region of the fluid with size ℓ that contains a velocity difference of u_ℓ . The rate of energy dissipated for unit mass and time by each eddy follows necessarily the law $\epsilon \sim u_\ell^3/\ell$. The Navier-Stokes equations for an incompressible fluid can be written as

$$\nabla \cdot \mathbf{u} = 0 \tag{1.1a}$$

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\rho^{-1} \nabla p + \nu \nabla^2 \mathbf{u}, \tag{1.1b}$$

where \mathbf{u} is the velocity vector, ∇ is the gradient vector operator, ρ is the density of the fluid, p is pressure, ν is the kinematic viscosity, and ∇^2 is the Laplace operator. An equation for $\epsilon = \partial(V^{-1} \int_V |\mathbf{u}|^2/2 d\mathbf{x}^3)/\partial t$ can be obtained multiplying (1.1b) by \mathbf{u} , and integrating over a volume V :

$$\epsilon = \frac{-\Delta p}{\rho V} \int_A \mathbf{u} d^2\mathbf{x} - \frac{\nu}{V} \int_V |\nabla \mathbf{u}|^2 d^3\mathbf{x}, \tag{1.2}$$

where A is the boundary of V , and \mathbf{x} is the coordinate vector. When ϵ becomes independent of viscosity in the limit of $\nu = 0$, $|\nabla \mathbf{u}|^2$ has to become almost singular. Turbulence is the mechanism for those strong gradients to appear far from a solid boundary.

Turbulence is a multiscale phenomenon that ranges from the largest possible scale present in the phenomenon to the probably microscopic scale where the energy is dissipated by viscosity. Assuming that the eddies mentioned earlier can develop without the influence of any boundary, they can be classified from their characteristic length.

The integral scale L represents the largest scale that can be defined in the phenomenon, the characteristic size of the forces that produce the motion of the fluid. In the case of a pipe, it could be the radius or the diameter. The Reynolds number associated to those eddies is usually very large. A pipe of radius $L = 0.1m$ that transports fresh water at a velocity of $U = 10m/s$ has a characteristic Reynolds number of $UL/\nu \simeq 10^6$.

The Kolmogorov's length η is the scale where the gradients of velocity are important, and the viscous dissipation of kinetic energy is relevant. The Reynolds number of the eddies of size η has to be of the order of $(\epsilon\eta^4/\nu^3)^{1/3} \sim 1$, and the Kolmogorov scale is therefore $\eta = (\nu^3/\epsilon)^{1/4}$. Note that the characteristic velocity associated with eddies of size η is $u_\eta = (\nu\epsilon)^{1/4}$, and the characteristic velocity gradient is $u_\eta/\eta = (\nu\epsilon)^{1/2}$.

If L and η are different enough, there is a range of intermediate scales l that are not influenced by the mechanism of energy injection nor dissipation, and the eddies have no characteristic length-scale, meaning that any length between L and η can be found on the flow, but none will be particularly relevant. Kolmogorov (1941) deduced from dimensional arguments that the associated one-dimension energy spectrum follows

$$E_{uu}(k) \sim \epsilon^{2/3} k^{-5/3}, \quad (1.3)$$

which is widely known as Kolmogorov's 5/3 law, with $k \sim 1/l$ the wavenumber. It quantifies the amount of energy that, on average, the eddies of size l transfer to those immediately smaller. Despite this is a relatively simplistic description, it is one of the most successful results in the field.

1.2 Turbulent wall bounded flows.

The previous classification, despite it roughly describes any turbulent flow, is not accurate in the presence of an inhomogeneity. Particularly in the case of a solid wall where a no slip ($\mathbf{u} = 0$) boundary condition is applied. The presence of a solid boundary introduces another fundamental length-scale in the flow, the distance along its perpendicular direction y . The other two fundamental coordinates are x , in the direction of the free stream

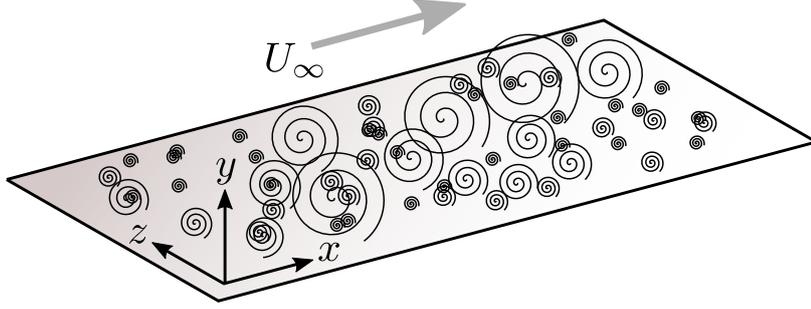


Figure 1.1: Sketch of the coordinates based on the position of the wall. The turbulent eddies are represented by whirls of different sizes

and z , perpendicular to the previous two. The components along each coordinate are u , v and w respectively, and the magnitude of the velocity of the free stream is U_∞ . A sketch of the previous definitions is presented in figure 1.1.

We can define the Reynolds average as the operation $\phi = \langle \phi \rangle + \tilde{\phi}$, where an arbitrary variable ϕ is decomposed as a mean term $\langle \phi \rangle$, and a fluctuations term $\tilde{\phi}$. In the particular case of a wall-bounded flow, the averaging of $\langle \cdot \rangle$ occurs along the three statistically homogeneous dimensions x , z and time. In consequence $\langle \phi \rangle$ is only a function of y . If we apply the Reynolds average to equation (1.1b), assuming the presence of the wall and homogeneity along the z coordinate, we obtain

$$\nabla \cdot \langle \mathbf{u} \rangle = 0 \quad (1.4a)$$

$$\langle \mathbf{u} \rangle \cdot \nabla \langle \mathbf{u} \rangle + \rho^{-1} \nabla \langle p \rangle = \nu \nabla^2 \langle \mathbf{u} \rangle - \nabla \cdot \langle \tilde{\mathbf{u}} \tilde{\mathbf{u}} \rangle, \quad (1.4b)$$

Very close to the wall, viscosity has to be dominant over the rest of the terms, and equation (1.4b) is trivially integrable to

$$\langle u \rangle = y \left. \frac{\partial \langle u \rangle}{\partial y} \right|_{y=0}. \quad (1.5)$$

With the expression of the shear stress at the wall

$$\tau_w = \nu \left. \frac{\partial \langle u \rangle}{\partial y} \right|_{y=0}, \quad (1.6)$$

and defining the friction velocity as $u_\tau = \sqrt{\tau_w}$, the non-dimensional identity

$$\frac{\langle u \rangle}{u_\tau} = y \frac{u_\tau}{\nu} \quad (1.7)$$

defines the characteristic scales for velocity and length in the near-wall region, the friction velocity u_τ and the viscous length ν/u_τ respectively. Equation (1.7) can be written as $\langle u \rangle^+ = y^+$, where the superscript $+$ denotes the use of *wall units*. Note that a characteristic time can also be defined as ν/u_τ^2 . The zone where the mean velocity profile fits this description, and viscosity is the dominant feature of the flow, is called viscous sublayer.

The separation of scales is a fundamental requirement for some properties of turbulence to emerge, as it was mentioned in §1.1. A similar argument can be done about the separation between the small scales close to the wall, of length ν/u_τ , and the characteristic length of the outer scales of the flow like the radius of the pipe in the example in §1.1, that it will be generally called δ . According to Townsend (1976), the eddies in a wall bounded flow have a characteristic velocity of u_τ and a size proportional to the distance to the wall y . In consequence, using the previous dimensional argument given by Kolmogorov, the energy-dissipation rate is $\epsilon \propto u_\tau^3/y$. In addition, the amount of kinetic energy injected by the largest scales into the turbulent motion, also called production of turbulent kinetic energy, is¹ $\Pi = u_\tau^2 \partial \langle u \rangle / \partial y$. If the scale separation is sufficiently large, meaning $\nu/u_\tau \ll y \ll \delta$, there is local equilibrium between the amount of energy introduced by the shear and the energy dissipated by the turbulent eddies, therefore

$$u_\tau^2 \frac{\partial \langle u \rangle}{\partial y} = \kappa \frac{u_\tau^3}{y}, \quad (1.8)$$

where κ is the Kármán constant, the ratio between the local production and dissipation. If equation (1.8) is integrated, one obtains a logarithmic solution for the mean velocity profile

$$\langle u \rangle^+ = \frac{1}{\kappa} \log y^+ + C, \quad (1.9)$$

where C is an additive constant that depends on the details of the flow near the wall. While the deduction of the presence of a logarithmic region in the mean profile of the streamwise component of velocity is subject to strong assumptions, it has been repeatedly observed in a wide range of wall bounded flows with sufficient separation between the viscous and the largest scales (see figure 1.2). The separation between ν/u_τ and δ is therefore crucial, and the ratio between these two quantities defines the Reynolds number of a wall-bounded flow $Re_\tau = \delta u_\tau / \nu$, called *friction Reynolds number*. The properties of the turbulent flow in this region were recently reviewed by Jiménez (2013).

¹The proof of this identity is omitted for brevity

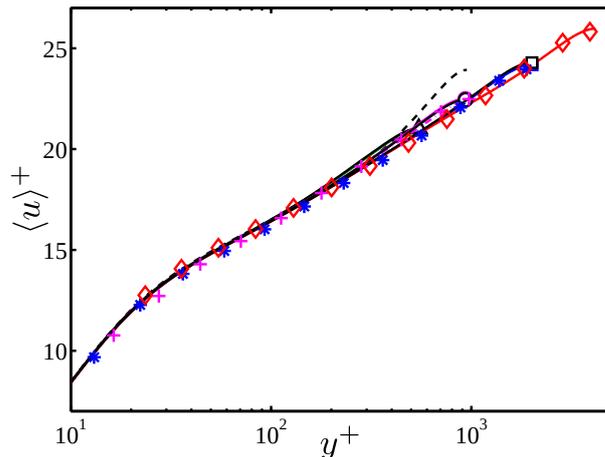


Figure 1.2: Evidence of the existence of the logarithmic layer described in equation (1.9) in turbulent channels within a range of Re_τ . Figure from Lozano-Durán & Jiménez (2014 a).

Below the logarithmic layer, where $y \gg \nu/u_\tau$, and above the viscous sublayer, where $y^+ \sim 1$ there is an overlap region where the two scales (y and ν/u_τ) are relevant. That region is frequently called *buffer layer*, and it is visible at $y^+ = 10 - 100$ in figure 1.2. The turbulent structures in the buffer layer are nonlinear and moderately chaotic, but their configuration is simple enough to be fully described by Jiménez & Moin (1991) and Jiménez & Pinelli (1999).

We have examined two of the three characteristic regions of a wall-bounded flow. The viscous sublayer and the logarithmic layer are relatively universal, since they are strongly influenced by the presence of a wall. They have been under intense research, and they have been modeled with relative success. These two zones correspond to roughly the 95% of the energy dissipation, but only the 25% of the volume. The region where the influence of δ is important is not particularly intense, but it is much less understood. The influence of large scales depend strongly on the configuration of the largest structures that carry most of the kinetic energy. In the case of the turbulent channels and pipes, those structures are at the center, and they are pushed by the pressure gradient. In the case of a boundary layer, the largest scales appear naturally as the byproduct of the interaction between the turbulent motion and the free stream. In any case, it has been repeatedly observed that these largest scales are not universal, and depend on the geometrical configuration of the problem.

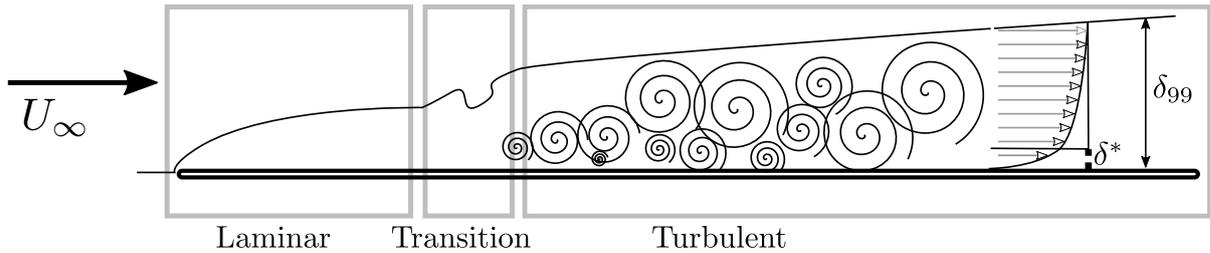


Figure 1.3: Regimes of the evolution of a boundary layer. The solid line that originates at the edge of the plate represents the far edge of the boundary layer. Beyond that line the free stream is not perturbed. At the right side of the figure there is a sketch of the average streamwise velocity profile, and the graphical interpretation of the displacement thickness δ^* , and the boundary layer thickness δ_{99} .

1.3 Turbulent boundary layers.

A turbulent boundary layer is a turbulent bounded flow that occurs when there is a relative motion between a solid boundary and a fluid at a sufficient speed. Turbulent boundary layers are particularly relevant in aeronautics. Roughly 20% of the drag of a plane flying at cruise speed is contributed by the skin drag, that occurs within the turbulent boundary layer. They are also important in meteorology. The interaction between the winds and the terrain often forms a turbulent boundary layer, that characterizes the lowest region of the troposphere.

Assume a free stream at a velocity U_∞ and a flat plate with infinitesimal thickness aligned with the direction of the current, like the ones sketched in figure 1.3. The interaction between the edge of the plate and the flow forms a stagnation point, followed by a laminar region that concentrates the gradients between the free stream (at a velocity of U_∞) and the solid wall (at $\mathbf{u} = 0$), a laminar boundary layer. This phenomenon was first described by Prandtl (1904), and an analytic solution in absence of external pressure gradients was found by Blasius (1950). If the boundary layer is let to evolve, it eventually transitions to a turbulent regime. The transitional behavior of boundary layers is one of the main topics in the field of linear stability analysis of fluids, but it is out of the scope of this work.

The turbulent regime of a boundary layer is usually characterized by the following parameters. A common measure for the thickness of the boundary layer is δ_{99} , which is also a good estimation of the size of the largest scales L . It is defined as the wall-normal coordinate where the average streamwise velocity is equal to the 99% of the free stream

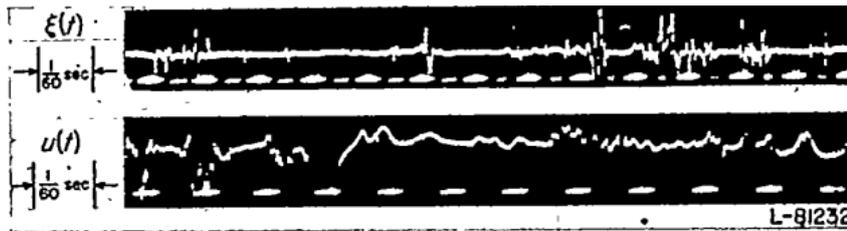


Figure 1.4: Example of intermittent signal obtained by a hot wire in the outer region of a turbulent boundary layer. From Corrsin & Kistler (1955)

$\langle u \rangle|_{y=\delta_{99}} = 0.99U_\infty$. The second parameter is the displacement thickness δ^* , defined as the distance by which the wall would have to be moved in the wall-normal direction to give the same flow rate in an unperturbed stream as occurs in the real fluid:

$$\delta^* = \int_0^\infty \left(1 - \frac{\langle u \rangle}{U_\infty}\right) dy. \quad (1.10)$$

A graphical interpretation of this equation is presented in the rightmost part of figure 1.3. The displacement thickness evaluates the streamwise growth of the boundary layer, since $d\delta^*/dx$ is the amount of mass that the boundary layer entrains per unit length. A similar quantity can be defined with the momentum of the free stream and of the real fluid, called momentum thickness θ , defined as

$$\theta = \int_0^\infty \frac{\langle u \rangle}{U_\infty} \left(1 - \frac{\langle u \rangle}{U_\infty}\right) dy. \quad (1.11)$$

This quantity is also relevant to characterize the streamwise evolution of the boundary layer. The integration of (1.1b) in absence of pressure gradients along the streamwise direction yields to the Kármán momentum integral equation

$$\frac{c_f}{2} = \frac{d\theta}{dx}, \quad (1.12)$$

where $c_f = u_\tau^2/\rho U_\infty^2$ is the friction coefficient. Equation (1.12) links the two phenomena studied in this thesis, the friction that occurs at the near-wall region, and the entrainment that occurs in the outer region of the boundary layer.

1.4 Turbulent entrainment of irrotational fluid.

The outer region of a boundary layer is also called *intermittent region*, because of the characteristic signal produced by a hot wire when it is placed there (see figure 1.4). A wall-parallel plane sufficiently far from the wall, say $y = \delta_{99}$, contains roughly the same

amount of turbulent flow than irrotational flow in the free stream. In consequence, a hot wire placed at that height would output an intermittent signal made from smooth and fast-changing sections. This opens an important question about the characteristics of boundary layers in comparison to other wall-bounded flows. Does the presence of irrotational flow affect the turbulent motion? In other words. Are the turbulent eddies in the outer region of the boundary layer aware that there is irrotational fluid nearby? If the answer is affirmative, the local dynamics of the interaction between the turbulent and the non-turbulent flow have to be described. If the answer is negative, the free stream is just a boundary-less condition for the turbulent eddies, that keep their collective behavior unaltered until they are dissipated.

This interaction does exist. The mass acquired by the boundary layer during its spatial evolution has to be obtained from the free stream. Entrainment is the process for which the turbulent motion propagates by giving vorticity (and any other fluctuation that is characteristic of the turbulent motion) to fluid that was previously irrotational. Entrainment is present in any external turbulent flow like jets and mixing layers, and its importance in the dynamics of turbulence in the outer region of the boundary layer are not yet determined.

The question is therefore how entrainment affects turbulent motion. One defining property of turbulence is vorticity $\boldsymbol{\omega} = \nabla \wedge \mathbf{u}$, where $\nabla \wedge$ is the rotor operator. Within the turbulent flow vorticity can be created, advected amplified or damped. On the other hand, the irrotational free stream has no mechanism to produce vorticity from scratch. It is clear that in the last term, the non-turbulent flow becomes turbulent by direct interaction, the question is whether an additional process promotes vorticity diffusion.

If vorticity propagation is purely local, a concept that has been frequently called nibbling, the interaction between the turbulent and the non-turbulent can be simple. On the other hand, if entrainment is a multi-scale process that involves the transmission of momentum from the large coherent structures to the non-turbulent flow, called engulfment by some authors, it is less probable that the interaction between the two states of the fluid can be given a simple description. This questions motivate a thorough analysis of the intermittent zone and the thin layer where vorticity diffuses out to the non-turbulent region, the turbulent/non-turbulent interface.

1.5 Boundary layers over rough walls.

Many boundary layers of practical interest happen over rough walls. The atmospheric boundary layer, driven by the geostrophic component of the wind, flows over a wide variety of roughness patterns. Large ships accumulate significant amounts of fouling, that can affect the properties of the boundary layer around the vessel and reduce its performance.

The most important effect of roughness is, with little exceptions, to increase the friction coefficient. At the present moment there is no widely accepted model to predict the friction coefficient from a given roughness pattern, an indication on how complex the interaction between a non-smooth solid wall and the structure of turbulent motion can be. There are, however, relevant theories about the mentioned interaction, like the wall similarity hypothesis by Townsend (1976). According to that hypothesis, with sufficient scale separation, and if the roughness does not interact directly with the logarithmic layer, the effect that the roughness has on the turbulent motion in boundary layers is confined within a region with a thickness of the order of the height of the roughness. Above that region of influence, the rest of the flow only perceives the change in the friction coefficient. Given the conditions, the range of application of this hypothesis may seem small, but a relevant amount of atmospheric and industrial flows occur within this parameter range. Townsend's hypothesis is relevant for the modeling of boundary layers over rough surfaces, and determining its limitations is of practical interest.

The consensus nowadays is that Townsend's hypothesis is valid, but only as a first approximation. Boundary layers with shallow roughness behave almost like boundary layers over smooth walls, but when friction is increased by a relevant amount, some of the basic statistics do not scale with the corresponding friction coefficient. It is significative that the similarity between rough and smooth walls is lost precisely far from the wall, in the intermittent region. If we remember equation (1.12), the friction coefficient is roughly proportional to the rate of growth of the boundary layer, that is related to the amount of mass that is entrained per unit length. This connection between roughness and intermittency suggests another explanation for the lack of accuracy of Townsend's similarity hypothesis, that was suggested by Robert Antonia in a private conversation with Javier Jiménez. The differences between smooth and rough-wall boundary layers in the intermittent region could be explained by the additional entrainment, and roughness would be only indirectly involved in the mentioned lack of agreement.

The characteristic scale of the intermittent region is δ (δ_{99} in this case), and there is no particular reason to think that the distance to the wall y is the most important scale of length past $y = 0.5\delta_{99}$. In consequence, the differences in the mean profiles between the two cases may be caused by y not being a suitable variable, neither $\langle \cdot \rangle$ a suitable ensemble average.

The detailed analysis of the intermittent region confronts us to a particularly hard geometrical problem. The direct observation of the region shows that the arrangement of eddies is complex, featuring overhangs with no particular orientation. It is not clear how to measure the geometrical properties of turbulence without a trivial reference frame, like a symmetry axis, a wall, or a simple surface. The study of the intermittent region will require the development of new measurement techniques.

1.6 Turbulence as a fractal.

Fractals appear in nature when two aspects of the same phenomenon have to balance, but an apparent singularity is required for that to happen. The amount of oxygen required by an organism scales with its volume, but the transport of mass between two fluids across a membrane scales with the area. To be able to grow to larger sizes, mammals evolved to have fractal lungs, and fish to have fractal gills. The same way, a fluid motion at a high Reynolds number with vanishing viscosity would dissipate almost no kinetic energy unless velocity gradients become almost singular. Two popular introductions to fractals are Mandelbrot (1983) and Schroeder (2012).

The early description of turbulence given by Richardson (1922) is compatible with the notion of fractal, but that does not mean that turbulence is a fractal in the strict sense. The velocity field in a turbulent flow is smooth at scales comparable to η , and large scales are seldom self-similar. However, some of the aspects of fractal surfaces in three-dimensional spaces are found in an enstrophy (the norm of the vorticity vector $\omega = |\boldsymbol{\omega}|$) isocontour like the one presented in figure 1.5.

The flow in the intermittent region involves multiple scales, and since we do not know which of those are important in entrainment, it is crucial that the analysis does not discard any. It is possible that entrainment is vorticity being diffused across a fractal surface, and that the turbulent/non-turbulent interface ingests irrotational flow similarly to how oxygen is diffused into the blood stream across gills. If that surface is simplified

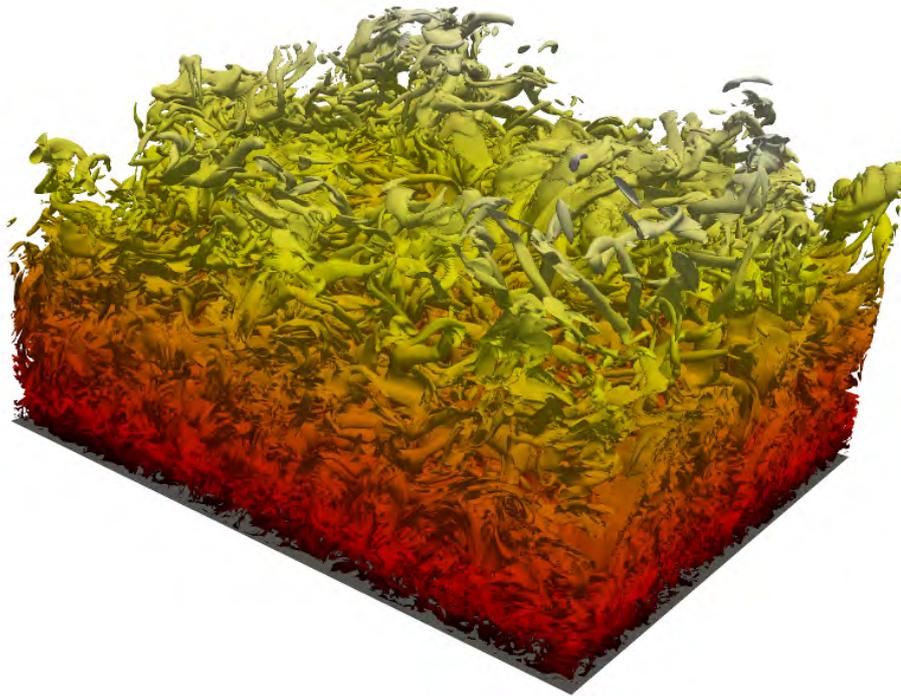


Figure 1.5: Enstrophy contour $\omega^+ = 0.4$ of a boundary layer at $Re_\tau \simeq 2000$. The contour is colored with the distance to the wall. The stream goes from bottom-right to top-left. The box corresponds to a small portion of the simulation at the mentioned Reynolds number, approximately $2\delta_{99}$ wide by $3\delta_{99}$ long.

and processed, the measurements of its properties may not be accurate. If the complexity of gills was ignored, the conclusion would be that all the fish that there is should be dead.

1.7 Direct Numerical Simulation of turbulent flows.

Equations (1.1a) and (1.1b) can be integrated numerically. Computational fluid dynamics is nowadays a relatively mature discipline, and some of its achievements have been relevant to the comprehension of the structure of turbulent flow. The field has steadily evolved since Moin & Mahesh (1998), and new simulation codes and more powerful supercomputers have provided a wealth of interesting results.

Direct Numerical Simulation (DNS) is the numerical resolution of the Navier-Stokes equations with sufficient spatial and temporal accuracy to obtain a solution that contains all the relevant scales of the flow. It is now widely known that these simulations require a significant amount of computational resources. It has been mentioned previously that the separation between the smallest and the largest scales of a turbulent flow is crucial to understand its properties. The verification of Townsend's wall similarity hypothesis

requires that $\delta \gg \nu/u_\tau$, and an even more demanding condition is required to study the intermittent region, $\delta \gg \eta$. Therefore, the target Reynolds number for a boundary layer simulation is *the largest possible*. The parameter that quantifies scale separation in wall-bounded flows is Re_τ , and the amount of information contained in a turbulent boundary layer scales as $Re_\tau^{9/4}$. This means that increasing the Reynolds number by a factor of 2 requires almost 5 times more memory and storage, and almost 8 times more computational resources.

The evolution of supercomputer resources has followed Moore's law (Schaller, 1997) quite consistently for the last two decades (see figure 1.6), but that does not mean that it is trivial for users to obtain the same speedup for their applications. The most powerful supercomputers have experimented a series of architectural changes during the last 20 years. For instance, the fastest supercomputer in the early 2000, the Earth Simulator, had a shared-memory architecture; the fastest supercomputer in 2010, Jaguar, was a distributed-memory CPU-only cluster; and the fastest at the time of writing this document, Tianhe-2, is a distributed-memory accelerated cluster. In consequence, the code for a DNS simulation has on average a longer life than the architecture of the most powerful supercomputer.

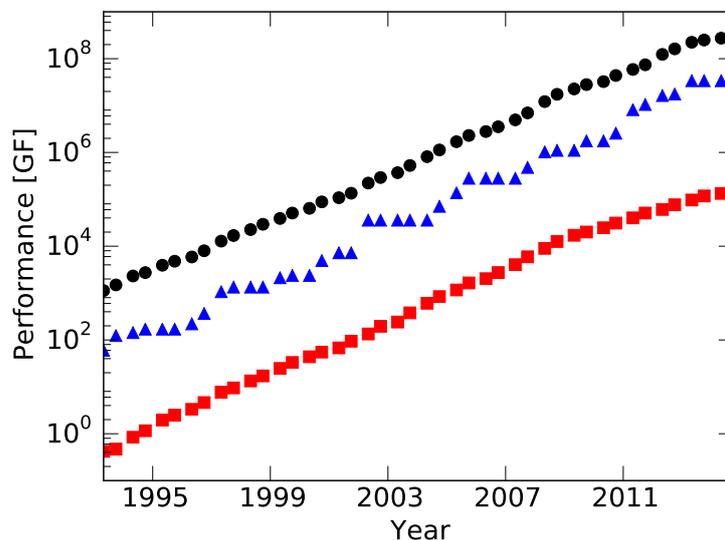


Figure 1.6: Evolution of the performance of the aggregation of the 500 most powerful supercomputers (dots), the most powerful supercomputer (triangles) and the 500th most powerful supercomputer (squares), with their respective temporal fits. Giga FLOPS (GF) is a unit of performance, and corresponds to the billions of floating point operations one supercomputer is able to perform per second. Source: (<http://top500.org>)

Porting to the current supercomputer architecture is not a trivial task. In the case of the Blue Gene, the family of supercomputers used in this work, the interconnection mesh is not conventional, and the amount of memory per per computation unit (usually a core) is a serious constraint for the design of the domain decomposition and the inter-process communication strategy.

Once the simulation has run, it usually generates a huge amount of data, of the order of hundred terabytes in the present case. A relatively simple operation, like moving those data out from the supercomputer center, can be a challenging exercise. In this case, it even required the use of fast wide area connections between supercomputer centers, and to send around 50 TiB storage physically across Spain.

1.8 Goals

The goals of the work described in this thesis are the following:

- To design and to implement the new generation of DNS code able to run turbulent boundary layers at $Re_\tau > 2000$ in the most powerful contemporary supercomputers.
- To simulate a turbulent boundary layer at a sufficient Re_τ , with a forcing able to increase entrainment by a significant amount. That forcing should be similar enough to the presence of a rough wall to put to test Townsend's wall similarity hypothesis.
- To develop the tools necessary to analyze the intermittent region of a boundary layer. Those tools should include as few assumptions about the geometrical aspects of the flow as possible.
- To determine the properties of turbulent motion in the intermittent region, with special emphasis on the interaction between the turbulent eddies in the outer part of the boundary layer and the free stream.
- To determine if the deviation from the collapse predicted by Townsend's wall similarity hypothesis is due to a change in the structure of the turbulent motion, or to a geometrical effect caused by the additional entrainment.

1.9 Structure of this thesis.

This document is separated in three relatively independent chapters. Chapter 2 describes the development of a novel turbulent boundary layer simulation code able to run and scale in the most recent supercomputer architectures. Chapter 3 describes the intermittent region of a boundary layer, presents a novel set of tools needed to analyze the turbulent/non-turbulent interface in full detail, and determines the fundamental scales present in the outer region of the boundary layer. Chapter 4 describes the simulation of a turbulent boundary layer with accelerated entrainment and its relationship with actual roughness. The structure of that forced boundary layer is then compared with an almost identical case without any forcing. Finally, the tools developed in Chapter 3 are applied to the forced boundary layer to analyze the validity of Townsend's wall similarity hypothesis in the intermittent region. Chapter 5 concludes and proposes some ideas for future work.

Direct numerical simulation of turbulent boundary layers at high Reynolds numbers

2.1 Introduction

Our current research is focused on understanding the flow in the turbulent regions that are further from the wall, where range of scales is wider, and the Reynolds number plays a significant role. Our approach is to analyze data obtained from Direct Numerical Simulation (DNS). While high Reynolds number simulations exist for other wall bounded flows (mainly channels), similar data sets were not available for boundary layers. The present code was developed to generate them.

We needed a high resolution code that is able to perform a DNS of a boundary layer that has good performance and excellent scalability. The starting point was the parallel Message Passing Interface (MPI)-only code described in detail in Simens *et al.* (2009). It needed severe modifications to satisfy the additional constraints that the BG/P architecture introduces. The most relevant change was to add a second level of parallelism with OpenMP, necessary to achieve the desired degree of scalability and performance; and an auxiliary domain to extend the computational box.

Two different simulations using this code were completed in two BG/P supercomputers

⁰Some content of this chapter was published as Borrell *et al.* (2013), *A code for direct numerical simulation of turbulent boundary layers in BG/P supercomputers*.

using 32768 cores, a zero pressure gradient boundary layer over a flat plate with a Reynolds number based on the momentum thickness of $Re_\theta = 1100 - 6800$ on Intrepid at the Advanced Leadership Computing Facility (ALCF) (Sillero *et al.*, 2013) and a forced boundary layer with artificial roughness with $Re_\theta = 300 - 4200$ (Borrell *et al.*, 2011) on Jugene at the Jülich Supercomputing Center (JSC).

Previous DNS of turbulent boundary layers at similar Reynolds numbers

This is a domain-specific code designed to solve a zero pressure gradient turbulent boundary layer in a rectangular domain. It is not comparable to codes like nek5000 (Fischer *et al.*, 2008) that, while also achieving excellent scalability in the same supercomputer architecture, are general purpose and are designed to handle more complex geometries.

One could classify previous simulations of turbulent boundary layers on how they deal with the inhomogeneity of the streamwise direction.

It is worth mentioning the pioneering work of Spalart (1988), despite covering a low range of Re_θ . Periodicity in the streamwise direction was enforced with a multiple-scaling transform of the coordinates as well as approximate treatment of the Navier-Stokes equations. This simplification is accurate as long as the streamwise growth of the boundary layer is small; therefore, it is only valid when simulating a short domain.

Another technique to deal with the inhomogeneity is to enforce periodicity by adding a fringe region at the end of the domain, where the flow is forced back to the laminar regime. This technique was applied, for instance, in Khujadze & Oberlack (2004), Wu & Moin (2009) and Schlatter & Örlü (2010), where the simulation ranges up to $Re_\theta = 4060$. While this method is useful to study the phenomenon of laminar-turbulent transition, it requires to start from a laminar flow condition. This may be a limiting factor when the region of interest is only the one further downstream. It also requires some perturbation to trigger turbulence, what makes the flow dependant on the tripping technique. On the other hand, it allows periodic treatment of the streamwise direction, and simplifies the algorithm significantly.

Finally, one can generate an inflow boundary condition that is already turbulent. While Simens *et al.* (2009), Lee *et al.* (2011) and Pirozzoli *et al.* (2010) rescale the flow with a

scheme similar to the one proposed by Lund *et al.* (1998), Ferrante & Elghobashi (2004) extended the cited method. The idea is to pick one cross-stream plane at an intermediate part of the domain, and to recycle it as inflow boundary condition. Two aspects must be taken into account: the separation between inflow and the recycled plane must be wide enough to ensure their independence, and the rescaling should take into account that the turbulent motion involves multiple scales, not only boundary layer thickness. In these simulations, the streamwise direction is non periodic and a finite-difference scheme has to be used. While this approach permits the simulation to start at almost any given value of Re_θ , the recycling process introduces an artificial inflow. All the scales have to evolve until they reach their asymptotic state; hence a portion of the simulation domain has to be discarded. A discussion about this accommodation length scale can be found in Sillero *et al.* (2010).

Compared to the previous related simulations, this code is focused on achieving the highest Reynolds number possible with the given computational resources. For example, the target for the smooth-wall case was to reach a friction Reynolds number $Re_\tau = 2000$, so that it could be compared with an existing simulation (Hoyas & Jiménez, 2006). That comparison introduces additional constraints regarding box size and resolution. At such Re_τ the flow is fully turbulent, there is no need to simulate the transition, and the recycling scheme of the previous implementation is kept.

Another key difference between the current code and the previous ones is that, despite running efficiently on any distributed memory supercomputer, it was tuned for a specific supercomputing architecture that imposes severe constraints on domain decomposition, communications and I/O.

It was also our intention to design an application as flexible as possible that was able to generate data sets at even higher Re_τ without introducing further design modifications in the next generation of supercomputers. Therefore, one of our goals is also to share implementation details that can be useful to design similar large-scale simulations.

The organization of this section is as follows. A basic description of the code is given in section 2, followed in sections 2.1 to 2.5 by the most relevant modifications to its previous version. Scalability is addressed in section 3; and parallel Input/Output, a new feature, is commented in section 4. Finally, validation and conclusions are in sections 4 and 5 respectively.

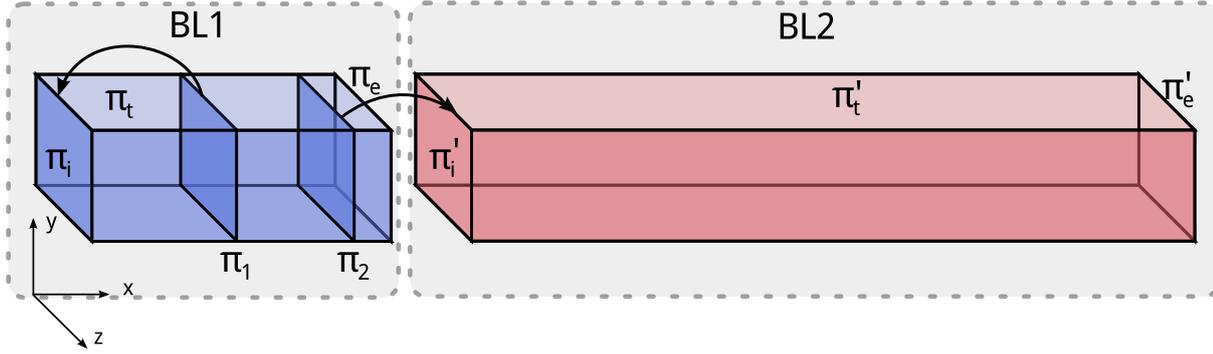


Figure 2.1: Schema of the computational domain and boundary conditions

2.2 The numerical code

The boundary layer is simulated in a parallelepiped over a smooth wall, with spatially periodic boundary conditions in the spanwise direction, but with non-periodic inflow and outflow in the streamwise direction. The code uses a well-established fractional-step method (Harlow & Welch, 1965; Perot, 1993) to solve the incompressible Navier-Stokes equations expressed in primitive variables, using spectral expansions in the spanwise direction, and compact finite differences (Lele, 1992) in the other two. A three sub-step, semi-implicit low storage Runge-Kutta scheme, in which wall-normal second derivative terms use a Crank-Nicholson scheme to increase the time step, is used to evolve the equations in time. A full description of the algorithm can be found in Simens *et al.* (2009).

For the problem considered, spatial derivatives are tightly coupled operations. Our code is constructed in such a way that only single data lines along one of the coordinate directions at a time have to be accessed globally. However, all the three directions have to be treated in every sub-step.

The code is single precision in the I/O operations and communications and double precision in the differentiation and interpolation operations where the implicit part of the compact finite differences and the fast Fourier transform can cause loss of significance.

Computational Setup

A schema of the computational domain can be seen in Figure 2.1. The x , y , and z axes correspond to the streamwise, wall-normal and spanwise directions, respectively. The simulation is split in two concatenated domains with different boundary conditions. The

planes π_i and π'_i are given inflow boundary conditions, and outflow boundary conditions are assigned to π_e and π'_e . The boundary conditions at the top of the boxes, π_t and π'_t , impose a zero pressure gradient on the domain. Finally, the spanwise direction is considered periodic. The purpose of the first boundary layer (Boundary Layer (BL)₁) is to provide accurate inflow boundary conditions to the second one (BL₂). The inflow of BL₁ is obtained from its own plane π_1 that is rescaled using a method based on the one proposed by Lund *et al.* (1998). The physical length of BL₁ is chosen to be long enough to let the large scales recover from an unrealistic initial condition, and once this asymptotic state has been reached, the plane π_2 is used to give BL₂ its inflow boundary condition. As a consequence, a small portion of the BL₁ simulation is thrown away.

Given that the goal of BL₁ is to allow the large scales to reach their asymptotic state and, given that the smaller scales reach a similar condition far more rapidly, BL₁ is run at a coarser resolution than BL₂. This setup permits computing a single boundary layer with significantly less computational work.

The separation between adjacent collocation points is determined by the resolution of the spatial discretization scheme and the local Kolmogorov scale. This scale changes depending on the distance to the wall, so using a non-uniform mesh in the wall normal direction is essential to save memory. To achieve a Reynolds number based on the friction velocity of up to $Re_\tau = 2000$ a computational box for the second domain of size $15360 \times 535 \times 4086$ for a total of 35×10^9 points per variable. In the case of the forced boundary layer of Borrell *et al.* (2011), the domain size of the main simulation is about half as big, $7860 \times 545 \times 4086$.

Domain decomposition and MPI communications

To take advantage of the distributed memory architectures, the computational domain must be partitioned. The only possible decomposition that guaranteed portability to the Blue Gene/P architecture was to use cross-stream planes schematized in Figure 2.2 as π_i .

To compute interpolations and derivatives over the x coordinate it is necessary to transpose the whole variable. This operation creates another elemental domain partition formed by lines in the streamwise direction, labeled in figure 2.2 as ϖ_i . Once these computations are finished the result is transposed back to planes π_i . A more traditional plane-to-plane transpose would be much simpler but it is not possible on the present su-



Figure 2.2: Elemental domains of the domain decomposition.

