UNIVERSIDAD POLITÉCNICA DE MADRID

ESCUELA TÉCNICA SUPERIOR DE INGENIERÍA AERONÁUTICA Y DEL ESPACIO

DETECTION OF INTENSE EVENTS IN WALL-BOUNDED TURBULENCE

TESIS DOCTORAL

Miguel Pérez Encinar

Ingeniero Aeronáutico

February 2020

DEPARTAMENTO DE MECÁNICA DE FLUÍDOS Y MOTOPROPULSIÓN

DETECTION OF INTENSE EVENTS IN WALL-BOUNDED
TURBULENCE

AUTOR

Miguel Pérez Encinar
Ingeniero Aeronáutico

DIRECTOR

Javier Jiménez Sendín
Doctor Ingeniero Aeronáutico

February 2020

CONTENTS

CONTENTS	v
ACRONYMS	vii
NOMENCLATURE	ix
ABSTRACT	xiii
RESUMEN	xv
ACKNOWLEDGMENTS	xvii
1 INTRODUCTION	1
1.1 What is ‘turbulence’?	1
1.2 Historical remarks	3
1.3 Wall-bounded turbulence	7
1.4 Coherent motions in wall-bounded flows	12
1.5 Aims	16
1.6 Contents	16
2 LINEARISED EDDIES IN CHANNEL FLOWS	19
2.1 The linear model	22
2.1.1 Optimal transient growth	23
2.2 Methods	25
2.2.1 The numerical data sets	25
2.2.2 Filtering	26
2.2.3 The band-pass filtered pseudo-spectrum	29
2.2.4 Filter parameters	31
2.3 Band-pass-filtered amplitude and inclination fields	33
2.3.1 Wall-parallel velocities	37
2.3.2 Conditional mean evolutions	39
2.4 Low-pass-filtered velocity fields	45
2.4.1 Space-time organisation of sweeps and ejections	55
2.5 Conclusions	58
3 ESTIMATING LOGARITHMIC LAYER TURBULENCE	61
3.1 Numerical experiments	65
3.2 The relation between p and Q	66
3.3 Linear Stochastic Estimation	70
3.4 Estimations of Q using LSE	72
3.4.1 Three-dimensional reconstructions of Q	74
3.5 Estimations of the velocity fields using LSE	76
3.5.1 Incomplete wall data	82
3.5.2 Reconstruction in physical space	87
3.5.3 Off-design behaviour	90
3.5.4 Practical application to limited input data	91

3.6	Conclusions	92
4	DETECTION OF LINEARISED BURSTS FROM THE WALL	95
4.1	The filtered-reconstruction estimator	95
4.2	Results	98
4.3	Conclusions	100
5	DISCUSSION AND CONCLUSIONS	103
5.1	Complex ‘velocity’ fields	103
5.2	Orr bursts in the logarithmic region of channel flows	104
5.3	Logarithmic layer turbulence: the view from the wall	106
5.4	Future work	107
5.4.1	Local wave approximation	107
5.4.2	What are the events seeding the Orr bursts?	108
5.4.3	Multi-scale interactions of Qs-Orr bursts	108
5.4.4	Towards a quantitative definition of wall-attached	109
	APPENDIX	111
A	SECOND-ORDER CONSISTENT DATABASES	113
A.1	Methods	118
A.1.1	Simulation algorithm	118
A.1.2	Enforcing boundary conditions at high Re	119
A.1.3	Code	120
A.2	A second-order consistent, time-resolved database	121
A.2.1	Methodology	121
A.2.2	Simulations	122
A.2.3	One-point statistics and energy spectra	123
A.2.4	Temporal statistics	125
A.3	Conclusions	127
	BIBLIOGRAPHY	131

ACRONYMS

NSE Navier–Stokes equations

RANS Reynolds-Averaged Navier–Stokes

DNS Direct Numerical Simulation

LES Large Eddy Simulation

LSE Linear Stochastic Estimation

HIT Homogeneous Isotropic Turbulence

UPO Unstable Periodic Orbit

POD Proper Orthogonal Decomposition

CMV Conditional Mean Velocity

NOMENCLATURE

CONSTANTS/ENTITIES

u_*	Characteristic velocity
ℓ	Characteristic length
h	Channel half-height
u_τ	Friction velocity
Re_τ	Friction Reynolds number
S^*	Corrsin shear parameter
κ	Kármán constant
$L_\bullet$	Channel box-length in the direction $\bullet$
$N_\bullet$	Number of points in the direction $\bullet$
Λ	Filter width or pseudo-wavelength
$\mathcal{E}$	Pseudo-spectrum
C	Convection velocity
Q_{2-4}	Pair of sweep and and ejection based on Reynolds stress quadrant analysis

OPERATORS

∇	Gradient operator
$\langle \bullet \rangle$	Ensemble average
Δ	Laplacian operator. Equivalently, ∇^2
$(\bullet)^*$	Complex conjugation (Hermitian transpose if $\bullet$ is a matrix)
$(\bullet)^+$	Wall (inner) units
$\widetilde{\bullet}$	Low-pass filtered magnitude
$\overline{\bullet}$	Band-pass filtered magnitude
$\frac{D\bullet}{Dt}$	(semi-Lagrangian) material derivative
$\langle \bullet \rangle_B$	Conditional evolution to the presence of a burst

$(\bullet)^\dagger$ Reconstruction of $\bullet$ from wall observables using LSE

VARIABLES

ε	Turbulent dissipation per unit volume and time
ν	Kinematic viscosity
$\mathbf{u}$	Velocity vector
$\mathbf{U}$	Mean velocity vector
$\mathbf{q}$	Fluctuating velocity vector
p	Modified (incompressible) pressure
η	Kolmogorov scale
S	Mean shear
k	Spatial wavenumber
τ	Total tangential stress
ϕ	Laplacian of the vertical velocity
ω	Vorticity vector ($\omega_\bullet$ is component in the direction $\bullet$)
$A_\bullet$	Amplitude of the band-pass filtered velocity component $\bullet$
$A_{\bullet M}$	Modal value of the band-pass-filtered amplitude $A_\bullet$
$\Psi_\bullet$	Inclination angle of the band-pass filtered velocity component $\bullet$
S_Λ	Effective shear for a given filter wavelength
ξ	Moving frame across the streamwise velocity that follows the mean convection

ABSTRACT

The objectives of this thesis are the detection and characterisation of intense events in turbulent wall-bounded flows, and to elucidate whether these events are controlled by quasi-linear dynamics or not. Specifically, to determine if ‘Orr bursts’, eddies whose dynamics can be described using the linear Orr–Sommerfeld equation, appear in fully nonlinear channel flows. Channel flows constitute the simplest wall-bounded flow and thus can be used as a model of the dynamics of turbulence in the vicinity of more general solid surfaces, like boundary layers or pipes.

The first part of this thesis focuses on the detection and characterisation of these events using full information from the flow. The velocity fields of time-resolved databases of channel flows are filtered using both complex band-pass filters, and real-valued low-pass filters. The former allow the computation of local amplitudes and local phases of the turbulent field within a range of wavenumbers, making direct comparisons with the wave-based linear theory possible. The latter bridge the gap between the localised wave-packets obtained from the band-pass filtered fields and turbulent eddies in physical space. The Orr bursts are found to populate the shear-dominated regions of the channel flow, with self-similar sizes respect to the wall, and time scales controlled by the mean shear. These events collaborate at sustaining the streamwise velocity streaks. The combination, and consistent results obtained from both filtering approaches, reveals striking similarities with the Q_{2-4} pairs responsible for momentum transfer in channel flows, strongly suggesting that the dynamics of the largest tall-attached eddies responsible for momentum transfer are well described by quasi-linear Orr bursts.

Second, the possibility of sensing the flow field far away from the wall is explored, using contemporaneous measurements of the three stresses at the wall. The equation for the pressure at the wall is investigated, uncovering that only a small portion of the flow is responsible for generating said pressure, and that the best possible estimation of the flow that can be inferred exclusively from the equation is unphysical, as it does not consider other constraints of the Navier–Stokes equations (NSE). The equation does however suggest that reasonable, data-driven reconstructions of the source term are possible by means of the cross-correlations of the source term and the pressure, which are also related to each other by the linear equation. The form of this reconstruction is found to be equivalent to Linear Stochastic Estimation (LSE), allowing the reconstruction of other quantities that, unlike the source term of the pressure, do not have a linear equation linking them to the wall, e. g. the velocities. Reconstructions of the source term and the velo-

cities are computed, and their properties show that only large eddies which scale with the distance to the wall are observable, suggesting a quantifiable definition of attached eddies.

Finally, the possibility of using [LSE](#) to detect Orr bursts from wall measurements is explored. The intense events can be detected using a complex band-pass filtered linear operator, similar to the one used to reconstruct the flow. The most intense events can be detected as they happen, but cannot be predicted with short anticipation, which requires full information from the flow. Even though the nature of their linearities differ, the increased performance of the linear reconstructions in the flow regions which correspond to linear Orr events reassures that this events are reasonably controlled by the mean shear.

RESUMEN

Los objetivos de esta tesis son la detección y caracterización de eventos intensos en flujos turbulentos delimitados por paredes, y elucidar si estos eventos están controlados por dinámicas cuasilineales o no. Específicamente, determinar si los “Orr bursts”, remolinos cuya dinámica se puede describir usando la ecuación lineal de Orr–Sommerfeld, aparecen en flujos de canales completamente no lineales. Los flujos de canal constituyen el flujo más simple limitado por paredes y, por lo tanto, se pueden emplear como modelo de la dinámica de la turbulencia en la vecindad de superficies sólidas más generales, como capas límite o tuberías.

La primera parte de esta tesis se centra en la detección y caracterización de estos eventos, utilizando información completa del flujo. Los campos de velocidad de flujos de canal almacenados en bases de datos resueltas en el tiempo se filtran, utilizando tanto filtros de pasa banda complejos como filtros de pasa bajos reales. Los primeros permiten el cálculo de amplitudes locales y fases locales del campo turbulento, a la vez que aíslan un rango de números de onda, haciendo posibles comparaciones directas con la teoría lineal, que está basada en ondas. Los filtros pasa bajos cierra la brecha existente entre los paquetes de ondas localizados obtenidos de los campos filtrados con el pasa banda y los remolinos turbulentos en el espacio físico. Se demuestra que los “Orr bursts” ocupan las regiones dominadas por la cortadura media del canal, con tamaños autosemejantes respecto de la distancia a la pared, y con escalas de tiempo controladas por la cortadura medio. Estos eventos colaboran para mantener las ‘streaks’ de la perturbación de la velocidad en el sentido del flujo. La combinación y los resultados consistentes obtenidos con ambos filtros revelan similitudes sorprendentes con las parejas Q_{2-4} , responsables de la transferencia de momento en los flujos de pared; lo que sugiere fuertemente que los remolinos más grandes asociados a la transferencia de momento están razonablemente bien descritos con la dinámica de los “Orr bursts” cuasilineales.

En segundo lugar, se explora la posibilidad de detectar el campo de velocidades lejos de la pared, utilizando mediciones contemporáneas de los tres esfuerzos limitados a esta. Se investiga la ecuación para la presión en la pared, descubriendo que solo una pequeña porción del flujo es responsable en generar dicha presión, y que la mejor estimación posible del flujo que puede inferirse exclusivamente de la ecuación no es física, al no considerar otras restricciones de las Navier–Stokes equations (NSE). Sin embargo, la ecuación sugiere que es posible obtener reconstrucciones razonables mediante las correlaciones cruzadas del término fuente y la presión, que también es-

tán relacionadas entre sí por la ecuación lineal. Se encuentra que la forma de esta reconstrucción es equivalente a Linear Stochastic Estimation ([LSE](#)), permitiendo así la reconstrucción de otras cantidades que, a diferencia del término fuente de la presión, no tienen una ecuación lineal ligándolas a la pared, como por ejemplo, las velocidades. Se calculan las reconstrucciones del término fuente y de las velocidades, y sus propiedades muestran que sólo los remolinos grandes que se escalan con la distancia a la pared son observables, lo que sugiere una definición cuantificable de remolinos ‘attached’ (ligados) a la pared.

Finalmente, se explora la posibilidad de usar [LSE](#) para detectar “Orr bursts” a partir de mediciones en la pared. Los eventos intensos pueden detectarse utilizando un operador lineal filtrado con el filtro pasa banda complejo, similar al utilizado para reconstruir el flujo. Los eventos más intensos se pueden detectar a medida que ocurren, pero no se pueden predecir con poca anticipación, debido a que esto último requiere información completa del flujo. Aunque la naturaleza de sus linealidades difiere, el rendimiento de las reconstrucciones lineales en las regiones de flujo que corresponden a eventos de Orr lineales es mejor que en el fondo turbulento, corroborando que estos eventos están razonablemente controlados por la cortadura media.

ACKNOWLEDGEMENTS

I would like to thank the people who are somehow reflected in this work, in no particular order.

Family and friends are never properly acknowledged. They are the only collective of people who are always there, unconditionally rooting for you, and offering invaluable support when it is needed. I'm indebted to all of them for respecting my aloofness through all of these years.

I would like to thank the people responsible for a tangible portion of my knowledge: my teachers. An excellent book is a poor substitute for a good teacher and, thankfully, I have fond memories of many of the latter. Special mentions are in order for my first teachers, my grandfather José María, who taught me how to read; and my sister Irene, whose tough lessons in algebra seeded my talent for maths at a tender age. I cannot thank my father Rafael for any lesson in particular, but he is without any doubts, the person I have learnt the most from.

This work would have not been the same without the labmates and colleagues I have meet through the years. I have shared with them many moments during this thesis, both personal and professional, and a piece of some of them can be found in the pages of this dissertation.

If my colleagues are in some of these pages, Marta is in all of them. If people are to be judged by who stands at their side, I can walk tall with good margins. Thanks for giving sense to otherwise useless thoughts.

I am grateful to my advisor, Javier Jiménez. The most valuable lesson I have learnt from him is that the wisdom of any individual should never be assumed, and that you can learn something new from anyone, regardless of their long career or lack thereof. It is remarkable that someone with such a long trajectory in research is still willing to learn new things from anyone, including undergrad students. I have also learnt wall turbulence from him, but that is far less valuable knowledge.

Finally, I would like to thank the European Research Council. This work was supported by the COTURB project of the European Research Council (ERC2014.AdG-669505), and without their grant, my experience during this thesis would have been completely different.

*To my crewmates, each individually, and as a whole.
To who they are, and to what they represent.*

It's the questions we can't answer that teach us the most. They teach us how to think. If you give a man an answer, all he gains is a little fact. But give him a question and he'll look for his own answers.

— *The Wise Man's Fear*, PATRICK ROTHFUSS

INTRODUCTION

In this chapter we provide a concise introduction to the turbulence phenomenon. Turbulence, as an ‘open’ centenary problem, is a sufficiently vast topic that allows for an introduction spanning several many-paged volumes; or at least a considerably large one [172]. Instead, here we offer a short description of the features of turbulence directly related to the matters discussed in the remaining chapters. We refer the reader to any of the many available monographs on the topic [15, 33, 95, 129, 166, 173] for a detailed description of turbulence as a phenomenon, and as an open problem.

1.1 WHAT IS ‘TURBULENCE’?

Most people, including those without any scientific background, have an intuitive notion of what turbulence is. For some, it evokes swirls in the oceans or in the atmosphere, whereas others may have observed it in waterfalls or in the smoke of a cigarette. More universal is perhaps the fear suffered by most when the term is invoked in a flight overseas. On the contrary, other flows are known to move smoothly or in the ‘laminar’ regime. For example, very viscous flows such as honey appear uniform when stirred, and barely opened taps produce a crystal-clear stream of water. If the latter are examples of flows moving ‘slowly’, the former show different aspects of the same phenomenon, which can be observed when fluids are moving fast enough, and may serve as an *intuitive* definition of turbulence.

However, a *scientific* definition of turbulence requires universal concepts capable of representing the reproducible characteristics of the previous examples. Even if turbulence research pioneered in the mid XIX century [32, 54], a satisfactory definition of the phenomenon, even if just qualitatively, has been elusive ever since and still remains a subject of controversy within the research community. Here we list some of features that are widely accepted as characteristics of turbulent flows, but the reader should be aware that some of them may be controversial for other scientists [166, 173].

DISSIPATION Turbulent flows are highly dissipative, and thus capable of thermalising more kinetic energy into heat than their laminar counterparts. Owing to their high dissipation, turbulent flows are responsible for approximately 25% of the world’s energy consumption [67], as most flows that are transported through pipes, such as water or oil, are in the turbulent regime, as well as the flows surrounding our vehicles

are mostly turbulent. Unlike in laminar flows, dissipation in turbulent flows is extremely robust and independent of most of the characteristics of the actual flow, when properly normalised [172].

CHAOS Turbulent flows are not smooth, but highly irregular or ‘chaotic’. Note that this is different from randomness, as turbulent flows are believed to satisfy the deterministic Navier–Stokes equations (NSE). In a turbulent flow, the same state will always evolve into the same flow, but a small perturbation of the initial conditions will give a completely different outcome [144]. It is also worth noting that chaos is an intrinsic property of the equations, i.e. no external stochastic forcing is required for chaos to develop. A direct consequence of their chaotic nature is the loss of dependence of the statistical properties of turbulent flows from their initial conditions. In a general sense, different realisations of the same flow, e.g. water through a pipe of a given diameter and mass flux, will have the same statistical properties even if the detail of velocity patterns are completely different across experiments.

DIFFUSIVITY Turbulence is transport. Turbulence strongly enhances diffusion of momentum, energy, temperature, or passive objects (mixing). This is a necessary but not sufficient property of turbulent flows, as many laminar flows can have turbulent-like mixing properties over passive scalars, e.g. different species or small particles can be mixed as fast by certain laminar flows as if they were turbulent [172].

SEPARATION OF SCALES In the atmospheric boundary layer, for example, turbulent perturbations range from hundreds of kilometres to several micrometres. Several of these different scales interact with each other, often exchanging momentum or energy among them. The large separation of interacting scales is one of the key aspects making the study of turbulence an extremely complicated subject.

NON-LINEARITY Turbulence is a highly non-linear phenomenon. Non-linear interactions are responsible for sustaining turbulent processes across scales. Noteworthy is the capability of the velocity gradients to amplify themselves and interact non-locally, which is the ultimate source of all the previous properties.

All of these characteristics are universal of fully developed turbulent flows, although specific differences exist for different flows, like boundary layers, jets or accretion disks.

I.2 HISTORICAL REMARKS

For a monograph on this topic, the reader is referred to Davidson et al. [34] and references therein.

Interest in turbulence predates modern science, and references as old as the work of Leonardo da Vinci (c. 1500) exist. However, a formal scientific study had to wait until the experiments of Hagen [54, 55] and Darcy [32], who independently found two different flow regimes in duct flows. In the first regime, the pressure drop was linear with the velocity and the kinematic viscosity of the fluid, while in the second the pressure drop was quadratic with the velocity and independent of the viscosity. Note that this second regime poses a striking implication, as the turbulent dissipation ε per unit volume and time in an incompressible flow is,

$$\varepsilon = \nu \|\nabla \mathbf{u}\|^2, \quad (1.1)$$

where ν is the kinematic viscosity, ∇ is the gradient operator, $\|\bullet\|^2$ is the squared L_2 -norm and $\mathbf{u}$ is the velocity vector. Equation (1.1) implies that if the dissipation remains constant in the limit of vanishing viscosity, the gradients have to become unboundedly large; a property which is often called the *dissipation anomaly* of turbulent flows [173]. Following the experiments of Hagen and Darcy, Boussinesq [22] distinguished for the first time the ‘turbulent’ (turbulent) and ‘smooth’ (laminar) regimes of fluid flows. He also observed that the former had enhanced momentum transport and dissipation, and introduced for the first time the concept of an *eddy viscosity*, which linearly accounted for the transport enhancement owing to the turbulent strain field. Towards the end of the century, it was Osborne Reynolds [136, 137] who made some important advances in the field, most still relevant today. His most important contribution is the identification of the critical parameter which controls the transition to turbulence from laminar regime in incompressible flows. This non-dimensional quantity that now bears his name, the *Reynolds number*, is the ratio between the characteristic time of inertial and viscous dynamics,

$$Re = \frac{u_* \ell}{\nu}, \quad (1.2)$$

where u_* is a characteristic velocity and ℓ is a characteristic length of the flow. When Re is large, inertial dynamics dominate the flow motion, and flows eventually become turbulent. Reynolds also realised that the turbulent regime was too complex to pursue a detailed description of the flow, and thereby a statistical description was more useful. Therefore, he introduced the *Reynolds decomposition* of the velocity field respect to the mean flow,

$$\mathbf{u} = \langle \mathbf{u} \rangle + (\mathbf{u} - \langle \mathbf{u} \rangle) = \mathbf{U} + \mathbf{q}, \quad (1.3)$$

where $\mathbf{U}$ is the mean velocity, $\mathbf{q}$ is the fluctuating velocity and $\langle \bullet \rangle$ is the ensemble average. When trying to pose equations for $\mathbf{U}$, known as the Reynolds-Averaged Navier–Stokes (RANS) equations, he realised that they are similar in nature to the NSE, but contain an additional term, the *Reynolds stress tensor* $\langle \mathbf{q}^T \mathbf{q} \rangle$. Because the Reynolds stresses cannot be determined solely from the mean components, a *closure problem* arises, and models are required to obtain an autonomous set of equations for the average flow. After the works of Reynolds, several authors, with Kármán [78], Prandtl [130] and Taylor [163] among them, proposed closure models inspired by the kinetic theory of gases, built upon an analogy between turbulent eddies in the flow and molecules in a gas. These *mixing length* models estimate the length and velocity scales which are required for an eddy viscosity, as proposed by Boussinesq [22].