Case	Re_θ	Nodes	$N_x \times N_y \times N_z$	DoF (Gp)	Time/DoF
BL_1	1100-3000	512	$3585 \times 315 \times 2560$	2.89	$13.98 \mu s$
BL_2	2800-6650	7680	$15361 \times 535 \times 4096$	33.66	$18.01 \mu s$

Table 2.1: Computational setup for the auxiliary BL_1 and main BL_2 boundary layers in Sillero *et al.* (2013): Degrees of freedom (DoF) in giga points; Time/DoF is the amount of total CPU (core) time spent to compute a degree of freedom for every step.

percomputer architecture. The low available amount of memory per node and the need for a large computational domain mandate that no essential domain decomposition based on planes that includes the streamwise direction can be stored as a whole. The ϖ_i pencils can be considered as a secondary partition of such plane.

Each of these two boundary layers is mapped to an MPI group. The first group runs the auxiliary simulation at coarse resolution and it consists of 512 nodes while the second MPI group comprises 7680 nodes and runs the main one in high resolution. The first MPI group is only about 8.5% of the total computational cost. This information is shown in Table 2.1.

The two computational domains communicate with each other only twice per sub-step, to send the π_2 plane from BL_1 to BL_2 and to synchronize the time step, using an additional MPI group that includes all the processes.

The work done by each group must be balanced since each MPI group has to wait for the other one in global operations, otherwise one group will slow down the other one that must remain idle waiting for the other group. The worst-case scenario occurs when the auxiliary simulation slows down the main one. The time taken by communication for the auxiliary simulation has been improved using a customized node topology described in section 2.2.

Global transposes and collective communications

Roughly 45% of the overall execution time is spent transposing the variables from planes to pencils and back; therefore, it was mandatory to optimize the global transpose as much as possible. Preliminary tests revealed that the most suitable communication strategy was to use the `MPI_ALLTOALLV` routine and the Blue Gene/P supercomputer (BG/P) torus network. This method is twice as fast than our previous custom transpose routine based on point-to-point communication over the same network implemented in Simens *et al.* (2009).

The global transpose is split into three sub-steps. The first one changes the alignment of the buffer containing a variable and casts the data from double to single precision to reduce the amount of information to be communicated. If more than one π plane is stored in every node then the buffer comprises the portion of contiguous data belonging to that node in order to keep message sizes as big as possible.

The second sub-step is a call to the `MPI_ALLTOALLV` routine. In this case the possibility of performing collective communications with derived data types would simplify the algorithm, but unfortunately it is not a feature of the present MPI standard. This is the reason why the global transpose is split into three sub-steps.

The third and last sub-step transpose the resulting buffer aligning the data ϖ -wise. This last transpose has been optimized using a blocking strategy because the array to be transposed has many times more rows than columns. The whole array is split into smaller and squarer arrays that are transposed separately. The aspect ratio of those smaller arrays is optimized for cache performance using collected data from a series of tests. Finally the data is cast to double precision again.

The procedure to transpose from ϖ_i pencils to π_i planes is similar and is split in three sub-steps too.

Node mapping in the Blue Gene family of supercomputers.

Mapping virtual processes onto physical processors is one of the essential issues in parallel computing, a field of intense study in the last decade. Proper mapping is critical to achieve sustainable and scalable performance in modern supercomputing systems.

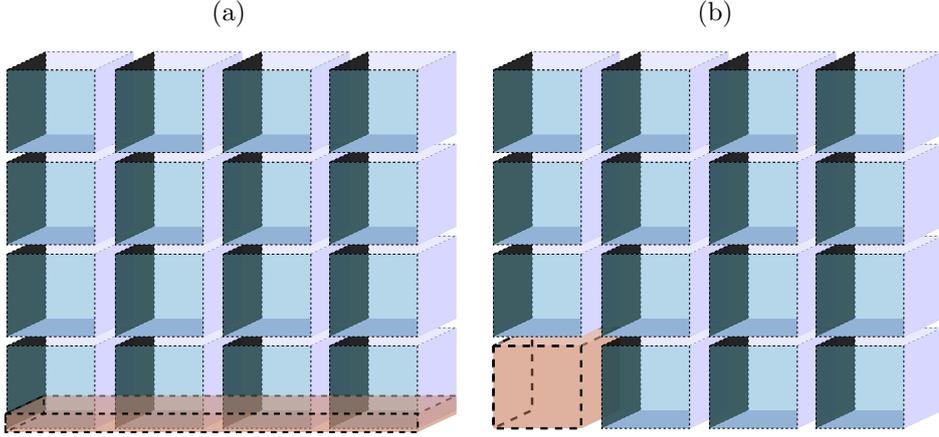


Figure 2.3: (a) Predefined and (b) custom node mapping for a 8192 node partition in a $[8, 32, 32]$ topology. The position of the nodes assigned to BL_1 is highlighted with a different color and a thicker line. The predefined mapping assigns to BL_1 the nodes in a $[8, 32, 2]$ sub-domain. Custom mapping assigns the nodes to a $[8, 8, 8]$ sub-domain and. BL_2 is mapped to the rest of the domain till complete the partition.

Blue Gene/P has a torus network topology except for allocations smaller than 512 nodes, in which the torus degenerates to a mesh. Therefore, each node is connected to six nodes by a direct link. The location of a node within the torus can be described by three coordinates $[X, Y, Z]$.

Different physical layouts of MPI tasks onto physical processors are predefined depending of the number of nodes to be allocated. The predefined mapping for a 512 node partition is a $[8, 8, 8]$ topology, while for 8192 nodes it is $[8, 32, 32]$ as it is shown in Figure 2.3.

Changing the node topology completely changes the graph embedding problem and the path in which the MPI message travels. This can increase or decrease the number of hops needed to connect one node to another, and as a result, alter the communication time to send a message. Fine tuning for specific problems can considerably improve the time spent in communications. Table 2.2 shows different mappings that have been evaluated for our specific problem size. The custom mapping reduces the communication time for BL_1 by a factor of two. The work load for BL_1 is estimated using this new communication time while the load for BL_2 is fixed. Balance is achieved minimizing the time in which BL_1 or BL_2 are idle in the global communications.

The choice of a user-defined mapping is motivated due to the particular distribution of nodes and MPI groups. The first boundary layer BL_1 runs in 512 MPI processes mapped onto the first 512 nodes, while BL_2 runs in 7680 MPI processes mapped onto

Topology	Nodes	Comm BL_1	Comm BL_2
Predefined [8, 8, 8]	512	27.77	—
Custom [32, 32, 8]	8192	79.59	86.09
Predefined [32, 32, 8]	8192	160.22	85.44

Table 2.2: Time spent in communication during global transposes. Different node topologies are presented for 10 time steps and for each boundary layer. Times are given in seconds.

the nodes ranging from 512 to 8192. The optimum topology for our particular problem would be the one in which the number of hops for each MPI group is minimum since collective communications occur locally for each group. For a single 512 node partition the optimum is the use of [8, 8, 8] topology, in which messages travel within a single communication switch. We have found the optimum mapping for BL_1 to be [8, 8, 8] sub-domain within the predefined [8, 32, 32], as shown in the right side of figure 2.3. BL_2 is mapped to the remaining nodes using the predefined topology and no other mappings have been further tested. Although a [8, 8, 8] topology is used for BL_1 by analogy with the single 512 node partition, communication time is nevertheless greater. This is due to the sub-optimal performance of using a 2D mesh instead of a 3D torus network, as already discussed. Finally, the reason can be found in the new hardware connection, since the 512 nodes and 8192 nodes of the 3D torus network are physically connected in a different way. This leads to the increase in the number of hops for BL_1 collective communications, since messages cannot travel within a single communication switch anymore.

The methodology to optimize communications for another size partitions would be similar to the one just described: mapping virtual processes to nodes that are physically as close as possible so the number of hops is minimized.

The network architecture was upgraded for the following generation, the Blue Gene/Q (Chen *et al.*, 2012). The new network, with a 5-D torus architecture, presumably avoids those kind of blockages because the routing algorithm has more physical connections available, but node mapping is a relatively simple technique that improves locality and reduces the amount of hops per message, and therefore improves performance. These issues are not present in other network topologies, like a full fat tree, since every node has a direct connection with the rest. Fat networks are seldom found in very large supercomputers, since the cost per node does not scale linearly like torus networks, but is the most common in small and middle sized supercomputers.

$N_{threads}$	$Comp. T.$	E	$Total T.$	η
1	60.820	1	70.528	1
2	30.895	0.984	38.951	0.905
4	16.470	0.923	24.438	0.721

Table 2.3: OpenMP scalability test performed on 512 nodes. Two efficiencies are given: E is based on the computation time only ($Comp. T.$) and η is based on the total time per step ($Total T.$) and is lower given that only one of the OpenMP threads is able to transfer data to other processes. Times are given in seconds.

The hybrid MPI-OpenMP approach

Introducing OpenMP adds a second domain decomposition to the π_i and ϖ_i used for MPI. The most important non trivial uses of OpenMP are the parallelization of the compact finite differences operators for wall-normal derivatives, that require a tridiagonal solver, and the Fast Fourier Transforms in the spanwise direction. While threaded versions of both band-diagonal solvers and FFT exist, our decision was to handle the OpenMP parallel regions by hand to ensure portability between the different available platforms.

It is important to state that the reason to mix concurrency and parallelism was not driven by the need for more performance but because the small memory capacity of the Blue Gene/P node, which does not allow a physically-significant block of data to be allocated to each core. For instance, in the forced boundary layer case, a single π plane is stored in every node that has been assigned to BL_2 . While very special attention was payed to the collective transpose, that takes almost half of the runtime, the goal of using OpenMP was to use all the available resources of the node. Once we achieved the required scalability and performance, no further tuning was explored.

Some tests were run in a 512 node configuration after porting the code to OpenMP. The results are shown in Table 2.3. These samples suggest that almost no penalty is paid when the computations are parallelized with OpenMP.

2.3 Scalability and portability

Extensive data about scalability was collected during the test runs in a BG/P system, and a later scalability system in a Blue Gene/Q supercomputer (BG/Q). The most relevant cases are listed in the Table 2.4.

All the simulations run keep a linear weak scaling up to 8192 nodes (32768 cores). There

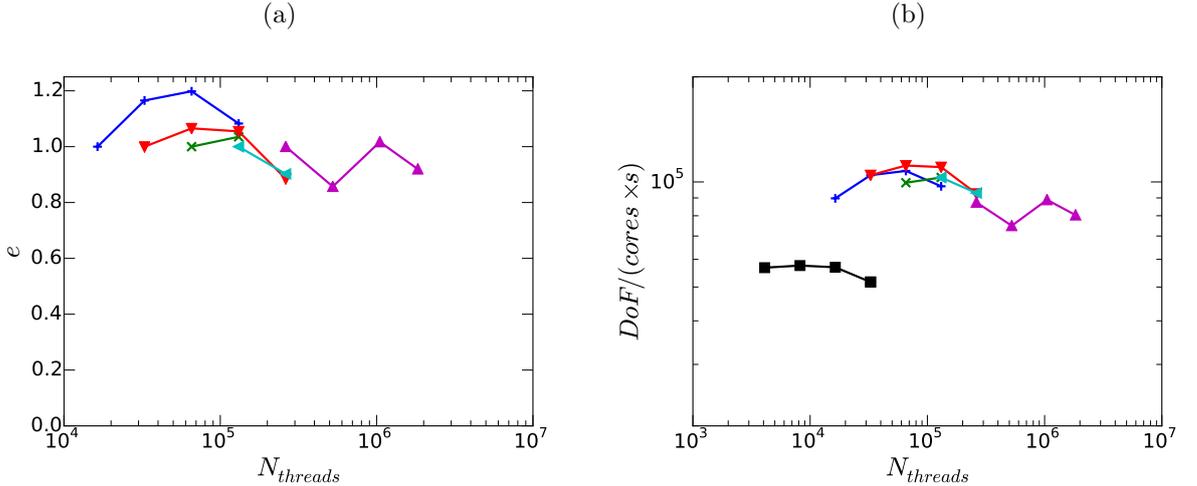


Figure 2.4: (a) Efficiency and (b) performance of the code in the BG/P (squares) and the BG/Q (triangles and circles) versus the number of computation threads. There are 4 threads per node in the BG/P, and 64 threads per node in the BG/Q. The data of this figure is presented in table 2.4. Each curve of efficiency represents one of the tests, that is presented in a different color. There are strong scalability tests (lines with crosses), and weak scalability tests (lines with triangles). The line with squares corresponds to the scalability tests in the BG/P, that follow no pattern in particular.

is only one supercomputer with a higher node count than the BG/Q used in these tests. We are pretty much confident that this code is as scalable as any code could be.

Communication time is typically 40% of the total run time, and that both computation and communication are scaling as expected. The global transpose implementation shows excellent scalability in all the test cases shown in Figure 2.4(a). It is important to mention that, in the BG/P supercomputer architecture, the linear scaling is kept even when the estimated message size is about 1 kB in size. All our previous implementations of the global transpose in more conventional high performance networks were not successfully scaling with an estimated message size below 3kiB. Some of the cases tested in the BG/Q are far beyond the limit of a realistic simulations with the present technology. A case of size $458753 \times 1024 \times 8192$ corresponds to $Re_\tau \sim 1000 - 30.000$, it would consume around 10^{10} CPU hours, and require 0.5 PiB of storage.

This code has also run successfully and with perfect scalability figures in small x86-64 clusters with a fat-tree Infiniband network; Hermit, a CRAY XE-6 system in HLRS (Stuttgart); SuperMUC, a large Beowulf cluster with x86-64 processors and an Infiniband network in Munich; Intrepid and Jugene, two BG/P systems in the Argonne National Laboratory and in Juelich Supercomputing Center; and Juqueen, a BG/Q system also

at Juelich. This makes us confident that all the performance optimizations are portable across different supercomputer architectures.

2.4 Parallel I/O

Intermediate stages of the simulation in the form of flow fields (velocities and pressure) are an important result and are saved even more often than would be required for checkpointing. Another mandatory feature to maintain the scalability with a large node count is the support for parallel collective I/O operations when a parallel file system is available. A handful of alternatives have been tested on parallel file systems, such as the use of raw POSIX calls enforcing the file system block size, sionlib (developed at Juelich) and parallel HDF5 (Group *et al.*, 2014).

HDF5 is a more convenient choice for storing and distributing scientific data than the alternatives tested because, despite having better performance (Frings *et al.*, 2009), they require translating the resulting files to a more useful format. Unfortunately, sufficient performance could not be achieved without tuning the I/O process. HDF5 performance depends on the availability of a cache in the file system. The observed behavior in the BG/P systems was that writing was one, and sometimes two, orders of magnitude slower than reading because in the GPFS used the write cache was turned off. To overcome this issue, when the MPI I/O driver for HDF5 is used, the sieve buffer size parameter of HDF5 can be set to the file system block size. The resulting write bandwidth for 8192 nodes in the Jugene BG/P system was increased up to 16GiB/s, which is similar to the read bandwidth 22GiB/s and closer to the estimated maximum.

2.5 Validation

The numerical scheme is identical to the previous version of the code, which was appropriately validated in Simens *et al.* (2009) and in Jiménez *et al.* (2010), where it was also compared with other experiments and simulations at comparable Reynolds numbers. However, some basic one-point statistics are presented for the present high Reynolds number simulation, showing excellent agreement with numerical and experimental data sets too.

In figures 2.5, a-d are shown the mean and fluctuations of the velocity profiles of the

present simulation compared with other available experimental (De Graaff & Eaton, 2000; Österlund *et al.*, 2000) and numerical (Schlatter & Örlü, 2010) data sets for two different Reynolds numbers. The agreement is excellent.

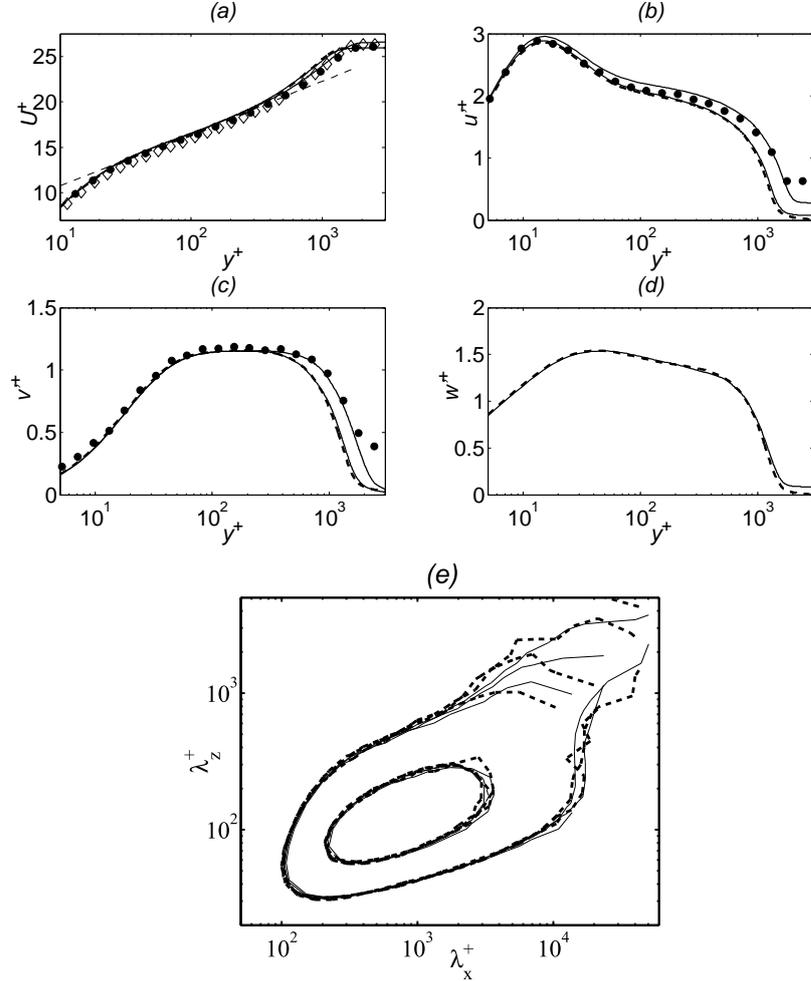


Figure 2.5: Experiments by De Graaff & Eaton (2000), \bullet , $Re_\theta = 5261$; and Österlund *et al.* (2000), \diamond , $Re_\theta = 5156$. Simulations by Schlatter & Örlü (2010), - -, $Re_\theta = 4060$; present, —, $Re_\theta = 4060, 5261$. The law $\log(y^+)/0.4 + 5$ is the discontinuous straight line in (a). (e) Two-dimensional premultiplied energy spectra $k_x k_z E_{uu}(k)$ at three Reynolds numbers and 15 wall units height for channels Hoyas & Jiménez (2006) (solid) and boundary layers (dashed), present and Jiménez *et al.* (2010)

Figure 5e presents the premultiplied energy spectrum $k_x k_z E_{uu}$ where k stands for the wavenumber $k = 2\pi/\lambda$ associated to the wavelength λ . Three different Reynolds numbers ($Re_\tau = 550, 980$ and 2000) at height 15 wall units have been considered. It is a more complete check than one-point statistics because it shows the energy contained in eddies of any size at a given distance to the wall.

The inner contour are the wavelengths whose energy is the 54% of the most energetic

modes, while the outer corresponds to the 14% of this peak. The energy spectrum contains information for all the scales, ranging from the smallest ones at the Kolmogorov scale, to the largest ones of the order of the edge of the boundary layer thickness. For example, the new simulation has enough resolution (fine mesh) to resolve the smallest energy-containing eddies as well as the largest ones where the energy reside. Close to the wall both channels and boundary layers (including the new simulation) are very similar at small scales, where they are Reynolds independent. The footprint of the largest structures, which depends on the Reynolds number, appears in the spectrum as a handle in the top-right corner, and the length of those eddies increases with the Reynolds number.

2.6 Conclusions and future work

A hybrid OpenMP-MPI code has been developed from its original MPI version to perform direct numerical simulations of boundary layers over smooth walls at high Reynolds numbers. The code has been tested in a Blue Gene/P computer using up to 8192 MPI processes, and four threads per process for OpenMP, showing good scalability for both MPI and OpenMP.

Some of the changes were necessary because of the architecture, like hybrid parallelism, all-collective communications and parallel I/O. Others were introduced to correct the somehow unpredictable influence of the inflow boundary conditions at large Reynolds numbers in turbulent boundary layers. This coupled the problem of defining the simulation and tuning the code. The solution here presented is the result at the end of this process.

The simultaneous use of OpenMP and MPI was relatively straightforward in our case, and is becoming a common feature in modern scalable codes. Collective communications are a similar case; once the global transposed was modified according to the suggestions of the Blue Gene handbook and system administrators, performance and scalability were improved significantly.

The approach of simulating two different computational domains, each at a different resolution, has proven to be effective and can be used in other spatially developing turbulent flows. However, it became an issue for our communications scheme. The solution was to separate the auxiliary low-resolution and the main high-resolution simulation in two

different MPI groups and to define a customized mapping of processes onto physical processors. While this particular kind of tuning is not necessary when the process count is low, it is crucial when one is using thousands of nodes over torus networks.

Parallel I/O had a large impact too, despite its performance changes depending on the particular hardware configuration of the platform.

At the time of publication of this thesis two simulations using it have been successfully completed, each one producing valuable data for the study of wall bounded turbulence and boundary layers in particular.

Some features of this new code are considered mandatory for the new generation of supercomputers. We hope that this experience can be a guideline for porting similar codes. Some implementation details that are described are particular to the Blue Gene/P, but the code has been tested successfully in many other supercomputer architectures, and it has been able to efficiently scale up to an entire Blue Gene/Q supercomputer with 458752 cores.

(a) Blue Gene/P (■)

Cores	$N_x \times N_y \times N_z$	N_t	Time/ DoF
4096	$3457 \times 646 \times 1536$	3.43	$17.6 \mu s$
8192	$6145 \times 646 \times 1536$	6.10	$17.4 \mu s$
16384	$8193 \times 711 \times 1536$	8.94	$17.6 \mu s$
32768	$8193 \times 711 \times 2048$	11.93	$19.4 \mu s$
32768	$16385 \times 801 \times 4608$	60.47	$19.3 \mu s$

(b) Blue Gene/Q. Test 1 (+)

Cores	$N_x \times N_y \times N_z$	N_t	Time/ DoF
4096	$8193 \times 315 \times 2048$	5.29	$11.14 \mu s$
8192	$8193 \times 315 \times 2048$	5.29	$9.56 \mu s$
16384	$8193 \times 315 \times 2048$	5.29	$9.30 \mu s$
32768	$8193 \times 315 \times 2048$	5.29	$10.29 \mu s$

(c) Blue Gene/Q. Test 2 (×)

Cores	$N_x \times N_y \times N_z$	N_t	Time/ DoF
16384	$4096 \times 711 \times 4096$	47.71	$10.05 \mu s$
32768	$4096 \times 711 \times 4096$	47.71	$9.71 \mu s$

(d) Blue Gene/Q. Test 3 (▼)

Cores	$N_x \times N_y \times N_z$	N_t	Time/ DoF
8192	$2048 \times 315 \times 2048$	5.29	$9.56 \mu s$
16384	$4096 \times 315 \times 2048$	10.57	$8.97 \mu s$
32768	$8192 \times 315 \times 2048$	21.13	$9.06 \mu s$
65536	$16384 \times 315 \times 2048$	42.27	$10.80 \mu s$

(e) Blue Gene/Q. Test 4 (◀)

Cores	$N_x \times N_y \times N_z$	N_t	Time/ DoF
32768	$8192 \times 711 \times 2048$	47.71	$9.71 \mu s$
65536	$16384 \times 711 \times 2048$	95.43	$10.76 \mu s$

(f) Blue Gene/Q. Test 5 (▲)

Cores	$N_x \times N_y \times N_z$	N_t	Time/ DoF
65536	$16384 \times 1024 \times 8192$	137.45	$11.44 \mu s$
131072	$32768 \times 1024 \times 8192$	274.89	$13.35 \mu s$
262144	$65536 \times 1024 \times 8192$	547.76	$11.25 \mu s$
458752	$114688 \times 1024 \times 8192$	962.08	$12.44 \mu s$

Table 2.4: Data collected from the profiled test cases. Time/ DoF is the amount of total CPU (core) time spent to compute a degree of freedom for every step; N_t is the size in GiB of a buffer of size $N_x \times N_y \times N_z$ (the lower the better). In the BG/Q, the number of threads is four times the number of cores.

Properties of the intermittent region and the turbulent/non-turbulent interface

3.1 Introduction

It has been known since the early days of turbulence research that the near-wall region of boundary layers follows the law of the wall, but that the outer region is influenced by the interaction between turbulence and the free stream (Klebanoff, 1955), whose most obvious consequence is the ‘wake’ component of the mean velocity profile (Coles, 1956; Jiménez *et al.*, 2010). Early work by Corrsin (1943) revealed the presence of a sharp but irregular boundary between turbulent and non-turbulent flow, and the intermittent character of the flow near that boundary. It is also known that, although the outer part of boundary layers has some similarities to a wake (Coles, 1956), intermittency does not behave identically in different flows (Gartshore, 1966). This is true even if the extent of the intermittent region, as quantified by the fraction of time during which the flow at a given point is turbulent (Townsend, 1948), is found to be similar in many flows.

Much of the research on the turbulent/non-turbulent (T/NT) interface has dealt with the entrainment process by which the irrotational flow acquires vorticity. An important early result was that the surface area of the T/NT interface is much larger than its projected area in the direction normal to the wall, and that it is intensely folded (Fiedler & Head, 1966). This observation was the origin of two conjectures summarized in Townsend (1976). The first one is that the interface itself has similar mass transfer per unit area in

⁰Most of the content of this chapter is drafted for publication as Borrell & Jiménez (2015), *Properties of the turbulent/non-turbulent interface in boundary layers*.

all turbulent flows, and that the different entrainment rates (stronger in jets and wakes, weaker in boundary layers and in plane mixing layers) are due to different folding intensities. The other conjecture deals with the mechanism of entrainment. It is clear that irrotational fluid can only acquire vorticity by viscous diffusion. But if the interface is folded enough, large pockets of irrotational flow can be trapped by large coherent structures and driven deep into the turbulent region before acquiring vorticity. To add some nomenclature to the argument, small-scale entrainment is usually called nibbling, while the process by which large blobs of irrotational fluid are swallowed by the turbulent flow before becoming vortical is called engulfment (Mathew & Basu, 2002). This engulfment hypothesis is alternative to the view by Corrsin & Kistler (1955), which assumed that entrainment is only a small scale process. This discussion made clear that understanding the geometry of the T/NT interface is an important step previous to clarifying what is the most relevant mechanism for entrainment.

Capturing the geometry T/NT interface is challenging in both experiments and simulations, partly because of its complexity. The thickness of the intermittent zone is comparable to the boundary layer thickness δ_{99} (Corrsin & Kistler, 1955), while we will see that the strong gradients present in the interface contain scales of the order of the Kolmogorov length η . The interface inherits the fractal-like properties of the underlying turbulent flow (Sreenivasan *et al.*, 1989) and, since turbulent flows typically contain eddies of all possible sizes between the smallest and largest scales, all of them have to be considered when the interface geometry is studied. As a result, important questions about entrainment in turbulent flows had to wait for the necessary data to be available.

Some experimental techniques can capture the interface with considerable detail, and the methods described in Prasad & Sreenivasan (1989) are still used today. The properties of the flow surrounding the interface were not measured until more recently, with the advent of Particle Image Velocimetry (PIV) (Westerweel *et al.*, 2002) and Particle Tracking Velocimetry (PTV) (Holzner *et al.*, 2008). However, experiments are typically restricted to two-dimensional sections of the flow, and the three-dimensional description of the field requires direct numerical simulations (DNS).

Just as experiments, simulations have their limitations. The range of available scales is the most obvious, and is crucial if the scaling properties of a phenomenon are to be studied. Direct numerical simulations at Reynolds numbers large enough to observe a

reasonable scale separation are a recent achievement. While there have been boundary layer simulations at moderate Reynolds numbers for some time (Jiménez *et al.*, 2010; Lee & Sung, 2013; Schlatter & Örlü, 2010), a domain size sufficiently large to obtain a deep T/NT interface requires state-of-the-art in DNSes (Pirozzoli & Bernardini, 2013; Sillero *et al.*, 2013). The Reynolds numbers of these newer simulations is comparable to that of most experiments for which the interface can be analysed in any detail.

These larger and more accurate representations of the flow field, and better analysis tools, have called into question the consensus of what is the dominant mechanism of entrainment. Dahm & Dimotakis (1987), Ferre *et al.* (1990), Mungal *et al.* (1991) and Dimotakis (2000) suggested that engulfment is the dominant process, but later works like Mathew & Basu (2002), and Westerweel *et al.* (2005) emphasized again the importance of nibbling. The dichotomy may have something to do with the level of description desired, since it is clear that viscosity is the ultimate mechanism for vorticity diffusion, but it is equally clear that the complex geometry of the interface has to be taken into account in determining the rate of diffusion.

To determine which scales are most relevant to entrainment requires the study of the turbulent structures present at the vicinity of the interface, which implies the analysis of the properties of the flow in a reference frame linked to the interface itself. Fiedler & Head (1966) presented results obtained from hot wires, but it was not until the work of Bisset *et al.* (2002), Westerweel *et al.* (2002), da Silva *et al.* (2011) and van Reeuwijk & Holzner (2014) that conditional profiles relevant to the scaling of the interface were shown. Bisset *et al.* (2002) mentioned that the T/NT interface could contain at least two layers with different scaling properties: a turbulent region where the major exchanges between the irrotational fluid and the fully turbulent core occur, and an extremely thin viscous superlayer at its outer boundary, already conjectured by Corrsin & Kistler (1955). A recent review of the state of the art is da Silva *et al.* (2014 *a*).

The length scales of the interface provide information about the configuration of the nearby eddies, and about how they are affected by the irrotational flow. An important question is whether the interface is different in any respect from the rest of the turbulent flow or, on the contrary, the eddies near the interface are representative of the bulk of the flow with no major influence from the outer stream. The main candidates for the scaling of the T/NT interface are the Kolmogorov viscous length η and the Taylor microscale λ .

The interface thickness of a DNS temporal jet was computed by da Silva & Taveira (2010) for a single Reynolds number. They found it to be of the order of the Taylor microscale. This result did not imply scaling, because the ratio λ/η depends on the Reynolds number, but Gampert *et al.* (2013) were able to scale with λ the average thickness of the interface of a passive scalar in a jet over the range $Re_\lambda = 61 - 141$. This would agree with the theory described in Hunt & Durbin (1999), based on the hypothesis that the interfaces are subject to relatively strong shear. As a consequence, eddies are blocked and squeezed instead of escaping to the irrotational side. The interface would then have different dynamics from the rest of the flow, and a characteristic length of λ , although it should be noted that the scalar and vorticity interfaces are not necessarily identical.

The goal of this paper is to obtain the properties of the T/NT interface, with emphasis on the relatively large-scale interactions across the fractal intermittent layer rather than on the thinner viscous superlayer. We also analyse the consequences of the threshold used for interface detection. New methods are developed for the geometric characterisation of large surfaces of arbitrary complexity in three-dimensional space, and for the conditional analysis of scalar fields with respect to those surfaces. These methods are used to describe the properties of the flow depending on its position relative to the T/NT interface, and to determine the characteristic thickness of the interface layer. The choice of the identification threshold is given special attention, as well as the choice of the variable being thresholded.

The paper is organized as follows. The next section is a short description of the data set used in this research. The characteristics of the intermittent zone that are relevant for the T/NT interface detection criteria, based on a vorticity isocontour, are presented in §3.3, followed in §3.4 by a quantitative analysis of the geometrical properties of the interface and its dependence on the threshold. Section 3.5 presents the conditional analysis of the flow using the interface as a reference frame. In particular, §3.5 contains the description of the structure of vorticity in the T/NT interface layer, and the determination of its thickness. Section 3.6 explores an alternative definition of the interface as an isosurface of the rate of strain, and the conditional properties of quantities other than the vorticity, such as the rate of strain and the mean velocity gradient. Finally, §3.7 concludes.

3.2 Description of the data.

The boundary layer is simulated in a rectangular box over a no-slip smooth wall. The spanwise boundary conditions are periodic, while inflow and outflow conditions are imposed in the streamwise direction. A transpiration velocity in the boundary opposite to the wall keeps the pressure gradient very close to zero. The simulation code and its implementation are thoroughly explained in Simens *et al.* (2009), and the modifications to achieve higher Reynolds numbers are presented in Borrell *et al.* (2013). The axes in the streamwise, wall-normal and spanwise directions are x , y , and z respectively. The velocity vector is \mathbf{u} , with components along each axis u , v , and w , respectively. Wall units are defined in terms of the friction velocity u_τ and of the kinematic viscosity ν , and are denoted by a '+' superscript. The brackets $\langle \cdot \rangle$ denote the average at a given wall-normal location, and primes denote root-mean-squared values. Both are functions of x and y . The boundary layer thickness is δ_{99} , defined as the distance to the wall at which $\langle u \rangle$ is 99% of the free-stream velocity. The Kolmogorov length is $\eta = (\nu^3 / \langle \varepsilon \rangle)^{1/4}$, where

$$\langle \varepsilon \rangle = \nu \langle |\nabla \mathbf{u}|^2 \rangle \quad (3.1)$$

is the turbulent kinetic energy dissipation rate. A third relevant length is the Taylor microscale λ

$$\lambda = \sqrt{\frac{15\nu |\mathbf{u}'|^2}{\langle \varepsilon \rangle}}. \quad (3.2)$$

Table 3.1 and figure 3.1 summarize the important parameters and characteristics of the simulation, which was designed to achieve convergence of all the scales of the flow in the domain labeled BL, over a range of Reynolds numbers as wide as possible. Two simulations are run simultaneously with a synchronized time step, but the purpose of the auxiliary simulation BL_{aux} is just to provide inflow boundary conditions for BL . Only data from BL are used in this paper. A detailed discussion of the effects of the inflow and the convergence of the flow properties to their asymptotic behaviour can be found in Sillero *et al.* (2013).

The simulation agrees excellently with previous experiments and direct numerical simulations (Sillero *et al.*, 2013). The Taylor microscale Reynolds number, $Re_\lambda = \lambda u' / \nu \simeq O(100)$, is comparable to those of most experiments and simulations used in the analysis of the T/NT interface in free shear flows, and higher than those in the boundary layers used for that purpose. The friction Reynolds number δ_{99}^+ ranges over a factor of two (see

3. PROPERTIES OF THE T/NT INTERFACE

Case	N_x, N_y, N_z	δ_{99}^+	Re_λ	δ_{99}/η	δ_{99}/λ	Tu_τ/δ_{99}
BL_{aux}	$3585 \times 315 \times 2560$	630 – 1100				
BL	$15361 \times 535 \times 4096$	1000 – 2000	75 – 109	242 – 440	14.2 – 21.4	11.5

Table 3.1: Summary of important parameters of the simulation. N_x , N_y and N_z are the size of the computational grid. The Taylor-microscale Reynolds number Re_λ is the maximum observed for the given station, and is attained at $y/\delta_{99} \simeq 0.4 - 0.6$. The thickness of the boundary layer is given using the Kolmogorov length and Taylor microscale as units. In both cases, η and λ are estimated at $y = 0.6\delta_{99}$. The running time T is normalised with properties at the middle of the BL box.

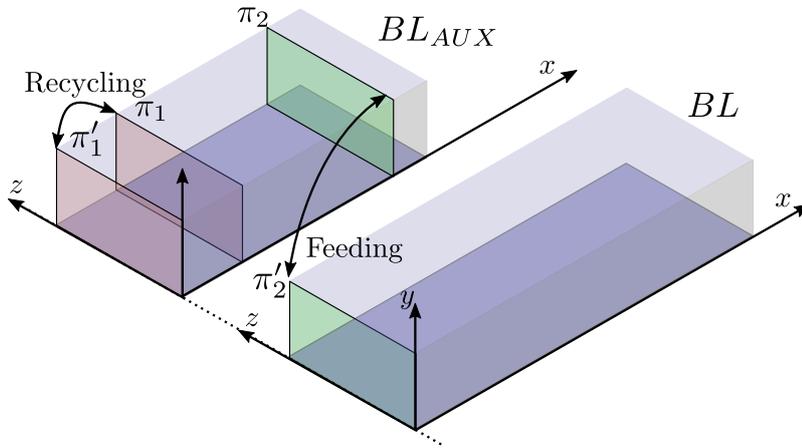


Figure 3.1: Sketch of the simulation and the boundary conditions. The inflow boundary conditions for BL are obtained from BL_{AUX} , copying the plane π_2 to to the first plane of BL at π'_2 . The streamwise location of π_2 is chosen so that the flow has recovered from the recycling scheme (π_1 is recycled to π'_1) used to start BL_{AUX} from a turbulent inflow condition.

table 3.1), allowing it to be used as a parameter in the analysis. The resulting ratio of δ_{99}/η ranges over a factor of 1.8, easily allowing the distinction between outer (δ_{99}) and viscous (η) scaling. The corresponding range of the ratio λ/η between the Taylor and Kolmogorov microscales is narrower (approximately 1.2), and the direct differentiation of scaling with the two quantities is harder to establish with the present data. The averaged properties of the data set have been accumulated during the complete history of the simulation, while some of the more detailed results have been obtained from at least eight flow snapshots, sufficiently separated to ensure statistical independence.

3.3 The turbulent/non-turbulent interface detection criterion.

The first step to study the T/NT interface is to define a criterion to discriminate between turbulent and non-turbulent flow. Unfortunately, we will see that different methods produce different interfaces, and that the criteria found in the literature are variable enough to be difficult to compare consistently. Historically, the first interface detections were based on a cut-off frequency for the one-point streamwise-velocity signal, in the expectation that turbulent fluctuations could be easily distinguished by their faster time scales. As better descriptions of the flow became available, the interface came to be defined by an indicator function with two components: a scalar field related to the turbulent fluctuations, and a threshold. Prasad & Sreenivasan (1989) use the concentration of a passive scalar injected in the turbulent side, and threshold it at the least probable value of the concentration. Da Silva *et al.* (2011) and Bisset *et al.* (2002), use the spanwise vorticity and a particularly low vorticity value as the threshold. In boundary layers, Jiménez *et al.* (2010) use the vorticity magnitude, and a threshold based on a sharp jump in the vorticity probability density function (PDF) at $y = \delta_{99}$. Chauhan *et al.* (2014) use the streamwise velocity as their scalar, and choose the highest velocity threshold for which the PDF of the height of the interface above the wall can be fitted by a gaussian.

Our criterion is based on the vorticity magnitude, $\omega = |\boldsymbol{\omega}|$, which has several desirable properties as a turbulence indicator. In the first place, the incompressible identity

$$\nabla^2 \mathbf{u} = -\nabla \wedge \boldsymbol{\omega} \quad (3.3)$$

implies that the characteristic turbulent dissipation of energy requires vorticity. Secondly, while velocity gradients can be created by pressure fluctuations in potential flow, there is no inviscid mechanism to create vorticity fluctuations. As a consequence, even if vorticity is not conserved, any vorticity in the boundary layer is ultimately linked to the wall. The vorticity magnitude can be easily obtained from DNS, and has been used in studies of the T/NT interface for boundary layers (Jiménez *et al.*, 2010), jets (da Silva *et al.*, 2011) and wakes (Bisset *et al.*, 2002). Note that some of these properties of the vorticity magnitude do not extend to its individual components.

We will define a point as turbulent if

$$\omega(x, y, z, t) > \omega_0, \quad (3.4)$$

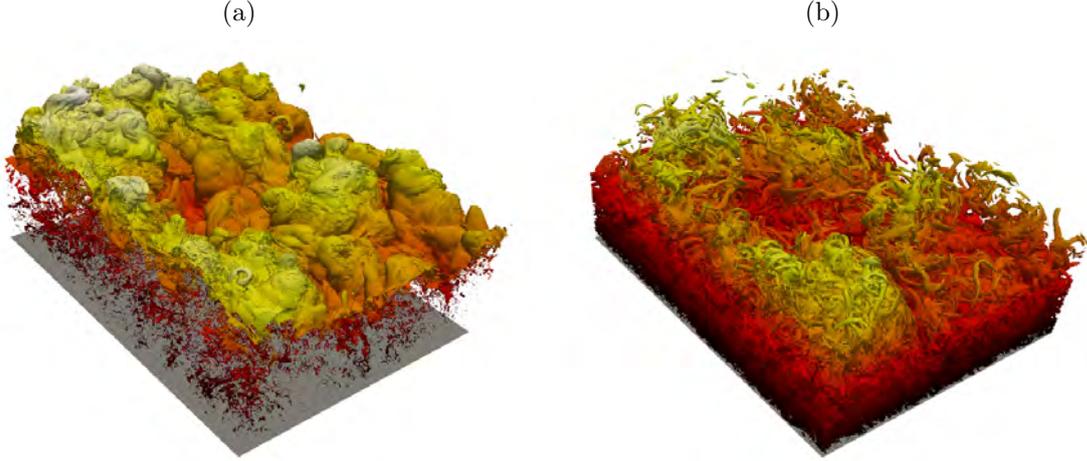


Figure 3.2: Vorticity magnitude isosurfaces of the same region of the present data at $\delta_{99}^+ \simeq 2000$, of size $3\delta_{99} \times 2\delta_{99}$ in the streamwise and spanwise directions, respectively. (a) $\omega_0^+ = 5 \times 10^{-4}$ ($\omega_0^* = 0.022$). (b) $\omega_0^+ = 5 \times 10^{-3}$ ($\omega_0^* = 0.22$). For the definition of ω^* , see (3.7). The flow is from top-left to bottom right, but note that the spanwise and streamwise directions are barely distinguishable.

so that the T/NT interface is defined by $\omega = \omega_0$. The remaining task is to determine an optimum threshold ω_0 , either from the properties of the resulting interface or from comparisons with previous investigations.