Almost contemporaneously, Rayleigh [134], Orr [124] and Sommerfeld [154] studied for the first time the linear stability of shear flows, in an attempt of linking potential instabilities of the mean flow to the transition to turbulence. The field of linear stability analysis has succeeded at explaining several properties of shear flows ever since, and we will see that some of their predictions are central to this thesis.

At the beginning of the xx century, Richardson [139] proposed a very successful multi-scale hierarchical model of turbulence. The model postulates that energy is introduced in the flow at large scales or in large eddies, which then break into smaller ones due to turbulent (non-linear) dynamics. This process is repeated successively until the eddies are small enough to be dissipated viscously. He famously summarised the model in *Weather prediction by numerical process* [140], in the form of a short rime,

*Big whirls have little whirls that feed on their velocity,
and little whirls have lesser whirls and so on to viscosity.*

This phenomenological picture of turbulence remains (after some formalisation) as one of the widest accepted models for macroscopic energy transfer in turbulent flows.

In the 30s, Taylor [164, 165] introduced several useful concepts, which are still relevant in turbulence research. One practical approximation, used in experiments nowadays, is the *Taylor’s frozen turbulent hypothesis*, that relates temporal to spatial derivatives, assuming that the passing time of turbulent eddies due to the advection by the mean flow is considerably faster than their internal dynamics. He also introduced the use of velocity correlations and Fourier transforms, and defined the concept of Homogeneous Isotropic Turbulence (HIT), showing that this flow is statistically similar to the flow in a wind tunnel after going through a grid. HIT is a turbulent flow where every statistic is translation, rotation and reflection invariant. Its importance in turbulent research is twofold. First, is the simplest form of turbulence

that carries all of the characteristics enumerated in §1.1. Second, for scales sufficiently small, is believed¹ that most turbulent flows are equivalent to HIT. From the concept of HIT, Kármán and Howarth [79] formulated an equation that describes the energy transfer between different wavelengths, known as the Kármán–Howarth equation.

In the following decade, and upon the concepts foreseen in Richardson’s model, and the Kármán–Howarth equation; the works of Kolmogorov [88, 89] (usually abbreviated as ‘K41’) produced one of the very few successful predictions in turbulence: The $-5/3$ slope of the turbulent energy spectra. Consider that Richardson’s model holds, and that we are dealing with an HIT flow in statistically steady state. In this flow, the energy is injected at (large) scales of size $\sim \ell_E$, e.g. in the aforementioned grid-in-a-wind-tunnel experiment, ℓ_E would be the characteristic length of the grid. A turbulent eddy with characteristic size ℓ_I , and characteristic velocity u_I , breaks in times of the order of its eddy turnover time, $t_I \sim \ell_I/u_I$. Following Richardson’s, this process is repeated each step, and the energy transferred in each step is $u_I^2/t_I = u_I^3/\ell_I$. Because the flow is in statistical equilibrium, the average energy transfer has to be balanced by the average dissipation (1.1),

$$\varepsilon \sim \frac{u_I^3}{\ell_I} \longrightarrow u_I \sim (\varepsilon \ell)^{1/3}, \quad (1.4)$$

which gives a relation between the characteristic size and velocities of these eddies. The break-up process cannot hold forever. Every step divides the characteristic length by some factor, and at some point the eddy turnover time t_I is comparable to the viscous time scale $t_v \sim \ell^2/\nu$; which implies that the eddy dissipates its energy before cascading forward and thus the cascade process ends. Equating both times gives,

$$t_\eta = t_v \rightarrow \frac{\eta}{u_\eta} = \frac{\eta^2}{\nu} \xrightarrow{(1.4)} \frac{\eta}{(u_\eta \varepsilon)^{1/3}} \rightarrow \eta = \left(\frac{\nu^3}{\varepsilon} \right), \quad (1.5)$$

where η is the *Kolmogorov scale*. For the scales ℓ_I to remain inviscid, $\ell_I \gg \eta$; and they are called *inertial* when they also far from the energy-injection (or *integral*) length scale, $\ell_E \gg \ell_I$. Using dimensional analysis, it can be seen that the energy spectra of the inertial range follows,

$$\Phi_{u_\ell u_\ell} \sim \varepsilon^{2/3} k^{-5/3}, \quad (1.6)$$

where $k \propto 1/\ell$ is the spatial wavenumber. As anticipated, (1.6) is one of the most sound predictions obtained in turbulence research, even if most of the assumptions preceding it are now known to be inaccurate [16]. Kolmogorov’s

¹ Although never proved, and controversial [92], is generally accepted that for sufficiently large scale separation, i.e. high enough Re , scales far away from the energy-injection range should be roughly isotropic [88].

model was later used by Onsager [123] to show that in the limit of vanishing viscosity, (1.1) gives constant dissipation, consistently with the observations of Hägen and Darcy among others.

Kolmogorov's theory reasonably predicts the scaling of the energy spectra across the 'inertial' range of scales, which is a second order moment, but experiments showed that higher moments are not predicted correctly [16]. The reason for the poor scaling of the high-order moments is that flows are *intermittent*, i.e. non-Gaussian, which is not accounted for in the original cascade model. More refined models of turbulent cascades have been proposed, including Kolmogorov [90] which modifies the original theory to account for intermittency. Another successful branch of models account for intermittency using *multiplicative cascades*, proposed by Frisch et al. [50] and refined by many [49, 63, 113]. Another successful family of models are the *shell models*, which sacrifice locality in space for simplicity [19].

During the 1960–70s, several important advances were made. First, the fast development of experimental techniques allowed the visualisation of *coherent motions* in both closed [87] and open flows [23]. Up to this point, the only descriptions of turbulence were statistical, disregarding the flow structure, and the realisation that organised motions exist within turbulent flows was paradigm-breaking.

In the 70-80s, the growth in computing power allowed Orszag and Patterson [125] to perform the first Direct Numerical Simulation (DNS). Later, Rogallo [142] increased the Reynolds number of Orszag and Patterson, providing the first simulation of turbulence with an incipient scale separation. In the same decade, Kim et al. [85] and Spalart [156] produced the first simulations of canonical wall-bounded flows, which clarified many of the observed features of wall-bounded flows despite their marginally turbulent Reynolds numbers. The quest for a larger Reynolds number (which requires larger simulations) lying in the 'experimental' range has not stopped ever since. In the context of HIT, Jiménez et al. [76] increased Rogallo's grid by a factor of 4, from 128^3 to 512^3 degrees of freedom, and a decade later Yokokawa et al. [185] elevated it up to 4096^3 . For wall-bounded flows, a similar story holds, with the works of Moser et al. [119], Álamo and Jiménez [6], Hoyas and Jiménez [57], Lee and Moser [93], in channel flows; and Jiménez et al. [72] and Sillero et al. [152] in boundary layers. Nowadays the Reynolds numbers that simulations can achieve are comparable to experiments in many cases, and the challenges of DNS are more related to domain sizes [102] or storing time-resolved data efficiently [103, 174].

The last decades have seen many fruitful advances, and many different approaches have been used to try to understand turbulence dynamics in a simpler way. Standing over the concepts of coherent motions, Cardesa et al. [26] studied intense energy events locally in time, scale and space—

5 dimensions—, showing that the energy transfer (cascade) processes are reasonably local across all the dimensions. A different approach to turbulent research, often called *dynamical systems approach* has been very active in the last decades. The basic idea consists in reinterpreting a turbulent flow as a highly dimensional system evolving in time as a point in the space of all the degrees of freedom. A famous early example exists in the works of Lorenz [98], which constructed a extremely reduced, forced, convective system of just 3 degrees of freedom with chaotic properties. One particularly successful application of this approach to transition-to-turbulence is the computation of exact periodic orbits, usually known as Unstable Periodic Orbit (UPO). Although these solutions exist outside of the turbulent attractor [41], researchers believe that the flow spends nontrivial time in the vicinity of these orbits [29]. Since the early computations by Kawahara and Kida [80], many UPOs have been computed for several flow configurations [12, 31, 35, 52, 73, 149], although their dynamical relevance is still a matter of debate.

I.3 WALL-BOUNDED TURBULENCE

In the previous section we intentionally cast aside the historical advances concerning turbulent wall-bounded flows.

Wall-bounded flows are central to this thesis, and we considered appropriated reserving a section for them.

Most turbulent flows in nature harvest their energy from the velocity difference between two regions, so they are *shear* flows. In many practical applications, the source of the shear is a solid, impermeable surface, making them fall in the category of *wall-bounded* flows. Owing to their importance in engineering, they have been studied from the beginning of turbulent research, but due to the intrinsic complications of turbulence many questions remain open. The most important theory of wall bounded flows consists on a collection of predictions for the first order statistics of wall-bounded flows, and is colloquially known as the *classical theory* of wall bounded turbulence. The details and the historical developments of the theory can be found in any monograph about wall-bounded flows [129, 166, 171], and here we will only show the predictions the theory makes about channel flows, even if most of them are shared with boundary layers and pipe flows. Channel flows, as sketched in figure 1.1, are arguably the simplest domain of the three, consisting in the flow between two parallel infinite plates, sustained by a constant pressure gradient. Its laminar counterpart is the classical Poiseuille flow, and some authors also prefer this name for the turbulent case.

The assumptions made in the remainder of this sections apply to channels whose half-height is h and the streamwise, wall-normal and spanwise

FIGURE I.1: Top: Sketch of the channel (Poiseuille) flow. Bottom: xy -section of a channel flow at large Re [174].

directions are x , y and z respectively. The streamwise direction is defined by the constant pressure gradient, and thus is the only direction with mean flow,

$$\mathbf{U} = (U(y), 0, 0). \quad (1.7)$$

Note that the mean streamwise profile only depends on the wall-normal coordinate as both the streamwise and spanwise directions are homogeneous. The perturbation velocities, $\mathbf{q} = (u, v, w)$, are defined following (1.3), and for v and w , they also constitute the total velocity. Due to the homogeneity of the wall-parallel directions and ergodicity² in time, any ensemble average in channel flows is equivalent to averages in the wall-parallel directions and time, $\langle \cdot \rangle \equiv \langle \cdot \rangle_{x,z,t}$, and will be used as our ensemble average herein, unless explicitly stated otherwise. Incompressible channel flows follow the incompressible NSE,

$$\partial_t \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p + \nu \Delta \mathbf{u}, \quad (1.8a)$$

$$\nabla \cdot \mathbf{u} = 0, \quad (1.8b)$$

² Although it is widely accepted that channel flows are ergodic, to the authors' knowledge this has never been proved.

where p is the quotient between the pressure and the constant density and Δ is the Laplacian operator acting on the velocity vector. Introducing (1.7) in (1.8a), and taking the ensemble average,

$$\partial_y \langle uv \rangle = -\partial_x P + \nu \partial_y^2 U, \quad (1.9a)$$

$$\partial_y \langle v^2 \rangle = -\partial_y P, \quad (1.9b)$$

where P is the mean modified pressure. The term $-\langle uv \rangle$ is usually known as the *tangential Reynolds stress*, and is responsible for the shape of the mean velocity profile, whereas the pressure gradient is responsible for the total mass flux. Integrating (1.9b),

$$P_*(x) = P + \langle v^2 \rangle, \quad (1.10)$$

where P_* is an integration constant depending only in the streamwise direction, as the term $\langle v^2 \rangle$ is homogeneous in x ,

$$d_x P_* = \partial_x P + \partial_x \langle v^2 \rangle, \quad (1.11)$$

and the average pressure gradient $\partial_x P$ is independent of y . Substituting (1.10) into (1.9a) and integrating across the wall-normal direction gives,

$$\langle uv \rangle = -y d_x P_* + \nu \partial_y U - \tau_*, \quad (1.12)$$

where τ_* is the integration constant. Particularising (1.12) at the wall $y = 0$ shows that the constant is equal to the average stress at the wall, $\tau_* = \nu(\partial_y U)|_{y=0}$. The constant $d_x P_*$ can be obtained by particularising the equation at $y = h$,

$$d_x P_* = \frac{\tau_*}{h}, \quad (1.13)$$

where the symmetry in y of the channel makes the velocity terms vanish. Removing the constant from (1.12) yields,

$$\tau = -\langle uv \rangle + \nu \partial_y U = \tau_* \left(1 - \frac{y}{h} \right). \quad (1.14)$$

The total tangential stress τ is linear with the distance to the wall. Because of the role of τ_* in (1.14) and the linearity of the equation, the square root of the stress, the *friction velocity* u_τ , serves as the natural velocity scale of channel flows. It follows from (1.14) that the other two scaling variables are h and ν . The three quantities define the parameter,

$$Re_\tau = \frac{u_\tau h}{\nu}, \quad (1.15)$$

known as the friction Reynolds number, which represents the ratio between the small friction-governed scales in the channel and the largest motions,

and serves as the primary Reynolds number in channel flows. From the three parameters, two different nondimensional groups can be formed. The first, with u_τ , ν/u_τ and ν/u_τ^2 as velocity, length and time scales respectively is known as *wall units* or inner units, and widely represented with a superindex $+$. These units are appropriated to scale motions whose dynamics are governed by viscosity or the friction at the wall. The second group is formed by u_τ , h and h/u_τ , and is called *outer units*. Outer units are the natural scales of motions not affected by the friction at the wall, and thus governed mostly by inertial dynamics. The ‘unit’ time scale in outer units also receives the name of *eddy turnover time*, and it represents the typical time scale of the largest motions found in channels.

The total stress balance (1.14) can be normalised with both sets of units in order to study the asymptotic behaviour of channels when $Re_\tau \gg 1$. Starting with the behaviour close to the wall, (1.14) can be nondimensionalised in inner units,

$$-\langle u^+ v^+ \rangle + \partial_{y^+} U^+ = 1 - \frac{y}{Re_\tau}. \quad (1.16)$$

At the wall, impermeability and no-slip imply that $\langle u^+ v^+ \rangle = 0$ and $U^+ = 0$, allowing the integration of (1.16) in its vicinity,

$$U^+ = f(y^+) + \mathcal{O}(1/Re_\tau), \quad (1.17a)$$

$$\langle u^+ v^+ \rangle = g(y^+) + \mathcal{O}(1/Re_\tau), \quad (1.17b)$$

and close to the wall, $\langle u^+ v^+ \rangle \approx 0$, which simplifies (1.17a) to,

$$U^+ \approx y^+. \quad (1.18)$$

These results are known as the *law of the wall* [130] and have been validated experimentally [121, 135]. The region where (1.18) holds, limited in practice to $y^+ \lesssim 5$, is known as the *viscous sublayer* and the region close to the wall, where viscosity is relevant (below $y^+ \lesssim 60$), as the *buffer layer*. The relations (1.17) are valid approximately up to $y/h \approx 0.2$.

In order to study the behaviour of the total stress far away from the wall, equation (1.14) can be normalised using outer units,

$$-\langle u^+ v^+ \rangle + \frac{1}{Re_\tau} \partial_y U^+ = 1 - y/h, \quad (1.19)$$

which in the limit $Re_\tau \rightarrow \infty$ simplifies to

$$-\langle u^+ v^+ \rangle \approx 1 - y/h, \quad (1.20)$$

which in practice is valid for $y^+ \gtrsim 100$. While (1.20) does not provide information about the mean profile, it follows from it that far away from the wall viscosity is irrelevant for the dynamics thereof. Kármán [78] showed that

$$U^+ - U_c^+ = F(y/h), \quad (1.21)$$

constructed using dimensional arguments. Equation (1.21) is known as the *velocity defect law*. The relation (1.17a) is valid when $y \ll h$, whereas (1.21) is valid if $y^+ \gg 1$. Regardless of the empirical values of the validity limits for these laws, at high Reynolds numbers there must exist a region far enough from the wall where the mean profile is dominated by inertial forces but small enough to be unaffected by the channel height. This region is known as the *intermediate region*. In this layer, both (1.17a) and (1.21) are valid, and hence,

$$U = u_\tau f(y^+), \quad (1.22a)$$

$$U = U_c + u_\tau F(y/h). \quad (1.22b)$$

Equations (1.22) can be matched using their gradients to avoid the constant U_c ,

$$y^+ d_{y^+} f(y^+) = \frac{y}{h} d_{y/h} F(y/h) = \kappa^{-1}, \quad (1.23)$$

where the equating constant κ^{-1} follows from the lack of common variables between both terms. This constant is known as the *Kármán constant*, as it was first obtained by Kármán [78] using similarity arguments. The method used here, matching the zeroth-order asymptotic expansions (1.17a) and (1.21) is due to Millikan [115]. The value κ is thought to be a universal constant for wall-bounded flows, but the values obtained from different DNS and experiments are scattered around $\kappa \approx 0.4$. Equation (1.23) can be integrated to,

$$U^+ = \frac{1}{\kappa} \log(y^+) + A, \quad (1.24a)$$

$$(U - U_c)^+ = \frac{1}{\kappa} \log(y/h) - B, \quad (1.24b)$$

where A is known as the *intercept constant* and its value depends on the roughness of the wall [64, 121]. Because equations (1.24) imply a logarithmic mean profile, the intermediate region is usually known as the *logarithmic region*, and its existence has seen agreement both experimentally and numerically at $Re_\tau \gtrsim 1000$. Nevertheless, other theories for the intermediate region exist, considering, for example, no velocity scale or a virtual origin [122] or the centerline velocity as the velocity scale [182]. The character of this region, as well as the value of the Kármán constant if considered logarithmic, are of great importance for several turbulence models [81, 157], and its accurate measurement is one important contribution of high- Re DNSs.

The validity of both inner and outer approximations, as well as the logarithmic region, are gathered in figure 1.2, which shows the mean profiles of a wide range of Reynolds numbers (550–4200).

FIGURE 1.2: Mean profiles of channel flows at $Re_\tau \in (550, 4200)$. Colours are different Re_τ simulations, 550 [7], 934 [7], 2003 [57], 4200 [102]; the data is available at the Jiménez group database <http://torroja.dmt.upm.es/>. (a) Scaling in inner units. The dotted line is (1.18), and the chained line its validity limit, $y^+ \approx 5$. The dashed line is the end of the buffer layer $y^+ \approx 100$. The short vertical colored lines are $y/h \approx 0.2$ for each simulation, the empirical validity limit of (1.17). The solid thick line is a logarithmic profile as (1.24a) with $\kappa^{-1} = 0.4$, $A = 5$. (b) Scaling in outer units. The short vertical lines fence the validity of (1.21) for each simulation, $y^+ \approx 100$. The dashed line marks the end of the logarithmic region, $y/h = 0.2$. The thick solid line is a logarithmic profile as (1.24b) with $\kappa^{-1} = 0.4$, $B = 0.2$.

1.4 COHERENT MOTIONS IN WALL-BOUNDED FLOWS

For a recent perspective on the state of the art of the description of wall-turbulence in terms of coherent motions, the reader is referred to Jiménez [69].

The asymptotic relations in §1.3 offer some information about the behaviour of the mean profile and about the mean Reynolds stress, but no information about the instantaneous flow structures whatsoever. As in any other turbulent flow, the velocity perturbations in channel flows are not random, and appear in the form of turbulent eddies. Owing to the inhomogeneity across the wall-normal direction, these eddies are expected to be reasonably different in the three different regions of the channel: the buffer, logarithmic and outer layers.

The buffer layer is the most studied and best understood region. This is partially due to the early availability of DNSs, but also because its dynamics are the simplest of the three regions. Because of the dominating effect of the viscosity, there is no separation of scales in the buffer layer, and production and dissipation occur essentially at the same length scales [65]. The most well-known example of coherent structure in wall-bounded flows are the *streamwise velocity streaks* (or simply streaks) that populate the buffer region.