The simplest tool is three-dimensional visualization, preferably of a relatively large part of the interface. Figures 3.2(a,b) show the interface in a domain whose wall-parallel size is several boundary layer thicknesses, for two thresholds separated by an order of magnitude. The two figures are clearly different. Figure 3.2(a) can be described as a moderately complex envelope with scattered small regions of low vorticity within the turbulent side, while figure 3.2(b) has a large number of handles and contortions that span a significant fraction of the boundary layer thickness.

Another useful tool is the joint PDF of the vorticity magnitude and of the vertical distance to the wall, which is presented in figure 3.3(a) as a premultiplied PDF, $\omega\Gamma_{\omega,y}$, to account for the logarithmic scale of the vorticity. It has two well-defined regions. The high-vorticity near-wall points of the turbulent core of the boundary layer are in the lower right-hand corner. Points far from the wall with very low vorticity, representing the ideally irrotational non-turbulent free stream, are in the top left corner. Their residual vorticity is due to the finite accuracy of the inflow condition, but it is about four orders of magnitude weaker than the turbulent values, and easily distinguished from them. In the present data set, the details of the joint PDF depend only weakly on the Reynolds

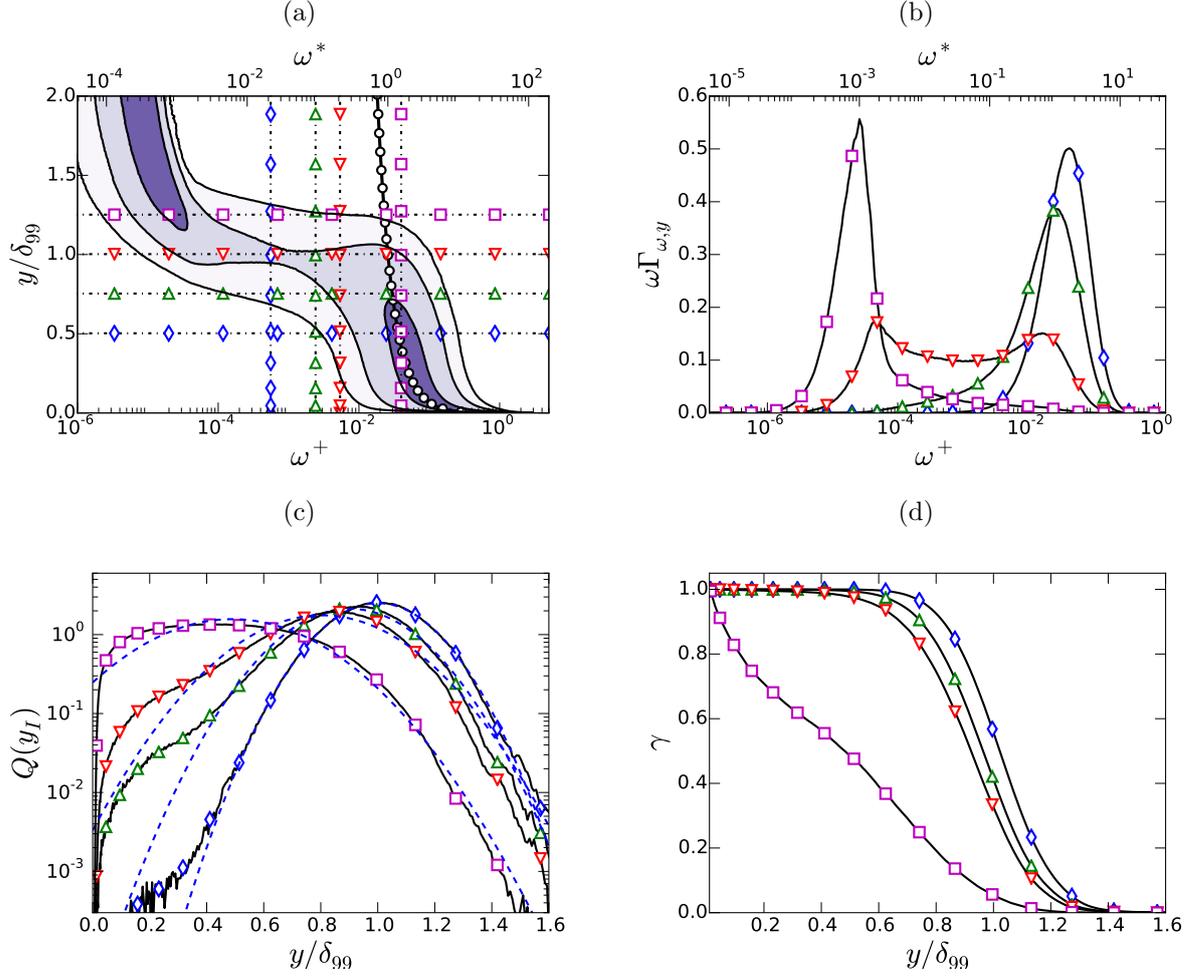


Figure 3.3: (a) Premultiplied joint PDF, $\omega\Gamma_{\omega,y}$, of the wall-normal distance and the vorticity magnitude. Contours contain 50%, 90%, and 99% of points, respectively. Two vorticity scales are provided, wall units ω^+ , and the ω^* defined in (3.7). The line with open circles is $\omega^+ = (y^+)^{-1/2}$. The horizontal and vertical lines correspond to the one-dimensional sections in (b, c), using the same markers. (b) Sections of $\omega\Gamma_{\omega,y}$ at four different distances to the wall: \circ , $y/\delta_{99} = 0.4$; \triangle , 0.7 ; \square , 1 ; ∇ , 1.3 . (c) One-dimensional PDF, $Q(y_I)$, of the vertical position of the interface when its average position is: \diamond , $\langle y_I \rangle / \delta_{99} = 1$ ($\omega_0^* = 0.022$); \triangle , 0.9 ($\omega_0^* = 0.09$); ∇ , 0.8 ($\omega_0^* = 0.2$); \square , $\omega_0^* = 2.0$. The vorticity threshold of the first and third curves are those of the isosurfaces in figures 3.2(a,b), respectively. The dashed line fitting each curve is the gaussian distribution with the same mean and standard deviation. (d) Intermittency factor for the four thresholds in (c). $\delta_{99}^+ \simeq 1500$.

number.

On the turbulent side of the PDF, the mode of the vorticity distribution follows closely the expected y -dependence of its root-mean-squared value, ω' , which can be estimated from the approximate balance between the production and the pseudo-dissipation of the turbulent kinetic energy,

$$\nu\omega'^2 \simeq -\langle uv \rangle \frac{\partial \langle u \rangle}{\partial y} \simeq \frac{u_\tau^3}{\kappa y}, \quad (3.5)$$

where $\kappa \simeq 0.4$ is the von Kármán constant. Equation (3.5) holds above $y^+ \simeq 50$ (Jiménez, 2013), and provides a characteristic magnitude for the vorticity fluctuations,

$$\langle \omega^+ \rangle \simeq (\kappa y^+)^{-1/2}. \quad (3.6)$$

We will use this value, particularised at the edge of the boundary layer, to define dimensionless *star* units for the vorticity,

$$\omega^* = \omega(\delta_{99}^+)^{1/2} \nu / u_\tau^2, \quad (3.7)$$

which are linked to the interface. The usual scaling $\omega^+ = \omega \nu / u_\tau^2$ is linked to the wall. The ratio ω^* / ω^+ varies in our data by a factor of 1.4, and we will see below that star units collapse most properties of the interface substantially better than wall units.

The definition of ω^* can be adapted to flows other than the boundary layer by normalising the vorticity with the root-mean-squared value of the enstrophy fluctuations just within the interface. We will occasionally do this for the purpose of comparison.

There is a band connecting the turbulent and non-turbulent regions of figure 3.3(a) that spans several orders of magnitude of the vorticity and extends over $y/\delta_{99} = 0.3 - 1.5$. Four horizontal sections of $\omega \Gamma_{\omega,y}$ are presented in figure 3.3(b) at different y . The PDF at $y = \delta_{99}$ is particularly interesting, because it shows the separation between the two regions of the flow. Its two mild peaks bracket a plateau three orders of magnitude wide, from the expected turbulent value $\omega^* \simeq \kappa^{-1/2}$ on the right, to the free-stream residual vorticity on the left. Any vorticity within this plateau could in principle be used as a threshold for the interface but, even with generous safety margins at both ends, this leaves a full order of magnitude of possible choices. This would not be a problem if thresholds within this range produced similar results, but they do not. The two isosurfaces in figure 3.2 are obtained with thresholds within the plateau. They correspond to the first and third left-most vertical lines in figure 3.3(a).

Other quantities frequently used to analyse the properties of the edge of boundary layers can be obtained from $\Gamma_{\omega,y}$. The intermittency parameter γ

$$\gamma(y; \omega_0) = \int_{\omega_0}^{\infty} \Gamma_{\omega,y} d\omega / \int_0^{\infty} \Gamma_{\omega,y} d\omega, \quad (3.8)$$

is the probability that a point at a given distance from the wall is turbulent according to (3.4). The sections of $\Gamma_{\omega,y}$ at constant ω provide the marginal probability distribution

Case	Symbol	$\langle y_I \rangle / \delta_{99}$	$\sigma(y_I) / \delta_{99}$	ω_0^*
Jiménez <i>et al.</i> (2010)	○	0.92	0.033	0.077
Kovaszny <i>et al.</i> (1970)	◇	0.83	0.044	0.155
Corrsin & Kistler (1955)	△	0.82	0.047	0.178
Murlis <i>et al.</i> (1982)	▽	0.8	0.052	0.221
Klebanoff (1955)	□	0.78	0.055	0.256
Chauhan <i>et al.</i> (2014)	◁	0.71	0.063	0.395

Table 3.2: Properties of $Q(y_I)$ for the different values of $\langle y_I \rangle$ found in the literature. The values of $\sigma(y_I)$ and ω_0^* are obtained from the present data set, and correspond to the threshold required to match $\langle y_I \rangle$ for each entry, and to its standard deviation.

$Q(y_I)$ of the height y_I of the interface above the wall,

$$Q(y_I; \omega_0) = -\partial\gamma/\partial\omega_0 = \Gamma_{\omega_0, y} / \int_0^\infty \Gamma_{\omega_0, y} dy. \quad (3.9)$$

Four examples of $Q(y_I)$ and $\gamma(y)$ are presented in figures 3.3(c,d). The thresholds of the first three are chosen so that the average height of the T/NT interface is $\langle y_I \rangle / \delta_{99} = 1$, 0.9, and 0.8, respectively, and are within the plateau in figure 3.3(a). We have already mentioned that the first and third ones are used in figure 3.2. This confirms that the threshold has an important effect on the geometrical properties of the T/NT interface, even for properties that are easily measurable. Note that, although $Q(y_I)$ and $\gamma(y)$ are linked by the first equality in (3.9), γ is not very sensitive to the changes in Q , and always looks approximately gaussian. The fourth line in figures 3.3(c,d), marked with open squares, is $\omega_0^* = 1.6$, and corresponds to the right-most end of the plateau in figure 3.3(a). It behaves differently from the other three PDFs, and neither $Q(y_I)$ nor $\gamma(y)$ can be approximated as gaussian. This threshold does not represent the interface any more, and can best be understood as describing the internal structure of the turbulent vorticity.

The mean value $\langle y_I \rangle$ and the standard deviation $\sigma(y_I)$ of the interface height are presented in figure 3.4 as functions of ω_0 . Three regimes can be distinguished. The first one, below $\omega_0^* = 2 \times 10^{-3}$, reflects the vorticity fluctuations in the free stream, and therefore is basically a numerical artifact. In the second one, between $\omega_0^* = 2 \times 10^{-3}$ and $\omega_0^* = 0.1$, the average position of the interface is $\langle y_I \rangle \simeq \delta_{99}$. The tails of $Q(y_I)$ are symmetric in this range (figure 3.3c), and the PDF is well approximated by a normal distribution. Above $\omega_0^* = 0.1$, the left tail of $Q(y_I)$ is influenced by the turbulent core flow, $\langle y_I \rangle$ drops faster with the threshold, and the standard deviation increases slightly.

The values of $\langle y_I \rangle$ available in the literature are compiled in table 3.2 and marked with their corresponding symbols in figure 3.4. They can be used as guides in choosing our

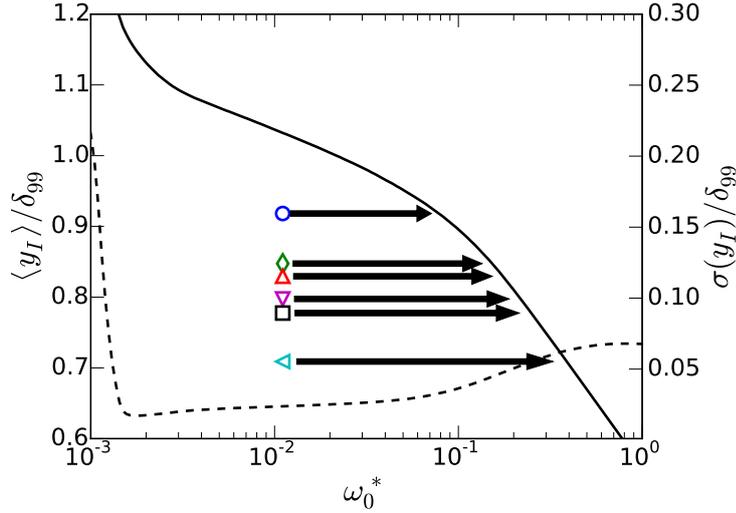


Figure 3.4: The solid line and left vertical axis are the average of the height, y_I , of the vorticity isosurface, and the dashed line and right vertical axis are its standard deviation, both as functions of ω_0^* . The values of $\langle y_I \rangle$ in table 3.2 are marked by their corresponding symbols, with an arrow pointing to the matching ω_0^* . Both curves change very little within our range of Reynolds numbers.

threshold. There is a fairly large spread between the choices of Jiménez *et al.* (2010) and of Chauhan *et al.* (2014) which, if translated to vorticity thresholds using figure 3.4, would imply half an order of magnitude in ω_0^* . The thresholds in figures 3.2(a,b) corresponds to $\langle y_I \rangle \approx \delta_{99}$ and $0.8\delta_{99}$, respectively.

Figure 3.3(b) suggests that $\omega_0^* = 0.022$, for which $\langle y_I \rangle = \delta_{99}$, should be a reasonable threshold, since it is at this height that the vorticity PDF is widest and bimodal. However, figure 3.4 and table 3.2 show that this threshold is an order of magnitude lower than most values used in previous works.

The definition in Prasad & Sreenivasan (1989) can be adapted to cases without a passive scalar, using the vorticity magnitude as a tracer (Gampert *et al.*, 2014; da Silva *et al.*, 2014 b). Applying this criterion to the present data would imply $\omega_0^* = 0.05$ and $\langle y_I \rangle = 0.95\delta_{99}$, which is again lower than the values found in the literature, and comparable to figure 3.2(a).

Chauhan *et al.* (2014) define the interface as the highest isocontour of the streamwise velocity field for which the PDF of y_I can be fitted by a gaussian. In trying to apply a similar criterion to the vorticity, we find that $Q(y_I)$ is approximately gaussian for $\omega_0^* \in (2 \times 10^{-3} - 0.1)$. This corresponds to $\langle y_I \rangle / \delta_{99} \in (1.1 - 0.9)$. Although the lowest end of this range agrees with the mean interface height in Jiménez *et al.* (2010), it is very far

from the value $\langle y_I \rangle = 0.71$ obtained by Chauhan *et al.* (2014). This shows that the vorticity and velocity interfaces are different, and that the criterion in Chauhan *et al.* (2014) should not be used for the vorticity.

In summary, since neither the intermittency properties of the interface nor previous studies provide guidance on a unique vorticity threshold, we defer our decision until we study the evolution of the interface over the rather wide range $\omega_0^* \in (0.001 - 10)$.

Alternative scalars other than vorticity magnitude

Another difference between this study and Chauhan *et al.* (2014) is that the latter uses a scalar quantity different than vorticity, equivalent to

$$\tilde{k} = \frac{(u - U_\infty)^2 + v^2}{U_\infty^2}, \quad (3.10)$$

to study the T/NT interface.

We have seen that interfaces obtained with different identification criteria do not necessarily match. Two different scalar quantities related with the turbulent fluctuations (like vorticity and \tilde{k}), also provide different descriptions of the structure of the intermittent region. A Joint PDF of ω and \tilde{k} is presented in figure 3.5(a), showing that the two quantities are not strongly correlated. One condition fulfilled by vorticity of practical importance for the study of the T/NT interface is the clear separation between the characteristic values of the turbulent and the non/turbulent sides. We can see that \tilde{k} does not fulfill that condition. The shape of the premultiplied joint PDF $\omega P_{\omega, \tilde{k}}$ shows a continuous change of \tilde{k} . While a threshold of vorticity (at around $\omega^* = 0.01$) can be used to separate between high and low \tilde{k} , the reciprocal is not true. A threshold of \tilde{k} at 0.05 corresponds to a range of vorticities almost five orders of magnitude wide.

This behaviour is not particular to boundary layers, and it is also present in other turbulent flows like temporal jets. The scalar

$$\hat{k} = (u - U_\infty)^2 + v^2 + w^2 \quad (3.11)$$

is defined to overcome the issue that \tilde{k} is singular if $U_\infty = 0$, which is one of the usual characteristics of temporal jets. In the premultiplied joint PDF of \hat{k} and vorticity, presented in figure 3.5(b), it is clear that the separation between the two states of the flow, turbulent and non-turbulent, is given by vorticity. On the other hand, the square of

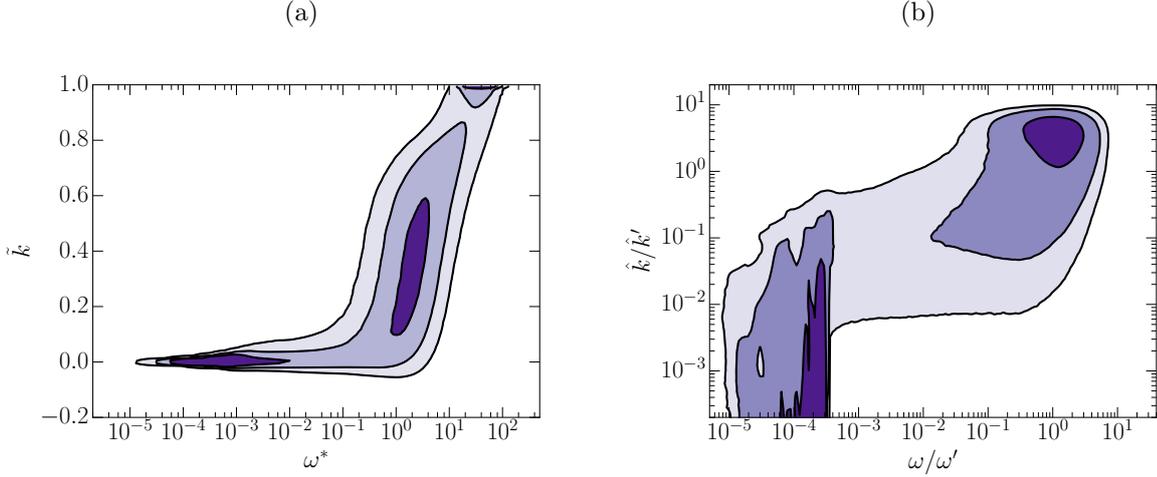


Figure 3.5: (a) Premultiplied joint PDF of vorticity magnitude and \tilde{k} (see equation (3.10)) $\omega P_{\omega, \tilde{k}}$ in the present boundary layer at $\delta_{99}^+ = 1900$. (b) Premultiplied joint PDF of vorticity magnitude and \hat{k} (see equation (3.11)) $\omega \hat{k} P_{\omega, \hat{k}}$ in a temporal jet. Contours contain 50%, 90% and 99% of points respectively.

the velocity fluctuations is much less intermittent. Any possible threshold of \hat{k} crosses the non-marginal contours of the premultiplied joint PDF of the boundary layer and the temporal jet.

This does not mean that velocity fluctuations are not meaningful to study the intermittent region or the T/NT interface, but it suggests that \tilde{k} and \hat{k} lack one particularly desirable property as turbulence indicators.

3.4 The geometry of the T/NT interface.

In this section we study directly the geometry of the T/NT interface. Since our interface is an isosurface separating the vortical turbulent fluid from the irrotational free stream, it inherits some of the geometrical characteristics of both. Moreover, since changing the vorticity threshold moves the isosurface towards one or the other of those two regions, we can use its dependence on ω_0 to explore the geometry of the two flow regimes and, incidentally, to decide which threshold is likely to behave best as an interface for a particular purpose. For example, figure 3.2(a) appears to represent better the free stream, while figure 3.2(b) is more representative of the interior of turbulence.

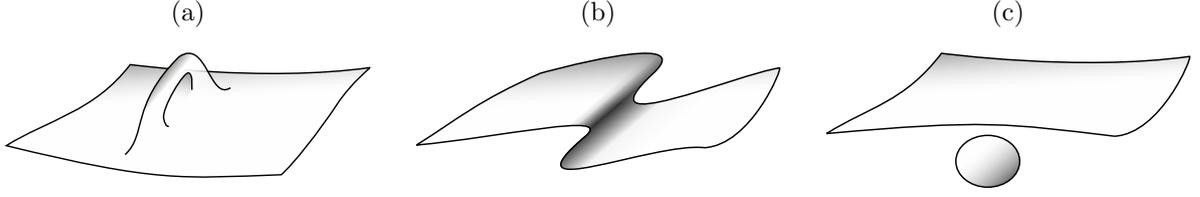


Figure 3.6: Sketch of the three basic geometrical features in the vorticity isosurface: (a) handles, (b) pockets and (c) bubbles.

Bubbles and drops

The first task is to precisely define the interface separating the flow into turbulent and non-turbulent regions. This is not as straightforward as the previous section may suggest. Figure 3.2(b) shows that the vorticity isocontour is not usually a singly connected surface. Depending on the threshold, there may be a few individual components, or several thousands, but one of them is typically much larger than the rest and spans the whole computational box. The smaller components can be classified as interior *bubbles* of low vorticity within the turbulent region (figure 3.6c), or *drops* of high vorticity in the free stream. It will be shown in §3.5 that, although there may be a large number of bubbles, they are too small to contribute significantly to most quantities related to the T/NT interface, and that there are generally very few drops. In consequence, the rest of the paper defines the interface as the largest singly connected component of the vorticity isosurface that separates ‘smoothed’ irrotational and vortical regions from which drops and bubbles have been eliminated.

The method to obtain this largest component is sketched in figure 3.7. We first decompose the computational domain into a structured set of computational cells (voxels). The flow properties are defined at the vertices of these cells. We next obtain the set $\Omega_{\omega>}$ of voxels for which at least one vertex has a vorticity higher than the threshold (figure 3.7a). This set contains the drops and the bulk of the turbulent flow. Similarly, we obtain the set $\Omega_{\omega<}$ for which at least one vertex has a vorticity lower than the threshold, containing the bubbles and the bulk of non-turbulent flow (figure 3.7b). Each of these sets has a connected component many times larger than the rest (about seven orders of magnitude larger in our data). In the case of $\Omega_{\omega>}$, this largest set represents the bulk of the turbulent flow, Ω_t . In the case of $\Omega_{\omega<}$, it represents the bulk of the free stream, Ω_n . The sets of voxels containing the drops, Ω_d , and the bubbles, Ω_b , are obtained by subtracting these largest components from their respective sets. Thus, $\Omega_d = \Omega_{\omega>} - \Omega_t$ and $\Omega_b = \Omega_{\omega<} - \Omega_n$.

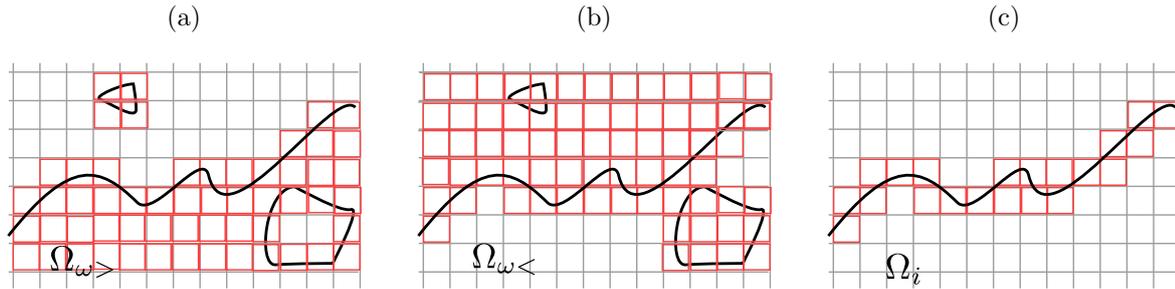


Figure 3.7: (a) Set $\Omega_{\omega>}$ of voxels where $\omega > \omega_0$. (b) Set $\Omega_{\omega<}$ of voxels where $\omega < \omega_0$. (c) The set Ω_i of voxels that contain the interface, obtained from (3.12).

The final step is to define the set of voxels of the cleaned T/NT interface as (figure 3.7c)

$$\Omega_i = \underbrace{(\Omega_t \cup \Omega_b)}_{\text{Turbulent side}} \cap \underbrace{(\Omega_n \cup \Omega_d)}_{\text{Non-turbulent side}} . \quad (3.12)$$

Note that each term of this equation is a practical definition of the smoothed turbulent region (the bulk of the turbulent flow plus the bubbles), and of the smoothed non-turbulent region (the bulk of the non-turbulent flow plus the drops).

Drops and bubbles should not be confused with other possible complications of the interface, also represented in figure 3.6(a,b) as handles, and overhangs or ‘pockets’. The former complicate the topology of the flow, and are unavoidable. The latter are topologically neutral, but may be important from the experimental or dynamical point of view. They hide part of the surface to some observational procedures, and may be precursors for large-scale engulfing. At this point, the interface is still a set of voxels that has to be converted into a surface, but this representation is sufficient for the analysis in the next two sections.

Fractal dimension

Mandelbrot (1975) was the first to suggest that the hierarchy of turbulent eddies can be approximated by a fractal when the Reynolds number is large enough. This was first verified by Sreenivasan & Meneveau (1986) for the bulk of the flow, and by Sreenivasan *et al.* (1989) for the T/NT interface. The latter also proposed a simple theory to relate both results. The fractal dimension of the vorticity isosurface measures how contorted it is, and is a useful statistical measure of its complexity. The most widely used definition is the box-counting Kolmogorov capacity: if N_b is the number of boxes of size r required to cover a set S , such as the interface, the fractal dimension D is defined by $N_b \propto r^{-D}$.

In practice, the computation of fractal dimensions is complicated because turbulence is only self-similar in a limited range of scales. Vorticity is smooth at scales of the order of the Kolmogorov microscale, and the largest eddies responsible for the energy input are not self-similar. In cases in which an extended power law is not immediately obvious, a reasonable redefinition of the box-counting dimension is

$$\dim_b = -\lim_{r \rightarrow \varsigma} \frac{\log N_b}{\log r}, \quad (3.13)$$

where ς stands for the smallest possible box size at which the data set remains self-similar or, in the present case, for the computational voxel size. This last requirement is difficult to define, and it is hard to speak of a fractal unless the self-similar range extends over a reasonably wide range.

Sreenivasan *et al.* (1989) found a clear power law from two-dimensional sections of the interface, and measured a constant dimension D away from the saturation caused by the shortest and longest scales. They concluded that the interface is a monofractal in that range. Moisy & Jiménez (2004) computed the fractal dimension of three-dimensional enstrophy isosurfaces in homogeneous turbulence, using the full three-dimensional field instead of cross sections, and three-dimensional boxes instead of two-dimensional ones. They found that the self-similar range observed by Sreenivasan *et al.* (1989) is only an approximation, and defined a local dimensional exponent to account for the dependence on the box size

$$D_b(r) = -\frac{d \log N_b}{d \log r}. \quad (3.14)$$

This definition includes the previous two. If $D_b(r)$ is constant and the T/NT interface is a monofractal, $\dim_b = D = D_b(r)$.

The local dimensional exponent (3.14) of our interface is presented in figure 3.8 as a function of ω_0^* for several Reynolds numbers. Figure 3.8(a) plots D_b for the smallest possible value of r , and tries to approximate (3.13). Figure 3.8(b) plots the maximum value D_b across r . The differences between the two figures quantify how far from a monofractal is the T/NT interface. All the data in figure 3.8 refer to the dimension of the ‘cleaned’ interface. The dimension of the full vorticity isosurface is dominated by the cloud of very small low-vorticity bubbles, and is usually $D < 2$. Note the good collapse of the different Reynolds numbers when parametrised with ω_0^* . The black horizontal bar near the peaks of both figures is the variation of ω^*/ω^+ in the range of Reynolds number of our data set. A similar bar is included in all later figures that make a Reynolds number

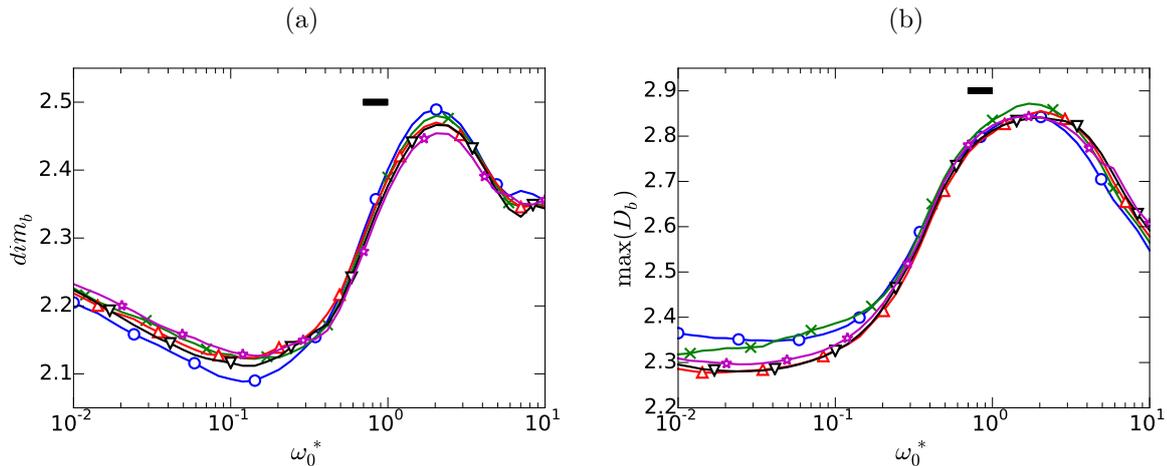


Figure 3.8: (a) Estimation of the box-counting fractal dimension obtained from the local fractal exponent in the limit of small box size, and (b) the maximum of the local exponent over r . Symbols are δ_{99}^+ : \circ , 1100; \times , 1300; ∇ , 1500; \triangle , 1700; \star , 1900. The horizontal bar is the variation of ω^*/ω^+ in our range of δ_{99}^+ .

comparison, and measures how much the collapse of the different curves would deteriorate if the data had been normalised with ω_0^+ instead of with ω_0^* .

The dependence of D_b on ω_0^* confirms the visual impression from figure 3.2 that the threshold has a dramatic effect on the interface. At low thresholds, the dimension approaches the smooth limit $D = 2$ but, at higher ones, the T/NT interface is significantly more convoluted. Sreenivasan *et al.* (1989) predict $D = 7/3$ for the T/NT interface, which is within the range of the present results; it would correspond to $\omega_0^* \simeq 1$ in figure 3.8(a), and to the lowest possible dimension in figure 3.8(b).

Regardless of the differences in their absolute values, the two estimations of the fractal dimension in figure 3.8 behave similarly with respect to ω_0^* . There is a transition between $\omega_0^* = 0.2$ and $\omega_0^* = 2$, across which the geometrical complexity of the interface increases significantly. The threshold $\omega_0^* \simeq 0.2$ already appeared in the analysis of figure 3.4, and corresponds to the vorticity where the average location of the interface $\langle y_I \rangle$ decreases fastest as the threshold increases. This suggests that for $\omega_0^* \gtrsim 0.2$, the surface begins to move inside the turbulent core, where it reflects the geometrical features of the turbulent vorticity itself.

The decrease of the dimension beyond $\omega_0^* \simeq 2$ was already observed by Moisy & Jiménez (2004), who used thresholds of the order of $\omega_0^* \simeq 2 - 12$ to study the geometry of the *volume* of the vorticity in isotropic turbulence. There is no simple relation between the

fractal dimensions of a set and of its surface, but Moisy & Jiménez (2004) noted that in the limit of very high thresholds the vorticity would be reduced to a discrete cloud of points for which $D \simeq 0$. A similar argument can be applied to the interface.

Genus

The geometric complexity of an object can also be characterized by its topological properties. The genus g is a topological invariant of any connected orientable surface, and measures the number of its ‘handles’ (figure 3.6a). A sphere has genus zero, a torus has genus one, and two connected tori have genus two. To our knowledge, the genus was first used to characterize turbulent structures in homogeneous turbulence by Leung *et al.* (2012), who cite instances of its earlier use in disciplines such as astrophysics. In most of those cases, the genus is obtained by integrating the mean and gaussian curvatures over the interface, which requires a careful triangulation of the surface. This step is time consuming and prone to errors, and we bypass it by computing the genus directly from the Euler characteristic χ of the numerically defined contour. The algorithm is described in Lozano-Durán & Borrell (2015), and is adapted to exploit the structure of discrete data in structured grids.

Any numerical isosurface in a cartesian grid is a polyhedron of stacked parallelepipeds. If V is the number of vertices, E the number of edges, and F the number of faces, its Euler characteristic is given by the Euler–Poincaré formula,

$$\chi = V - E + F, \quad (3.15)$$

and the genus is

$$g = 1 - \frac{\chi}{2}. \quad (3.16)$$

The genus is a measure of complexity, like the fractal dimension, but the two are not equivalent. A wrinkled piece of paper has genus zero, independently of the amount of wrinkling. A regular Brownian surface is defined as a fractal single-valued map on the real plane. It has a fractal dimension $D = 2.5$, but no handles (Russ, 1994).

As in the previous section, we compute the genus for the largest connected component of the vorticity isosurface. It is shown in figure 3.9(a) normalized by its maximum over ω_0 . There is a topological transition in which handles begin to appear over roughly the same range, $\omega_0^* \simeq (0.2 - 2)$, as the growth of the fractal dimension. Around $\omega_0^* \simeq 1$,

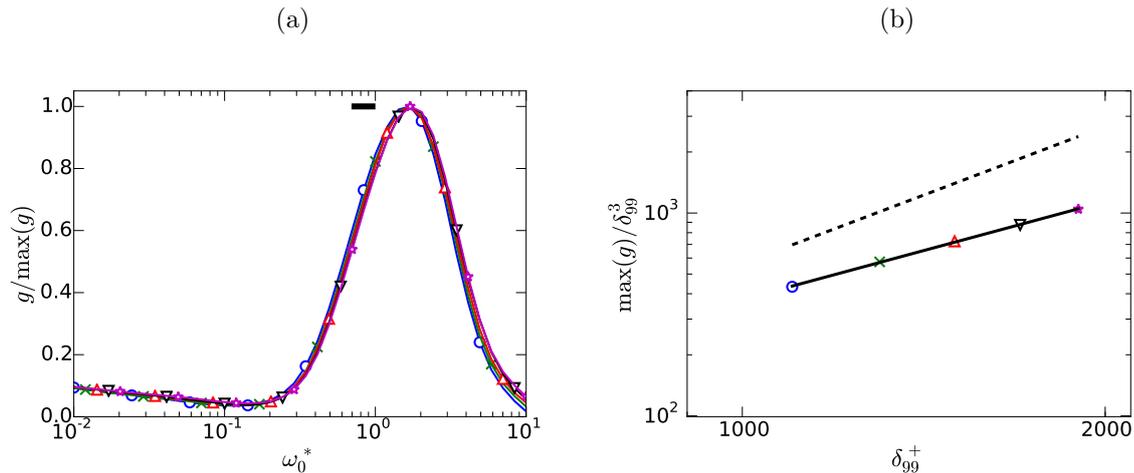


Figure 3.9: (a) Genus normalized with its maximum over ω_0 . The horizontal bar is the variation of ω^*/ω^+ in our range of δ_{99}^+ . The horizontal bar is the ratio of ω^*/ω^+ over the present range of δ_{99}^+ . (b) Maximum genus per cubic boundary-layer thickness, occurring in all cases at $\omega_0^* \simeq 2$. Markers as in (a). Both axes are logarithmic. The solid line is a power law fit, $\max(g)/\delta_{99}^+ \propto \delta_{99}^{+1.6}$. The dashed one is the Kolmogorov limit, $\max(g)/\delta_{99}^+ \propto \delta_{99}^{+9/4}$. Symbols are δ_{99}^+ , as in figure 3.8.

handles are the dominant feature of the surface, and there are hundreds or thousands of them in a volume $O(\delta_{99}^3)$. We suggested in the discussion of the fractal dimension that the T/NT interface at these high thresholds is basically a reflection of the internal geometry of the turbulent vorticity, and the reasons for the decrease of the dimension beyond the end of the transition also apply here. Some turbulent features disappear for very large thresholds, causing the genus to decrease. The maximum genus occurs at the end of the topological transition $\omega_0^* \simeq 2$, and figure 3.9(b) shows that it increases with the Reynolds number as $\max(g)/\delta_{99}^+ \propto \delta_{99}^{+1.6}$. This exponent is somewhat smaller than for the number of Kolmogorov-size structures per cubic integral scale $\delta_{99}^{+9/4}$, which sets an upper bound for the scaling of the possible complexity. Note again the good collapse provided by ω_0^* for the Reynolds number dependence of the genus.

This predominance of handles will become important for the conditional analysis of the flow in the next section. When the analysis of a surface with handles is carried out using a lower-dimensional section, such as a two-dimensional plane or a line, the results can be subject to interpretation artefacts. For example, the planar section of a torus across its principal axis is two circles, giving the impression of two disconnected geometrical objects. Up to a point, the same is true for pockets such as those in figure 3.6(b). For example, the interface shown below in figure 3.13(b) is a section of a singly-connected isosurface, although it appears to contain many unconnected irrotational bubbles within

the turbulent region. Another effect of the handles has to do with values conditioned to the direction normal to the interface. The usual assumption in this case is that a normal defined from high towards low vorticity values points into the the free stream. In a handle, or in a narrow pocket, this is only true over distances of the order of the feature thickness, and becomes an issue if handles and pockets are dominant. The problem is less pressing when the threshold is chosen below $\omega_0^* \simeq 0.2$, where the T/NT interface is smoother, but figure 3.4 shows that a lot of the published work uses thresholds within the topological transition, characterised by non-trivial fractal dimensions and, presumably, large genera.

The main conclusion from this section is that the properties of the fully turbulent flow appear gradually in the geometry of the interface as the threshold traverses the topological transition, and that the handles, folds, and high fractal dimensions are probably the reflection of the internal structure of the flow.

3.5 Conditional analysis of the vorticity field.

In this section we study the properties of the vorticity field as a function of the distance to the T/NT interface. Given the geometrical complexity of the interface, it is to be expected that different definitions of distance produce different conditional results. To allow us to differentiate between genuine flow properties and possible measurement artefacts, we will study the cases in which the results of two alternative distance definitions that are not equivalent.