They had been observed in early experiments [20, 83, 87], and in the first DNS [85]. It was shown by Jiménez and Moin [74] that the streaks take part in a cycle together with quasi-streamwise vortices. The cycle is thereof capable of sustaining turbulent activity in the buffer layer, independently of the dynamics of the outer regions. Many other works have enriched this view of the buffer layer [56, 73, 75, 150, 177] using very different approaches, from linear stability of a streaky turbulent background to exact coherent motions (UPOs). The overall picture consist in a streaky background that gets destroyed due to (quasi-linear) instabilities, seeding turbulent kinetic energy to non-linear processes that generate quasi-streamwise vortices which in turn regenerate the streak due to the ‘lift-up’ of the mean profile. At large Re , structures from the outer layer extend to the wall, affecting the structures there. Some models exist for this effect [108], which is known to be more than a simple superposition of the large (outer) scales over the small scales in the buffer layer.

Several theories coexist in the logarithmic region. Despite their differences, most of them are consistent with the celebrated model of Townsend [171], known as the *attached eddy model*. The general idea is that in the logarithmic region there is no scale other than u_τ , as neither ν nor h are appropriated for this region. The missing scale is the distance to the wall y , which implies that the dynamics of the eddies that contribute to the Reynolds stress are self-similar with the distance from the wall. This suggests that, aside from *inactive motions* [171] which do not contribute to the mean Reynolds stress, the largest scales at every height within the logarithmic region should be self-similar. Up to now, there is ample experimental evidence of self-similarity in the logarithmic region [101, 109, 111], although it is seen below that different authors propose different eddies as self-similar ‘building blocks’ for the logarithmic layer.

Note that there is a different scale that can be used to complement the friction velocity in the logarithmic layer, and is the time scale S^{-1} imposed by the mean shear S . Even if for an exact logarithmic profile, both nondimensional groups are equivalent ($y^+ S^+ = \kappa^{-1}$), the physical interpretation is slightly different. In the former interpretation, the group suggests that the structures are constrained in size by the distance from the wall, and observations from transition inspired some authors to believe that the structures were created in the buffer layer and grew in size until they could not grow any bigger. In the later scenario, the largest eddies are generated by taking energy from the mean shear, and their length scale is dictated by inertial dynamics. Because the mean profile is a product of the non-linearity of the NSEs, both views are equivalent, except for the mentioned growth mechanism, which is not directly implied by the self-similarity. Considering recent contributions, there is growing evidence [66, 68, 100, 116, present work] that the turbulent eddies

in the logarithmic region are generated by dynamics controlled by the mean shear.

The oldest ‘building block’ that was proposed to describe wall-bounded flows is due to Theodorsen [167], and is known as the *horseshoe vortex*. Inspired by the horseshoe, and the self-similarity of Townsend [171], Perry and Chong [127, 128] proposed a model consisting of self-similar *lambda* wall-attached vortices. Their characteristic velocity is u_τ and their length y , and following Townsend [171] their density is inversely proportional to y in order to recover a logarithmic profile. Similar to the lambda and horseshoe vortices, a refined model was proposed by Adrian et al. [3], known as the *hairpin packet* model. The model consists in a concatenation along the streamwise direction of horseshoe/hairpin vortices, which grow from the wall until they reach outer-scale sizes, in times much longer than their typical eddy turnover [186]. The packets are associated to low streamwise momentum regions, and rely on several growth mechanisms [169]. Recent DNS simulations have shown that hairpin packets populate transitional boundary layers [183], but they are absent in other simulations at large Re [146, 152]. Recent works, such as Dennis and Nickels [36] recover only pieces of hairpins, e. g. ‘legs’, ‘heads’, when analysing high- Re databases, questioning their relevance in turbulence far from transition.

A different family of structures can be traced back to the early experiments by Kline et al. [87]. In these experiment, the authors identified streamwise streaks as well as *bursts* of wall-normal velocity, that moved low streamwise momentum far away from the wall. This qualitative realisation was later formalised in Wallace et al. [178] and Lu and Willmarth [104] in the *quadrant analysis*, which is used to detected intense regions of local intense Reynolds stress $-uv$ which in turn can be used to detect the ejections of Kline *et al.* [21]. A review of these early experiments, emphasising the ambiguities and controversies of structures in wall turbulence is Robinson [141]. The quadrant analysis was extended to three dimensions using DNS data by Lozano-Durán et al. [101] which uncovered that two different families of *uvsters* exist, the ones attached to the wall, which extend from it up to a certain distance; and the detached. The attached family has self-similar sizes and aspect ratios across the logarithmic layer of channels. Their instantaneous structure is far from smooth, with fractal-like topologies. Lozano-Durán and Jiménez [103] extended the three dimensional analysis to time, showing that the time scales are also self-similar. Moreover, they found that the ejections (Q_2 , positive v , negative u) and sweeps (Q_4 , negative v , positive u) are not unrelated, and usually come in pairs of sweep and ejection. This pair constitutes an skew-symmetrical object in the spanwise direction, in contrast with the symmetrical horseshoe/lambda/hairpin vortices model. Similar structures can be found in homogeneous shear flow simulations [38], although these are

FIGURE 1.3: Top: Sketch of the hairpin packet model as devised by Adrian et al. [3] © Cambridge University Press. Bottom: Sketch of a pair of sweep-ejection pair and their time evolution, courtesy of [99].

completely symmetrical in y whereas in channels the sweeps tend to be bigger than the ejections. The symmetry across the spanwise direction is an important difference because the instabilities capable of creating skew-symmetrical perturbations in channels are stronger than their symmetrical counterparts [177]. A sketch of the structural model is showcased in figure 1.3.

A different approach used to obtain characteristic ‘eddies’ consist on computing the eigenfunctions (or singular value decomposition) of a tensor of interest, believed to capture some relevant piece of information about the flow. The most prominent example in fluid mechanics are the Proper Orthogonal Decomposition (POD), which are the eigenfunctions of the two-point (two-time) cross-correlation tensor,

$$C_{ij}(\mathbf{x}, \mathbf{x}', t, t') = u_i(\mathbf{x}, t) u_j(\mathbf{x}', t'). \quad (1.25)$$

It was introduced by Lumley [105] and its known in other fields as the Karhunen–Loève decomposition. Due to the translation invariance across the stream and spanwise directions, these eigenfunctions are Fourier modes for x and z . Nevertheless, they can be ‘compacted’ [117] to construct eddies representative of the energy-containing scales of the flow. Different tensor decompositions are possible. The origins of Adrian’s hairpins can be traced back to a similar approach but using the LSE operator. Other tensors used to identify characteristic motions are the resolvent operator of McKeon and Sharma [110], which contains full information of the linearised NSE; the Dynamic Mode Decomposition [147], which identifies temporal frequencies and their associated spatial structure; or the somewhat exotic product of the controllability and observability Grammians [143], which balance modes according to a mixed criterion. One drawback of these methods is that is usually hard to identify (visualise) similar features in instantaneous flow fields.

1.5 AIMS

There are two independent but related goals to this thesis. The first and most important is trying to provide a better dynamical description of the mostly kinematic Q_{2-4} pairs of Lozano-Durán and Jiménez [103]. The second, more practical, is studying to which extent is detecting these (and other) flow structures from limited measurements of the flow fields.

Attempts to improve the knowledge about the dynamics of these objects are not new. In two works that can be considered roots for the first part of this thesis [66, 68], Jiménez linked a linearised transient growth mechanism [124] to the dynamics of fully nonlinear channel flows. However, these works were limited to unlocalised structures in minimal channels, which made very difficult relating them to the localised objects of other authors.

For the second part, the relation of the different variables at the wall (pressure, and the two shear stresses) with the velocities far away from it is analysed. Inspired by the results obtained in the first part, a simple linear estimator is chosen to estimate the flow far away from the wall, and the capabilities of this method to detect coherent objects are emphasised.

Finally, the possibility of detecting bursts solely from wall measurements is explored, as well as the differences this entails with respect to detection from full data.

1.6 CONTENTS

The remainder of this manuscript is organised in four chapters. Chapter 2 is devoted to the characterisation of quasi-linear dynamics in channel flows.

For this purpose, a novel method which locally approximates the flow fields as streamwise wave-packets is proposed, bridging this way the gap between the linear mechanisms and localised objects. These methods are applied to time-resolved databases of channel flows, in an attempt to elucidate the dynamics controlling the momentum-carrying eddies of the logarithmic region. These quasi-linear eddies share most of their properties (if not all) with the (attached) momentum carrying eddies of [101, 103], strongly suggesting that the dynamics controlling the later are mostly linear. The contents of this chapter are under consideration for publication in the Journal of Fluid Mechanics, under the title “Momentum transfer by linearised eddies in turbulent channel flows”, with J. Jiménez as a coauthor.

Chapter 3 analyses the relation of the different observables at the wall with the flow field far from it. The flow fields are found to be observable at relatively far distances even with a relatively simple linear method. The quality of the reconstructions is analysed and found to be self-similar respect to the distance to the wall, suggesting that the observables eddies are those attached to the wall. It is also shown that the reconstructions of the velocity fields contain a significant amount of the streamwise velocity perturbations and of the tangential Reynolds stress. The contents of this chapter are published in Physical Review Fluids (4, 114603), under the title “Logarithmic-layer Turbulence: a view from the wall”, with J. Jiménez as a coauthor; and in the Proceedings of the CTR Summer Program (2018, pp. 217–226), with the different title “Reconstructing channel turbulence from wall observations”, and A. Lozano-Durán and J. Jiménez as coauthors.

In chapter 4, the results from previous chapters are used to develop a method capable of detect intense turbulent events from the wall. The analysis shows that the events that can be detected at a given wall distance are larger than the most energetic. It is argued that the detected structures are the ‘feet’ of the Orr bursts.

In the fifth and last chapter we discuss the significance of the results of this thesis and conclude.

The contents of this chapter are under consideration for publication in the Journal of Fluid Dynamics, in a paper titled “Momentum transfer by linearised eddies in channel flows”, with J. Jiménez as a coauthor.

Nonlinearity is an essential characteristic of turbulent phenomena [166]. Linear models of statistically steady turbulent flow neither retain its chaotic behaviour [144] nor its self-similar multi-scale organisation [89]. In addition, linear models cannot predict the intensity of turbulent perturbations. Nevertheless, some features of shear flows, such as the structure of the velocity perturbations at the energy-injection scale, can be described reasonably well by the linearised Navier–Stokes equations. In fact, since the source of turbulent kinetic energy in a shear flow is the mean shear, it is unavoidable that at least part of the energy production mechanism should be linear, in the sense of involving interactions of the perturbations with the mean flow, rather than of the perturbations among themselves.

The best-known example is the Kelvin–Helmholtz instability of mean profiles with inflection points, which controls many of the properties of the large-scale structures in fully turbulent free-shear flows [23]. On the contrary, the mean profiles of wall-bounded flows lack an inflection point, and are linearly stable [138]. Still, attempts to derive properties of wall-bounded turbulence by assuming that the mean profile is marginally stable to perturbations have a long history [107], and it has been known for some time [24] that modally stable linearised perturbations can transiently grow by drawing energy from the mean shear. Many linear models reproduce some of the large-scale features of channel flow well from linearised dynamics [11, 44, 58].

One such mechanism was proposed by Orr [124]. The cross-shear velocity component (wall-normal in wall-bounded flows) is amplified over times of the order of the inverse of the shear when backwards-leaning perturbations are tilted forward by the mean velocity profile. The perturbations grow until they are normal to the shear and are damped past that point, because continuity requires that the change in vertical scale due to the tilting is balanced by the velocity amplitude. Since perturbations eventually decay, there is no net production of kinetic energy for the wall-normal velocity, but, for three-dimensional perturbations whose wavefronts are oblique to the mean stream, some of the energy is transferred to the other two velocity components. Moreover, since continuity does not interact with the

wall-normal variation of the wall-parallel velocities, the effect of this ‘lift up’ remains after the wall-normal velocity has vanished [66], forming the characteristic streamwise-velocity streaks that populate wall-bounded flows [87]. Because the growth of the wall-normal velocity is much shorter-lived than the resulting streaks, we use the term ‘burst’ to describe it. The term was originally coined for the ejections of low-speed streaks close to the wall in boundary layers [83]. It was abandoned for a while due to controversies about the nature of the ejections, but eventually reclaimed, in the sense introduced above, to denote the intermittent behaviour of wall-bounded flows [48, 73, 80].

The connection between linearised Orr amplification and bursting was studied by Jiménez [66], who showed that the length and time scales predicted by the linear model were in agreement with the bursting of the logarithmic layer in minimal channels. This work was extended by Jiménez [68], who tracked the temporal evolution of the largest Fourier modes in channels designed to be minimal for the logarithmic layer. The strongest wall-normal velocity events were found to closely follow the predictions of the linearised equations when the mode was initially ‘coherent’ across the wall-normal direction (i.e. could be represented by a wavetrain). A key aspect of these results is the relation between the ‘tilting’ angle of the perturbations and their amplitude, which is predicted by the linearised Orr-Sommerfeld equations, and serves as an identification criterion for Orr-like bursts.

Several problems remain. The works just cited treat Orr bursts as individual Fourier modes. This is no problem for linearised theory over spatially homogeneous directions, since modes can eventually be added to form more complicated structures, but it becomes questionable when bursts are to be identified in (nonlinear) extended flow fields. The bursts observed in experiments and simulations are localised in space as well as in time, and cannot be described as infinite wavetrains [2]. Extensions of the classical quadrant analysis [178, 180] to three-dimensional Reynolds-stress structures [39, 101], and later to their temporal evolution [103], show that much of the momentum transfer in channels and in other shear flows is carried by transient eddies with geometric aspect ratios of order unity, rather than by infinite wavetrains. The reason why Jiménez [68] looked for Orr-like events in a minimal channel was precisely to be able to use Fourier methods while maintaining the aspect ratios of experimental observations, and the analysis failed when applied to wavelengths much shorter than the size of the simulation box. It is not obvious whether the results from uniform wavetrains can be generalised to a population of localised bursts of different sizes in a large simulation box, and many of the tools used for the former, such as the wavetrain inclination angle, need to be redefined for the latter.

Another important question is to what extent Orr bursting explains the momentum transfer in full-sized channels. For example, it was shown by Jiménez [68] that the largest-scale field of the wall-normal velocity in a minimal channel can be described as a linearised Orr burst over 65% of the time, but it is unclear which is the equivalent fraction in large boxes. For example, the Reynolds stress in minimal boxes is estimated to burst 20% of the time [69], but the stress structures identified by [101] only cover 8% of the wall.

In the present work, we introduce a band-pass filtering method that allows us to extend the analysis of bursting beyond minimal channels, detecting structures of arbitrary size in full-sized simulations. Using large boxes is specially important when studying the characteristic time and length scales of the turbulent structures, because the largest scales analysed in minimal computational boxes are also the most prone to suffer simulation artefacts. We study multiple scales across different Reynolds numbers and simulation boxes, uncovering a range of self-similar Orr-like events in the logarithmic region, with sizes of the order of their distance from the wall.

On the other hand, it is important to understand the limitations of our approach. The main one has to do with scale. The reason why Fourier modes are useful in analysing the linearised equations is that harmonic functions are eigenfunctions of uniform translations in space and time. Any linear mechanism without an explicit dependence on the homogeneous coordinates cannot change the wall-parallel length scales, including the Fourier wavenumbers. The range of scales mentioned in the previous paragraph cannot be generated by the linearised equations, and remains an empirical input parameter. The effect of the band-pass filters used in this thesis is also to isolate a narrow range of scales, which remains constant throughout the observed evolutions. In this respect, both linearisation and the band-pass analysis exclude the multiscale aspect of turbulence dynamics. To remedy this, we also analyse the effect of low-pass filters that allow us to include the longer streaks of the streamwise velocity, and some multiscale effects.

The remainder of this chapter is organised as follows. The linear theory for the Orr mechanism in turbulent channel flows is briefly reviewed in §2.1. Section §2.2.1 covers the domain definition and the database that will be used in the paper, and the rest of §2.2 describes the methodology used to detect and characterise Orr events. Results are presented in §2.3 and §2.4. Section §2.3 discusses the statistics and spatio-temporal properties of band-pass filtered flows, which isolate the bursts, and §2.4 extends the results to the low-pass filtered velocity. Conclusions are gathered in §2.5.

2.1 THE LINEAR MODEL

Let $\mathbf{x} = (x, y, z)$ be the streamwise, cross-shear and spanwise directions, being $\mathbf{q} = (u, v, w)$ the corresponding velocity perturbation components. We restrict ourselves to shear flows with mean velocity profiles $\mathbf{U} = (U(y), 0, 0)$ with no inflection points, where the average is taken over the two homogeneous directions (x, z) and time. The kinematic pressure and viscosity are p and ν , respectively, defined as the quotients between the dynamic quantities and the constant fluid density, which does not appear explicitly in the equations. The Navier–Stokes equations can be written as

$$\partial_t \mathbf{q} = -(\mathbf{U} + \mathbf{q}) \cdot \nabla (\mathbf{U} + \mathbf{q}) - \nabla (P + p) + \nu \nabla^2 (\mathbf{U} + \mathbf{q}), \quad (2.1)$$

$$\nabla \cdot \mathbf{q} = 0; \quad (2.2)$$

where $\mathbf{u} = \mathbf{U} + \mathbf{q}$ is the standard Reynolds decomposition, and P is the mean pressure. In general, capital letters denote ensemble averages and lower-case ones are fluctuations. Primes are reserved for root-mean-square (rms) intensities. The linearised version of (2.1) is

$$\partial_t \mathbf{q}_L = -\mathbf{U} \cdot \nabla \mathbf{q}_L - \mathbf{q}_L \nabla \mathbf{U} - \nabla p_L + \nu \nabla^2 \mathbf{q}_L, \quad (2.3)$$

where p_L is the fast pressure [84], although we suppress the sub-index ‘ L ’ when referring to linearised variables for the remainder of the section. The linearised equations can be manipulated into the Orr–Sommerfeld–Squire system [124, 159],

$$\partial_t \phi = -U \partial_x \phi + d_y S \partial_x v + \nu \nabla^2 \phi, \quad (2.4)$$

$$\partial_t \omega_y = -U \partial_x \omega_y - S \partial_z v + \nu \nabla^2 \omega_y; \quad (2.5)$$

where $\phi = \nabla^2 v$, $\boldsymbol{\omega} = (\omega_x, \omega_y, \omega_z)$ is the vorticity, and $S = d_y U$ is the mean shear. Although (2.4) is autonomous in v , continuity and (2.5) connect it with the wall-parallel velocities. Continuity and the linear pressure do not appear in (2.4), which is a vorticity equation, but their effect is hidden in the definition of ϕ . For example, the term $d_y S \partial_x v$, which is responsible for the Kelvin–Helmholtz instability of free shear flows [23], vanishes identically in a homogeneous shear, $U = Sy$, and (2.4) reduces to a simple advection–diffusion equation. Since such flows share many of their bursting properties with boundary layers and channels, it was argued by [66] that the amplification of the bursts in (2.4) has to be mediated by the pressure through the Laplacian operator in ϕ . In fact, for a homogeneous shear in the limit $\nu \rightarrow 0$, and a pure Fourier mode of streamwise wavenumber k_x , and spanwise wavenumber $k_z = 0$, the solutions to (2.4) are ‘Orr’ bursts of the form

$$v = \cos^2(\Psi) \exp [ik_x(x - Syt)], \quad (2.6)$$

FIGURE 2.1: Evolution of the amplitude and inclination angles of an optimal transient growth mode in a channel flow. $\lambda_x = \lambda_z$. (a) Snapshots of the velocity components (top to bottom, u, v, w), spaced by $St \approx 0.52$, from left to right. The dashed lines show the inclination angle computed as in [68]. The contours levels are $u = 4, v = 1, w = 3$ in uniform arbitrary units. (b) Amplitude as a function of the inclination angle. The markers are, $\bullet$, u ; $\blacksquare$, v ; $\blacktriangle$, w ; and correspond to the snapshots in (a)

where $\Psi = \arctan(St)$ is the inclination angle of the wavefront with respect to the y axis. The inclination angle is a reinterpretation of time, and (2.6) relates the amplitude of the v perturbations to the inclination of their wave fronts, being maximum when $\Psi = 0$ and vanishing towards $\Psi = \pm\pi/2$. The burst starts from a weak backwards-leaning perturbation, which amplifies as it is tilted forward by the shear, is strongest when it is normal to the shear, and is damped again after this moment. The whole evolution unfolds during times of the order of the shear time, S^{-1} , and the linearisation is only valid if this time is shorter than other evolution times in the equation.