Consider first the vertical distance Δ_v . Given a surface Ω , Δ_v is the distance between a point p and the topmost intersection with Ω of a line normal to the wall going through p . A sketch is given in figure 3.10(a), emphasizing that even if the line used to measure distance crosses the interface multiple times, only the highest intersection is taken into account. Note that discarding the lower intersections hides part of the complexity of the interface, and that some handles and pockets might not be captured. This criterion has been used to study the T/NT interface in boundary layers by Chauhan *et al.* (2014) using normals to the wall, and in jets by Westerweel *et al.* (2009) and da Silva & Taveira (2010) using normals to their symmetry plane.

Our second definition of distance is the separation between the point p and its closest point in Ω . We will call it ball (or minimum) distance Δ_b , and has a simple geometrical interpretation as the radius of the sphere tangent to the interface and centred at p . It

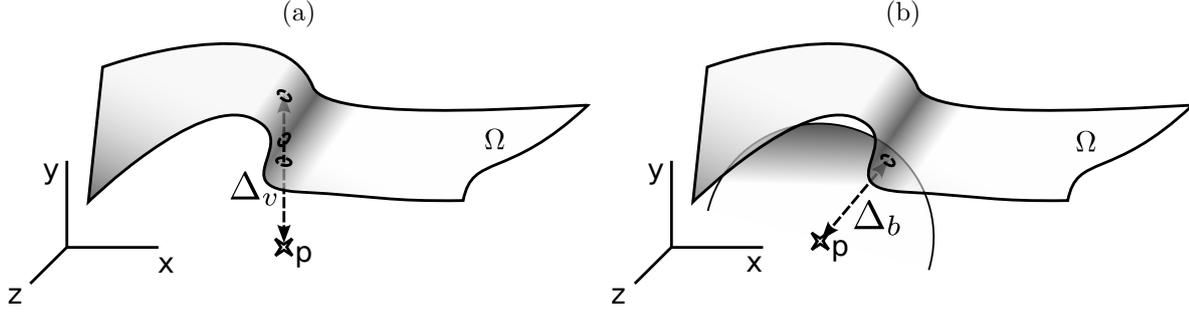


Figure 3.10: (a) Sketch of the vertical distance Δ_v , and (b) the ball distance Δ_b , between a point p and a surface Ω . In the case of Δ_v , the wall-normal line may intersect Ω multiple times, but only the top one is kept. Here, the surface has a pocket and the line crosses it three times. In the case of the ball distance, there is usually only one point where the sphere centred at p with radius Δ_b is tangent to Ω , marked here with a small circle.

is sketched in figure 3.10(b). Some properties of this distance are particularly convenient for a conditional analysis. Regardless of the complexity of the surface, there is always a closest surface point to any point in space, and the ball distance is always uniquely defined. If the point p is close to the interface, Δ_b is equivalent to the distance measured along the local normal. It also has a relatively simple mathematical formulation, since Δ_b satisfies the Eikonal equation $|\nabla(\Delta_b)| = 1$ with $\Delta_b = 0$ at the interface. This equation has a solution regardless of the complexity of the boundary condition, and can be integrated by several fast methods (Jones *et al.*, 2006).

The relation between the two distance definitions depends on the local orientation and complexity of the surface, particularly in the case of handles and pockets. Referring to figure 3.11(a), when the T/NT interface is mostly horizontal, simple and smooth, the two definitions produce similar results. When the interface is more complex or not parallel to the wall, as in figure 3.11(b), their results are different. For example, point p in figure 3.11(b) is very close to the interface in terms of Δ_b , but relatively deep into the turbulent side in terms of Δ_v .

Our algorithm to obtain the ball distance starts from the set Ω_i of interface voxels defined in (3.12). The vorticity within each voxel is approximated by a trilinear interpolation of the values at the vertices, so that the T/NT interface is approximated by a polyhedron of which each interface voxel contains a planar face. Finally, the point of each face closest to the center of the voxel is picked, and the interface is approximated by the set Ω_p of these points. The sets Ω_i and Ω_p are illustrated in figure 3.12.

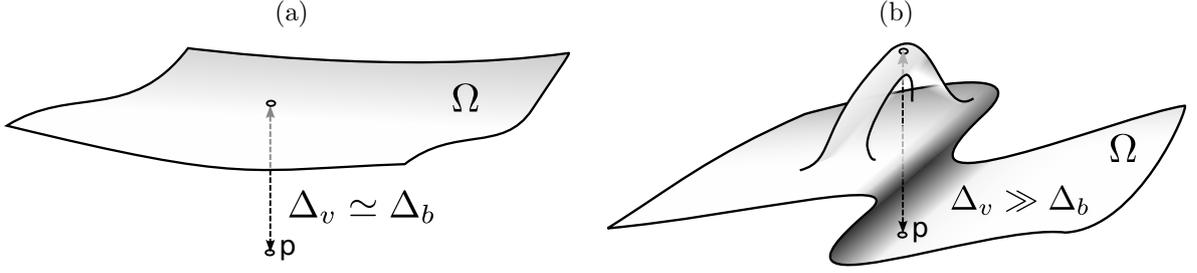


Figure 3.11: (a) Sketch of simple almost horizontal surface for which $\Delta_v \simeq \Delta_b$. (b) Example of a case in which both distances are very different. Here, point p is very close to the interface and $\Delta_b \simeq 0$, but it lies underneath a pocket and a handle, and $\Delta_v \gg \Delta_b$

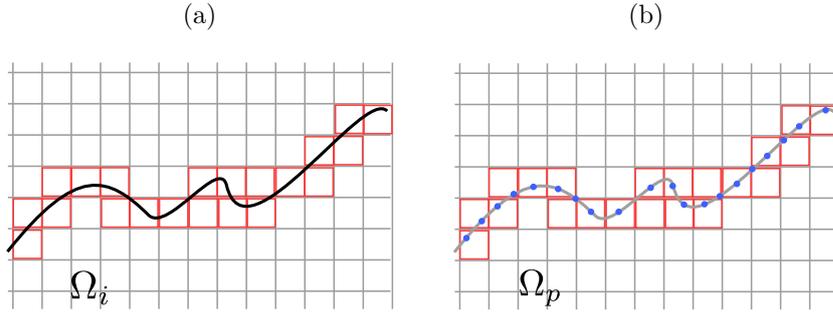


Figure 3.12: (a) Set Ω_i of voxels that contain the interface. (b) Set Ω_p of points used to approximate the T/NT interface.

The ball distance between p and Ω is approximated by the distance between p and its nearest neighbour in Ω_p . The nearest-neighbour search (NNS) is a common problem in optimization. If N_p is number of elements in Ω_p , a fast solution requiring $O(\log N_p)$ computations was found by Arya *et al.* (1998). Most data analysis packages and toolkits provide implementations of some variant of NNS, and free libraries are available (Muja & Lowe, 2014).

In our analysis, the distance to the interface is treated as a field, and computed for all the points in the computational domain. Assuming a total number N of field points, obtaining the field of ball distances requires $O(N \log N_p)$ operations. For our data, N_p is of the order of 10^8 , and N of the order of 10^9 for each snapshot.

The signed distance field

The two distance criteria are evaluated for every collocation point of the computational domain. The discrete fields obtained with the minimum and vertical distance are called the ball-distance field $\Delta_b(x, y, z)$, and the vertical-distance field $\Delta_v(x, y, z)$, respectively.

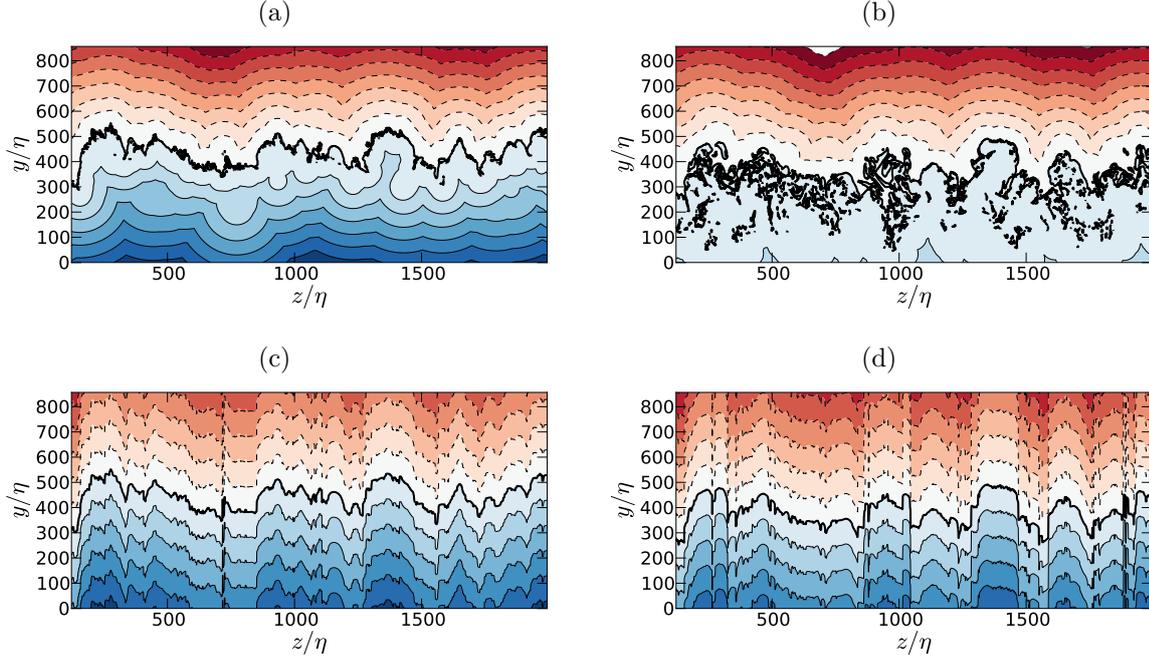


Figure 3.13: Cross-stream sections of the signed distance fields for the lowest and highest available thresholds, and for the two definitions of distance. (a) Δ_b and $\omega_0^* = 0.01$; (b) Δ_b and $\omega_0^* = 0.5$; (c) Δ_v and $\omega_0^* = 0.01$; (d) Δ_v and $\omega_0^* = 0.5$. All sections correspond to the same flow field at $\delta_{99}^+ = 1500$. The thicker solid line represents the T/NT interface for each distance definition, and always corresponds to a single connected surface. The isolated spots are due to three-dimensional contortions. Other contour levels are separated by 50η for Δ_b , and by 100η for the Δ_v . Negative contours are dashed.

The symbol Δ denotes distance regardless of a particular definition. The isosurface $\Delta = 0$ is the effective representation of the interface, but the two definitions of distance generate different isosurfaces. By convention, the distance to the interface is defined as positive or negative depending on whether the point is in the turbulent or in the non-turbulent region. Note that, because the distance is computed with respect to the cleaned interface defined in §3.4, turbulent and non-turbulent points refer to the smoothed flow regions. Bubbles are counted as turbulent, and drops as non-turbulent.

The concept of a distance field is also found in the study of the T/NT interface by Mellado *et al.* (2009), who use the length of the trajectories along lines of maximum gradient of an advected scalar to measure the distance with respect to the interface. While their definition can also be used regardless of the complexity of the surface, the gradient lines of the vorticity magnitude are very contorted in the turbulent side, and less suitable for conditional analysis than any of the definitions mentioned above.

Sections of the two distance fields of the same snapshot of the flow are shown in figure

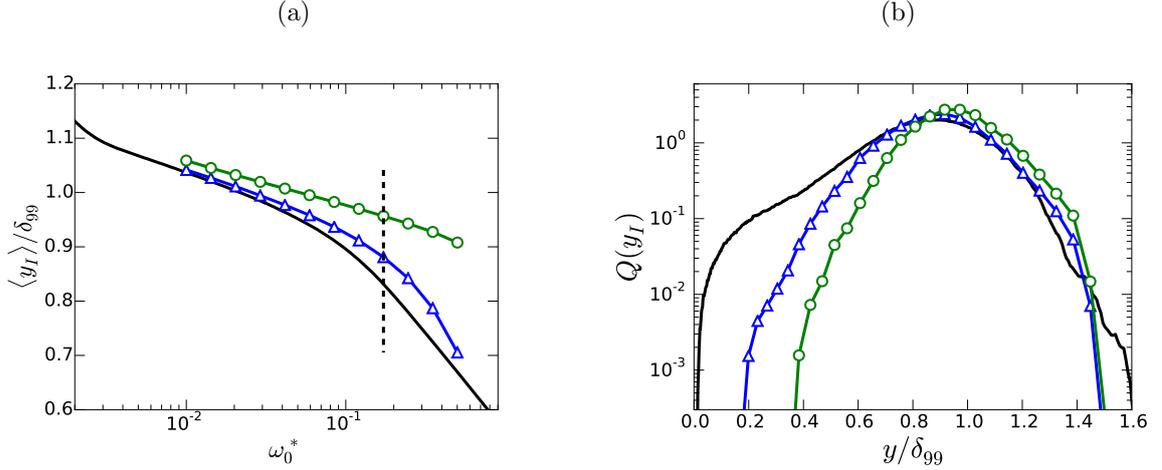


Figure 3.14: (a) Mean position of the T/NT interface as a function of the vorticity threshold. No symbols, vorticity isosurface as in figure 3.4; \triangle , interface defined as $\Delta_b = 0$; \circ , $\Delta_v = 0$. The threshold in (b) is the vertical dashed line. (b) PDFs of the vertical position of the three interfaces for $\omega_0^* = 0.19$. Lines as in (a). $\delta^+ = 1500$.

3.13, each one computed for two different thresholds. They are normalized with the Kolmogorov length η at $y = 0.6\delta_{99}$, which will be used as a reference length scale from now on. The first observation is that the two distances give fairly different results in the turbulent side, particularly for the higher vorticity thresholds. In the non-turbulent side, where the interface is more convex, the differences are not as important. When the threshold is within the topological transition, such as $\omega_0^* = 0.5$ in figures 3.13(b,d), the contortions of the ball-distance interface are so intense that there are very few points in the turbulent side for which $\Delta_b > 100\eta$. We emphasize that $\Delta_b = 0$ in figure 3.13(b) corresponds to a single connected surface from which bubbles have been removed, and that the apparently isolated contours within the turbulent side are artefacts of the two-dimensional section. Comparison of the results of the two thresholds for each distance definition shows that the vertical distance field in figures 3.13(c,d) is less sensitive to the contortions than the ball distance in figures 3.13(a,b), and also less sensitive to the choice of the threshold. Because of this, it misses most of the interface complexity and the existence of a topological transition.

Note that none of the interfaces defined by the distance criteria just mentioned exactly coincides with a vorticity isosurface. In the case of Δ_b the only difference is the absence of the bubbles and drops discarded in the smoothing step, and the deviations are relatively small. The vertical distance misses substantial parts of the isosurface, and may deviate a lot from it. For example, figure 3.14(a) shows the mean position of the two interfaces

as a function of ω_0^* , compared with the mean position of the vorticity isosurface. The mean $\langle y_I(\Delta_b) \rangle$ deviates little from the position $\langle y_I \rangle$ of the vorticity isosurface (figure 3.4), but $\langle y_I(\Delta_v) \rangle$ remains close to the edge of the boundary layer even when the vorticity isosurface moves closer to the wall. This is confirmed by the PDFs of the height of the three isosurfaces, given in figure 3.14(b). For low thresholds (not shown), the PDFs of the two interfaces and of the vorticity isosurface roughly coincide, and are approximately gaussian (Corrsin & Kistler, 1955). But for the higher threshold in figure 3.14(b), $y_I(\Delta_b)$ follows the isosurface into the core of the boundary layer substantially better than $y_I(\Delta_v)$. As a consequence, $y_I(\Delta_b)$ results in a much better representation of the intermittency parameters of the boundary layer, such as γ . Note that the vorticity threshold used in figure 3.14(b), $\omega_0^* = 0.09$, although relatively high, is below the beginning of the topological transition, and in the range of most of the studies collected in table 3.2.

Conditional analysis of distance and vorticity.

The properties of the vorticity conditioned to its position with respect to the interface can be analysed using the joint PDF of the vorticity magnitude and of the distance, $F_{\omega, \Delta}$. Figure 3.15 shows four examples corresponding to the thresholds and distance definitions in figure 3.13. They are part of a more complete set of five Reynolds numbers in the range $\delta_{99}^+ \in (1100 - 1900)$, and ten thresholds in $\omega_0^* \in (0.01 - 0.5)$, each of them computed for the two distance definitions mentioned above. Similar PDFs were obtained in Taveira & da Silva (2014) for planar jets and Δ_v , but only for a limited range of distances.

The joint PDF can be divided into four quadrants, separated by the axes $\Delta = 0$ and $\omega = \omega_0$, marked with dashed lines in figure 3.15. Given that the flow field is the same in the four figures, the differences in $F_{\omega, \Delta}$ are due to the different distance definitions and thresholds. We will order the quadrants in the usual counterclockwise way.

The first quadrant, which contains turbulent points of relatively high vorticity, represents the core turbulent flow. As already seen in figure 3.13, the minimum and vertical distances behave similarly for low thresholds (figures 3.15a,c), but very differently for thresholds within the topological transition. The field of vertical distances depends only slightly on the threshold (figures 3.15c,d), but there are few points at distances beyond $\Delta_b = 100\eta$ for the higher threshold in figure 3.15(b).

The second quadrant contains different geometrical objects depending on the distance

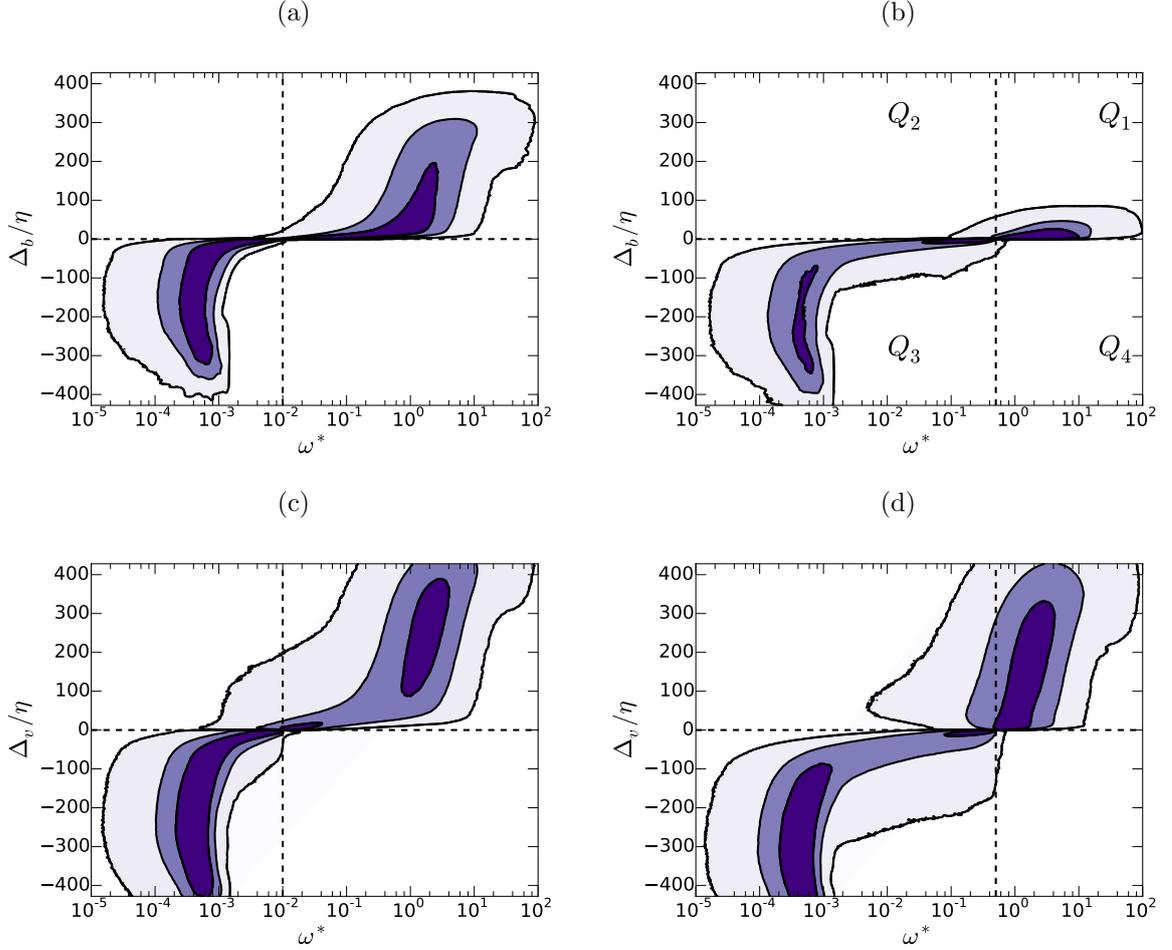


Figure 3.15: Pre-multiplied joint probability density function of vorticity and distance, $\omega F_{\omega, \Delta}$. Each subplot corresponds to the thresholds, Reynolds number, and distance definitions in figure 3.13 (top row Δ_b ; bottom row Δ_v ; left column $\omega_0^* = 0.01$; right column $\omega_0^* = 0.5$). Contours contain 50%, 90%, and 99% of points, respectively.

definition. It contains bubbles for Δ_b , and bubbles, handles, and pockets for Δ_v . For the ball distance, the weight of the second quadrant is always small compared with the first one, and contributes little to the averaged vorticity in the free-stream side of the interface (figures 3.15a,b). This is not the case for the vertical distance, and it is clear from figures 3.15(c,d) that the weight of this quadrant increases as the threshold increases and the interface becomes more complex. This quadrant, with especial reference to the properties of the pockets, will be studied in more detail in §3.5.

The third quadrant contains points of low vorticity classified as non-turbulent. It represents the bulk of the free stream which, in the case of Δ_b , also includes the irrotational pockets. It depends only weakly on the threshold and on the distance definition, except for $\omega \approx \omega_0$.

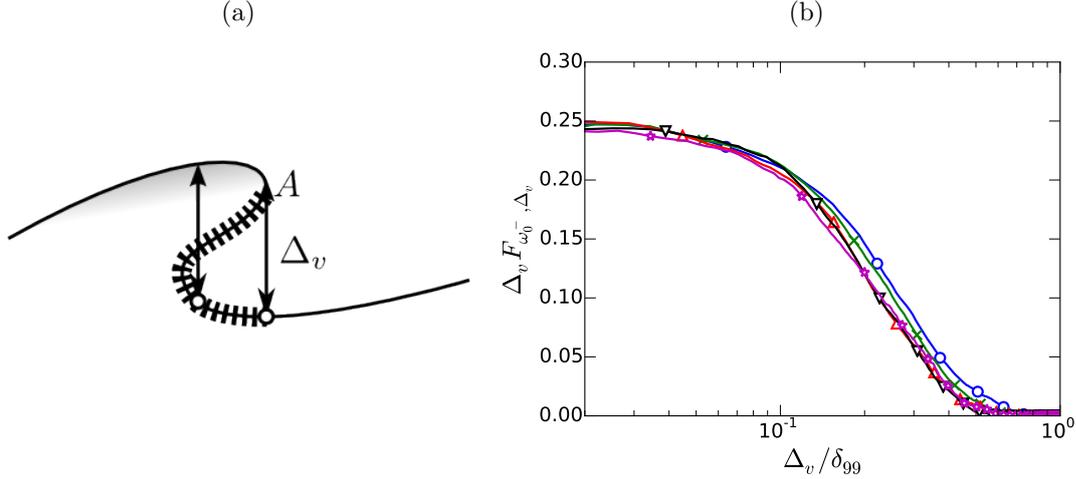


Figure 3.16: (a) Sketch of the discontinuity of the vertical distance at the edge of a pocket. The hatched line represents points where $\omega = \omega_0$ and $\Delta_b = 0$, but $\Delta_v > 0$. The non-turbulent region just outside A has $\Delta_b \simeq 0$ but $\Delta_v < 0$. (b) Premultiplied PDF, $\Delta_v F_{\omega_0^-, \Delta_v}$, of the vertical distance of the non-turbulent points with $\omega^* = 0.25\text{--}0.5$, whose vorticity is close to the threshold $\omega_0^* = 0.5$. \circ (blue), $\delta_{99}^+ = 1100$; \times (green), 1300; \triangle (red), 1500; ∇ (black), 1700; \star (magenta), 1900 .

The fourth quadrant, with $\omega > \omega_0$ and negative distances, corresponds to the objects defined in §3.4 as drops. It is almost empty for all the cases considered in this study, confirming that the smoothing of the free stream described in §3.4 does not affect the results of the conditional analysis.

The influence of the distance definition on the joint PDF is most visible in the neighbourhood of the vertical axis, $\omega = \omega_0$. These are points in which the vorticity is close to the identification isosurface, but that may be incorrectly identified as being far from the interface. The range of possible ball distances for $\omega \simeq \omega_0$ (figure 3.15a,b) is very narrow, $|\Delta_b| < 30\eta$, especially in the second quadrant, and can be interpreted as a typical position of the irrotational bubbles that have been mislabelled as turbulent by the smoothing process. On the other hand, the vertical distances in the same region can be as large as 200η to both sides of the interface (figure 3.15c,d). Denote by ω_0^- the vorticities just below the threshold. The wide Δ_v tails of $F_{\omega_0^-, \Delta_v}$ have several causes, sketched in figure 3.16(a). On the positive side, $\Delta_v > 0$ in Q_2 , all the points represented with a hatched line in that figure are in the vorticity and Δ_b isosurfaces, but not in the Δ_v interface, which is only the top of the overhang. Points near the hatched line have vorticities close to ω_0 , but they are counted as being deep within the turbulent region by Δ_v . The $\Delta_v < 0$ tail of $F_{\omega_0^-, \Delta_v}$ in Q_3 contains points whose vorticity is slightly below than the threshold, but

which are classified by Δ_v as being far within the irrotational region. They correspond to points such as A in figure 3.16(a), in which the orientation of the interface is vertical and induces a discontinuity in the height of the Δ_v interface. Such discontinuities are clearly visible in figures 3.13(c,d). These tangencies are less common than the overhung surfaces, and the mass in the negative tail of $F_{\omega_0^-, \Delta_v}$ is typically smaller than in the positive one, especially in the convoluted interfaces at the higher thresholds (15% in the case of figure 3.13d).

It is clear from 3.16(a) that the negative tail of $F_{\omega_0^-, \Delta_v}$ contains information about the ‘depth’ of the pockets, rather than about the thickness of the interface. The premultiplied probability distribution $\Delta_v F_{\omega_0^-, \Delta_v}$, integrated over the band $\omega \in (\omega_0/2, \omega_0)$, is presented in figure 3.16(b) for a relatively high threshold. It is well approximated by a power law $F_{\omega_0^-, \Delta_v} \propto \Delta_v^{-1}$ for $\Delta_v \lesssim 0.2\delta_{99}$. Although the reason for this particular power is not completely clear, it suggests a regular structure for the Δ_v interface. That interface has no overhangs, and represents pockets as holes with steep sides. If we assume pockets of size Δ , the contribution of each hole to the PDF in figure 3.16(b) would be proportional to the $O(\Delta)$ length of its lip. Their number would be proportional to Δ^{-2} and the total lip length would be proportional to Δ^{-1} , as in the figure.

In any case, the fact that the distribution of pocket heights satisfies a power law is consistent with the fractal nature of the interface, and suggests that the discontinuities are a consequence of a self-similar hierarchy of overhangs. For the threshold in figure 3.16(b), the self-similar range ends around $\Delta_v \approx 0.2\delta_{99}$, and the probability of finding pockets deeper than that limit is very low. This is about three times the standard deviation of the position of the vorticity isosurface for this threshold (figure 3.4). At lower thresholds, such as those in figures 3.15(a,c), the self similar range disappears, and the ‘pocket’ distribution is concentrated into a peak near $\Delta_v = 10\eta$.

Conditional averages

The averaged vorticity conditioned to the distance to the interface can be computed from $F_{\omega, \Delta}$ as

$$\bar{\omega}(\Delta) = \frac{\int_0^\infty \omega F_{\omega, \Delta} d\omega}{\int_0^\infty F_{\omega, \Delta} d\omega}. \quad (3.17)$$

It is given by the solid lines with squares in figures 3.17(a-d), and is equivalent to the conditional vorticity profiles in Bisset *et al.* (2002), Westerweel *et al.* (2002) and da Silva

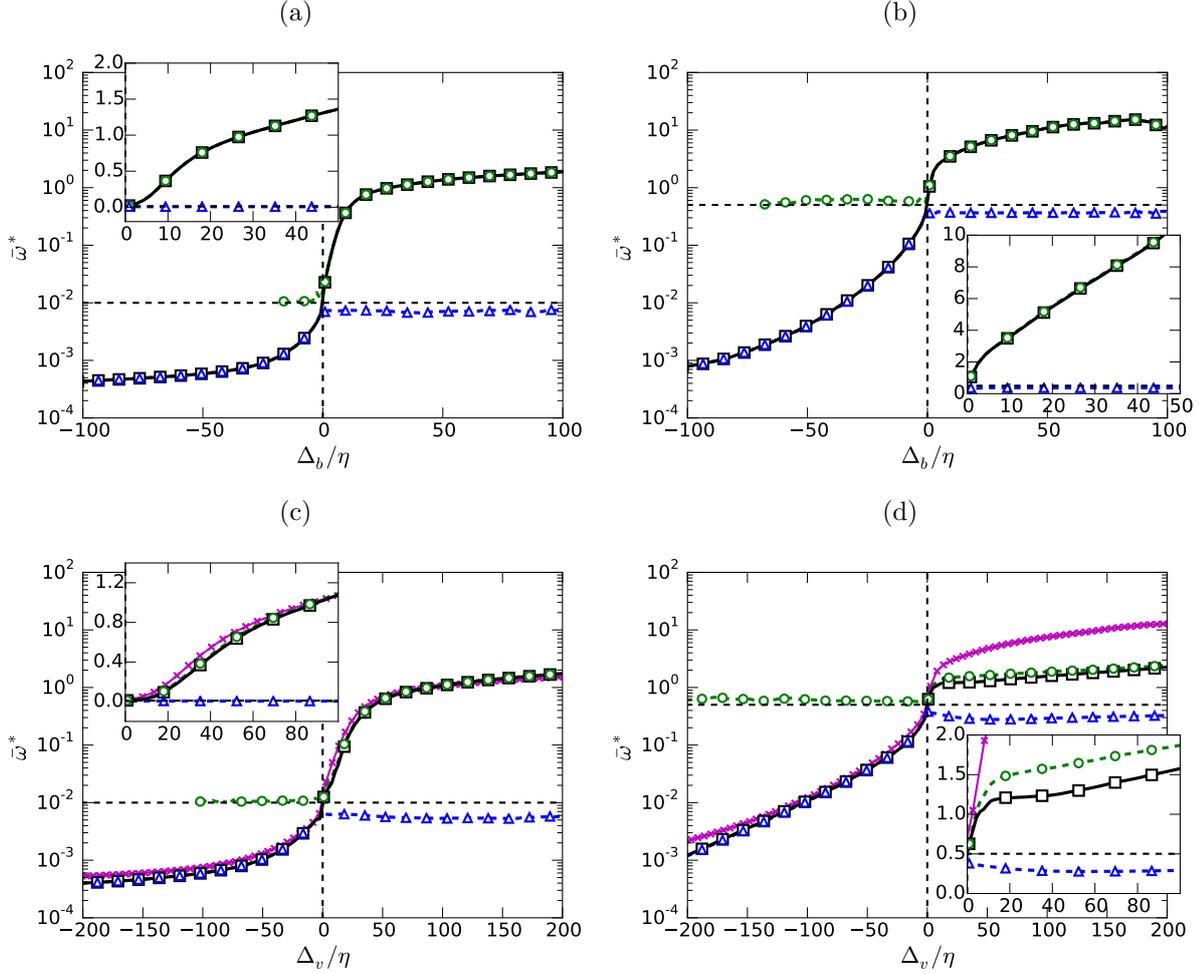


Figure 3.17: Conditional profiles: $-\square-$, $\bar{\omega}_1$; $-\circ-$, $\bar{\omega}_2$; $-\triangle-$, $\bar{\omega}_3$. The figures follow the same arrangement as in figure 3.13 (top row Δ_b ; bottom row Δ_v ; left column $\omega_0^* = 0.01$; right column $\omega_0^* = 0.5$). The black dashed lines correspond to the value of the threshold (vertical) and zero distance (horizontal). The inset in each figure correspond to the same plot, using linear coordinates for the vorticity magnitude.

et al. (2011). Note the use of the bar over the symbol to distinguish (3.17) from the more usual mean profile $\langle \omega \rangle$ at a given distance from the wall, defined as

$$\langle \omega \rangle(y) = \frac{\int_0^\infty \omega \Gamma_{\omega,y} d\omega}{\int_0^\infty \Gamma_{\omega,y} d\omega}. \quad (3.18)$$

We will use the notation $\bar{\omega}(\Delta_b)$ and $\bar{\omega}(\Delta_v)$ to distinguish between conditional profiles obtained with each definition of distance.

The conditional vorticity in all the panels of figure 3.17 increases to its expected fully turbulent level, $\omega^* = O(1)$, within a few Kolmogorov lengths from the interface. This rise is monotonic, except perhaps for the plateau at $\Delta_v/\eta = 15\text{--}40$ in figure 3.17(d), which starts at a distance to the interface of the order of the diameter of a single turbulent

vortex (10η , see Jiménez *et al.*, 1993).

The existence of a plateau or of a maximum in the conditional vorticity profile near the T/NT interface has been mentioned in wakes (Bisset *et al.*, 2002; Townsend, 1976) and reported in jets (da Silva *et al.*, 2011; Westerweel *et al.*, 2009). Its presence has sometimes been used to define the thickness of the interface layer (da Silva & Taveira, 2010), and taken as the basis for theoretical models in which the interface is maintained by the presence of a strong localised shear (Hunt & Durbin, 1999). Similar models have been used to suggest similarities between the T/NT interface in jets (Westerweel *et al.*, 2009) and strong internal vortex layers in homogeneous turbulence (Ishihara *et al.*, 2013). Chauhan *et al.* (2014) report a strong conditional vorticity peak in boundary layers, but their interface is defined in terms of the streamwise velocity, and is probably unrelated to the one discussed here. Moreover, not all these papers use the same definition of the interface or even the same thresholded scalar. In fact, when da Silva *et al.* (2014 *a*) compile conditional vorticity statistics for a variety of flows, the only obvious peak is found at the early stages of the evolution of a shearless mixing layer (da Silva & Taveira, 2010). Bisset *et al.* (2002) also find strong vorticity peaks for some high vorticity thresholds in their wake, but attribute them to the presence of isolated vorticity patches, and discard them in favour of a lower threshold ($\omega^* \approx 0.1$) for which the maximum is barely noticeable. Note that, if the vorticity were particularly intense close to the interface, a plateau analogous to the one in figure 3.17(d) should also appear in the $\bar{\omega}(\Delta_b)$ profile in figure 3.17(d), but this is not the case. An alternative explanation is that the vorticity close to the interface is not really particularly intense but that, when the conditional profiles are obtained as a function of Δ_v at a sufficiently high threshold, some non-turbulent flow is counted as being turbulent within the inner part of the interface, lowering the local average vorticity.

To differentiate between the two hypotheses we split the conditional profile $\bar{\omega}(\Delta_v)$ into contributions from the high-vorticity first quadrant, Q_1 , and the mislabeled non-turbulent points in Q_2 . Equation (3.17) can be split into

$$\bar{\omega} = W_1\bar{\omega}_1 + W_2\bar{\omega}_2, \quad (3.19)$$

where

$$\bar{\omega}_1 = \frac{\int_{\omega_0}^{\infty} \omega F_{\omega,\Delta} d\omega}{\int_{\omega_0}^{\infty} F_{\omega,\Delta} d\omega}, \quad \bar{\omega}_2 = \frac{\int_0^{\omega_0} \omega F_{\omega,\Delta} d\omega}{\int_0^{\omega_0} F_{\omega,\Delta} d\omega}, \quad (3.20)$$

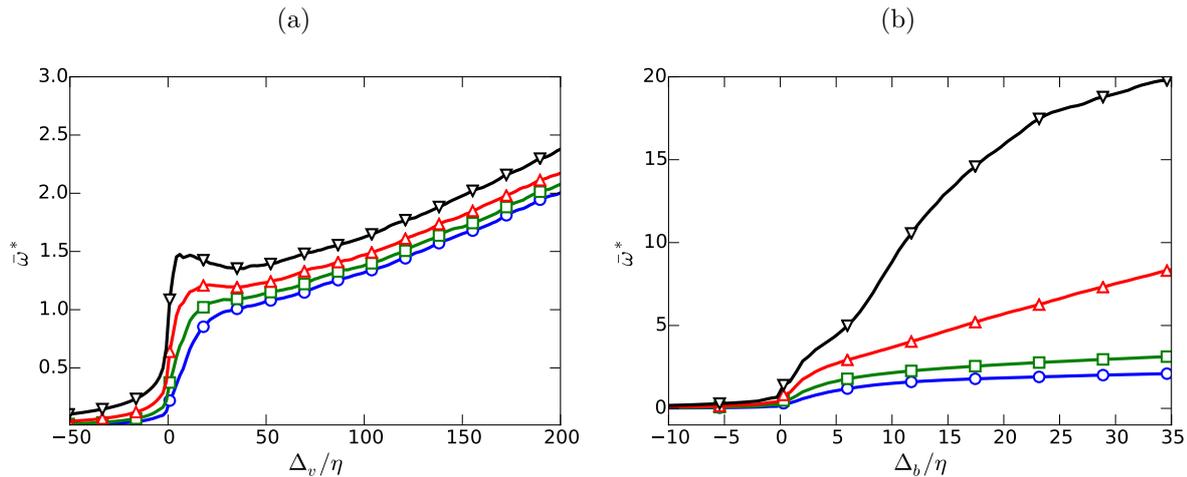


Figure 3.18: (a) Conditional vorticity profiles for $\delta_{99}^+ = 1900$, computed as in figure 3.17, as functions of the threshold. \circ , $\omega_0^* = 0.17$; \square , 0.29; \triangle , 0.52; ∇ , 0.88. (a) Δ_v . (b) Δ_b .

are the conditional averages for Q_1 and Q_2 , and

$$W_1 = \frac{\int_{\omega_0}^{\infty} F_{\omega,\Delta} d\omega}{\int_0^{\infty} F_{\omega,\Delta} d\omega}, \quad W_2 = \frac{\int_0^{\omega_0} F_{\omega,\Delta} d\omega}{\int_0^{\infty} F_{\omega,\Delta} d\omega}, \quad (3.21)$$

are the corresponding weights. The profiles of $\bar{\omega}$, $\bar{\omega}_1$, and $\bar{\omega}_2$ are given in figure 3.17. In the case of low thresholds (left column of the figure), $\bar{\omega}_1 \simeq \bar{\omega}$, and the contribution of the second quadrant is small, regardless of the distance definition.

The only case in which $\bar{\omega}_1$ is clearly different from the overall average is figure 3.17(d), in which the contribution of the handles and pockets is significant. In this figure, the maximum relative weight of Q_2 is $W_2 \simeq W_1/4$ at $\Delta_v = 20\eta$. At the even higher thresholds at which the interface reaches its maximum geometrical complexity near the end of the topological transition, the weights of the two quadrants are comparable. This has a noticeable effect on the conditional profiles at the higher thresholds, and it is clear from figure 3.17(d) that the plateau is a consequence of the negative contribution from $\bar{\omega}_2$. If we consider this contribution as a spurious effect of Δ_v , the ‘true’ conditional vorticity $\bar{\omega}_1$ in figure 3.17(d) increases monotonically near the the interface. In essence, the conditional vorticity remains constant or decreases away from the interface because Δ_v misclassifies some weakly vortical pockets as part of the turbulent flow.

This effect is clearer in figure 3.18, which presents conditional vorticities for several interface thresholds. Figure 3.18(a) is computed with Δ_v , and develops a plateau and eventually a peak as the threshold increases. As in figure 3.17(d), it can be shown that this is a due to the increasingly negative contribution from the pockets as the complexity

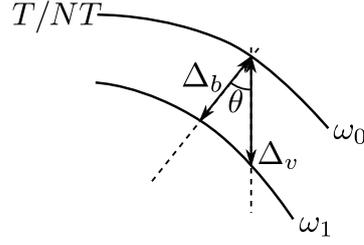


Figure 3.19: Sketch of the geometrical meaning of equation (3.22). The ball distance between any point with vorticity ω_1 and the T/NT interface corresponds to the distance along the local normal. The angle theta θ is measured from the local normal to the vertical direction. This figure is a simplification of the three-dimensional case.

of the interface increases. Figure 3.18(b) presents the same cases computed for Δ_b , and shows no trace of an interface peak.