For non-zero viscosity, there is a perturbation Reynolds number, $Re = S/k_x^2 \nu$ [68], and (2.6) decays in times of $O(S^{-1} Re^{1/3})$, which are typically long for large structures. Nonlinearity is represented by the local eddy turnover time, $\|\mathbf{q}\|^2/\varepsilon$, where $\|\mathbf{q}\|^2/2$ is the turbulent kinetic energy and ε is the energy dissipation rate. Its ratio to the shear time is the Corrsin shear parameter, $S^* = S\|\mathbf{q}\|^2/\varepsilon$. In wall-bounded flows, $S^* \approx 10$ in the range $50 \nu/u_\tau < y < 0.7 h$, where u_τ is the friction velocity [66], and (2.3) is therefore valid at least for the energy-containing structures in this range.

2.1.1 Optimal transient growth

For more complex mean velocity profiles, Orr-like solutions can be computed numerically. We use the algorithm in [148] to find solutions that optimally amplify the total kinetic energy, including a y -dependent eddy viscosity to account for the nonlinear turbulent dissipation, and a velocity profile and kinematic viscosity corresponding to a turbulent channel at

$Re_\tau = u_\tau h / \nu = 2000$, as in [11] and [132]. Álamo and Jiménez [11] found that the maximum energy amplification occurs for $\lambda_x \gg \lambda_z$, but [69] remarked that these wavenumbers are associated with the slow amplification of the streamwise-velocity streaks through (2.5), and to very long evolution times that may be limited by the nonlinear or viscous effects mentioned above. We are here interested in situations such as (2.6), where the energy amplification is associated with the transient bursting of v over a few shear times. These were shown by [69] to be strongest at $\lambda_x \approx \lambda_z$.

Figure 2.1(a) shows as shaded contours several snapshots of the most amplified transient mode with $\lambda_x = \lambda_z \approx 0.7h$, plotted at intervals of $S(y_g)\Delta t \approx 0.23$, where $S(y_g)$ is the shear at the centre of gravity of the v^2 profile at maximum amplification. It was shown in [66] that the dimensions of the most amplified inviscid v -modes in the logarithmic layer are proportional to $\lambda_0 = 2\pi/(k_x^2 + k_z^2)^{1/2}$, and that their centre of gravity is at $y_g \approx 0.22\lambda_0$, or $y_g \approx 0.16\lambda_x$ for equilateral wavevectors. The same is true in figure 2.1(a), even if this case differs from [66] in using an eddy viscosity. The temporal evolution of the perturbations in [66] is controlled by $\lambda_0 S(y_g)t / \lambda_x$. In the range of wavelengths $250\nu/u_\tau \lesssim \lambda_x \lesssim 4h$, corresponding to $40\nu/u_\tau \lesssim y_g \lesssim 0.64h$, we can assume that $S(y_g) \propto u_\tau/y_g \propto u_\tau/\lambda_0$, and that the scaling of the time is $\lambda_0 S(y_g)t / \lambda_x \propto u_\tau t / \lambda_x$. The result is that both the dimensions and the temporal evolution of the amplified structures scale with λ_x . The relation between the inclination angle Ψ and the amplitude of v was shown in [66] to be fairly independent of the wavenumbers, and even of the type of flow. These results are extended here to the three velocity components, and to the eddy-viscosity equations, and we indeed find that the optimal modes in the neighbourhood of $\lambda_x \approx \lambda_z$ scale well with λ_x in the range of wavenumbers mentioned above. For example, the shape and inclination angles of the solution in figure 2.1(a) are remarkably similar to those in figure 19 of [69], where the wavelengths are much larger, $\lambda_x = 3.9h$, $\lambda_z = 5.55h$, and the solid contour lines superimposed on the second snapshot in figure 2.1(a), which show the moment of maximum amplitude of a mode with $\lambda_x = \lambda_z \approx 0.35h$, also agree very well with those at the longer wavelengths. The self-similar behaviour ceases to hold at heights corresponding to the buffer layer, where the shear is no longer inversely proportional to y .

The dashed lines in figure 2.1(a) are the vertically averaged inclination angles at each time, computed using the formula for single modes in [68]; for a more general one see (2.18) below. They represent well the instantaneous structure of v , but they are biased towards the lower half of the structures for u and w because the averaging in (2.18) is weighted with the energy, which tends to be concentrated near the wall for these two components. This is also where the shear is highest, leading to flatter inclinations. Figure 2.1(b) shows

the amplitude-inclination relation for the three velocity components in figure 2.1(a). The trajectory for v is reminiscent of (2.6), but the wall-parallel components tilt forward much faster, as shown by the larger separation of their first few symbols, which are uniformly spaced in time. Inspection of figure 2.1(a) shows that this is also due to the weighting of the angle close to the wall, where the higher shear results in faster characteristic times. The upper part of the u and w structures does not tilt much faster than that of v , while the root of the latter tilts as fast as those of the other two components. The last half of the angle-amplitude evolution is also different for the three velocity components. The spanwise velocity has an amplitude peak at approximately the same time as v , which is not present for u . At the end of the evolution, v decays at a rate dictated by the shear, while u and w only decay under the effect of the eddy viscosity. For inviscid perturbations in a uniform shear, the three velocity components have well differentiated behaviours. During the burst, v grows and decays in $t = O(S^{-1})$, as in (2.6), but u and w never decay [66]. In all cases, continuity requires that the two wall-parallel have to be proportional to each other once v has died. An important message of figure 2.1 is that optimum growth solutions should only be taken as indicative. The properties enumerated in the last few sentences are common to any initial condition, while optimum growth is meaningless in an inviscid flow. The behaviour of v , whose equation (2.4) is autonomous, is essentially independent of the initial conditions, which mostly define the origin of time, but u and w satisfy the forced equation (2.5), whose evolution depends of v and of their initial conditions. The optimum transient growth criterion selects the initial condition with a highest local maximum, but the v burst and the v -less final decay are robust properties of most initial conditions involving backwards-leaning perturbations. Note also that the energy growth in figure 2.1(b) is not particularly large, and that the same is true of all the cases mapped in figure 29 of [69].

2.2 METHODS

The remainder of this chapter seeks to relate the behaviour of fully nonlinear channel flow simulations to the linearised dynamics sketched above.

2.2.1 The numerical data sets

We use direct simulations of canonical incompressible turbulent channel flows whose half-height is h . Normalisation with v and u_τ is represented by a ‘+’ superscript, and our primary Reynolds number is $Re_\tau = h^+$. The simulations use periodic boundary conditions in the wall-parallel directions, with periods L_x and L_z , and are summarised in table 2.1. The Reyn-

CASE	Re_τ	L_x/h	L_z/h	N_x	N_y	N_z	D_y	N_f	REF.
S1000	932	2π	π	512	385	512	CH	1800	[103]
L1000	934	8π	3π	2048	385	1536	CH	72	[9]
S2000	2009	2π	π	1024	633	1024	FD	590	[103]
F2000	2000	8π	3π	512	512	512	FD	1146	[A, 174]
L2000	2003	8π	3π	4096	633	3072	FD	30	[57]
S4000	4164	2π	π	2048	1081	2048	FD	40	[103]

TABLE 2.1: Parameters of the simulations. $L_{x,z}$ are the streamwise and spanwise periods of the numerical box. $N_{x,y,z}$ are the number of collocation grid points in y , or Fourier modes in x, z , except for F2000, which was computed at full resolution but stored at the reduced one given in the table. D_y represents the discretisation used for the wall-normal direction: ‘CB’ refers to Chebyshev polynomials and ‘FD’ to compact finite differences. N_f is the number of fields used to compute statistics. The snapshots of S1000 and F2000 are close enough in time to compute temporal derivatives.

olds numbers range from $Re_\tau \approx 900$ to 4000, and the simulations include both medium-sized ($L_x = 2\pi h, L_z = \pi h$) and large computational boxes ($L_x = 8\pi h, L_z = 3\pi h$). The equations of motion are written as evolution equations for the (x, z) averages of the streamwise and spanwise velocities, and for ϕ and ω_y , as in (2.4)–(2.5) without linearisation [85]. The spatial discretisation is Fourier spectral in the two periodic directions, dealiased using the pseudo-spectral 3/2 rule, but the discretisation along y varies among simulations, as indicated in table 2.1. Some cases use Chebyshev polynomials collocated at the Gauss–Lobatto–Chebyshev nodes. Others use spectral-like compact finite differences [47, 94], up to 12th-order consistent, in grids whose spacing is adjusted to keep the resolution approximately constant in terms of the local isotropic Kolmogorov scale $\eta = (\nu^3/\varepsilon)^{1/4}$. Full details can be found in the original publications. The data stored are snapshots of the three velocity components, which in most cases are approximately statistically independent and can only be used to compile instantaneous statistics. However, the snapshots in S1000 and F2000 are stored closely enough in time to provide temporal derivatives and histories without recomputing the flow evolution.

2.2.2 Filtering

As discussed in §2.1, the relation between the instantaneous tilting angle of the wall-normal velocity perturbations and their amplitude can be used as a diagnostic property for Orr bursts. [68] defined the inclination angle of the wavefronts of a pure Fourier mode in terms of the wall-normal derivative

of its complex phase, but this definition, as well as that of the amplitude, needs to be generalised for spatially localised objects. Consider the low-pass-filtered field

$$\tilde{\mathbf{q}}(x, y, z, t; \Lambda) = \int_{\mathcal{B}} \Gamma(x' - x, z' - z, \Lambda) \mathbf{q}(x', y, z', t) dx' dz', \quad (2.7)$$

where $\mathcal{B}$ stands for the full computational box, and

$$\Gamma(x, z; \Lambda_{0L}, \Lambda_{0z}) = \gamma(x, \Lambda_{0L}) \gamma(z, \Lambda_{0z}), \quad (2.8)$$

with

$$\gamma(\xi, \Lambda_0) = \frac{1}{\Lambda_0 \sqrt{2\pi}} \exp\left[-\frac{1}{2}(\xi/\Lambda_0)^2\right], \quad (2.9)$$

and the band-pass filtered one

$$\bar{\mathbf{q}}(x, y, z, t; \Lambda) = \int_{\mathcal{B}} G(x' - x, z' - z, \Lambda) \mathbf{q}(x', y, z') dx' dz', \quad (2.10)$$

with

$$G(x, z; \Lambda, \Lambda_{0x}, \Lambda_{0z}) = \Gamma(x, z; \Lambda_{0x}, \Lambda_{0z}) \exp(i2\pi x/\Lambda). \quad (2.11)$$

The spectral transfer functions of the two kernels are

$$|\hat{\Gamma}|^2 = \exp\left[-(k_z \Lambda_{0z})^2 - (k_x \Lambda_{0L})^2\right], \quad (2.12)$$

$$|\hat{G}|^2 = \exp\left[-(k_z \Lambda_{0z})^2 - [(k_x - 2\pi/\Lambda) \Lambda_{0x}]^2\right], \quad (2.13)$$

where the carat stands for the (x, z) Fourier transform. Both filters have spectral width $\Delta k = O(\Lambda_0^{-1})$, but the band-pass filter G is centred at the wavenumber $k_x = 2\pi/\Lambda$ (and $k_z = 0$), while the low-pass filter Γ is centred at $k_x = k_z = 0$. Thus, if Λ and the Λ_0 's are chosen to be of the same order, Γ smooths the flow by damping everything shorter and narrower than $O(\Lambda_0)$, while G selects only structures which may be wide, but whose length is of the order of Λ . The latter is useful to isolate localised bursts, while the former also retains the larger structures of the flow around them.