Note that the distances in figure 3.18(b) are much lower than in figure 3.18(a), while the conditional vorticities are higher. In fact, similar conditional vorticities are found when the horizontal axis of figure 3.18(a) is extended to $\Delta_v \simeq 400$, carrying the plot to the neighbourhood of the wall. The plot of $\bar{\omega}(\Delta_v)$ for these large distances is very similar to a shifted version of $\omega'(y)$ (figure 3.3a). The vorticity isosurface at these high thresholds permeates the whole boundary layer, and occasionally comes very close to the wall. The ball distance recognises this fact and brings the strong near-wall vorticity closer to the interface, while the vertical distance misses that complexity.

This discussion suggests that the apparent strongly vortical interface layer observed in some of the studies mentioned above is an artefact of how a one-dimensional definition of distance interacts with a fully three-dimensional geometry, whether that definition is linked to the normal to the wall or to some other arbitrary direction. In particular, it is due to the neglect of the effect of irrotational pockets on the conditional quantities. We will discuss the relevance of these pockets in the entrainment process in §sec:pockets.

Practical similarity between Δ_b and Δ_v

Under some assumptions, minimum and vertical distance can be related. At short distances from the interface, and assuming that the ball distance is equivalent to the distance along the local normal to the surface,

$$\Delta_b \sim \Delta_v \cos(\theta), \quad (3.22)$$

where θ is the angle between the local normal and the vertical direction. The geometrical meaning of equation (3.22) is sketched in figure 3.19. The average ratio between the two distance measurements can be estimated with

$$\frac{\overline{\Delta_b}}{\overline{\Delta_v}} \sim \overline{\cos \theta} = \int_0^{\pi/2} \cos \theta P(\theta) d\theta. \quad (3.23)$$

In the previous equation, $P(\theta)$ is the PDF of the local orientation of the surface, the angle between the local normal and the vertical direction. Looking at the magenta line with crosses in figure 3.17(c), $\overline{\Delta_b/\Delta_v} = 1/\pi$ fits the conditional profiles for a low threshold with remarkable accuracy, even for relatively long distances respect to the interface. This resemblance is lost when the threshold is chosen within the topological transition (figure 3.17(d)), where the profiles are no longer similar and the approximation of the equation (3.22) is not valid. This ratio of $1/\pi$ cannot be obtained with a simple shape. A cylinder with the axis parallel to the wall is oriented as $P(\theta) = 2/\pi$, hence $\overline{\cos \theta} = 2/\pi$. The orientation of a sphere corresponds to $P(\theta) = \sin \theta$, and $\overline{\cos \theta} = 1/2$. The relatively low ratio suggests that the interface is seldom oriented parallel to the wall, and in consequence, the simplified picture of the T/NT interface as a wavy, slightly contorted surface, would be conceptually wrong.

The relevance of pockets

We saw in figure 3.16 that pockets form a self-similar hierarchy of many different sizes, and it has been conjectured that their formation signals the large-scale engulfment of irrotational fluid before it is finally entrained by small-scale ‘nibbling’. Their abundance has been used to quantify the relative importance of the two processes (Mathew & Basu, 2002; Sandham *et al.*, 1988).

We can define pockets as regions identified by the ball distance as part of the free stream, $\Delta_b < 0$, and by the vertical distance as turbulent, $\Delta_v > 0$. For the purpose of this section, they include the underside of handles as well as simple folds of the interface. Figure 3.20(a) shows the joint PDF of the two distances in the range corresponding to pockets. The figure is drawn for the relatively high vorticity threshold of figures 3.13(b,d), guaranteeing both the presence of abundant pockets and the possibility of observing how the vorticity diffuses into the irrotational flow. It includes the two extreme Reynolds numbers in our simulation, allowing some scaling comparisons. It turns out that the size of pockets, as measured by the maximum Δ_b , scales best in terms of the Kolmogorov

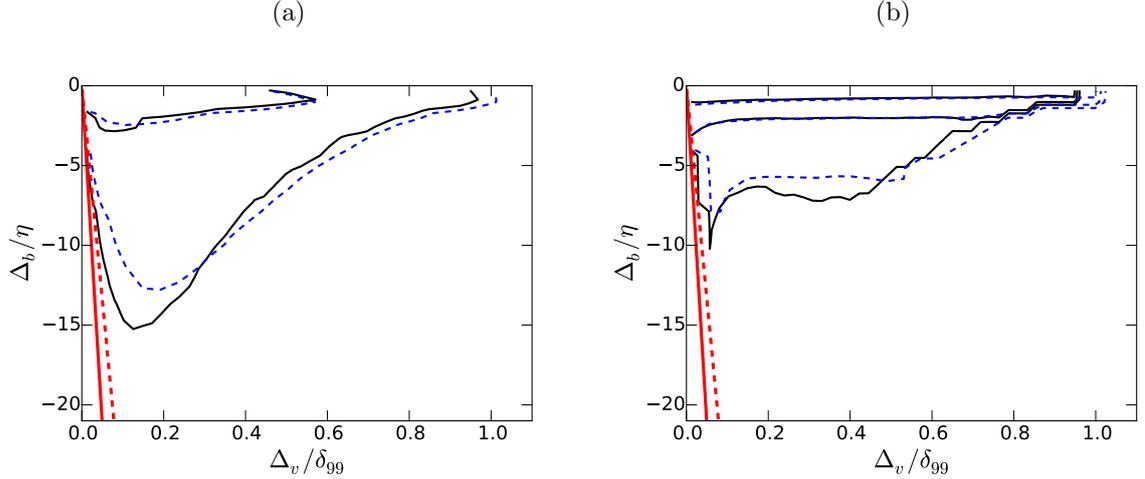


Figure 3.20: Flow within pockets of the interface at $\omega_0^* = 0.5$. ---, $\delta_{99}^+ = 1100$; —, $\delta_{99}^+ = 1900$ (a) Joint PDF of Δ_v and Δ_b within the region. The two contours for each Reynolds number contain 60% and 98% of the points, respectively. (b) Average vorticity within the pockets as a function of Δ_v and Δ_b within the region that contains 98% of the points. Contours are $\omega^* = 0.15, 0.25, 0.35$. The two diagonal red lines are $\Delta_v = \Delta_b$.

viscous length, while their depth within the layer, as measured by Δ_v , scales better with the boundary-layer thickness. The joint PDF is roughly triangular. It is bounded on the left by the trivial limit $\Delta_b \leq \Delta_v$, plotted for each Reynolds number as a thick inclined straight line, and on the right by a roughly hyperbolic curve that can be interpreted to mean that deeper pockets (large Δ_v) tend to be smaller (small Δ_b), presumably because they have been broken down by the turbulence while being entrained into the layer.

The question of whether being entrained into a pocket also promotes the diffusion of vorticity is tested in figure 3.20(b), which shows the distribution of the conditionally averaged vorticity in the same parameter space as figure 3.20(a). Note that all the vorticity levels in this figure are below the interface threshold, so that the band of higher vorticities along the top of figure 3.20(b) portrays how vorticity diffuses into the irrotational fluid. Its width, approximately $5\text{--}10\eta$, strongly suggest a viscous origin (van Reeuwijk & Holzner, 2014), and it is clear from the figure that the vorticity is correlated with the ball distance, but not with the vertical position with respect to the interface. The only exceptions are points near the line $\Delta_v = \Delta_b$, where both measures coincide.

The implication is that the fluid within pockets is sensitive to how close it is to the interface, but not to how deep it is within the turbulent layer. If engulfment were an important mechanism to promote the diffusion of vorticity into the irrotational fluid, for example by preferentially straining it, one would expect some correlation between Δ_v and

the width of the diffusion band at the top of figure 3.20(b), but there is little evidence for that. Apparently, whether the fluid is within a pocket or not is immaterial to its behaviour, although the break-up of the deeper pockets into smaller sizes should enhance the overall effect of viscous diffusion. We will only use Δ_b from now on in our analysis.

The thickness of the interface layer.

While the previous sections deal with the properties of the interface *surface*, it is also interesting to characterise the properties of the interface *layer*, understood as the part of the turbulent flow that is directly influenced by its proximity of the free stream. As a first step, figures 3.21(a,b) reproduce the first and second (turbulent) quadrants of the joint PDFs of the vorticity and distance in figures 3.15(a,b). The distance axis is now logarithmic, to emphasize the region close to the interface, and each figure includes the two extreme Reynolds numbers in our data set.

Three regions can be distinguished in order of increasing distance from the interface. The first and closest to the interface contains the strongest vorticity gradients. If we define the limit of this layer by the intersection of two straight lines tangent to the probability isocontours near and far from the interface, its thickness scales well in Kolmogorov units for the different Reynolds numbers, implying a viscous origin. It is approximately 10η in figure 3.21(a) ($\omega_0^* = 0.01$), 5η in figure 3.21(b) ($\omega_0^* = 0.09$), and almost vanishes at the beginning of the topological transition, $\omega_0^* = 0.2$ (not shown). In the cases in which this region can be identified in the joint PDF, its limit is roughly $\bar{\omega}^* = 1$, which we have seen above to be the level of fully developed turbulence.

The viscosity-dominated region just outside the interface has been recently studied by van Reeuwijk & Holzner (2014) and Taveira & da Silva (2014) in temporally evolving turbulence fronts. They identify it with the ‘superlayer’ conjectured by Corrsin & Kistler (1955), and find that its characteristic thickness is the Kolmogorov microscale computed with the energy dissipation rate of the core flow. The enstrophy level in this viscous layer depends somewhat on the definition, but is typically very low. The viscous region in figure 3.21 is probably not the superlayer, whose observation requires a higher numerical resolution and a quieter free stream than those in our simulation (van Reeuwijk & Holzner, 2014). We will see later that both the rate of strain and the vortex stretching remain high in the viscous layer of figure 3.21, and that region is probably best interpreted as part of the ‘buffer layer’ defined by van Reeuwijk & Holzner (2014) in the range $\omega^* \in (0.1 - 1)$.

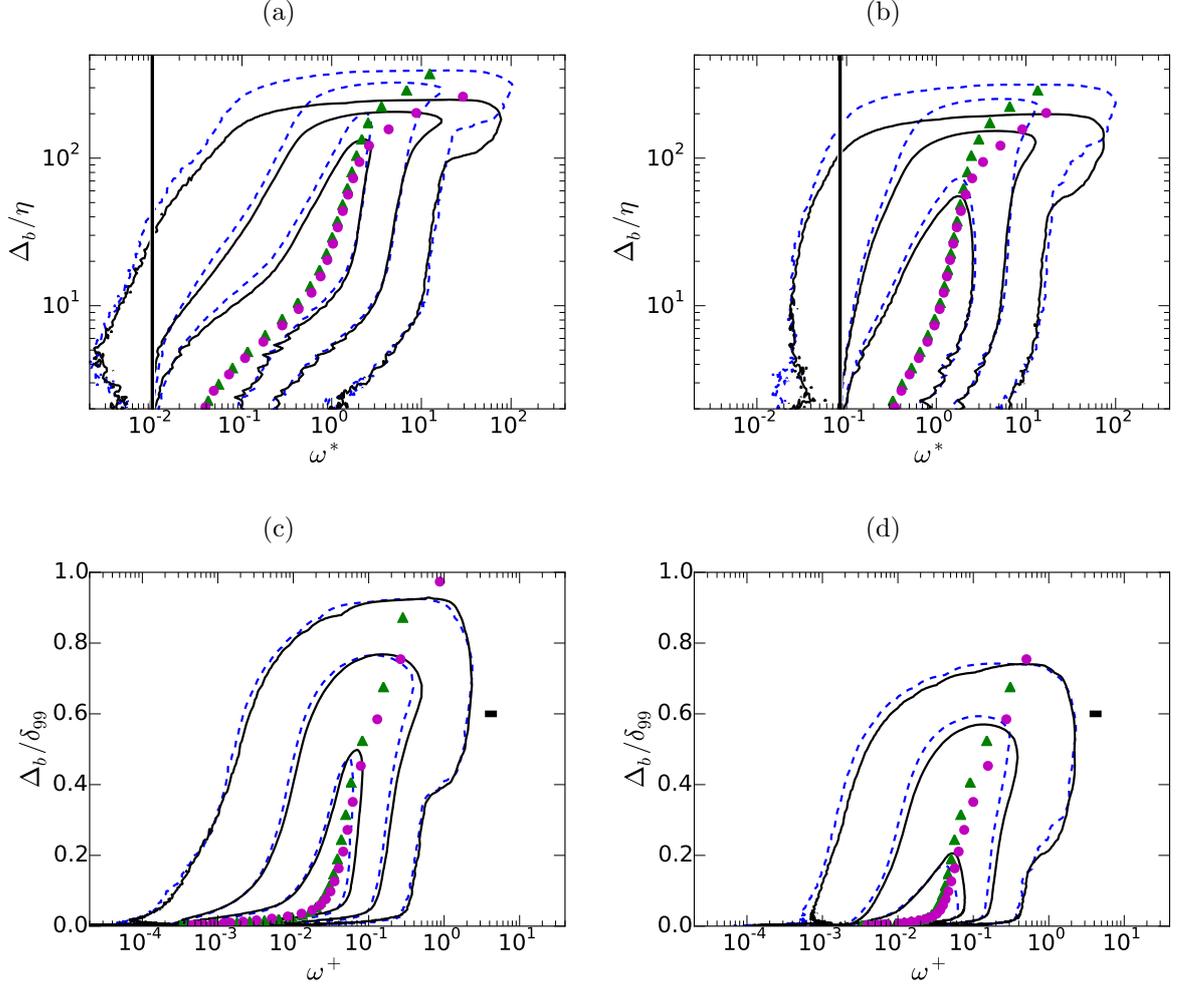


Figure 3.21: Premultiplied joint probability density function of the vorticity and ball distance in the turbulent side of the interface ($\omega F_{\omega, \Delta_b}$) for: (a) a low threshold, $\omega_0^* = 0.01$, and (b) a moderate one at the beginning of the topological transition, $\omega_0^* = 0.09$. Two Reynolds numbers are presented in each figure, $\delta_{99}^+ = 1100$ (— black), and $\delta_{99}^+ = 1900$ (--- blue). The vertical solid line is ω_0^* . (c,d) Same as (a,b), but with the vorticity in wall units and the distance normalized with the boundary layer thickness. The horizontal bar is the variation of ω^*/ω^+ in our range of δ_{99}^+ . The curves with markers correspond to the average vorticity magnitudes for each Reynolds number, $\delta_{99}^+ = 1900$ (green Δ), and $\delta_{99}^+ = 1100$ (magenta \circ). Contours contain 50%, 90%, and 99% of points, respectively.

In analogy to the similarly named layer in wall-bounded turbulence, both nonlinear and viscous effect are important that region. It is interesting that such a hybrid mechanism was proposed by Townsend (1976), who noted that viscous diffusion of vorticity and its tangential transport should be of the same order close to the interface. The predicted result of this mechanism is also a thickness $O(\eta)$. Note that both viscous mechanisms, but especially the latter one, extend to the neighborhood of the interface the intuitive idea that the small scales of turbulence are universal (Kolmogorov, 1941).

The region beyond the viscous interface layer is self similar, in the sense that both the conditionally averaged vorticity and the probability isocontours follow power laws in Δ_b . That implies that there is no intrinsic length or vorticity scale, which is consistent with the intuitive idea of an inertial layer but is a little difficult to interpret dynamically. In the first place, different parts of the PDF scale differently. The strong-vorticity isocontours to the right of figure 3.21(a) are almost vertical ($\omega \propto \Delta_b^0$), but those corresponding to weak vorticity on the left of the figure follow $\omega \propto \Delta_b^1$, and the conditional mean enstrophy approximately satisfies $\omega \propto \Delta_b^{1/2}$. In the second place, those slopes change with the detection threshold, and it is hard to distinguish any power law in figure 3.21(b) or at higher thresholds. We mentioned in (3.6)–(3.5) that $\langle \omega \rangle \propto y^{-1/2}$ is a consequence of the self-similarity of the logarithmic layer, but the same argument cannot be used here. The interface is not an impermeable boundary that limits the size of eddies as the wall does, although it could be argued that the size of the eddies defines the position of the interface. The trend in figure 3.21(a), that the larger eddies have more intense vorticity, is contrary to the inertial relation of homogeneous turbulence, $\omega \propto \Delta^{-2/3}$ (Kolmogorov, 1941), and the most plausible explanation is that larger eddies reach closer to the wall and are therefore stronger. The self-similarity in the figure may be coincidental.

The width of this intermediate region depends on the identification threshold, but scales with the boundary layer thickness. It extends to the hockey-stick at the top of the PDFs, which contains the points with the highest vorticity and farthest from the interface. This last region is mostly formed by points near the wall. When ω is scaled in wall units and Δ_b is normalized with the boundary-layer thickness, as in figures 3.21(c, d) the two Reynolds numbers collapse well for long distances and high vorticities. At the two Reynolds numbers in figure 3.21, $\delta_{99}/\eta \approx 250$ and 450, respectively.

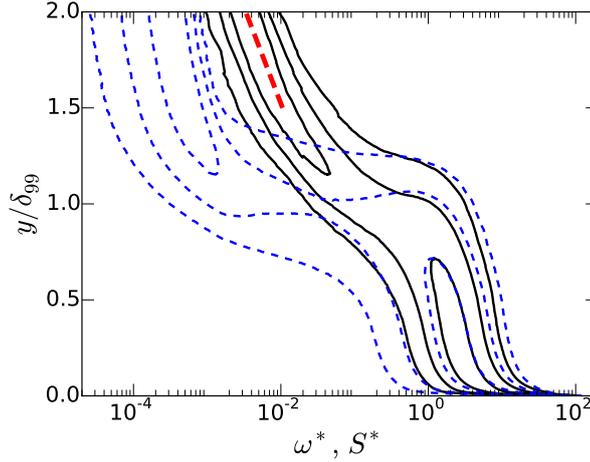


Figure 3.22: Premultiplied joint PDFs: —(black), $S\Gamma_{y,S}$; ---(blue), $\omega\Gamma_{y,\omega}$. Contours contain 50%, 90%, and 99% of points, respectively. The dashed diagonal is the exponential decay of the Fourier modes of irrotational strain with a wall-parallel wavelength $2\delta_{99}$.

3.6 Other velocity gradients

In the previous sections, we have discussed the properties of the vorticity field near a vorticity isosurface, and it is perhaps not surprising that they may be special. For example, an interesting question is whether the vorticity within the interface layer has different properties from the core of the turbulent flow, such as perhaps being weaker because it is less strained, but such questions are hard to answer if the interface is defined by the magnitude of the vorticity itself. It is useful for that purpose to determine the conditional properties of quantities other than the one being thresholded. In this section we will study the properties of the strain rate tensor \mathbf{S} in the neighbourhood of the vorticity interface, as well as the behaviour of the vorticity in the neighbourhood of an interface defined in terms of the strain. Define S as the euclidean norm of the rate-of-strain tensor, $S = \|\mathbf{S}\|$. In analogy to equation (3.7), and taking into account that

$$\langle \omega^2 \rangle = 2\langle S^2 \rangle \quad (3.24)$$

in homogeneous flows, the star units for the S are defined as

$$S^* = S \frac{\nu \sqrt{2\delta_{99}^+}}{u_\tau^2}. \quad (3.25)$$

Equation (3.24) then becomes $\langle \omega^2 \rangle^* = \langle S^2 \rangle^*$, and suggests that ω^* and S^* should be of the same order. The joint PDFs of S and y , and of ω and y , are presented in figure 3.22.

Both PDFs agree within the turbulent region in the right-lower corner of figure 3.22, supporting the normalisation (3.25), but the vorticity in the free stream on the left-

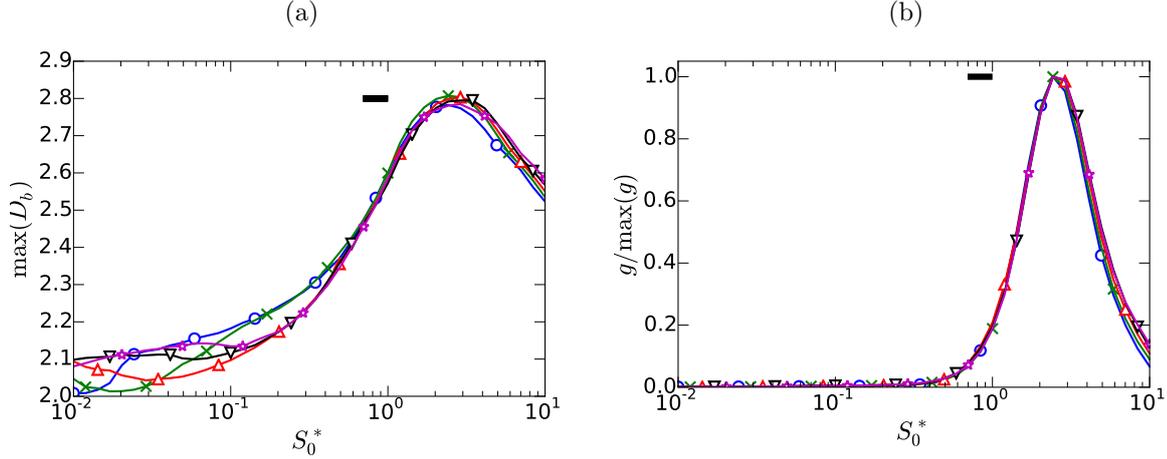


Figure 3.23: (a) Fractal dimension and (b) genus of an interface defined by thresholding the norm of the strain-rate tensor. Five different values of δ_{99}^+ are presented: \circ , 1100; \times , 1300; ∇ , 1500; \triangle , 1700; and \star , 1900. The horizontal bar is the variation of ω^*/ω^+ in our range of δ_{99}^+ .

hand side of the figure is almost two orders of magnitude lower than the rate of strain. This is not unexpected in a nominally irrotational part of the flow but, since (3.24) is a kinematic relation whose only condition is spacial homogeneity, the mismatch between the two magnitudes implies that the strain in the free stream has to be an inhomogeneous residual effect of the vortical flow within the boundary layer.

Any solenoidal velocity field can be written as

$$\mathbf{u} = \nabla \wedge \mathbf{B} + \nabla \phi, \quad (3.26)$$

where the potentials satisfy, $\nabla^2 \phi = 0$ and $\nabla^2 \mathbf{B} = -\boldsymbol{\omega}$ (Batchelor, 1967). In the irrotational free stream, both potentials satisfy Laplace's equation and, if they are expanded in terms of wall-parallel Fourier harmonics, decay away from the wall as $\exp(-ky)$, where $k^2 = k_x^2 + k_z^2$ is the magnitude of the wall-parallel wave vector. All the velocity components and the rate-of-strain tensor decay exponentially at the same rate, and the slowest decay corresponds to the largest horizontal wavelengths. The thick dashed diagonal in figure 3.22 is $S \propto \exp(-\pi y/\delta_{99})$, corresponding to the decay of irrotational velocity fluctuations due to structures within the boundary layer whose shortest dimension is $O(2\delta_{99})$. This is the order of magnitude of the largest structures in boundary layers (Sillero *et al.*, 2014).

The vorticity is unrelated to the velocity potentials, and decays much faster than the rate of strain as it enters the free stream. In fact, this was one reason why we originally chose vorticity over other quantities to characterise the T/NT interface.

Note that the vorticity also decays exponentially with y in the free stream, although at a much lower level than S . This is not a kinematic result, but a consequence of the numerical inflow conditions, which are used to determine the three velocities at the inflow but cannot control their derivatives with respect to x . As a consequence there is a residual vorticity in the free stream due to terms like $\partial_x v$, which inherits the exponential decay of the velocity potentials at the inflow plane.

A consequence of the relatively high strain levels in the free-stream is that the separation between its characteristic values in the turbulent and the non-turbulent sides is not as clear-cut as in the case of the vorticity. Even so, the topological transition happens for comparable thresholds. The geometrical properties of the strain interface are presented in figure 3.23. The fractal dimension in figure 3.23(a) should be compared to figure 3.8(b) for the vorticity interface. The strain isosurface is smoother, with minimum values close to the non-fractal value, $D \approx 2$. The maximum dimension is also somewhat lower than for the vorticity, in agreement with the observation by Moisy & Jiménez (2004) that strong dissipation structures are less fractal (plate-like) than those of vorticity (string-like). The evolution of the genus in figure 3.23(b) is also similar to the case of the vorticity, although the maximum genus and fractal dimension are reached for slightly higher thresholds than for the vorticity, $S_0^* \approx 2.5$ instead of $\omega_0^* \approx 1.5$. The topological transition is also narrower for the strain interface, especially for the genus in figure 3.23(b), which starts to increase at $S_0^* \approx 1$ instead of at $\omega_0^* \approx 0.3$, as it did in figure 3.9(a). The reason is probably that while the maximum dimension and genus mark the threshold for which the interface has fully moved into the core turbulent flow, the slower decay of the strain fluctuations with y means that low-strain isosurfaces are much farther from the wall than similar enstrophy thresholds, and the corresponding interfaces becomes regular much faster.

The evolution of the conditional statistics of the flow across the vorticity and strain isosurfaces are compared in figure 3.24. The two thresholds chosen are $S_0^* = 0.1$ and $\omega_0^* = 0.01$, both of which are within the plateau that separates the values of turbulent and non-turbulent flow in their joint PDFs, and well below the beginning of the respective topological transitions. The average height of the resulting interfaces is similar, $\langle y_I \rangle \approx 1$. Figure 3.24 shows the conditionally averaged enstrophy and strain for each of the two interfaces. They are plotted as functions of the respective ball distances, which we will denote by Δ_b^ω and Δ_b^S , respectively.

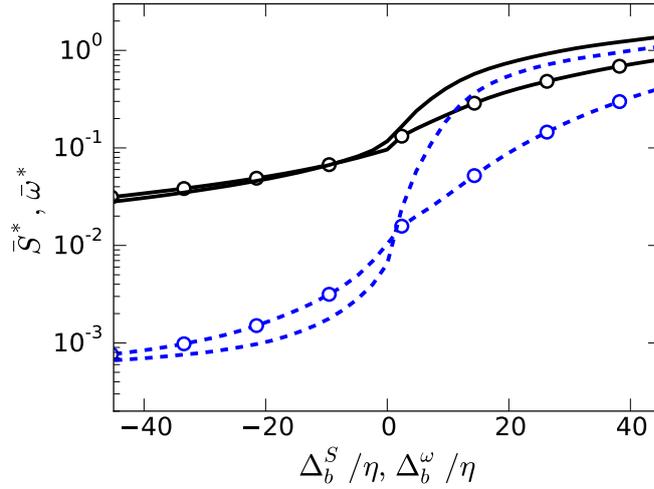


Figure 3.24: Conditional mean profiles of the vorticity and rate of strain, as functions of the distance Δ_b^ω to the $\omega_0^* = 0.01$ interface, or of the distance Δ_b^S to $S_0^* = 0.1$. $\delta_{99}^+ = 1900$. —, \bar{S} ; ---, $\bar{\omega}$. Lines without symbols are with respect to Δ_b^ω . Those with symbols are with respect to Δ_b^S .

The interesting question to be answered is whether the sharp vorticity gradient across the vorticity interface is a statistical artefact of the thresholding procedure, or a true physical effect. The former is a possibility, because enstrophy is fixed at the interface while moving slightly away from the geometrically complex isosurface could sample flow regions that are unrelated to it and representative of the bulk of the turbulent and irrotational regions. In a related example, Chauhan *et al.* (2014) find a sharp velocity jump across an interface defined in terms of the velocity magnitude, raising similar questions. In both cases, the sharp jump at the interface is what makes the criterion useful, and the reason to believe that an interface can be defined at all (Corrsin & Kistler, 1955). The lines without symbols in figure 3.24 refer to the vorticity interface. Within $\Delta_b = O(10\eta)$ of $\Delta_b^\omega = 0$ the conditional vorticity (dashed) drops by three orders of magnitude, and a similarly sharp gradient is seen for the conditional strain (solid).

The behaviour is different for the strain interface, represented by the lines with circles in figure 3.24. Both the vorticity and the strain cross relatively smoothly the level $\Delta_b^S = 0$. The difference between the two behaviours strongly suggest that while a sharp vorticity jump is a dynamically significant feature separating distinct regions of the flow, that of the strain is not. By inference, it is also probably true that the sharp velocity jumps in Chauhan *et al.* (2014) are true features of the flow, although it is unclear whether they are related to the enstrophy interfaces discussed here.

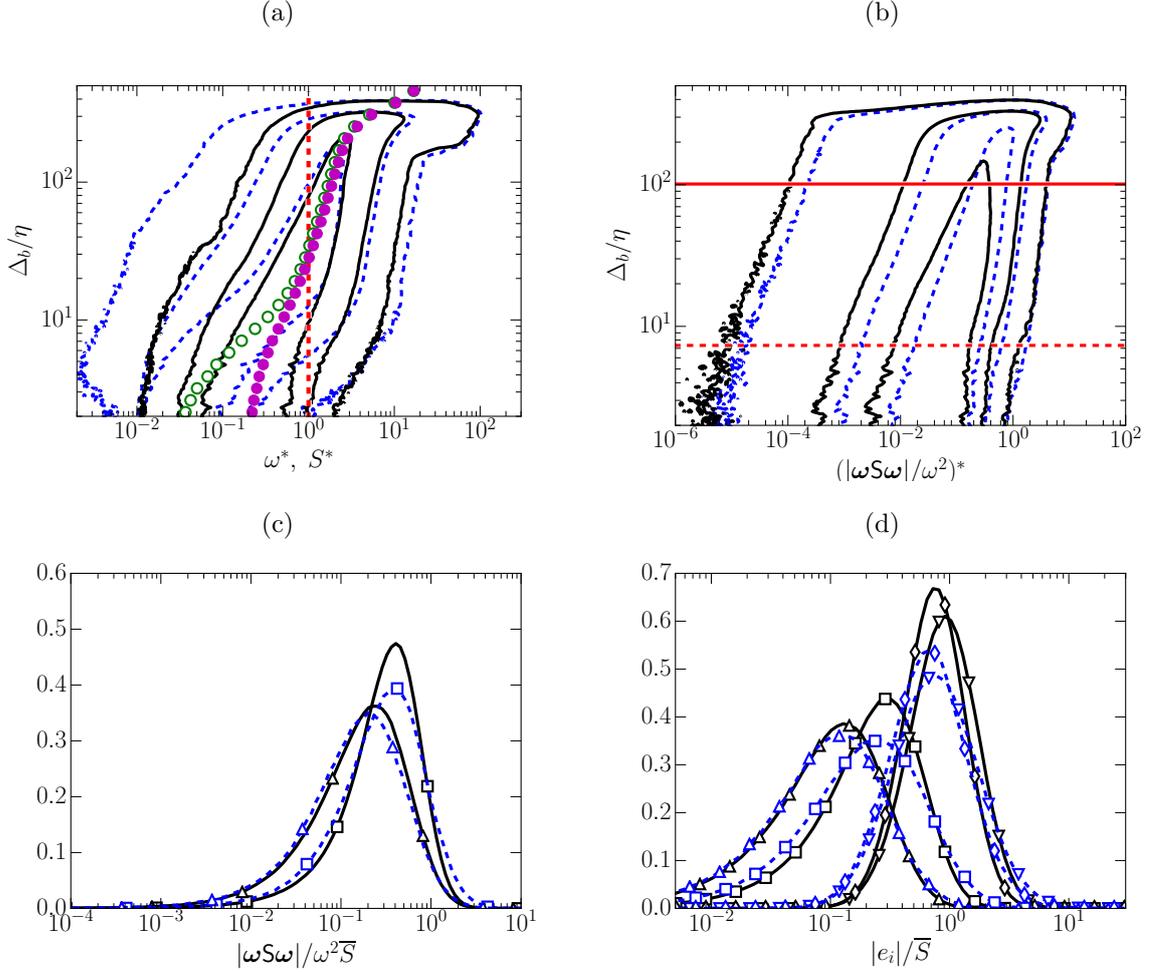


Figure 3.25: (a) Premultiplied joint PDFs at $\delta_{99}^+ = 1900$, with respect to the interface $\omega_0^* = 0.01$: --- (blue), $\omega F_{\omega, \Delta_b}$; — (black) $S F_{S, \Delta_b}$. The two symbol lines are: \circ , $\bar{\omega}^*$; \bullet , \bar{S}^* . The vertical line is $\omega^* = S^* = 1$. (b) Premultiplied joint PDF of Δ_b and: --- (blue), vortex stretching; — (black), vortex compression. The two horizontal lines: —, $\Delta_b = 100\eta$; ---, $\Delta_b = 7\eta$ are used in (c,d). Contours in (a,b) contain 50%, 90%, and 99% of points, respectively. (c) Premultiplied PDF of: Δ , normalised vortex stretching; \square , compression. (d) Premultiplied PDFs of the absolute values of the normalised individual eigenvalues of the rate-of-strain tensor: \diamond , largest; ∇ , smallest; Δ , positive intermediate; \square , negative intermediate. The abscissae in (c,d) are normalised with \bar{S} . In both cases, the PDFs are compiled at: ---, $\Delta_b = 7\eta$; —, $\Delta_b = 100\eta$.

As a conclusion from the previous results we are now ready to define a ‘natural’ interface as an enstrophy isosurface below the topological transition, such as $\omega_0^* \approx 0.01$. This threshold is somewhat lower than most of those compiled in table 3.2, with the result that the turbulent region contains part of the buffer and viscous superlayers defined by van Reeuwijk & Holzner (2014).

The structure of the flow with respect to this interface is displayed in figure 3.25. The joint PDFs of the vorticity and rate of strain with Δ_b are shown in figure 3.25(a). The

sharp decay of the vorticity below $\bar{\omega}^* = 1$ is clearly visible, spanning a thickness of about 20η . The strain decays slowly as it gets closer to the interface and far from the wall, but shows no especial behaviour within the buffer region. The vorticity in the buffer layer, even if subject to viscous effects, lives in a straining environment that is essentially the same as in the core turbulent flow.

This is seen more clearly in figure 3.25(b) which shows the PDF of the vortex stretching component of the strain, $\boldsymbol{\omega}\mathbf{S}\boldsymbol{\omega}/\omega^2$. The positive (stretching) and negative (compression) PDFs are plotted separately to allow a logarithmic representation. Both decays slowly and apparently self-similarly as they approach the interface, but do not change appreciably as they enter the buffer layer. Figure 3.25(c) shows the same result more quantitatively in the form of one-dimensional PDFs of the stretching term at two distances from the interface, one within the buffer region and another one in the core of the flow. The stretching has been normalised with the corresponding \bar{S} , which acts as a scaling parameter and absorbs most of the differences between the two levels. An even more detailed comparison is figure 3.25(d), which shows the PDFs of the individual eigenvalues of the rate of strain tensor. The PDFs at the two distances also collapse well. In both cases, the implication is that the straining environment within the buffer layer is essentially the same as in the core of the flow. Enstrophy is viscously diffused into the free stream, but it is stretched into core turbulence as it does so, in agreement with the model proposed in Townsend (1976). Note that the thickness of the region in which this takes place ($\Delta_b = 10 - 20\eta$) is comparable to the diameter of the individual Kolmogorov vorticity worms (Jiménez *et al.*, 1993).

3.7 Conclusions.

The T/NT interface of a turbulent boundary layer has been studied in the range of Reynolds numbers $\delta_{99}^+ = 1000 - 2000$. The emphasis is on the statistical description of the relatively large-scale interactions between turbulent and non-turbulent fluid in the fractal intermittent zone, rather than on the details of the small-scale structures where the interface can be considered smooth. This requires the introduction of a new definition of (ball) distance between a point and a general surface, which is compared with the more usual wall-normal (vertical) distance from the top of the interface. It is shown that the correct scale for the interface vorticity is the root-mean-squared magnitude of the enstrophy fluctuations at the edge of the boundary layer, $u_\tau^2(\delta_{99}^+)^{-1/2}/\nu$, rather than

wall units. The properties of an enstrophy isosurfaces scaled in this way, ω_0^* , are relatively independent of the Reynolds number.

The geometric complexity of the interface is explored as a function of the vorticity threshold. The ball and vertical distance generate different interfaces. Only the former captures the increase of interface complexity between $\omega_0^* \approx 0.1 - 2$, as the interface gradually moves from the free stream into the core turbulence. This transition manifests itself as an increase of the fractal dimension and of the topological genus of the interface, and reflects the appearance of a large number of surface corrugations, handles and pockets. The vertical distance is relatively insensitive to such features, and misses the transition. It is argued that this limitation of the vertical distance is responsible for some of the previously reported properties of the T/NT interface, such as the proposed layer of localised high vorticity at the edge of the turbulent region. That feature disappears with the new definition of distance, and can be reproduced with the old one.

What remains is a narrow layer of thickness $O(10\eta)$ in which the enstrophy decays from its core value, $\omega^* \approx 1$, to that of the free stream. To ascertain whether this sharp transition is a statistical artefact of the thresholding procedure or a true physical feature, an interface based on thresholding the norm of the rate-of-strain tensor is also studied. It is found that, whereas both the enstrophy and the strain change sharply across the vorticity interface, neither of them does so across the strain interface. It is concluded that enstrophy thresholding represents a physical feature, while thresholding the strain does not. The conditionally averaged properties in the neighbourhood of an enstrophy interface with a threshold ($\omega_0^* = 0.01$) below the topological transition is studied in some detail. It is found that even within the $O(10\eta)$ layer in which the vorticity decays sharply, the properly scaled straining structure of the flow is essentially identical to core turbulence. This fractal ‘buffer’ layer, even if controlled by viscous diffusion, retains most of the structure of the interior of the flow. There is no evidence of an intermediate length scale between η and δ_{99} , such as the Taylor microscale conjectured by Hunt & Durbin (1999) and found by Gampert *et al.* (2013). However, the range of Reynolds numbers in our simulation is not wide enough to completely discard that possibility.

The analysis of the conditional vorticity distribution with respect to the two definitions of distance throws some light on the controversy between entrainment by engulfing or by nibbling. The rate at which vorticity diffuses into irrotational pockets within the

turbulent region is independent of their position within the layer, but pockets closer to the wall tend to be smaller than those closer to the edge of the layer, presumably because they are broken down as they are entrained. This should enhance entrainment. The size of the entrained pockets scales in viscous units, but they are found at depths that scale with the boundary layer thickness.

The effect of increased entrainment

4.1 Introduction

Turbulent boundary layers over rough walls have an undeniable technological interest. Many engineering flows, and the majority of geophysical flows, happen over rough surfaces. It has been studied since the beginning of research in turbulence (Colebrook & White, 1937; Nikuradse, 1933), and it is still a subject of intense research. This chapter is focused on a particular topic related with wall-bounded flows over rough walls, the effect that roughness has on the structure of the turbulent motion. This problem has been reviewed by Raupach *et al.* (1991), Jiménez (2004), and Flack & Schultz (2014).

Turbulent boundary layers over smooth walls can be characterized with a limited set of length scales. Above 100 wall units from the wall, the majority of the properties of the turbulent motion can be described only with the distance to the wall y , and the boundary layer thickness δ_{99} . The analysis in Chapter 3 shows that the influence of the presence of the irrotational free stream is contained within a thin interface layer, while the rest of the turbulent flow remains apparently not affected. This does not mean that the geometrical characteristics of the intermittent region are simple or not relevant. The turbulent/non-turbulent interface (T/NT) interface is a fractal-like surface that contains all the possible length-scales from the Kolmogorov microscale, where turbulence is smooth, to the boundary layer thickness. But none of those intermediate scales seems to be particularly relevant.

The presence of a rough wall may introduce other length scales, like the average height of the roughness elements k , or other features of the roughness pattern L_i . One particularly

relevant simplification of the problem is to assume that $\delta_{99} \gg (\nu/u_\tau, k, L_i)$, meaning that there is sufficient separation between the boundary layer thickness and the wall-related scales for a log-layer to appear (Raupach *et al.*, 1991). This simplification is one of the fundamental assumptions of this chapter.

Clauser (1954) showed that the primary effect of surface roughness is to cause a downward shift of the mean streamwise velocity profile. Equation 1.9 becomes

$$\langle u \rangle^+ = \kappa^{-1} \log y^+ + C_0 - \Delta U^+(k^+, L_i^+), \quad (4.1)$$

where ΔU^+ is the roughness function, equal to zero for a smooth wall, y is the wall-normal coordinate, and κ is the Kármán constant. There are two important limits in the roughness function depending on k^+ . When the height of the roughness elements is small, and for some particular roughness geometries, $\Delta U = 0$. This particular case is called hydrodynamically smooth roughness, and the flow around the roughness elements is dominated by viscous effects. When $k^+ \rightarrow \infty$, the skin friction is independent of the Reynolds number $\delta u_\tau/\nu$, where δ can be the boundary layer thickness δ_{99} in boundary layers, or the semi-height h in channels. When that condition is fulfilled, additional shear stress is dominated by the form drag on the roughness elements, and the roughness function reaches an asymptote. This is the fully rough regime, that introduces another simplification to the list of relevant lengths. According to Flack & Schultz (2014), there is a mapping between the roughness function ΔU^+ and the equivalent sand roughness k_s^+ , defined as

$$\langle u \rangle^+ = \kappa^{-1} \log \frac{y^+}{k_s^+} + C_1, \quad (4.2)$$

regardless of the geometrical details of the roughness pattern. Note that equation (4.2) is a rearrangement of equation (4.1). This suggests that the set of parameters L_i are not relevant in the limit of large k_s^+ . This reduces the previous set of lengths in the case of the fully rough regime to Δ , ν/u_τ , and k_s . The upper limit of the hydrodynamically smooth regime is $k_s^+ \sim 5$, while the lower limit of the fully rough regime is $k_s^+ \sim 70$.

Between those two limits there is the transitionally rough regime, where the Reynolds number of the roughness elements k^+ is moderate, and the flow around them has no particularly dominating term. In consequence, the macroscopic effect of a particular roughness geometry, the additional friction, is hard to predict. In this regime, the presence of the roughness may be contained within a region with a thickness of the order of k , called roughness sublayer, while the rest of the flow may not be sensitive to the details of the

geometrical characteristics of the rough wall. There is one particularly relevant hypothesis about the characteristics of the turbulent motion above the roughness sublayer (in case it exists as described) due to Townsend (1976). Outside the roughness sublayer, the turbulent motions in a boundary layer at high Reynolds number is independent of the wall roughness and the viscosity, except for the role of the wall in setting the velocity scale u_τ . Considering the previous list of the possibly relevant length scales, ν/u_τ , k , and L_i are important in the roughness sublayer, while the rest of the boundary layer can be described only with δ and ν/u_τ . This means that under the described conditions, and if the wall similarity hypothesis is valid, u_τ is a similarity parameter for roughness.

This hypothesis is supported by several facts. The roughness sublayer occurs below the lower bound of the logarithmic layer at $y^+ \simeq 100$, and may not affect its dynamics that are characterized by the friction velocity and the distance to the wall. The Kármán constant κ is not affected by the presence of roughness in almost every experiment and simulation of boundary layers over rough walls, which is an indirect but strong evidence that the turbulent structure in the logarithmic layer is roughly independent of the nature of the wall. This independence was studied by Flores & Jiménez (2006), and Mizuno & Jiménez (2013), that confirmed that the dynamics of the near-wall region can be modified without changing the properties of the logarithmic layer.

One challenge to the wall similarity hypothesis is the existence of *d-roughness* (Perry *et al.*, 1969), where the height of the roughness layer observed in the average streamwise velocity profile is proportional to the boundary layer thickness. It has been observed only in the limit of what could be considered large k/δ_{99} , for a specific kind of roughness, transverse square rods, and with a given separation between rods (Bandyopadhyay, 1987). This geometry has the particular property of causing one of the highest amounts of additional drag for a given k . *d-roughness* could be caused by the strong interaction between the flow in the grooves that separate the rods, and the large-scale coherent structures of velocity of the flow. This mechanism was already sketched by Townsend (1976). However, the experiments where this behavior has been observed have not been successfully reproduced.

There is another case where the wall similarity hypothesis has been challenged. Krogstad & Antonia (1994) set up a boundary layer at a reasonably high Reynolds number, and a roughness pattern made with transverse square rods with $\delta_{99} \gg k$. This particular

roughness geometry is “strong”, in the sense that for a given k , the ratio k_s/k is the largest known. They found that the one-point correlation times for all velocity components are almost twice shorter for rough than for smooth boundary layers below $y/\delta_{99} = 0.5$, where y is the wall-normal coordinate. While this effect was not strong enough to be visible in the scaling of the average profile like in *d-roughness*, the structure of the velocity field seemed to be affected in all the thickness of the boundary layer. Interestingly, other roughness patterns had a weaker effect, and one-point correlation times were similar to the smooth wall’s. The same case was explored by Lee *et al.* (2011); Lee & Sung (2007); Schultz & Flack (2005); Volino *et al.* (2011); Wu & Christensen (2007), and a weaker but measurable change in the streamwise section of the correlations was found. More recently, Krogstad & Efros (2012) concluded that the strong differences in the temporal correlations of the Krogstad & Antonia (1994) experiment may be due to the small Reynolds number, and one of the conditions of applicability of the wall similarity hypothesis was therefore not fulfilled. The general consensus is that the effect of the roughness exists and it is measurable, but that it is not strong enough to discard the validity of the wall similarity hypothesis. On the other hand, there is no explanation to what causes the differences in the one- and two-point statistics. Similar experiments have been carried out in internal flows by Flores & Jiménez (2006); Krogstad *et al.* (2005), where the influence of roughness is even smaller, suggesting that roughness may affect boundary layers differently than channels and pipes.

The main difference between internal and external wall-bounded flows is entrainment, the ingestion of irrotational fluid during the streamwise evolution of the turbulent flow. Entrainment is related to the streamwise growth of the thickness of the boundary layer by the Kármán equation

$$\frac{d\theta}{dx} = \frac{u_\tau^2}{U_\infty^2}, \quad (4.3)$$

where x is the streamwise coordinate, θ is the momentum thickness, U_∞ is the velocity of the free stream, and u_τ is the friction velocity. It is also known that the ingestion of mass involves a thin interface layer that separates the turbulent flow from the non-turbulent free stream, the turbulent/non-turbulent interface, that was thoroughly described in Chapter 3. The goal is to test if the wall similarity hypothesis is valid in the T/NT interface too, which would suggest that it is valid across the whole thickness of the turbulent boundary layer, and whether the differences in the one- and two-point statistics caused by the presence of roughness are due to a different arrangement of the turbulent

eddies in the intermittent region, or to the fact that the structure of turbulent motion is affected by some additional friction.

The strategy is to simulate two very similar turbulent boundary layers without external pressure gradients and with the usual no-slip boundary condition at the wall. The first, also used in Sillero *et al.* (2013) and Sillero *et al.* (2014), can be considered a *canonical* simulation, while the second adds a forcing that increases the effective friction coefficient. The aim of the forcing is not to mimic actual roughness, but to increase the rate of entrainment without introducing any secondary effects that may cause a geometrical footprint on the flow. The comparison of these two cases should provide some quantitative evidence of the validity of the wall similarity hypothesis. At the same time, if the volumetric force affects the characteristics of the boundary layer similarly to the wealth of experiments with actual rough walls, the results will argue in favor of the validity of Townsend’s similarity hypothesis.

This chapter is structured as follows. The next section contains a description of the numerical experiment designed to test the wall similarity hypothesis. Section 4.3 contains the one-point statistics of the experiment, that are compared with other relevant data sets. The following §4.4 discusses in detail the two-point three-dimensional spatial correlations of the velocity components, and pressure. Section 4.5 analyzes in detail the properties of the intermittent region, and discusses the validity of the wall similarity hypothesis in the T/NT interface. Finally, §4.6 summarizes the most relevant results, and concludes.

4.2 The numerical experiment

A boundary layer is simulated in a parallelepiped over a flat plate with periodic spanwise boundary conditions and non-periodic streamwise and wall-normal directions. The velocity components in the streamwise (x), wall-normal (y), and spanwise (z) directions are u , v , and w respectively. The velocity vector is \mathbf{u} , brackets $\langle \cdot \rangle$ denote the average profile along the wall-normal direction, and primes denote root-mean-squared values. The $+$ superscript is used to express wall units, where u_τ is used as the unit for velocity, and ν/u_τ as the unit for length, where ν is the kinematic viscosity. Another important length is the boundary layer thickness δ , that corresponds to δ_{99} in the case of boundary layers, and the semi-height h in the case of channels. The Reynolds number used to compare the two type of flows is $\delta_{99}^+ = \delta_{99}u_\tau/\nu$ and $h^+ = hu_\tau/\nu$ in boundary layers and channels

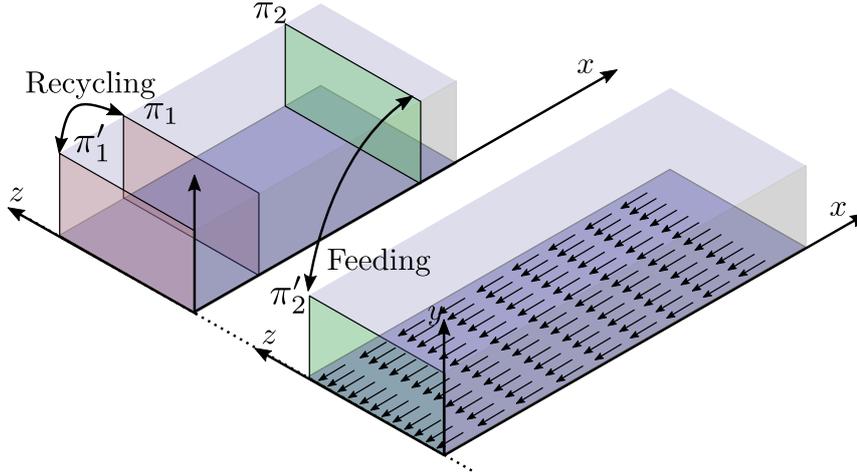


Figure 4.1: Sketch of the twin box configuration. The auxiliary boundary layer, shorter and with coarser resolution, is sufficiently long to recover from the incorrect initial boundary conditions generated by the recycling of the plane π_1 to the plane π'_1 . The plane π_2 is copied as the inflow boundary conditions of the main simulation, that introduces a volumetric force, represented as the arrows pointing towards the inflow, close to the wall.

Case	δ^+	δ^+ profiles	Label	Symbol
present	460-1675	1500	BL_f	— (red)
Sillero <i>et al.</i> (2013)	980-2025	1500	BL_s	- - (blue)
Del Alamo <i>et al.</i> (2004)	934	934	C_1	□ (magenta)
Hoyas & Jiménez (2006)	2003	2003	C_2	▽ (green)

Table 4.1: Summary of the important parameters of the four simulations used more often in this study.

respectively.

The numerical code used in this simulation is described in Simens *et al.* (2009), and in Chapter 2. The configuration, practically identical to Sillero *et al.* (2013), is sketched in figure 4.1. It exploits the fact that the code is able to run two synchronized simulations in a tandem configuration, where an auxiliary low-resolution simulation provides the inflow boundary conditions to a second domain by copying the plane π_2 to π'_2 each time step. The main different between the present simulation and Sillero *et al.* (2013) is the addition of a volumetric force in the near-wall region of the second domain, expressed in figure 4.1 as an array of arrows at the wall, and the intensity of the suction at the top face of the domain to preserve the zero-pressure-gradient condition. The twin domain configuration is of particular importance in this case, since the forcing is not allowed interact with the recycling process that changes the plane π_1 to build π'_1 , creating the turbulent inflow for the auxiliary domain.

The simulation has been run for approximately $14u_\tau/\delta_{99}$, where both u_τ and δ_{99} are evaluated at the end of the computational domain. This unit of time is equivalent to the characteristic turnover time of the largest scale in a boundary layer, and it is commonly called eddy turnover (Simens *et al.*, 2009). In this case, an eddy turnover is roughly equivalent to the time a virtual particle in the outer free stream needs to cross through the complete domain with a velocity of $u = U_\infty$. This simulation consumed about 40 million CPU hours in the Jugene BG/P supercomputer at the Juelich Supercomputing Center, and produced around 70 TiB of raw data.

The force is introduced as an additional term in the streamwise component of the momentum equation

$$\frac{\partial u}{\partial t} + \mathbf{u} \cdot \nabla u = -\frac{\partial p}{\partial x} + \nu \nabla^2 \mathbf{u} - u_0 g / T, \quad (4.4)$$

that also contributes to the total shear stress

$$\tau_w = \nu \frac{\partial \langle u \rangle}{\partial y} + \frac{1}{T} \int_0^\delta u_0 g \, dy, \quad (4.5)$$

and the friction velocity $u_\tau = \sqrt{\tau_w}$. This near-wall forcing has three components. The velocity u_0 is the streamwise component of the velocity averaged in the spanwise direction, and it is a function of the streamwise, and wall-normal directions. The mean profile $\langle u \rangle$ is the temporal average of u_0 . The weight g is a function of the distance to the wall

$$g = \frac{1}{2} \left(-\tanh \left(\frac{y - y_0}{y_1} \right) + 1 \right), \quad (4.6)$$

where the additional parameters y_0 , and y_1 control the thickness of the forced region, and the smoothness of the transition to the outer, non forced flow. Since the goal is to restrict the forcing to a layer very close to the wall, those parameters are set to $y_0 = 20^+$, and $y_1 = 5^+$. The third component T has units of time, and it is adjusted to keep the overall effect of the forcing equal to the average wall shear stress of an unforced boundary layer τ_w^0

$$T = \frac{1}{\tau_w^0} \int_0^\delta u_0 g \, dy \quad (4.7)$$

for each location in the streamwise direction. The determination of T is a calibration process, that requires u_0 , and τ_w^0 to be known previously. Fortunately, two simulations are available within this range of Reynolds numbers, Jiménez *et al.* (2010), and Sillero *et al.* (2013). The forcing has the projected effect of doubling the wall shear stress, but since its presence changes the flow, the practical effect is $\tau_w \simeq 1.75\tau_w^0$. This forcing does

not alter the no-slip boundary condition at $y = 0$, and does not require the definition of a virtual wall location.

A similar strategy is often used in meteorology to model the effect of an arbitrary surface roughness like the vegetal canopy (Belcher *et al.*, 2003), where the forcing is proportional to the square of the streamwise velocity, and the magnitude is controlled by a constant with units of length (Wood & Mason, 1993). This allows us to emphasize that the main goal of the present forcing scheme is not to model any particular kind of roughness, but to increase entrainment avoiding any kind of geometrical feature. This forcing fulfills several design goals. It is computationally cheap, and simple to implement in the numerical simulation, since u_0 corresponds to the zero mode of the Fourier decomposition in the spanwise direction. In consequence, this simulation has practically the same computational cost as its unforced counterpart, and the performance and scalability figures given in Chapter 2 are valid as well. The fact that the forcing is proportional to the averaged instantaneous velocity also means that detachment is very unlikely. If a significant portion of the boundary layer detaches, $u_0 < 0$ at that section, and the forcing will favor the local reattachment regardless of the value of T . The eventual detachment is also taken into account in the various models of the vegetal canopy, and the quadratic dependence with the streamwise component of velocity is often introduced as $u|u|$.

If this forcing is compared with actual roughness, the following parameters are obtained. The height of the roughness elements k can be approximated as the wall-normal coordinate where the weight $g = 0.5$, therefore $k = 25^+$. This value is not arbitrary, and it is discussed in §4.3. The ratio between the boundary layer thickness and the roughness height δ_{99}/k is comprised between 50 and 70, fulfilling the condition for application of the wall similarity hypothesis. The equivalent sand roughness, defined by equation (4.2), is $k_s \simeq 70$, at the limit between the transitionally and the fully rough regimes. Another interesting quantity is k_s/k , which was previously described as the strength of a given roughness geometry. The present forcing achieves $k_s/k \simeq 2.8$. Packed graded sand has a $k_s/k \simeq 1$. This implementation of forcing has a value of k_s/k similar to the 3-D staggered cubes of Volino *et al.* (2011). The transverse square rods, a particularly strong kind of roughness, achieve values of k_s/k between 8 and 14 (Efros & Krogstad, 2011; Krogstad & Antonia, 1994; Volino *et al.*, 2011).

4.3 Average properties.

Streamwise evolution of a boundary layer after a step change in the friction coefficient.

The forcing described in equations (4.4)-(4.7) is introduced only in the second domain, meaning and the flow overcomes a sudden change in the friction coefficient. The transient caused in a turbulent boundary layer by the sudden change of surface roughness is an interesting phenomenon, since it is relatively common in the case of the atmospheric boundary layer, and has been studied extensively since Colebrook (1939). The main tool of this study is the comparison of the present simulation with Sillero *et al.* (2013) and Sillero *et al.* (2014), therefore it is important that the forced boundary layer has reached its equilibrium configuration in all its thickness.

Sillero *et al.* (2013) analyzed the recovery of a turbulent boundary layer from incorrect inflow boundary conditions, and they used the eddy turnover distance \hat{x}

$$\hat{x} = \int_0^x \frac{u_\tau}{\delta_{99} U_\infty} dx, \quad (4.8)$$

that measures how far eddies are advected during a large-scale turnover time, to study the streamwise evolution of quantities such as the friction coefficient, the shape factor H or the wake intensity Π . They also mention that the larger outer scales may take a long time to converge, and that the recovery distance depends on Reynolds number. One important feature of the turnover distance is that it is able to collapse the evolution of the previously mentioned parameters regardless of the initial Reynolds number. If it is compared to other units of length, like the initial boundary layer thickness, the eddy turnover distance has a clearer physical interpretation. Assume an imaginary particle originating from the wall $y = 0$ with a wall-normal velocity of $v = u_\tau$ and a streamwise velocity of $u = U_\infty$. When the particle has reached the outer edge of the boundary layer $y = \delta_{99}$, it has been advected one eddy turnover length $\hat{x} = 1$.

The next step to analyze the streamwise evolution of the boundary layer is to choose a meaningful quantity. Efros & Krogstad (2011) studied the transient caused by the presence of cross-stream square rods in a comparable yet higher Reynolds number, and proposed the Clauser wake parameter G

$$G = \int_0^\infty \left(\frac{U_\infty - \langle u \rangle}{u_\tau} \right)^2 dy \Big/ \int_0^\infty \frac{U_\infty - \langle u \rangle}{u_\tau} dy \quad (4.9)$$

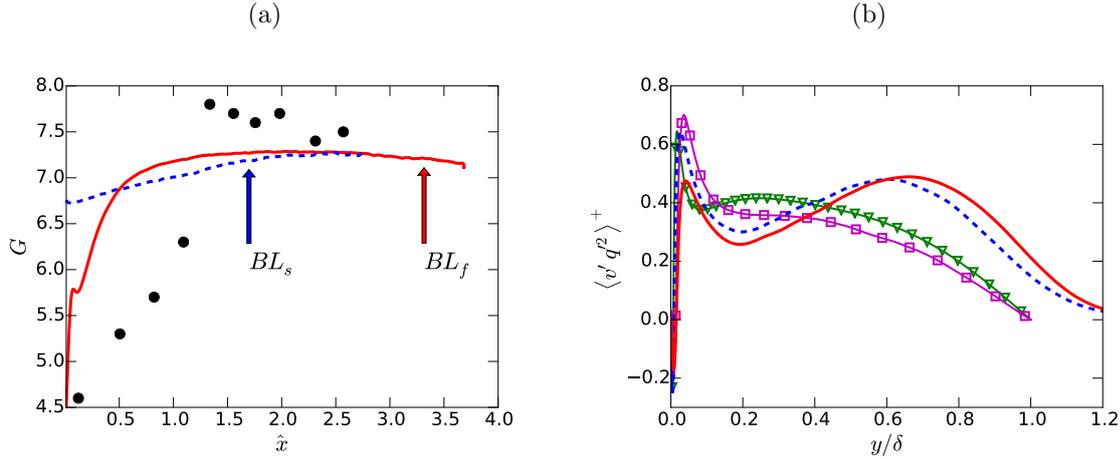


Figure 4.2: (a) Spatial evolution of the Clauser wake parameter G . The two arrows mark the stages of BL_s and BL_f where $\delta_{99}^+ = h^+ = 1500$. (b) Average profile of the wall-normal flux of turbulent kinetic energy $\langle v'q'^2 \rangle^+$ for the three reference cases in table 4.1.

Following Coles (1956), the shape of the wake should be more or less constant in equilibrium flows. It is plausible that once the Clauser wake parameter is no longer changing, the whole extent of the flow is in equilibrium. The turnover length that seems to collapse reasonably well the evolution of a variety of properties of the mean profile for a wide range of initial Reynolds numbers.

The streamwise evolution of the Clauser wake parameter G for three different boundary layers is presented in figure 4.2(a). The eddy-turnover distance in Efros & Krogstad (2011) has been estimated from the data provided. The smooth-wall boundary layer used as the baseline for the comparison can be considered converged, but the wake parameter grows slightly from 6.6 and 7.2. This suggests that the hypothesis that G is constant for boundary layers in equilibrium is only an approximation. In all the cases the final value for the wake parameter after $\hat{x} = 2.5$ is $G = (7 - 7.5)$ regardless of the fact that the wake of rough boundary layers is stronger. The effect of roughness and the forcing is similar, reducing G almost by a half. The rough and the forced boundary layers converge at different rates, but according to this criterion, both reach some level of convergence around $\hat{x} = 1$. At the point where BL_f is compared with BL_s (where $\delta^+ = 1500$ in both cases), the boundary layer has evolved for more than three eddy turnovers, and it is plausible that the comparison is consistent at that point.

Given the amount of data provided by the present DNS, it should be possible to characterize the effect of a sudden change in the friction coefficient with more detail. The effect has

been described by many experiments (Antonia & Luxton, 1971; Bradley, 1968; Cheng & Castro, 2002; Townsend, 1966), as an internal boundary layer that propagates within the bulk of the turbulent flow with a given wall-normal speed until a new equilibrium configuration is reached. It should be possible to design a tracking method for the internal boundary layer from its known characteristics. It is reasonable to think that the turbulent kinetic energy contained within the internal boundary layer is higher due to the additional friction. It is also known by the Kármán equation that the internal boundary layer has to grow faster than the original, outer one. In consequence, a suitable variable to track the internal boundary layer should be the wall-normal turbulent flux of streamwise velocity fluctuations $\langle v'q'^2 \rangle$, where $q'^2 = (u'^2 + v'^2 + w'^2)/2$, and it is presented in figure 4.2(b) for the cases listed in table 4.1. This quantity was used by Jiménez & Simens (2001) to characterize the independence of the small and the large scales in turbulent channels and boundary layers, and a similar quantity was used by Flores & Jiménez (2006) to show that the mentioned independence exists in channels regardless of the presence of roughness. A similar quantity is also available for experimental studies of boundary layers like Fernholz & Finley (1996).

The profile of $\langle v'q'^2 \rangle$ depends on the flow and the Reynolds number. In the case of turbulent channels at a low Reynolds number C_1 , the profile features only a maximum close to the wall at the same location as the peak of turbulent kinetic energy $y^+ = 15$. A second maximum emerges closer to the center of the channel of C_2 , at $h^+ \simeq 2000$ number, while the near-wall peak lowers its intensity. The profile of $\langle v'q'^2 \rangle$ for boundary layers features a similar near-wall peak, but a second and stronger maximum appears in both BL_f and BL_s closer to the edge of the boundary layer. According to the data of Fernholz & Finley (1996) shown in figure 1(b) of Jiménez & Simens (2001), the magnitude of the outer peak grows with the Reynolds number. The streamwise evolution of $\langle v'q'^2 \rangle$ in boundary layers, normalized with the value of the second maximum, is shown in figure 4.3.

The profile $\langle v'q'^2 \rangle$ of BL_s , in figure 4.3(a), evolves only slightly given the limited range of Reynolds numbers. The most remarkable characteristic is the widening of the separation between the inner and the outer peak, already described in Jiménez & Simens (2001). Note that, given the normalization used in figure 4.3, the value of the map at the second peak is 1. The same quantity in the case of BL_f , in figure 4.3(b). The evolution of the internal boundary layer is clearly visible. The additional drag creates a third maximum as

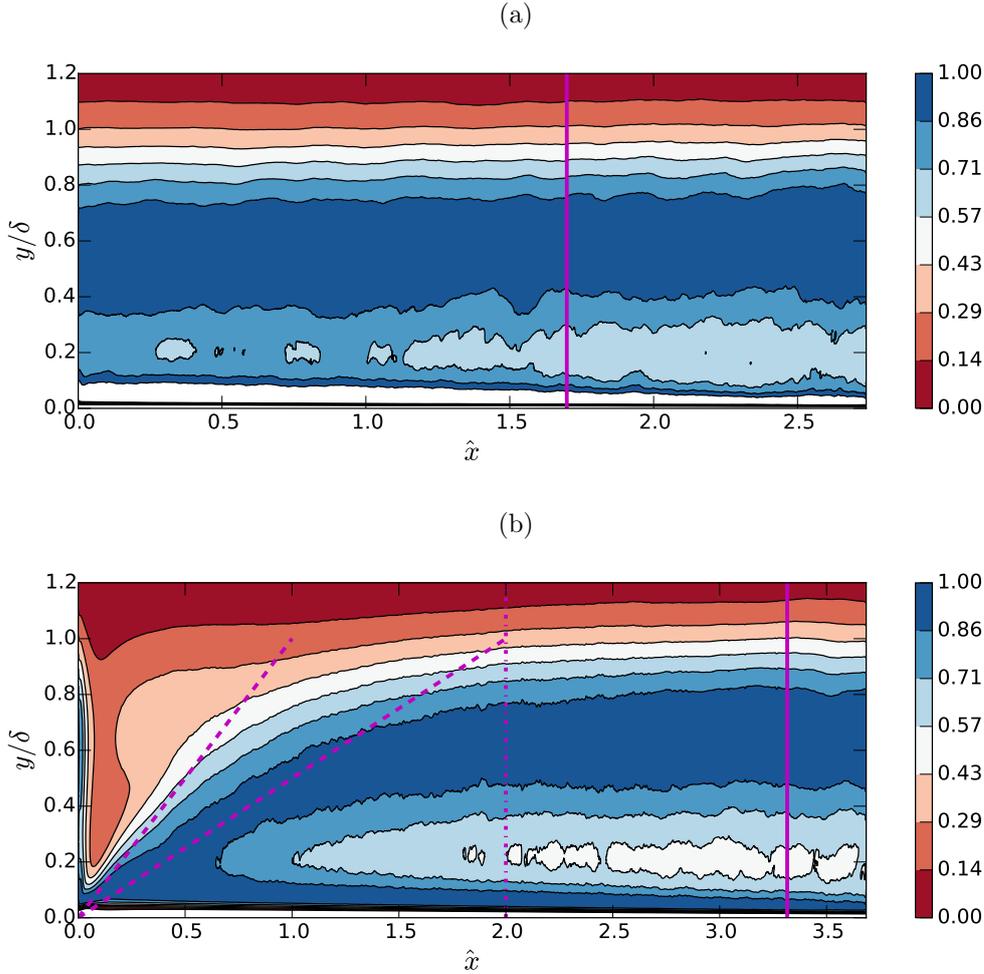


Figure 4.3: Spatial evolution of the wall-normal term of the turbulent kinetic energy velocity budget for (a) BL_f , and (b) BL_s , using the local turnover length as a distance unit. The two dashed lines in (a) correspond to a trajectory of an ideal particle traveling a wall-normal speed of u_τ (left), and $0.5u_\tau$ (right). The dash-dotted vertical line in (a) corresponds to $\hat{x} = 2$, the approximate end of the transient caused by the forcing. The solid vertical lines in (a), and (b) mark the streamwise location of $\delta_{99}^+ = 1500$, where the two boundary layers are compared.

soon as it is introduced, stronger than the one coming from upstream. The fact that only the internal boundary layer is visible is caused by the chosen normalization of $\langle v'q^2 \rangle$. The peak coming from the layer is still present, but it is hidden by the fact that the energy flux of the internal boundary layer is much stronger. The two peaks have merged at $\hat{x} = 1$, which is approximately the stage at which the wake is comparable according to the previous study of the G parameter.

One useful feature of the (\hat{x}, y) plane in figure 4.3 is that a theoretical particle with streamwise velocity $u = U_\infty$, and wall-normal velocity $v = cu_\tau$, where c is an arbitrary constant, follows a straight line. The left and right diagonal lines in figure 4.3(b), that

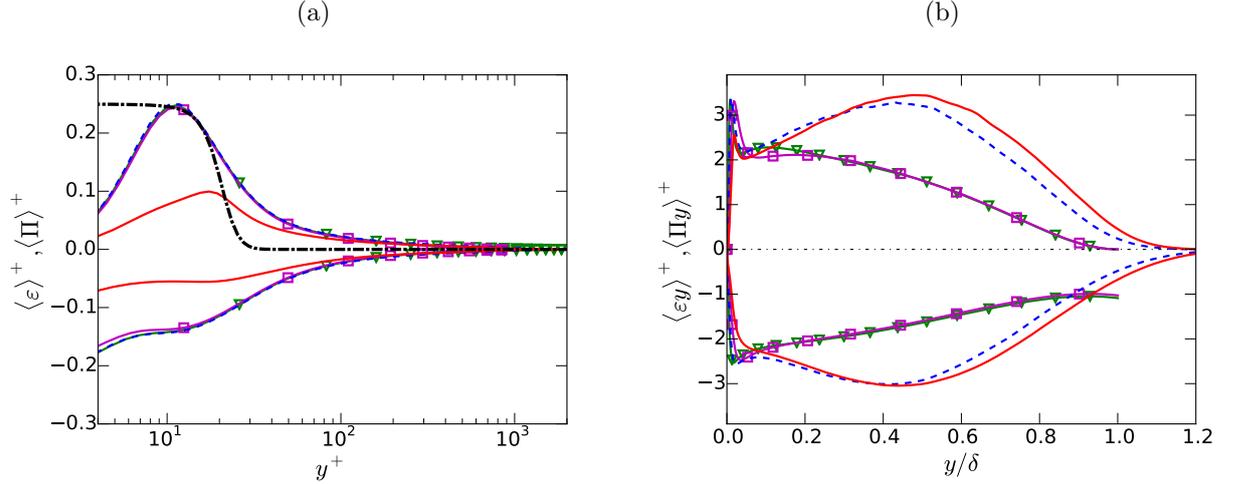


Figure 4.4: (a) Profiles of turbulent kinetic energy production and dissipation for the cases specified in table 4.1. The additional black dash-dotted line is the shape of the weighting function g normalized with the maximum value of Π . (b) Profiles of turbulent kinetic energy production and dissipation premultiplied by the distance to the wall y .

correspond to wall-normal velocities of u_τ and $u_\tau/2$ respectively, are particularly useful to quantify the growth of the internal boundary layer. Given that u_τ is precisely the velocity at which coherent structures are advected from and towards the wall (Lozano-Durán & Jiménez, 2014 *b*), it is reasonable to expect that the convergence after a change of the conditions at the wall is of the order of a few turnover length units. We can roughly estimate that after $\hat{x} = 2$, the majority of eddies that existed before the sudden change of the conditions at the wall have either decayed or have directly interacted with the forcing (or the roughness). The estimation of $\hat{x} = 2$ is purely qualitative, and the exact length most likely depends on the details of the sudden addition of drag.

Mean velocity and turbulence profiles.

We now examine the one-point statistics of the two boundary layers, BL_f and BL_s , and compare them with channels at similar Reynolds numbers C_1 and C_2 . Despite the outer zone of boundary layers and channels is not strictly comparable, channels are useful baseline to evaluate the effects of the additional entrainment, given that channels have no entrainment at all.

The profiles of production and dissipation of turbulent kinetic energy, $\langle \Pi \rangle$, and $\langle \varepsilon \rangle$, presented in figure 4.4(a). There is an important deficit of production and dissipation below $y^+ = 20$, precisely at the upper limit of the forcing $y = y_0$. Despite only one set of values

of y_0 and y_1 have been tested, this correlation between the point of maximum production and the limit of the forcing suggests that the that $y_0 + y_1$ is roughly comparable to k . The peak of production of turbulent kinetic energy at $y^+ = 15$, one of the important features of turbulent flows over smooth walls, is not present in BL_f . This means that the dynamics of the turbulent motion close to the wall are strongly affected by the forcing.

The same quantities in figure 4.4(a) are presented in figure 4.4(b) premultiplied by the distance to the wall y . This technique was already used in Jiménez *et al.* (2010) with the same purpose, to zoom into the outer region of the boundary layer. The profiles of production and dissipation show that the effect of the forcing is confined within a relatively thin region close to the wall. The shape of the profile is recovered roughly past $y_0 + y_1$, still within the buffer layer. The profiles of BL_s and BL_f are similar in the logarithmic layer. This was somewhat expected, given that the Kármán constant κ (see figure 4.5a), the ratio between production and dissipation, and one of the arguments that support the validity of the wall similarity hypothesis, is almost identical in both cases. Production and dissipation of BL_f are higher from the lower limit of the intermittent region (roughly at $y/\delta_{99} = 0.4$) to the edge of the boundary layer. The other terms in the turbulent kinetic energy equation that are important in the intermittent region (not shown), like pressure strain, turbulent transport, and turbulent convection, are also more intense in the forced case.

The average profile of the streamwise component of velocity (figure 4.5a) suffers the usual offset reported in practically all boundary layers over rough walls. The Kármán constant used to fit approximately the logarithmic profile is 0.41. When the mean profiles of BL_s and BL_f are compared in defect coordinates (not shown) BL_f presents a slightly stronger wake. The most perceivable difference in the streamwise velocity fluctuations (figure 4.5b) is the height of the peak, considerably weaker in BL_f . The forcing is acting on the streamwise component velocity, and it is proportional to it. It is to expect that the most visible effect happens in u . The other two components of the velocity vector do not show significant changes close to the wall. The spanwise component w' even shows its characteristic peak close to the wall, slightly damped and shifted at $y = 20^+$.

Apart from the near-wall peak of u' , the differences between BL_s and BL_f intensify above $y/\delta_{99} = 0.2$, the approximate upper limit of the logarithmic layer. Beyond that point, and particularly past $y/\delta_{99} = 0.6$, the differences between the two boundary layers are

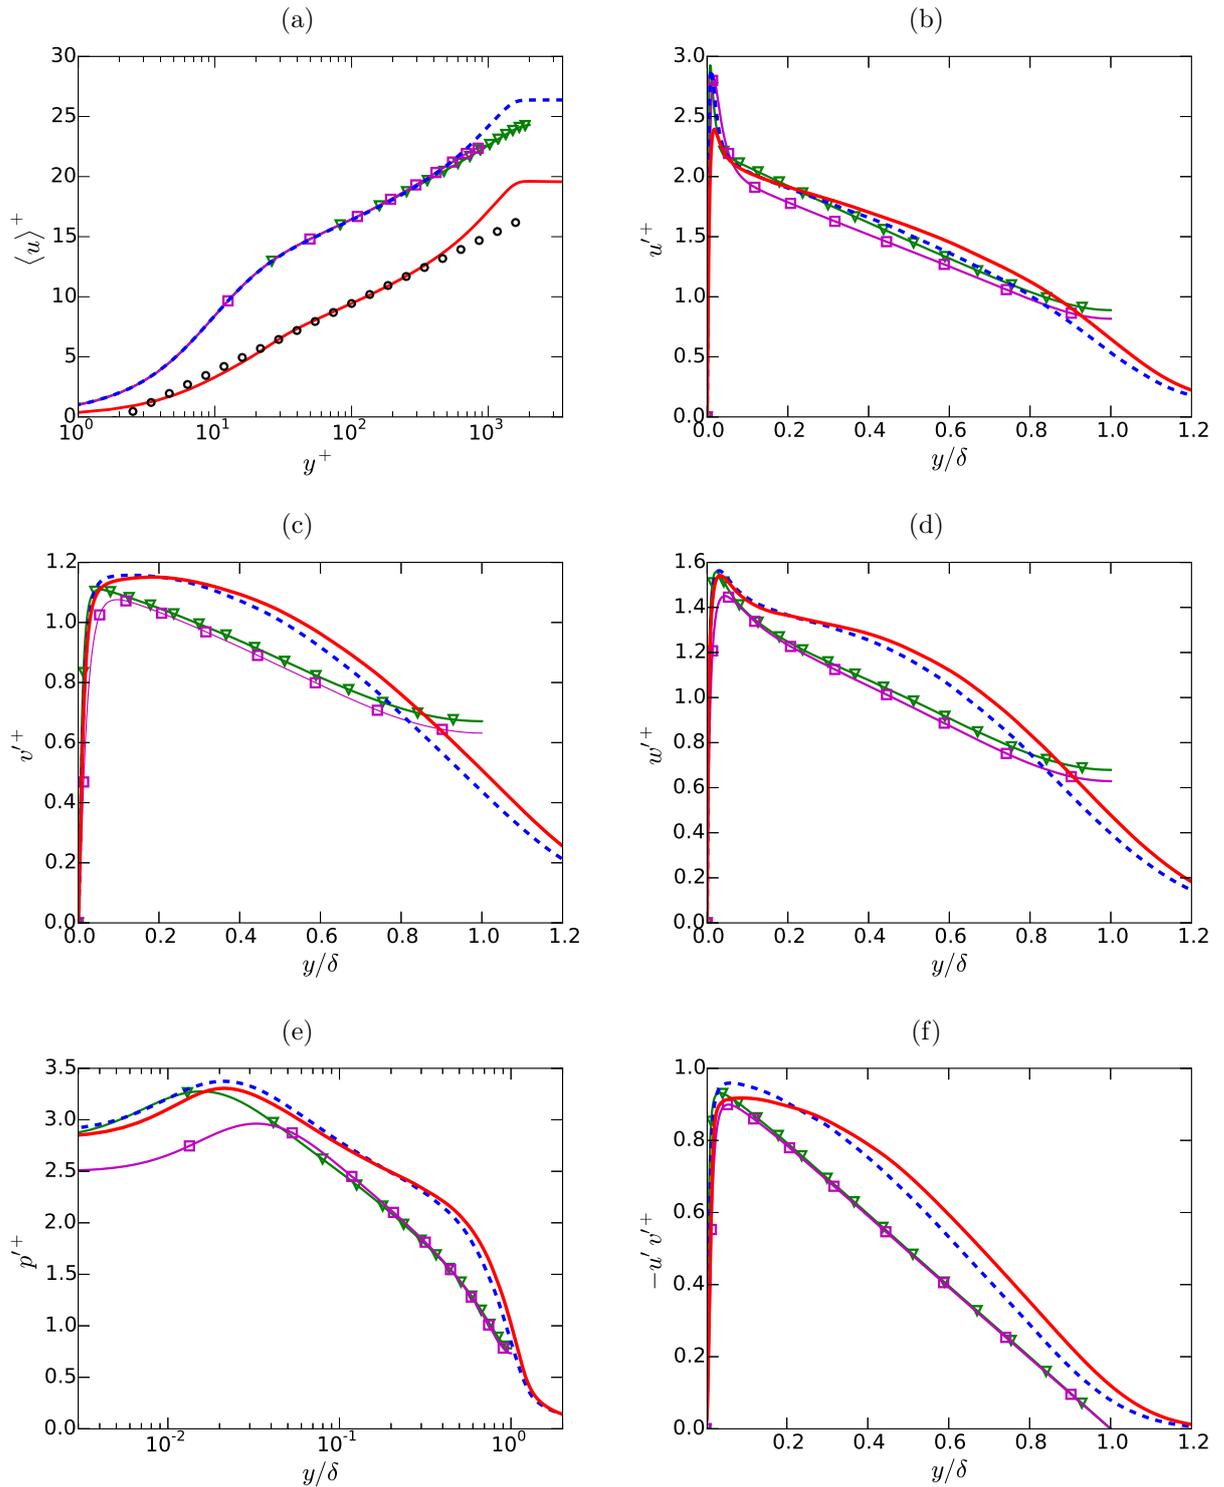


Figure 4.5: Comparison of the wall-normal profiles for the cases given in table 4.1. (a) Average streamwise velocity $\langle u \rangle^+$. The straight line of dots corresponds to a fit using equation (4.2). (b) Streamwise velocity fluctuations u'^+ . (c) Wall-normal velocity fluctuations v'^+ . (d) Spanwise velocity fluctuations w'^+ . (e) Pressure fluctuations p'^+ . (f) Reynolds shear stresses $-u'v'^+$

of the order of the deviation between BL_s and the two channels. This one of the main motivations of this work. The effect of the forcing is contained within the buffer layer, and the logarithmic layer seems to forget about its presence. Since those differences are within the intermittent region, entrainment is probably the cause. But we saw in Chapter 3 that the T/NT interface layer is also a very thin region that faces the irrotational free stream, and that the bulk of the turbulent flow is not noticeably affected by its presence. These two observations cannot be trivially connected. The T/NT interface is microscopic (happens at the Kolmogorov scale), while the differences in the mean profiles are macroscopic.

Pressure is a particularly interesting variable because it is likely to be relevant in the entrainment of non/turbulent fluid. According to §3, the turbulent/non-turbulent interface layer, the region where entrainment actually takes place, is characterized by strong vorticity gradients. In consequence, the non-turbulent flow carries almost no vorticity, but relatively intense pressure gradients caused by the relatively high level of strain. In addition, most data sets of boundary layers over rough walls at a moderate or high Reynolds number are obtained with experimental techniques, where the determination of the pressure field is very hard or not possible.

The pressure fluctuations profile is shown in 4.5(e), where a logarithmic axis is used to emphasize the existence of a logarithmic profile above $y^+ = 30$. A logarithmic profile is usually interpreted as the footprint of the additive contribution of elements that scale with the inverse of the distance to the wall. For instance, the logarithmic profile that appears in the velocity fluctuations profile is commonly interpreted as the aggregated contribution of self-similar eddies with the same characteristic velocity, a characteristic size of y , and a population proportional to $1/y$. This is the Townsend's attached eddy hypothesis. A similar argument can be applied to pressure in channels. Knowing that $\log y = \int 1/y dy$, the logarithmic profile could be caused for example by the addition of structures with equal intensity and size y , or by structures with intensity $1/y$ and constant size. Either way, the logarithmic profile involves a single length-scale, the distance to the wall. The turbulent pressure intensity profile in channels is well approximated by a logarithm between $y^+ = 50$ to $y/\delta_{99} = 0.5$. Above that height, the profile deviates from the trend in a more subtle way in boundary layers.

Jiménez *et al.* (2010) observed that boundary layers present a logarithmic pressure fluctuations profile with the same exponent as in channels, but consistently higher, and

concluded that the additional intensity is caused by the presence of intermittency and entrainment, and propagated down to the wall by the global character of pressure. Sillero *et al.* (2013) found that, with a Reynolds number $\delta_{99}^+ > 1000$, a hump appears between $y/\delta_{99} = (0.3 - 0.8)$. If the additive process suggested by the logarithmic profile exists, the hump indicates the presence of a second length scale. This length is most likely the boundary layer thickness, given that the deviation from the logarithmic trend is only clearly visible when the Reynolds number is large enough, and also given that the extent of the hump seems scales precisely with δ_{99} .

Comparing the pressure intensities of the two boundary layers, the fluctuations of the forced case close to the wall are lower, and more intense at the hump. This change of tendency occurs at the same distance from the wall for every velocity fluctuations profile, as well as pressure, approximately at the end of the logarithmic layer $y/\delta_{99} = 0.2$. The shape of the profiles, together with the fact that the stronger fluctuations appear past the logarithmic layer, favors the hypothesis that the forcing is only indirectly the cause of the stronger turbulence levels on the intermittent region.

The logarithmic region in the pressure profile of BL_s is incipient at $\delta_{99}^+ = 1500$, and more evident around $\delta_{99}^+ = 2000$. The same profile in BL_f , is slightly more tilted, and seems to have a different slope where the logarithmic region should appear. Given that this region appears once the separation of scales is sufficient, a possible explanation for this difference is that the ratio between ν/u_τ and δ must be higher in the forced case to observe an equivalent scale separation than in the smooth case. Unfortunately, $\delta_{99}^+ = 1500$ is the largest Reynolds number available in this case, and the available experiments at a higher Reynolds number were not able to measure the pressure fluctuations profile.

The conditions for scale separation were also studied by Krogstad & Efros (2012), where a boundary layer with the same roughness geometry than Krogstad & Antonia (1994) (cross-stream rods), but at a significantly higher Reynolds number, was analyzed. They found that the effect of the roughness on the turbulent statistics diminished with the Reynolds number and δ_{99}/k , and mentioned that the significantly shorter u coherent structures in Krogstad & Antonia (1994) could be a low Reynolds number effect.

The previous results can be interpreted as follows. The intensity of the pressure and velocity fluctuations at the outer region become more intense with the presence of roughness or forcing. The range of influence of the intermittent region, with a characteristic length

of δ_{99} , increases reducing the extent of the region that is characterized by the distance to the wall. The mechanism that increments the rate of spread of the boundary layer, with a characteristic length of k_s , does not leave a direct footprint on the flow, that can be described with the same two characteristic lengths as a smooth boundary layer. In consequence, roughness could be a tool to amplify the effect of the intermittent region over the rest of the flow.

Comparison between the present forcing and actual roughness.

Another important question whether this particular forcing is similar to actual roughness. The most similar cases are Lee *et al.* (2011) and Lee & Sung (2007), two direct numerical simulations of two common roughness patterns at a lower Reynolds number, probably too low to fulfill the previously mentioned conditions for scale separation, but with a precise description of the characteristics of the flow close to the roughness elements. Their staggered array of cubes has a $k_s^+ \simeq 60$, similar than the present case, while the cross-stream rods was designed to reproduce Krogstad & Antonia (1994), with $k_s^+ = 211$.

The effect on the turbulent profiles is similar to the observed in figure 4.5, with good support for the wall similarity hypothesis. Both roughness patterns lower considerably the peak of u' at $y^+ = 15$ similarly to the present case. The effect on the other two components is similar too, the wall-normal component is only slightly damped, while the peak in the spanwise component is unaltered or even enhanced with the presence of a three-dimensional roughness pattern. The most remarkable difference can be seen in the Reynolds shear stress profile, flatter in the present case, and without the inflection point that appears in the present case, that does not seem to affect the other components of the Reynolds stress tensor. This suggests that the effect of the forcing is comparable to actual roughness. The described behavior is analogous to other forcing schemes like the suction pattern of Flores & Jiménez (2006).

The comparison with experiments is not straightforward because to measuring the properties of the flow very close to the roughness elements is particularly hard. The measurements obtained far from the wall, above $y/\delta_{99} = 0.2$, should be accurate for hot wires and PIV. For cases with $k_s > 40$, a similar deviation from the behavior of smooth walls is seen in the majority of experiments for the three components of velocity, as it can be seen in the review by Flack & Schultz (2014). It is therefore reasonable to think that,

regardless of the validity of Townsend's wall similarity hypothesis, the effect of the present volumetric forcing is similar to the effect of actual roughness.

4.4 Velocity field structure

The goal of the two-point statistics is to determine the spatial organization of the flow. All the results shown previously have been obtained as an ensemble average of a quantity ϕ in a given point \mathbf{x} . This provides little information about how the value of ϕ in \mathbf{x} depends on its environment. Important events in wall bounded flows are known to be coherent, like ejections and sweeps. The most frequent tool to quantify that coherency is the autocorrelation coefficient $C_{\phi\phi}$ defined as

$$C_{\phi\phi} = \frac{\langle(\phi - \langle\phi\rangle)(\phi_1 - \langle\phi_1\rangle)\rangle}{\phi'\phi'_1}, \quad (4.10)$$

where ϕ is the value of the scalar in the point \mathbf{x} , and ϕ_1 is the value of the same field in an alternative position \mathbf{x}_1 . The brackets denote the ensemble average, and the primes are the standard deviations. Note that the autocorrelation coefficient is a real scalar quantity defined for every point in space \mathbf{x}_1 , and that the point from which the autocorrelation is computed \mathbf{x} is a parameter. Therefore the complete parameter space for the autocorrelation $C_{\phi\phi}$ has five dimensions, since only the spanwise direction is homogeneous. A single autocorrelation for a single variable for a given position \mathbf{x} has a size of the order of a complete flow field. Computing, managing and processing the results of the autocorrelations is a relatively challenging task for the present data sets due to their volume.

The autocorrelation measures how much the value of ϕ in \mathbf{x} is correlated with its environment. If a strong and frequent event is very localized its imprint in the autocorrelation will be small. On the other hand, if a strong and frequent fluctuation is coherent in space, it has to be visible in $C_{\phi\phi}$, and the extension of this coherence should appear in the correlation levels. This quantity does not capture coherence between that is very weak or intermittent. Assume that the value of ϕ at \mathbf{x} is 1, and the value ϕ_1 in \mathbf{x}_1 is 1 or -1 intermittently, each one half of the time. The value of $C_{\phi\phi}$ at \mathbf{x}_1 will be 0, because the average is smoothing out that intermittency. On the other hand, if the correlation is conditioned to positive or negative values, $C_{\phi\phi}$ at \mathbf{x}_1 will be 0.5. Sillero *et al.* (2014) obtained conditional correlations of the velocity field, but that result is considered beyond the goals of this work.

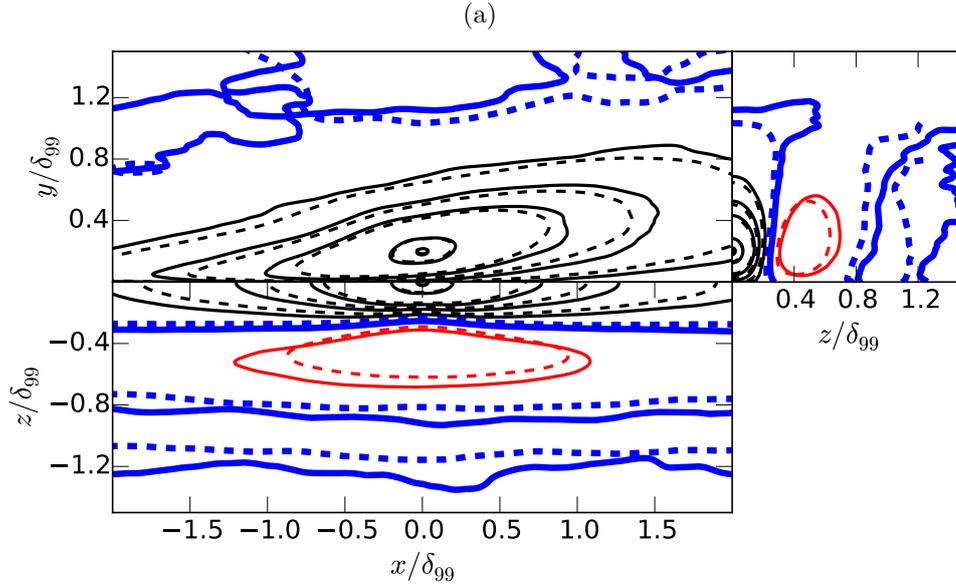


Figure 4.6: Sections of C_{uu} ; left top, section over the XY plane; left bottom, section over the XZ plane; right, section over the YZ plane. Three different distances from the wall are presented, (a) $y/\delta_{99} = 0.2$, (b) 0.6 and (c) 1.0. The two sets of contours correspond to BL_f (solid), and BL_s (dashed). Five black contours for positive correlation are presented, regularly spaced from 0.9 to 0.1. The thicker blue contour is line of zero correlation. Contours of negative correlation are red.

Townsend's wall similarity hypothesis is often formulated as an expression of the autocorrelation coefficient. According to Raupach *et al.* (1991), the hypothesis is valid if the shape of $C_{u_i u_i}$ (where u_i represents any of the three components of \mathbf{u}) is not affected by the presence of roughness. Note that the definition is given with the covariance R_{ij} , but it is equivalent because u_i^2 is contained in the standard deviations of equation (4.10).

Autocorrelation coefficients for each component of the velocity vector are presented in figures 4.6, 4.7 and 4.8. The two available boundary layers are used at the usual streamwise location where $\delta_{99}^+ = 1500$. A comprehensive study of this quantity, and a thorough comparison of BL_s to C_2 can be found in Sillero *et al.* (2014). The following discussion tries to determine if this anomaly in the shape of the autocorrelations invalidates Townsend's hypothesis.

The three heights used to compute the autocorrelation ($y/\delta_{99} = 0.2, 0.6, 1.0$) are not arbitrarily chosen. We saw in Chapter 3 that at the first stage $y/\delta_{99} = 0.2$, the presence of irrotational flow is almost impossible, and the average profiles are still dominated by the properties of the logarithmic layer. In the second stage, the level of intermittency is important, but it is not dominant, and the mode of vorticity that corresponds to the fully

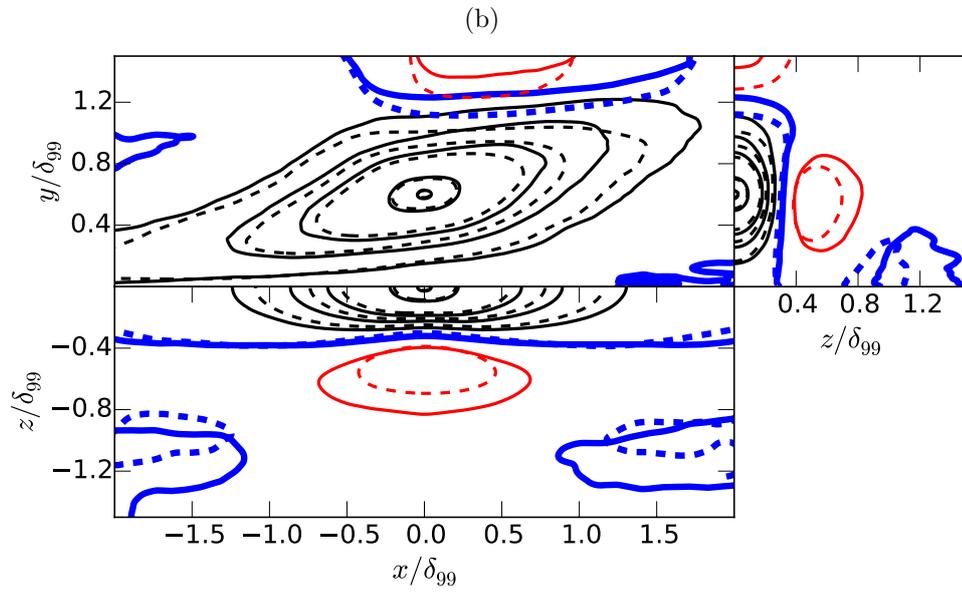


Figure 4.6: (continued)

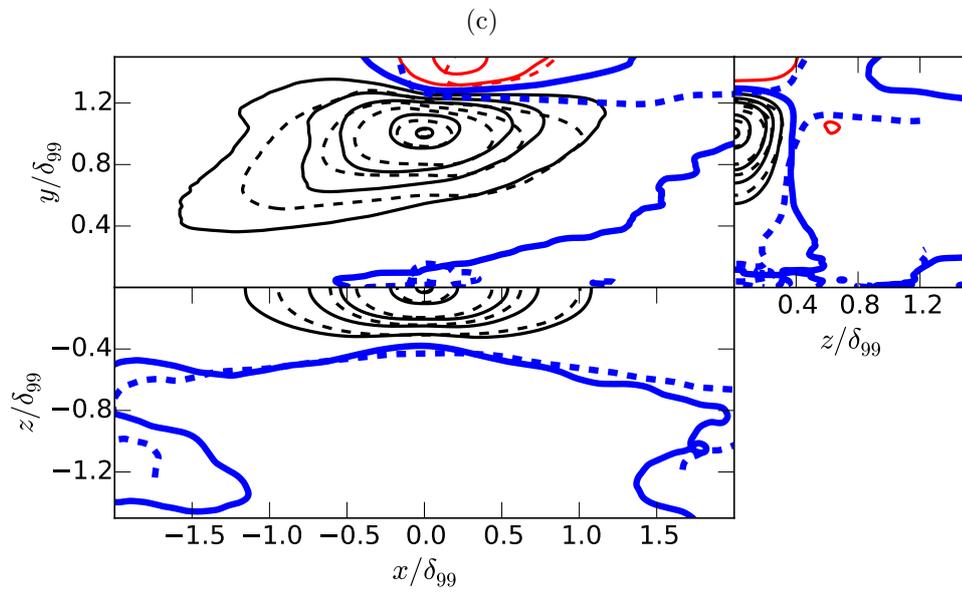


Figure 4.6: (continued)

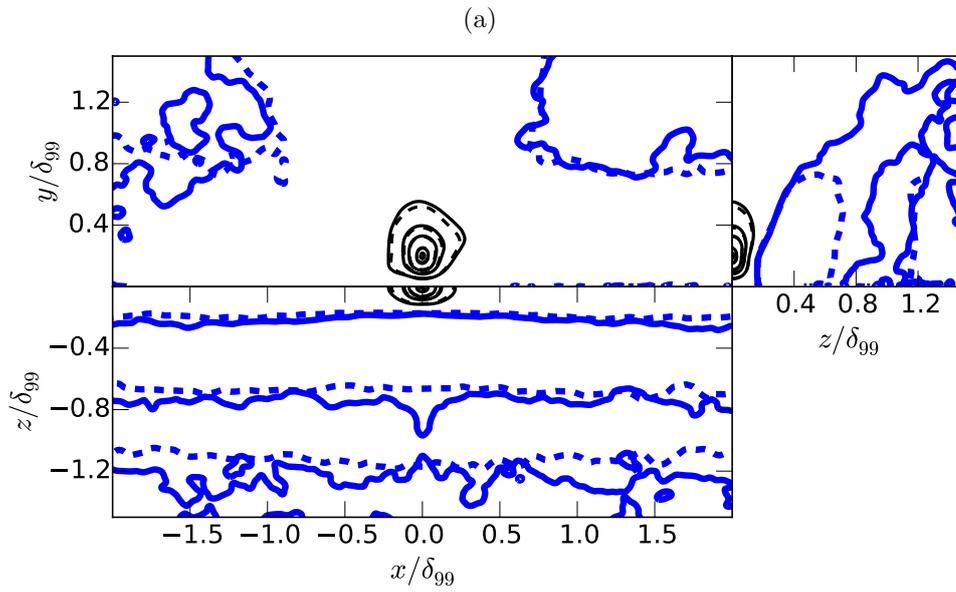


Figure 4.7: Sections of C_{vv} . The definitions are the same as figure 4.6.

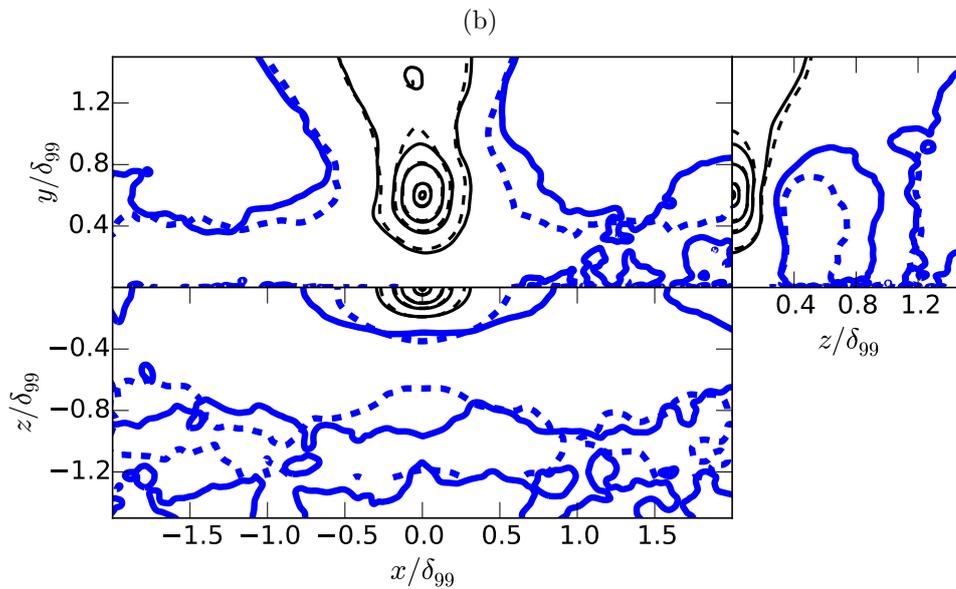


Figure 4.7: (continued)

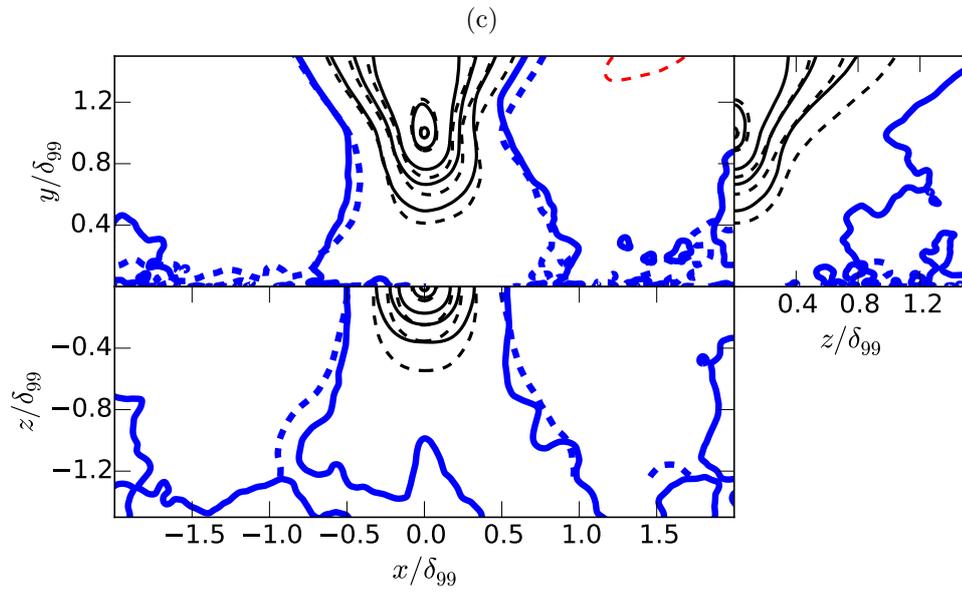
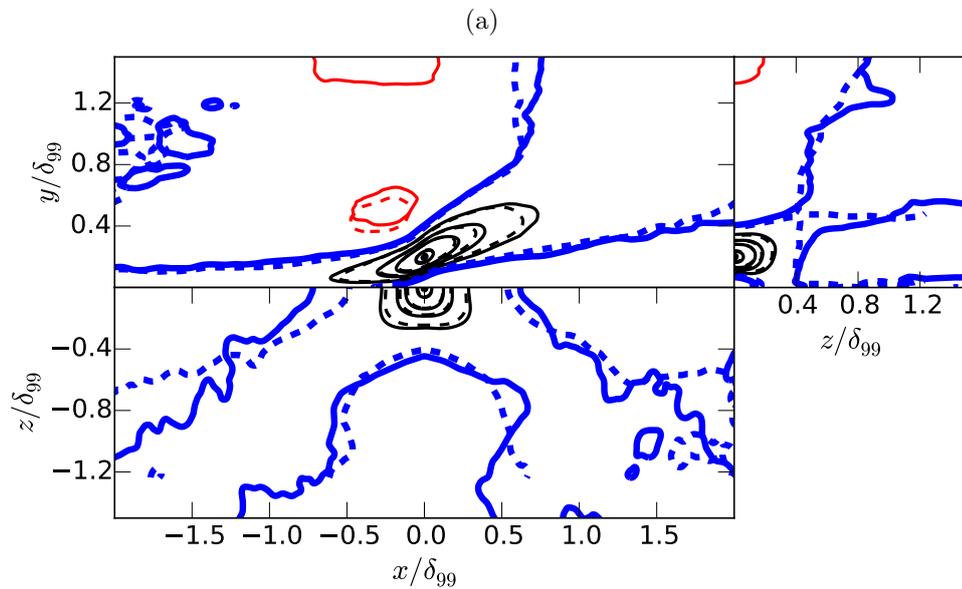


Figure 4.7: (continued)

Figure 4.8: Sections of C_{wv} . The definitions are the same as figure 4.6.

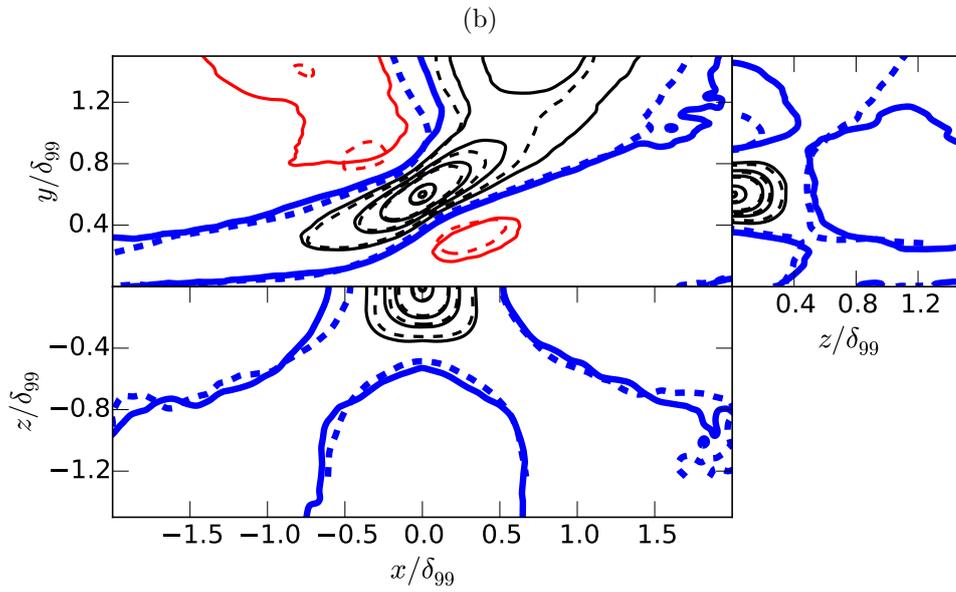


Figure 4.8: (continued)

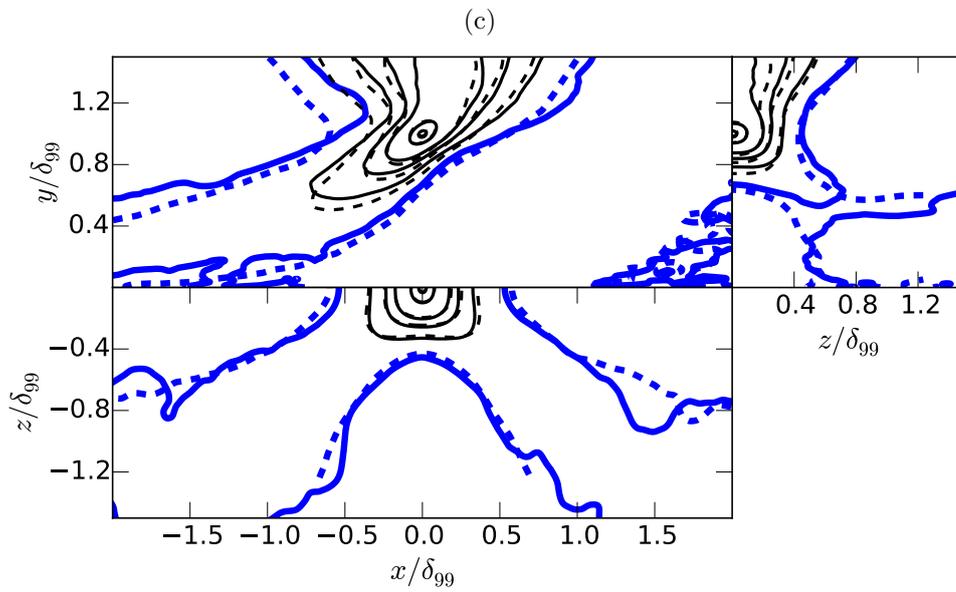


Figure 4.8: (continued)

turbulent flow is still clearly visible (see figure 3.3). At the edge of the boundary layer the level of intermittency is the highest, and the amount of turbulent and non-turbulent flow is similar. Note that the autocorrelation coefficient has not been conditioned in any way, and the results of $C_{\phi\phi}$ also contain the influence that the turbulent flow may have on the non-turbulent region and vice-versa.

The autocorrelation coefficient in BL_f has a very similar shape than BL_s , but it is consistently larger. This means that strong and coherent events in the velocity field are larger in BL_f . Many other cases of slightly longer autocorrelations with the presence of roughness can be found, like Lee *et al.* (2011), Volino *et al.* (2011) and Wu & Christensen (2010). In this case, and like many other experiments and simulations, the correlations are much less affected by the additional friction than in Krogstad & Antonia (1994). The differences between BL_s and BL_f are similar to the differences between BL_s and channels at the same Reynolds number, and those differences are interpreted as an effect of the large scales more than as an indication that the two flows are entirely different. It is also important to note that the shape of the contours of C_{uu} that mark relatively high correlation levels are similar, while the contours of C_{vv} and C_{ww} are equal within a reasonable range. The differences in the three components of the velocity are more visible when the autocorrelation is centered at $y/\delta_{99} = 1$, where the geometrical properties of the intermittent region have to be more relevant.

This result does not confirm Townsend's wall similarity hypothesis, but suggests a similar explanation for the differences in the mean profiles. The influence of the large scales is distinguishable in BL_f and BL_s , the same way it is distinguishable between boundary layers and channels. The differences can be explained by the fact that entrainment influences the layout of the turbulent structures in the intermittent region. The next step is to find if turbulence, regardless of having a slightly different organization in the forced case, is measurably different. The entrainment rate in BL_f is almost twice higher than BL_s , and it is important to determine if turbulence is different in any way close to the T/NT interface.

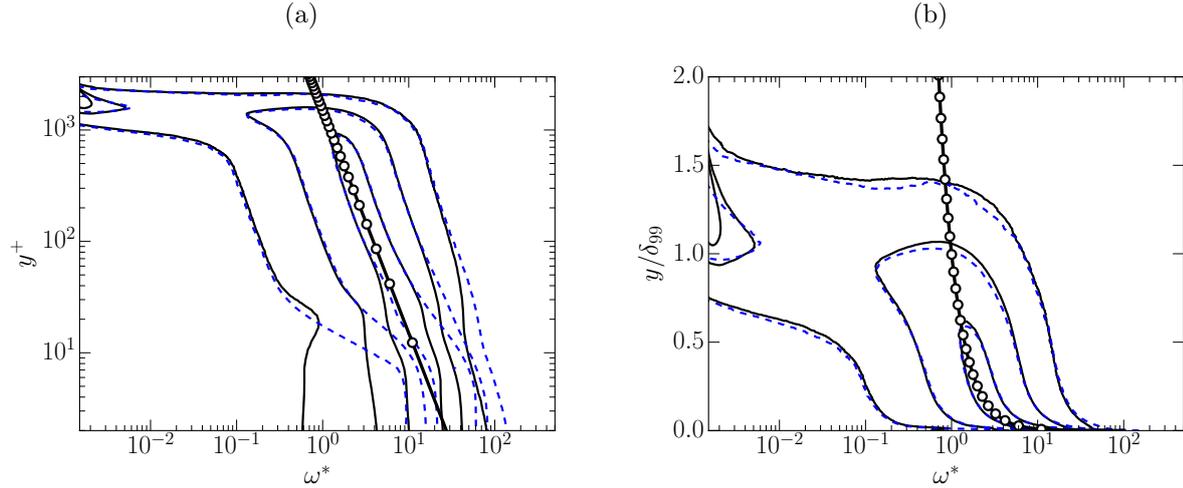


Figure 4.9: Premultiplied PDF of vorticity magnitude ω^* and distance to the wall y/δ_{99} , $\omega\Gamma_{\omega,y}$. Two different scales have been used in the vertical axis: (a) logarithmic to zoom into the logarithmic layer, and (b) linear to zoom into the intermittent region. The two contour sets correspond to BL_f (— black), and BL_s (--- blue). Each contour level contains 50%, 90% and 99% of points.

4.5 The effects of the additional entrainment in the intermittent zone

What follows is the study of the intermittent region analogous to §3.3-3.5, with the tools described and tuned for the data in Sillero *et al.* (2013), which is the base of the comparison.

We saw that most of the information about the intermittent properties of the flow could be summarized in the joint PDF of vorticity ω^* and wall-normal distance y ($\Gamma_{\omega,y}$ in figure 4.9). Other results that are more often found in the literature, like the intermittency parameter γ , are implicitly included in this analysis since they can be obtained from the joint PDF.

Representing the premultiplied joint PDF $\omega\Gamma_{\omega,y}$ with a logarithmic vertical axis (figure 4.9a), one can clearly see the effect of the forcing below $y^+ = 20$. Vorticity is significantly damped by the forcing, but the most probable values recover below $y^+ = 40$, and the complete shape of the joint PDF is recovered before reaching the logarithmic layer. Above $y^+ = 100$ the values of vorticity for the two boundary layers are almost identical, following the expected trend until the beginning of the intermittent region. This trend is derived from the approximation that at a given y dissipation is proportional to u_τ^3/y , which

assumes that it exists a logarithmic layer. The fact that both boundary layers follow the expected trend is not an additional validation of Townsend's wall similarity hypothesis because it is redundant with the fact that both cases have the same Kármán constant κ , but it confirms that star units $u_\tau^2/\nu\sqrt{\delta_{99}}$ are also suitable to compare the two boundary layers. Note that star units include the dependence with the Reynolds number and the fundamental scaling parameter u_τ . The wall similarity hypothesis will be valid if the properties of the intermittent region and the T/NT interface are identical in star units.

The agreement is worse far from the wall, and differences between BL_f and BL_s can be seen again when $\omega\Gamma_{\omega,y}$ is represented with a linear vertical axis (figure 4.9b). The fully turbulent flow with the vorticity characteristic of the logarithmic layer reaches deeper into the free stream, and farther from the wall. The differences in the intermittent region are measurable, but not important. We saw in §3.3 that the geometrical details are hidden by operating using the wall-normal distance as a parameter, and that the analysis of the T/NT interface is necessary to determine scales associated to entrainment. This difference in the intermittent region is consistent with the conclusions of the previous section §4.4. The properties of the turbulent motion are not necessarily different due to the additional friction, but its geometrical configuration beyond the logarithmic layer is.

The geometrical properties of the T/NT interface of BL_f are shown in figure 4.10. The interfaces of BL_f and BL_s similar complexity and topological properties. They experiment the same topological transition at a similar range of thresholds $\omega_0^* = 0.2 - 2$, affecting both estimations of the fractal dimension (figures 4.10 a,b) and the genus (figure 4.10c). The joint PDF $\Gamma_{\omega,y}$ shows that there is some more vorticity at the edge of the boundary layer in BL_f , which is consistent with the fact that the fractal dimensions and the maximum genus per unit volume (figure 4.10 d) are measurably higher. However, the precise value of these two quantities in the forced case is difficult to interpret.

The next step is the conditional analysis of the vorticity field, analogous to the one in §3.5 for BL_s . We have repeatedly mentioned that one possible explanation for the lack of agreement of the wall-normal profiles is the fact that y is probably not a suitable parameter to study the intermittent zone. The characteristic scale for the intermittent region is δ_{99} , and the geometrical configuration of the turbulent structures is not necessarily oriented with the wall, one condition that is fulfilled indeed in the logarithmic layer. In addition, the analysis of the two-point statistics suggest that the geometrical configuration of the

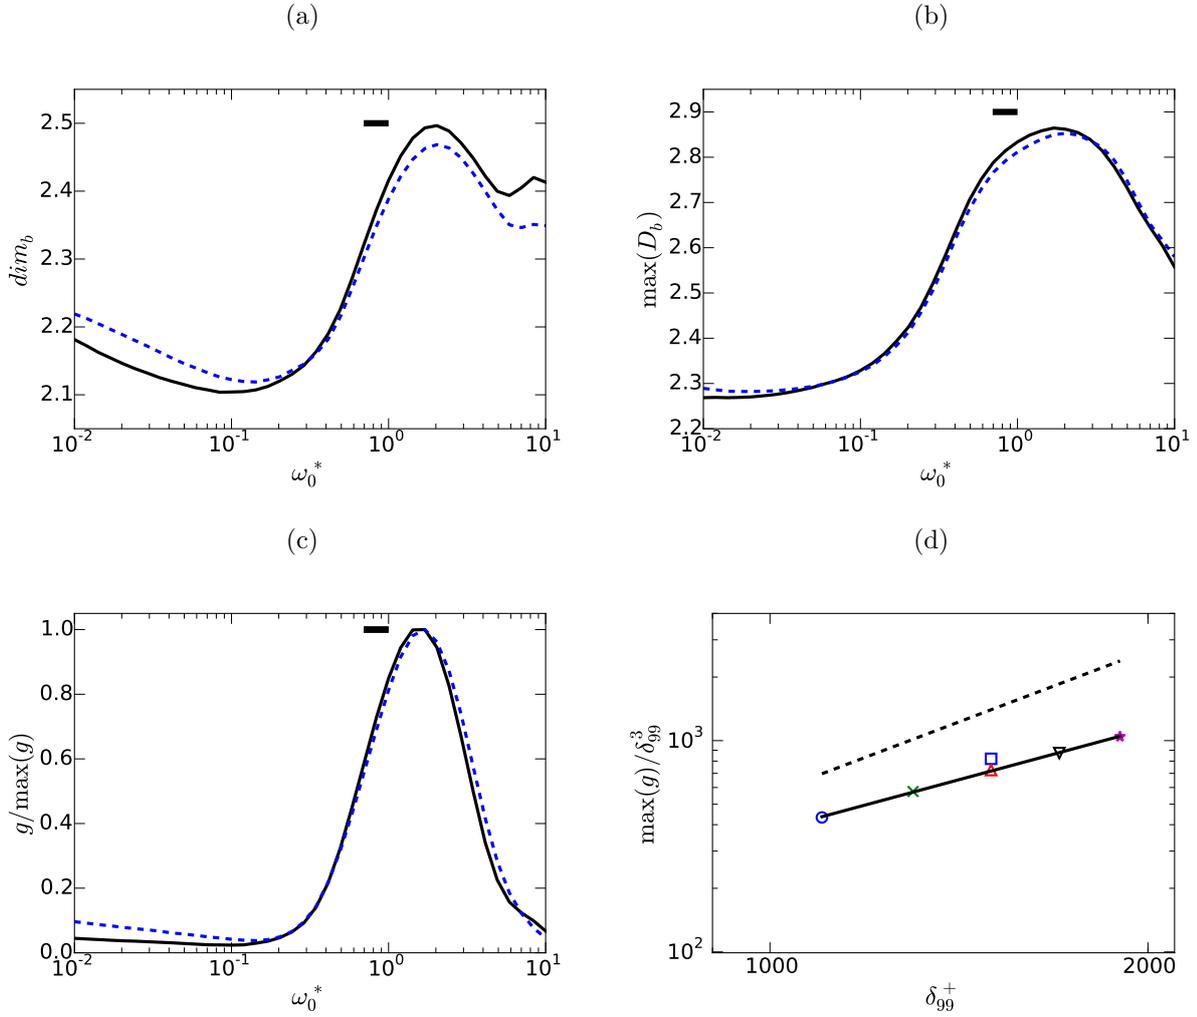


Figure 4.10: (a) Lower estimation of the fractal dimension (see equation (3.13)). (b) Higher estimation of the fractal dimension (see equation (3.14)). (c) Genus of the interface surface normalized with its maximum value for any threshold $g/\max(g)$. (d) Maximum genus per unit volume δ_{99}^3 , for $\omega_0^* \simeq 2$ for the rough interface surface (square). The additional points and the regression correspond to the smooth boundary layer (see figure 3.9). The thick lines in figures (a-c) correspond to the ratio ω^*/ω^+ for the smooth case.

turbulent flow of BL_f is different from BL_s .

The premultiplied joint PDF of vorticity and minimum distance to the T/NT interface $\omega F_{\omega, \Delta_b}$, presented in figure 4.11.

It is useful to recall the expressions for the wall normal average

$$\langle \omega \rangle(y) = \frac{\int_0^\infty \omega \Gamma_{\omega, y} d\omega}{\int_0^\infty \Gamma_{\omega, y} d\omega}, \quad (4.11)$$

the conditional average

$$\bar{\omega}(\Delta_b) = \frac{\int_0^\infty \omega F_{\omega, \Delta_b} d\omega}{\int_0^\infty F_{\omega, \Delta_b} d\omega} \quad (4.12)$$

despite these quantities will not be commented in detail in this section. F and Γ are the result of an identical computation set with a different reference frame. While Γ uses the distance to the wall y , F uses the distance to the interface Δ_b . They have indeed complementary properties. The wall is a suitable frame of reference to study the events that occur close to the wall. Above the logarithmic layer, where intermittency is relevant, the scaling respect to y is not necessary the most suitable.

The T/NT interface is a suitable reference frame to study the intermittent region because it is able to abstract the geometrical complexity, and the conditional analysis is easier to interpret. In other words, F is just a way to look at the boundary layer from the top to the bottom.

The premultiplied joint PDF $\omega F_{\omega, \Delta_b}$ for BL_f and BL_s are presented in figure 4.11 for two different thresholds at the usual Reynolds number. The most evident feature of the joint PDF is the good agreement between the two cases regardless of the threshold ω_0^* . The same three regions described in §3.5 are found: an interface layer close to the T/NT interface with a thickness that scales in Kolmogorov units, the bulk of the turbulent flow that is not affected by the presence of the free stream, and the near-wall region at the end of the hockey-stick-shaped region. The agreement is almost perfect in the regions where it is expected to occur, any point that is relatively far from the wall or has a vorticity $\omega^* > 1$. The conditional averages $\bar{\omega}^*$ of BL_f and BL_s are hard to distinguish below $\Delta_b/\eta < 100$.

The value of the slope of the contours of the premultiplied joint PDF and the conditional profile is particularly hard to explain, as it was mentioned in §3.5, but the fact that the power law is observed is consistent with the geometrical properties of the T/NT interface,

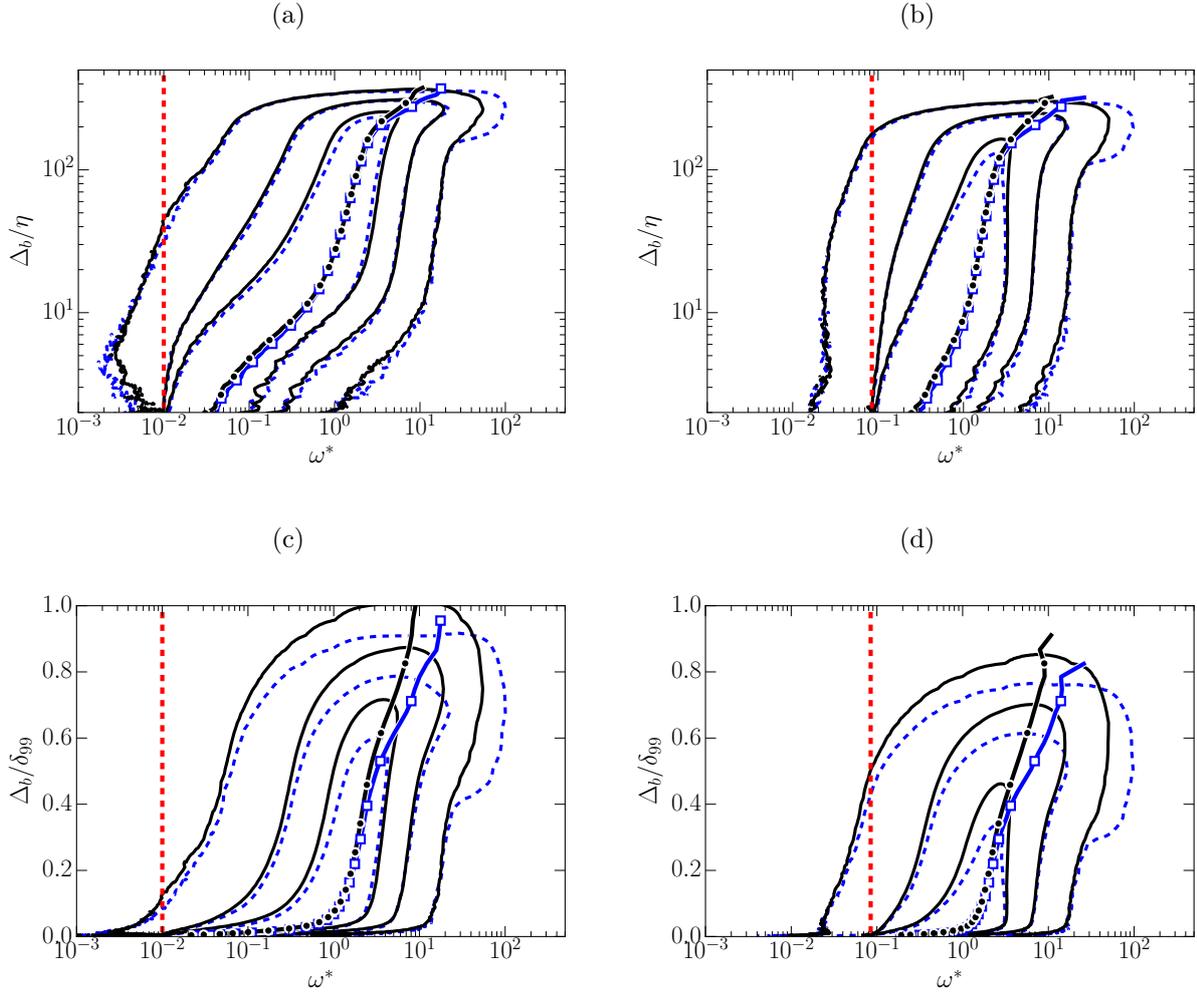


Figure 4.11: Premultiplied joint PDF of vorticity and minimum distance $\omega F_{\omega, \Delta_b}$ at (a) $\omega_0^+ = 0.01$, and (b) $\omega_0^+ = 0.09$. Each contour set corresponds to BL_f (solid) and BL_s (dashed). The vertical red dashed line highlights the value of the threshold. The two dotted lines are the conditional average $\bar{\omega}^*$ for BL_f (empty blue squares) and BL_s (solid black circles). (c,d) correspond to (a,b) with δ_{99} as the unit of length. Each contour line contains 50%, 90% and the majority (99.9%) of points respectively.

the surface $\Delta_b = 0$. It has been determined that the vorticity magnitude is proportional to $y^{-1/2}$. Furthermore, there is practically no relationship between the two reference frames $y = 0$ and $\Delta_b = 0$. $\Delta_b = 0$, is a fractal-like surface that contains all the possible lengths between δ_{99} and η , while the distance to the wall y has only a weak influence on the average orientation of the interface. In consequence the fully turbulent flow (with vorticity higher than $\omega^* = 1$) should be scale-independent of Δ_b .

When F_{ω, Δ_b} is scaled with the boundary layer thickness (figure 4.11c-d), we see that the differences between the BL_f and BL_s arise in the hockey-stick shaped region that corresponds to points that are closer to the wall. Vorticity is lower in the forced case, which is caused by the forcing, and the maximum Δ_b is also slightly longer, which is consistent with the fact that the arrangement of the turbulent eddies in the intermittent region is different. The fact that the maximum distance from the interface is longer in the forced case adds some more information about how turbulence is arranged. The ball distance is the minimum distance, therefore, the maximum Δ_b is the minimum distance between the wall and $\Delta_b = 0$, which is obtained after cleaning all the bubbles and drops (see §3.4) of a vorticity isocontour $\omega(x, y, z) = \omega_0$. The maximum ball distance is the distance between the lowest edges and overhangs and the wall, and it has little to do with the average location of the interface. The average vertical location of BL_s and BL_f is very similar, like the differences of the contours at the intermittent region in figure 4.9(b), while the maximum ball distance of the two cases differs by more than 10%. Handles and pockets, that were studied in detail in §3.5, are one of the consequences of the geometrical complexity of the interface. The T/NT interface of BL_f is therefore more complex (with more handles and a higher fractal dimension), but it is farther from the wall than in BL_s . At the same time, the T/NT interface layer seems to be identical in both boundary layers.

One of the consequences of this result is that the force boundary layer is apparently less intermittent, in the sense that the fully turbulent flow fills more volume between $y/\delta_{99} = 0 - 1$. There are more handles and pockets in BL_f than in BL_s , but the ones in the former do not reach as close to the wall as the latter. Another consequence is that it strengthens the conclusions of §3.5. The vertical location of handles and pockets is immaterial to the amount of non-turbulent fluid they entrain. The forced boundary layer entrains 75% more fluid per unit length in the streamwise direction, but handles and pockets are further from the wall. If engulfment was an important mechanism in entrainment, the case should be the opposite.

The mean conditional profiles of BL_s and BL_f are also different beyond $\Delta_b = 0.3\delta_{99}$, although this limit changes with the threshold. The differences are due to the combined effect of the forcing, that shrinks the tip of the hockey-stick-shaped corner of $\omega F_{\omega, \Delta_b}$, and the differences in the location of handles and pockets, that extends $\omega F_{\omega, \Delta_b}$ towards higher ball distances.

This result recovers the scaling predicted by Townsend's wall similarity hypothesis above $y/\delta_{99} = 0.4$. Once the geometrical effects in the intermittent region are subtracted, the scaling with u_τ is found in the turbulent/non-turbulent interface layer as well.

4.6 Conclusions

A zero-pressure-gradient turbulent boundary layer with a volumetric near-wall forcing has been simulated. The forcing has been designed to increase the wall friction without introducing any obvious geometrical effect, and to keep the performance and scalability figures of the DNS code. The actual simulation is split in two domains, a smaller one in charge of the generation of correct inflow boundary conditions, and a second and larger one where the forcing is applied.

The streamwise evolution of the transient caused by the sudden addition of the forcing is analyzed. The evolution of the Clauser parameter G as a function of the eddy turnover distance \hat{x} of the present case is comparable to a similar experiment. The boundary layer reaches its final state in about $\hat{x} = 2$. While the precise quantity is difficult to interpret, and it may depend on the details of the forcing or the roughness pattern, the order of magnitude is consistent with the current knowledge about streamwise convergence of boundary layers. The stage at which the present boundary layer is compared with its smooth-wall counterpart, at $\delta_{99}^+ = 1500$, is far enough from the end of the transient, at $\hat{x} = 3.3$, and the comparison between the two cases is very likely to be correct.

The one-point statistics of the forced boundary layer show that the effect of the forcing is similar to actual roughness. The most visible effect is that the peak of production of turbulent kinetic energy $y^+ = 15$ is damped and shifted further from the wall, where the effect of the forcing is significantly weaker. The streamwise component of the velocity fluctuations is also strongly affected, while the other two components preserve their qualitative features. The agreement between the forced and the smooth-wall boundary layers is relatively good at the logarithmic layer, but one-point statistics above $y/\delta_{99} = 0.4$

present differences that are of the order of the ones between boundary layers and channels. These differences are caused by entrainment, that affects the intermittent region and the geometrical configuration of eddies in the outer part of the boundary layers.

The two-point statistics of the velocity components support the previous conclusion. The shape of the three dimension autocorrelation of the three velocity components of the forced boundary layer is similar than the smooth-wall case, but the correlation lengths are consistently longer, particularly in the case of the streamwise component. Similar differences in the autocorrelation coefficient have been reported in practically every experiment and simulation of boundary layers over rough walls. Those differences are stronger when the point at which the autocorrelation is computed is further from the wall.

The geometrical complexity of the T/NT interface has an important role in the intermittent region and in entrainment. The results also suggest that the differences in the one-point statistics between the forced and the smooth boundary layers may be caused by the mentioned geometrical complexity. After changing the frame of reference from the wall (the $y = 0$ surface) to the T/NT interface (the $\Delta_b = 0$ surface), the scaling predicted by Townsend's wall similarity hypothesis is recovered. Despite the agreement in the one-point statistics is excellent in the T/NT interface layer, some geometrical details of the interface are different. The genus and the fractal dimension are higher in the forced boundary layer, suggesting that the T/NT is a more complex surface. The maximum ball distance is also higher, meaning that the forced boundary layer is slightly less intermittent.

This result also strengthens the conclusions derived from the analysis of the turbulent/non-turbulent interface layer in Chapter 3. A different boundary layer with an entrainment rate 75% higher exhibits exactly the same properties regarding its interaction with the non-turbulent free-stream, suggesting that entrainment is mostly a local phenomenon that occurs at the edge of the most external vortices of the boundary layer. It also confirms that engulfment, defined as the process for which large patches of non-turbulent flow are digested by the coherent structures of the turbulent motion, is not particularly relevant.

Finally, we conclude that the validity of Townsend's wall similarity hypothesis has been confirmed for one-point statistics covering the entire thickness of the boundary layer, but not for two-point statistics, where it can only be considered a good first approximation.

Chapter 5

Conclusions and future work

Townsend's wall similarity hypothesis is useful and relatively successful at modeling rough-wall turbulent boundary layers at sufficiently high Reynolds numbers, but this hypothesis was not fully backed by the available experimental data beyond the logarithmic layer. There was the belief that entrainment could play an important role in an accurate description of boundary layers, and that the additional friction caused by the presence of roughness was indirectly causing those differences. It was necessary to understand entrainment to study Townsend's hypothesis.

Average entrainment and friction are coupled by the Kármán momentum equation, but the actual mechanism for which entrainment occurs was not clear, neither its connection with what happens at the wall. It was not that the coupling between two models was particularly hard, but that those two models did not even exist.

The first step has been to understand entrainment and the region where it takes place, the turbulent/non-turbulent interface. One of the significant findings of this thesis is that the intermittent region is geometrically complex, and that complex plays an important role in the description of the properties of the outer region of boundary layers. One facet of this role is purely mathematical. The ensemble average operation requires that the direction across it is computed is statistically homogeneous. This is the reason why homogeneous and isotropic turbulence, and turbulent channels and pipes, have been such a successful tool. The geometrical complexity is significantly reduced when the large-scale configuration is simple. Intermittency is an indication that statistical homogeneity is probably lost, but in the case of turbulent boundary layers, the tools used in its analysis have been the same as in channels and pipes. Maybe because of the lack of a better

alternatives.

This is the reason why most of this thesis deals with the development of new tools. Simulations at higher Reynolds numbers, and algorithms to deal with the turbulent motion regardless of its geometrical complexity. Developing new tools is always hard, and interpreting results that no one has ever seen before is sometimes harder. In addition, some tools can be better than others if their building blocks are already available, are particularly efficient to compute, or easier to interpret. If we could build a computer that gave us the ultimate answer to Turbulence, and its answer was 42 , our efforts would be practically useless.

A new DNS code for the simulation of turbulent boundary layers has been developed. It is based on a previous code that provided the fundamental tools and algorithms, but the changes needed for the required levels of scalability and performance were substantial. Its scalability is particularly remarkable. Being able to use almost two million threads while keeping almost constant performance levels is a remarkable achievement, even when the supercomputing architecture is the best available¹. The implementation is also relatively simple, given the complexity of the algorithms. The twin box configuration, used in the two main data sets of this thesis, was developed in less than a month.

We have made an important effort to characterize on a simple way the geometrical features of the intermittent region. A fractal-like highly contorted surface has been characterized with three simple components: bubbles, handles and pockets of different sizes. This simplification is one of the achievements of this work, the identification of abstract quantities, like the fractal dimension and the genus, with actual structures that exist in the flow.

The ball distance field (called minimum distance field in other applications) is also another tool that has been proven itself useful. We have shown repeatedly that any measurement involving a distance includes its reference frame as a fundamental assumption. The vertical distance, one of the most common tools in the study of the T/NT interface, is not able to capture the geometrical complexity of the T/NT interface. As a consequence, some of the important properties of the interface that are relevant to understand entrainment are not properly measured. This was not known before this new tool, the ball distance, was

¹The present OpenTBL code is in the HighQ club, a series of featured tools that were able to use the whole Juqueen supercomputer: http://www.fz-juelich.de/ias/jsc/EN/Expertise/High-Q-Club/_node.html. We are grateful with the Juelich staff, that offered us the chance to look for the limits of this code.

developed. The lack of accuracy of the predictions made by Townsend's wall similarity hypothesis in the outer region of the boundary layer is a similar case. If the outer region of the boundary layer is measured using a more intrinsic reference frame than the wall, the agreement is recovered with great precision.

At the time of the elaboration of this thesis we can say that the behavior of the different coherent structures in the logarithmic layer is more or less understood. This study completes a relatively simple portrait of the dynamics of turbulence close to the edge of the boundary layer. The fully turbulent flow far from the wall behaves like in the logarithmic layer, dissipating energy at a rate of u_τ^3/y . The environment that vortices find at the edge of the boundary layer is the same as in the rest of the flow, there is no local effect caused by the presence of the free stream like some impact pressure or some sort of local activity, and vortices have a peaceful death. They are still stretched, since strain within the T/NT interface layer is relatively intense, but not more than any other vortex. Entrainment is mostly a local effect, caused by strong vorticity gradients, and the rest of the turbulent flow evolves ignoring its existence. If roughness is added to the equation the mechanism is essentially the same. The vorticity gradients are stronger by a factor that scales with the friction coefficient, and in consequence, entrainment is amplified.

The connection between roughness (or additional friction) and entrainment is that the spatial arrangement of eddies in the outer region of the boundary layer is affected by the changes at the wall. Boundary layers over rough walls are less intermittent, and correlation lengths are longer for any velocity component. But the properties at each particular point fulfill Townsend's wall similarity hypothesis. This is the final statement about the validity of Townsend's hypothesis. The average properties of any point within the boundary layer, from the lower bound of the logarithmic layer to the T/NT interface, follow the scaling predicted by the hypothesis. On the other hand, the structure of the turbulent motion is affected by the additional friction.

There are sound arguments to support these conclusions, that are a model for the properties of the intermittent region of boundary layers, and of the effect of rough walls within in the transitionally rough regime.

This portrait has implications on LES modeling. The dynamics of the turbulent flow are not fundamentally different in the intermittent region, and most formulations should obtain a flow that is correct considering the limitations of the models. The T/NT interface

would be of course filtered, and in consequence thickened, but mostly correct, since LES is theoretically capable to capture the differences in the large-scale features of the turbulent motion.

The tools developed in this thesis are already having a meaningful impact within the turbulence research community. The OpenTBL boundary layer code is now used by other research groups. At the time this thesis is written, we know of two other versions of this OpenTBL, one aimed to simulate boundary layers with adverse pressure gradients after laminar-turbulent transition, and a second one designed to simulate boundary layers with strong adverse pressure gradients. When we implemented this code, we tried to make it as future-proof as possible. It seems that, at least in some extent, we succeeded in our goal. Ball distance can be useful to study other aspects of turbulent flows. It is efficient, fast, and allows to formulate approximate gradients close to complex surfaces. It can be applied to study the T/NT interface of other external turbulent flows with close to no modification.

There are some questions that are left unanswered. In boundary layers, friction controls entrainment like a throttle. At this point we understand how entrainment works, and why Townsend's hypothesis is valid also at the T/NT interface. But the link between the wall and the T/NT interface, the spatial arrangement of eddies beyond the logarithmic layer, is far from being understood. There is also a resilient question in any work based on results from DNS. The practical scale separation found in the T/NT interface and in the forced boundary layer is marginal, and these conclusions have to be verified for boundary layers at higher Reynolds numbers.

Finally, I'd like to say that the time and effort invested in the development of new tools is never properly rewarded. Results are the fundamental outcome of any research and tools only pay off from the moment in which they are finished and correctly tuned. The OpenTBL code took almost two years of development. Some parts of it, like the I/O system that is responsible to write the raw results, had to be written four times from scratch. Some bugs took weeks to catch, and sometimes threatened the whole project. Babysitting the simulation was also an interesting experience too.

Data management was a constant issue. If developing tools is never properly rewarded, moving data around is not rewarded at all. The tools required to compute and to analyze the ball distance required a similar amount of effort. One of the most challenging bits of

this thesis was to interpret the results of a brand new tool, and to discover its potential. One of the personal conclusions of this thesis is that tools are never properly rewarded by themselves, but the development of things that are completely brand new is probably what has taught me the most during the time I have devoted to this work.

Bibliography

- ANTONIA, R. A. & LUXTON, R. E. 1971 The response of a turbulent boundary layer to a step change in surface roughness part 1. smooth to rough. *J Fluid Mech* **48** (04), 721–761.
- ARYA, S., MOUNT, D. M., NETANYAHU, N. S., SILVERMAN, R. & WU, A. Y. 1998 An optimal algorithm for approximate nearest neighbor searching fixed dimensions. *J. ACM* **45** (6), 891–923.
- BANDYOPADHYAY, P. R. 1987 Rough-wall turbulent boundary layers in the transition regime. *J Fluid Mech* **180**, 231–266.
- BATCHELOR, G. K. 1967 *An introduction to fluid dynamics*. Cambridge U. Press.
- BELCHER, S. E., JERRAM, N. & HUNT, J. C. R. 2003 Adjustment of a turbulent boundary layer to a canopy of roughness elements. *J. of Fluid Mech.* **488**, 369–398.
- BISSET, D. K., HUNT, J. C. R. & ROGERS, M. M. 2002 The turbulent/non-turbulent interface bounding a far wake. *J. Fluid Mech.* **451**, 383–410.
- BLASIUS, H. 1950 *The boundary layers in fluids with little friction. Translation of Grenzschichten in Flüssigkeiten mit kleiner Reibung.* Z Math Phys. v. 56 H. 1 (1908).. National Advisory Committee for Aeronautics.
- BORRELL, B., SILLERO, J. A., GUNGOR, A.G. & JIMÉNEZ, J. 2011 Direct numerical simulation of a high-entrainment turbulent boundary layer. EB-4. College Park, MD: Am. Phys. Soc.
- BORRELL, G. & JIMÉNEZ, J. 2015 Properties of the turbulent/non-turbulent interface in boundary layers. *J Fluid Mech* **Submitted** (-), -.

- BORRELL, G., SILLERO, J. A. & JIMÉNEZ, J. 2013 A code for direct numerical simulation of turbulent boundary layers at high reynolds numbers in bg/p supercomputers. *Comp. Fluids* **80**, 37–43.
- BRADLEY, E. F. 1968 A micrometeorological study of velocity profiles and surface drag in the region modified by a change in surface roughness. *Q. J. Roy. Meteor. Soc.* **94** (401), 361–379.
- CHAUHAN, K., PHILIP, J., DE SILVA, C. M., HUTCHINS, N. & MARUSIC, I. 2014 The turbulent/non-turbulent interface and entrainment in a boundary layer. *J. Fluid Mech.* **742**, 119–151.
- CHEN, D., CHOUDHURY, A., EISLEY, N. *et al.* 2012 Looking under the hood of the blue gene/q network. *Proc. IEEE/ACM Supercomputing (SC 2012), Salt Lake City* .
- CHENG, H. & CASTRO, I. P. 2002 Near-wall flow development after a step change in surface roughness. *Bound-Lay Meteorol* **105** (3), 411–432.
- CLAUSER, F. H. 1954 Turbulent boundary layers in adverse pressure gradients. *J. Aeronaut Sci* **21**, 91–108.
- COLEBROOK, CF & WHITE, CM 1937 Experiments with fluid friction in roughened pipes. *P. R. Soc. A* pp. 367–381.
- COLEBROOK, CYRIL FRANK 1939 Turbulent flow in pipes, with particular reference to the transition region between the smooth and rough pipe laws. *J. Inst. Civil Engrs* **11** (4), 133–156.
- COLES, D. 1956 The law of the wake in the turbulent boundary layer. *J. Fluid Mech.* **1**, 191–226.
- CORRSIN, S. 1943 Investigation of flow in an axially symmetrical heated jet of air. *NACA WR* **W-94**.
- CORRSIN, S. & KISTLER, A. L. 1955 Free-stream boundaries of turbulent flows. *NACA TR* **1244**.
- DAHM, W. J. A. & DIMOTAKIS, P. E 1987 Measurements of entrainment and mixing in turbulent jets. *AIAA J.* **25**, 1216–1223.

- DARCY, H. 1857 *Recherches expérimentales relatives au mouvement de l'eau dans les tuyaux*, , vol. 1. Mallet-Bachelier.
- DE GRAAFF, D. B. & EATON, J. K. 2000 Reynolds-number scaling of the flat-plate turbulent boundary layer. *J Fluid Mech* **422**, 319–346.
- DEL ALAMO, J. C., JIMÉNEZ, J., ZANDONADE, P. & MOSER, R. D. 2004 Scaling of the energy spectra of turbulent channels. *J. Fluid Mech.* **500**, 135–144.
- DIMOTAKIS, P. E. 2000 The mixing transition in turbulent flows. *J. Fluid Mech.* **409**, 69–98.
- EFROS, V. & KROGSTAD, P-Å. 2011 Development of a turbulent boundary layer after a step from smooth to rough surface. *Exp Fluids* **51** (6), 1563–1575.
- FERNHOLZ, H. H. & FINLEY, P. J. 1996 The incompressible zero-pressure-gradient turbulent boundary layer: an assessment of the data. *Prog Aerosp Sci* **32** (4), 245–311.
- FERRANTE, A. & ELGHOBASHI, S. E. 2004 A robust method for generating inflow conditions for direct simulations of spatially-developing turbulent boundary layers. *J Comput Phys* **198** (1), 372–387.
- FERRE, J. A., MUMFORD, J. C., SAVILL, A. M & GIRALT, F. 1990 Three-dimensional large-eddy motions and fine-scale activity in a plane turbulent wake. *J. Fluid Mech.* **210**, 371–414.
- FIEDLER, H. E. & HEAD, M. R. 1966 Intermittency measurements in the turbulent boundary layer. *J. Fluid Mech.* **25**, 719–735.
- FISCHER, P., KRUSE, J., MULLEN, J., TUFO, H., LOTTES, J. & KERKEMEIER, S. 2008 Nek5000—open source spectral element cfd solver. *Argonne National Laboratory, Mathematics and Computer Science Division, Argonne, IL, see <https://nek5000.mcs.anl.gov/index.php/MainPage>* .
- FLACK, K. A. & SCHULTZ, M. P. 2014 Roughness effects on wall-bounded turbulent flowsa). *Phys Fluids* **26** (10), 101305.
- FLORES, O. & JIMÉNEZ, J. 2006 Effect of wall-boundary disturbances on turbulent channel flows. *J Fluid Mech* **566**, 357–376.

- FRINGS, W., WOLF, F. & PETKOV, V. 2009 Scalable massively parallel i/o to task-local files. In *High Performance Computing Networking, Storage and Analysis, Proceedings of the Conference on*, pp. 1–11. IEEE.
- GAMPERT, M., BOSCHUNG, J., HENNIG, F., GAUDING, M. & PETERS, N. 2014 The vorticity versus the scalar criterion for the detection of the turbulent/non-turbulent interface. *J. Fluid Mech.* **750**, 578–596.
- GAMPERT, M., NARAYANASWAMY, V., SCHAEFER, P. & PETERS, N. 2013 Conditional statistics of the turbulent/non-turbulent interface in a jet flow. *J. Fluid Mech.* **731**, 615–638.
- GARTSHORE, I. S. 1966 An experimental examination of the large-eddy equilibrium hypothesis. *J. Fluid Mech.* **24**, 89–98.
- GROUP, HDF *et al.* 2014 Hierarchical data format, version 5 .
- HAGEN, G. H. L. 1854 *Über den einfluss der temperatur auf die bewegung des wassers in röhren*-. Druckerei der Königl. akademie der wissenschaften.
- HARLOW, F. H. & WELCH, J. E. 1965 Numerical calculation of time-dependent viscous incompressible flow of fluid with free surface. *Phys Fluids* **8** (12), 2182.
- HOLZNER, M., LIBERZON, A., NIKITIN, N., LÜTHI, B., KINZELBACH, W. & TSI-NOBER, A. 2008 A lagrangian investigation of the small-scale features of turbulent entrainment through particle tracking and direct numerical simulation. *J. Fluid Mech.* **598**, 465–475.
- HOYAS, SERGIO & JIMÉNEZ, JAVIER 2006 Scaling of the velocity fluctuations in turbulent channels up to $re\tau = 2003$. *Phys Fluids* **18** (1), 011702.
- HUNT, J. C. R. & DURBIN, P. A 1999 Perturbed vortical layers and shear sheltering. *Fluid Dyn. Res.* **24**, 375–404.
- ISHIHARA, T., KANEDA, Y. & HUNT, J.C.R. 2013 Thin shear layers in high reynolds number turbulence—dns results. *Flow Turb. and Comb.* **91** (4), 895–929.
- JIMÉNEZ, J. 2004 Turbulent flows over rough walls. *Annu Rev Fluid Mech* **36** (1), 173–196.
- JIMÉNEZ, J. 2013 Near-wall turbulence. *Phys. Fluids* **25** (10), 101302.

- JIMÉNEZ, J., HOYAS, S., SIMENS, M. P. & MIZUNO, Y. 2010 Turbulent boundary layers and channels at moderate reynolds numbers. *J Fluid Mech* **657**, 335–360.
- JIMÉNEZ, J. & MOIN, P. 1991 The minimal flow unit in near-wall turbulence. *Journal of Fluid Mechanics* **225**, 213–240.
- JIMÉNEZ, J. & PINELLI, A. 1999 The autonomous cycle of near-wall turbulence. *Journal of Fluid Mechanics* **389**, 335–359.
- JIMÉNEZ, J. & SIMENS, M. 2001 The largest scales in turbulent flow: The structures of the wall layer. In *Coherent Structures in Complex Systems* (ed. D. Reguera, J. M. Rubí & L. L. Bonilla), *Lect Notes Phys*, vol. 567, pp. 39–57. Springer Berlin Heidelberg.
- JIMÉNEZ, J., WRAY, A. A., SAFFMAN, P. G. & ROGALLO, R. S. 1993 The structure of intense vorticity in isotropic turbulence. *J. Fluid Mech* **255**, 65–90.
- JONES, M., BAERENTZEN, J. A. & SRAMEK, M. 2006 3d distance fields: A survey of techniques and applications. *IEEE T. Vis. Comput. Gr.* **12** (4), 581–599.
- KHUJADZE, G. & OBERLACK, M. 2004 Dns and scaling laws from new symmetry groups of zpg turbulent boundary layer flow. *Theor comp fluid dyn* **18** (5), 391–411.
- KLEBANOFF, P. S. 1955 Characteristics of turbulence in a boundary layer with zero pressure gradient. *NACA TR* **1247**.
- KOLMOGOROV, A. N. 1941 The local structure of turbulence in incompressible viscous fluid for very large reynolds numbers. In *Dokl. Akad. Nauk SSSR*, , vol. 30, pp. 299–303.
- KOVASZNAY, L. S. G., KIBENS, V. & BLACKWELDER, R. F. 1970 Large-scale motion in the intermittent region of a turbulent boundary layer. *J. Fluid Mech.* **41**, 283–325.
- KROGSTAD, P-Å., ANDERSSON, H. I., BAKKEN, O. M. & ASHRAFIAN, A. 2005 An experimental and numerical study of channel flow with rough walls. *J Fluid Mech* **530**, 327–352.
- KROGSTAD, P-Å. & ANTONIA, R. A. 1994 Structure of turbulent boundary layers on smooth and rough walls. *J Fluid Mech* **277**, 1–21.
- KROGSTAD, P-Å. & EFROS, V. 2012 About turbulence statistics in the outer part of a boundary layer developing over two-dimensional surface roughness. *Phys Fluids* **24** (7), 075112.

- LEE, J. H. & SUNG, H. J. 2013 Comparison of very-large-scale motions of turbulent pipe and boundary layer simulations. *Phys. Fluids* **25** (4), 045103.
- LEE, J. H., SUNG, H. J. & KROGSTAD, P.-Å. 2011 Direct numerical simulation of the turbulent boundary layer over a cube-roughened wall. *J Fluid Mech* **669**, 397–431.
- LEE, S. H. & SUNG, H. J. 2007 Direct numerical simulation of the turbulent boundary layer over a rod-roughened wall. *J Fluid Mech* **584**, 125–146.
- LELE, S. K. 1992 Compact finite difference schemes with spectral-like resolution. *J Comput Phys* **103** (1), 16–42.
- LEUNG, T, SWAMINATHAN, N. & DAVIDSON, P. A. 2012 Geometry and interaction of structures in homogeneous isotropic turbulence. *J. Fluid Mech.* **710**, 453–481.
- LOZANO-DURÁN, A. & BORRELL, G. 2015 An efficient algorithm to compute the genus of discrete surfaces and applications to turbulent flows. *TOMS* p. (Submitted).
- LOZANO-DURÁN, A. & JIMÉNEZ, J. 2014 *a* Effect of the computational domain on direct simulations of turbulent channels up to $re\tau = 4200$. *Phys Fluids* **26** (1), 011702.
- LOZANO-DURÁN, A. & JIMÉNEZ, J. 2014 *b* Time-resolved evolution of coherent structures in turbulent channels: characterization of eddies and cascades. *J. Fluid Mech* **759**, 432–471.
- LUND, T. S., WU, X. & SQUIRES, K. D. 1998 Generation of turbulent inflow data for spatially-developing boundary layer simulations. *J Comput Phys* **140** (2), 233–258.
- MANDELBROT, B. B. 1975 On the geometry of homogeneous turbulence, with stress on the fractal dimension of the iso-surfaces of scalars. *J. Fluid Mech.* **72**, 401–416.
- MANDELBROT, B. B. 1983 *The fractal geometry of nature*, , vol. 173. Macmillan.
- MATHEW, J. & BASU, A. J. 2002 Some characteristics of entrainment at a cylindrical turbulence boundary. *Phys. Fluids* **14** (7), 2065–2072.
- MELLADO, J. P., WANG, L. & PETERS, N. 2009 Gradient trajectory analysis of a scalar field with external intermittency. *J. Fluid Mech.* **626**, 333–365.
- MIZUNO, Y. & JIMÉNEZ, J. 2013 Wall turbulence without walls. *J Fluid Mech* **723**, 429–455.

- MOIN, PARVIZ & MAHESH, KRISHNAN 1998 Direct numerical simulation: a tool in turbulence research. *Annual review of fluid mechanics* **30** (1), 539–578.
- MOISY, F & JIMÉNEZ, J. 2004 Geometry and clustering of intense structures in isotropic turbulence. *J. Fluid Mech.* **513**, 111–133.
- MUJA, M. & LOWE, D. G. 2014 Scalable nearest neighbor algorithms for high dimensional data. *IEEE T. Pattern Anal.* **11**, 2227–2240.
- MUNGAL, M. G, KARASSO, P. S. & LOZANO, A. 1991 The visible structure of turbulent jet diffusion flames: large-scale organization and flame tip oscillation. *Comb. Sci. Tech.* **76**, 165–185.
- MURLIS, J., TSAI, J. & BRADSHAW, P 1982 The structure of turbulent boundary layers at low reynolds numbers. *J. Fluid Mech.* **122**, 13–56.
- NIKURADSE, J. 1933 Laws of flow in rough pipes. *NACA TM* **1292**.
- ÖSTERLUND, J. M., JOHANSSON, A. V., NAGIB, H. M. & HITES, M. H. 2000 A note on the overlap region in turbulent boundary layers. *Phys Fluids* **12** (1), 1–4.
- PEROT, J. B. 1993 An analysis of the fractional step method. *J Comput Phys* **108** (1), 51–58.
- PERRY, A. E., SCHOFIELD, W. H. & JOUBERT, P. N. 1969 Rough wall turbulent boundary layers. *J Fluid Mech* **37** (02), 383–413.
- PIROZZOLI, S. & BERNARDINI, M. 2013 Probing high-reynolds-number effects in numerical boundary layers. *Phys. Fluids* **25** (2), 021704.
- PIROZZOLI, S., BERNARDINI, M. & GRASSO, F. 2010 Direct numerical simulation of transonic shock/boundary layer interaction under conditions of incipient separation. *J Fluid Mech* **657**, 361–393.
- PRANDTL, L 1904 Verhandlungen des dritten internationalen mathematiker-kongresses. *Heidelberg, Leipeizig* pp. 484–491.
- PRASAD, R.R. & SREENIVASAN, K.R. 1989 Scalar interfaces in digital images of turbulent flows. *Exp. Fluids* **7** (4), 259–264.

- RAUPACH, MR, ANTONIA, RA & RAJAGOPALAN, S 1991 Rough-wall turbulent boundary layers. *Applied Mechanics Reviews* **44** (1), 1–25.
- VAN REEUWIJK, M. & HOLZNER, M. 2014 The turbulence boundary of a temporal jet. *J. Fluid Mech.* **739**, 254–275.
- REYNOLDS, O. 1895 On the dynamical theory of incompressible viscous fluids and the determination of the criterion. *Phil Trans Roy Soc Lon A* pp. 123–164.
- RICHARDSON, L. F. 1922 *Weather prediction by numerical process*. Cambridge University Press.
- RUSS, J. C 1994 *Fractal Surfaces*. 233 Spring Street, New York, N.Y. 10013: Plenum Press.
- SANDHAM, N. D., MUNGAL, M. G., BROADWELL, J. E. & REYNOLDS, W. C. 1988 Scalar entrainment in the mixing layer. In *Proc. Summ.*, pp. 69–76. Center for Turbulence Research.
- SCHALLER, ROBERT R 1997 Moore’s law: past, present and future. *Spectrum, IEEE* **34** (6), 52–59.
- SCHLATTER, P. & ÖRLÜ, R. 2010 Assessment of direct numerical simulation data of turbulent boundary layers. *J. Fluid Mech.* **659**, 116–126.
- SCHROEDER, M. R. 2012 *Fractals, chaos, power laws: Minutes from an infinite paradise*. Courier Corporation.
- SCHULTZ, M. P. & FLACK, K. A. 2005 Outer layer similarity in fully rough turbulent boundary layers. *Exp Fluids* **38** (3), 328–340.
- SILLERO, J.A., BORRELL, B., GUNGOR, A.G., JIMÉNEZ, J., MOSER, R. D. & A., OLIVER T. 2010 Direct simulation of a zero-pressure-gradient turbulent boundary layer up to $re\theta=6650$. EB-4. College Park, MD: Am. Phys. Soc.
- SILLERO, J. A., JIMÉNEZ, J. & MOSER, R. D. 2013 One-point statistics for turbulent wall-bounded flows at reynolds numbers up to $\delta^+ \simeq 2000$. *Phys Fluids* **25** (10), 105102.
- SILLERO, J. A., JIMÉNEZ, J. & MOSER, R. D. 2014 Two-point statistics for turbulent boundary layers and channels at reynolds numbers up to $\delta^+ \simeq 2000$. *Phys Fluids* **26** (10), 105109.

- DA SILVA, C. B., HUNT, J. C. R., EAMES, I. & WESTERWEEL, J. 2014 *a* Interfacial layers between regions of different turbulent intensity. *Ann. Rev. Fluid Mech.* **46**, 567–590.
- DA SILVA, C. B., DOS REIS, R. J. N. & PEREIRA, J. C. F. 2011 The intense vorticity structures near the turbulent/non-turbulent interface in a jet. *J. Fluid Mech.* **685**, 165–190.
- DA SILVA, C. B. & TAVEIRA, R. R. 2010 The thickness of the turbulent/nonturbulent interface is equal to the radius of the large vorticity structures near the edge of the shear layer. *Phys. Fluids* **22**, 121702.
- DA SILVA, C. B., TAVEIRA, R. R. & BORRELL, G. 2014 *b* Characteristics of the turbulent/nonturbulent interface in boundary layers, jets and shear-free turbulence. *J. Phys.* **506** (1), 012015.
- SIMENS, M., JIMÉNEZ, J., HOYAS, S. & MIZUNO, Y. 2009 A high-resolution code for turbulent boundary layers. *J. Comp. Phys.* **228**, 4218–4231.
- SPALART, P. R. 1988 Direct simulation of a turbulent boundary layer up to $\text{Re}_\theta = 1410$. *J. Fluid Mech.* **187**, 61–98.
- SREENIVASAN, K. R. & MENEVEAU, C. 1986 The fractal facets of turbulence. *J. Fluid Mech.* **173**, 357–386.
- SREENIVASAN, K. R., RAMSHANKAR, R. & MENEVEAU, C. 1989 Mixing, entrainment and fractal dimensions of surfaces in turbulent flows. *Proc. Royal Soc. London A* **421** (1860), 79–108.
- TAVEIRA, R. R. & DA SILVA, C. B. 2014 Characteristics of the viscous superlayer in shear free turbulence and in planar turbulent jets. *Phys. Fluids* **26** (2), 021702.
- TOWNSEND, A. A. 1948 Local isotropy in the turbulent wake of a cylinder. *Austral. J. Sci. Res. A* **1**, 161–174.
- TOWNSEND, A. A. 1966 The flow in a turbulent boundary layer after a change in surface roughness. *J. Fluid Mech.* **26** (02), 255–266.
- TOWNSEND, A. A. 1976 *The structure of turbulent shear flow*. The Pitt Building, Trumpington Street, Cambridge CB2 1RP: Cambridge University Press.

- VOLINO, R. J., SCHULTZ, M. P. & FLACK, K. A. 2011 Turbulence structure in boundary layers over periodic two-and three-dimensional roughness. *J Fluid Mech* **676**, 172–190.
- WESTERWEEL, J., FUKUSHIMA, C., PEDERSEN, J. M. & HUNT, J. C. R. 2005 Mechanics of the turbulent-nonturbulent interface of a jet. *Phys. Rev. Lett.* **95**, 174501.
- WESTERWEEL, J., FUKUSHIMA, C., PEDERSEN, J. M. & HUNT, J. C. R. 2009 Momentum and scalar transport at the turbulent/non-turbulent interface of a jet. *J. Fluid Mech.* **631**, 199–230.
- WESTERWEEL, J., HOFMANN, T., FUKUSHIMA, C. & HUNT, J. 2002 The turbulent/non-turbulent interface at the outer boundary of a self-similar turbulent jet. *Exp. Fluids* **33** (6), 873–878.
- WOOD, N. & MASON, P. 1993 The pressure force induced by neutral, turbulent flow over hills. *Q. J. Roy. Meteor. Soc.* **119** (514), 1233–1267.
- WU, X. & MOIN, P. 2009 Direct numerical simulation of turbulence in a nominally zero-pressure-gradient flat-plate boundary layer. *J Fluid Mech* **630**, 5–41.
- WU, Y. & CHRISTENSEN, K.T. 2010 Spatial structure of a turbulent boundary layer with irregular surface roughness. *Journal of Fluid Mechanics* **655**, 380–418.
- WU, Y. & CHRISTENSEN, K. T. 2007 Outer-layer similarity in the presence of a practical rough-wall topography. *Phys Fluids* **19** (8), 85108–85108.