The low-pass filter Γ presents few conceptual problems, and its results are described in §2.4, but the band-pass filter (2.10) generates a complex field which is not easily interpreted as a flow, and requires some discussion. Its kernel G is a monochromatic complex wave in the streamwise direction whose amplitude is modulated by the Gaussian (2.9) in the streamwise and spanwise directions. It is essentially a complex-valued continuous Morlet or Gabor wavelet [43], whose real and imaginary parts are represented in figure 2.2(a) as functions of x . The number of oscillations contained within the Gaussian envelope is approximately $2\Lambda_{0x}/\Lambda$. To keep the filters self-similar we choose $\Lambda_{0x} = \Lambda/2$ and $\Lambda_{0z} = \Lambda/6$, implying an aspect ratio $\Lambda_{0x}:\Lambda_{0z} \approx 3:1$,

FIGURE 2.2: (a) The wider filter on the left is the band-pass kernel $G(x, 0)$ for $\Lambda = 4\pi/9$. The solid darker line is $\text{Re}(G)$, and the lighter one is $\text{Im}(G)$. The two envelopes are: $---$, γ_x ; $-\cdot-$, γ_z . The narrower filter on the right corresponds to $\Lambda = \pi/18$, although the amplitude has been divided by four for the purpose of plotting. (b) Sample signal, $g = \sin(12x)$ if $x < 0$, and $\sin(3x)$ otherwise. (c) Its amplitude after band-pass filtering with $\Lambda = \pi/6$. The solid line in (c) is the amplitude recovered by $\text{Re}(G)$, and the dashed one corresponds to the full complex G .

which captures the strongest wall-normal velocity perturbations (see figure 2.3b). Changing this aspect ratio to 2:1 or 4:1 did not qualitatively change the results.

The main advantage of filtering a real-valued function with a complex wavelet is that the real and imaginary part of the resulting field preserve phase (positional) information, while its absolute value hides the oscillations of the wavelet [43, 160]. A visualisation for synthetic data is presented in figure 2.2(b).

In fact, the band-passed field, $\bar{a}$, of a generic velocity component can be interpreted as the coefficient of a local Fourier expansion that optimally approximates the flow within the Gaussian envelope (2.8). It minimises the weighted L_2 error

$$J[\bar{a}](x) = \int_{\mathcal{B}} \Gamma(x' - x, z' - z) |\mathcal{J}|^2 dx' dz', \quad (2.14)$$

$$\mathcal{J} = a(x', y, z') - \bar{a}^*(x, y, z) \exp[2\pi i(x' - x)/\Lambda], \quad (2.15)$$

where the asterisk denotes complex conjugation and a is either u , v or w . Note that a^* in this integral is evaluated at the centre of the filtering kernel, x , rather than at the integration variable x' , and that the Fourier wavetrain, $\exp[2\pi i(x' - x)/\Lambda]$, is centred in each case at the filter position. This implies that a^* can be interpreted as a field of coefficients of ‘local’ Fourier wavetrains, whose absolute value is $|a(x, y, z)|$, and whose argument is $-\arctan(\text{Im}(a)/\text{Re}(a))$. Thus, although $\bar{a}$ cannot be used as a filtered velo-

FIGURE 2.3: (a) Transfer functions of the: ---, band-pass filter G ; —, low-pass filters F . The contours are $(0.2 : 0.2 : 0.8)$. (b) Sample filtered premultiplied spectra of v in L2000, at $y/h = 0.15$ and $\Lambda/h = 0.175$. The dashed contour is $|\widehat{G}|^2 = 0.2$. The shaded contours are the full $k_x k_z E_{vv}$; —, low-pass filtered; — · —, band-pass filtered. The spectral contours are $k_x k_z E_{vv}/u_\tau^2 = (0.014, 0.029, 0.071, 0.114)$.

city field, its absolute value is the local wavetrain amplitude [160], and the wall-normal derivative of its argument is a local inclination angle, as in [68].

It is convenient to choose the short-wave limits of the low-pass and band-pass filters to be of the same order, so that they isolate related structures. The $(1/e)$ limit of (2.13) is $k = 2\pi/\Lambda + 1/\Lambda_{0x}$, while that of (2.12) is $k = 1/\Lambda_{0L}$. Equating them results in $\Lambda_{0L} = \Lambda/(2\pi + \Lambda/\Lambda_{0x})$. For our previous choice of $\Lambda_{0x} = \Lambda/2$, this implies $\Lambda_{0L} = \Lambda/8$. The transfer function of the two filters is represented in figure 2.3(a). The cospectrum of any two filtered variables can be computed from the unfiltered cospectrum as $E_{\bar{a}\bar{c}} = E_{ac}|\widehat{G}|^2$, with an equivalent formula for $E_{\widetilde{a}\widetilde{c}}$. An example of the effect of the band- and low-pass filters on the spectrum of a representative plane of v is shown in figure 2.3(b).

Note that the computation of the filtered fields can be done economically because the streamwise and spanwise directions are periodic, and fast Fourier transforms can be used to compute the convolutions.

2.2.3 The band-pass filtered pseudo-spectrum

The perturbation energy of the filtered flow fields can be directly computed from their spectrum,

$$\mathcal{E}_{ac}(y; \Lambda) = \sum_{\forall k_x, k_z} E_{\bar{a}\bar{c}}(k_x, y, k_z; \Lambda) = \sum_{\forall k_x, k_z} |\widehat{G}|^2(k_x, k_z; \Lambda) E_{ac}(k_x, y, k_z), \quad (2.16)$$

and is represented in figure 2.4 for the three velocity components and for the tangential Reynolds stress, as a function of y and of the filter wavelength

FIGURE 2.4: Pseudo-spectra of the band-pass filtered velocity components and of their tangential Reynolds stress, as a function of the filter wavelength and of the distance from the wall. $Re_\tau = 5200$ [93]. The contours are 90% and 50% of the maximum for each filter wavelength. The solid and dashed lines bracket the most energetic part of $\mathcal{E}_{vv}$ in (b). The upper (solid) limit is $y = \Lambda$. The lower (dashed) one is $y = 4\Lambda/11$. The chain-dotted line is the wall-normal location of the maximum for each filter width. (a) $\mathcal{E}_{uu}$. (b) $\mathcal{E}_{vv}$. (c) $\mathcal{E}_{wv}$. (d) $-\mathcal{E}_{uv}$.

Λ . Because the filter wavelength represents the size of the scales retained by the filter, we can think of $\mathcal{E}_{ac}(y; \Lambda)$ as a ‘pseudo-spectrum’ with pseudo-wavelength Λ . Figures 2.4(a–d) show that the wall-normal location of the maximum energy is proportional to the filter width for a range of heights corresponding to an extended logarithmic layer, $y \in (100\nu/u_\tau, 0.4h)$. This is also approximately the region where a self-similar hierarchy of wall-attached velocity structures with sizes proportional to their distance from the wall is found to exist [171], and we will refer to it as our definition of the logarithmic layer from now on.

Figure 2.4 allows us to determine which distances from the wall should be used to characterise eddies of different scales. The band-pass filter (2.11) produces a flow field that is locally averaged within each wall-parallel plane, but eddies and bursts also have a vertical dimension, which figure 2.4 shows to be proportional to their horizontal ones [101]. It was shown by [66] that bursts may be characterised by a vertically averaged amplitude,

$$A_a^2 = \frac{1}{y_1 - y_0} \int_{y_0}^{y_1} |\bar{a}|^2 dy, \quad (2.17)$$

and by a mean inclination angle,

$$\Psi_a = \arctan \frac{\Lambda \int_{y_0}^{y_1} \text{Im}(\bar{a}^* \partial_y \bar{a}) dy}{2\pi \int_{y_0}^{y_1} |\bar{a}|^2 dy}, \quad (2.18)$$

where a stands for either u , v or w , and the integrals extend over the intense part of the eddy. The average in (2.18) is weighted with the square of the

amplitude because strong perturbations typically coexist in the filtering window with weaker ones, whose inclination angle is not necessarily small, nor relevant.

The limits y_0 and y_1 are chosen to bracket the intense band of $\mathcal{E}_{vv}$ in figure 2.4(b), and the same limits are used for all variables. The upper limit is very close to the Corrsin spectral scale, $y_1 = \Lambda$, above which the shear is too weak to interact with eddies of size Λ [69]. The lower limit, $y_0 = 4\Lambda/11$, also scales with Λ , and represents the point below which impermeability damps v . Because these limits are adjusted for the wall-normal velocity, while the u or w structures are larger than those of v , figure 2.4(a,c) shows that they are too far from the wall to capture more than a fraction of the energy of the wall-parallel components. On the other hand, the tangential Reynolds stress in figure 2.4(d) is captured well, supporting the classical classification into ‘active’ and ‘inactive’ variables in [171].

From now on, we will use the quantities in (2.17)–(2.18) to represent the local amplitude and inclination angle of our flow fields. Their main practical advantage is that they reduce the dimensionality of our data set, because each filter has an associated three-dimensional data space (the two wall-parallel positions and time), instead of the full four-dimensional one. In addition, they act as a filter in y , excluding structures outside the band (y_0, y_1) , or with very different vertical dimensions from the band thickness. They are centred at $[x, (y_0 + y_1)/2, z]$, and the vertical average is only meaningful whenever a velocity perturbation of size comparable to Λ happens to be within the vertical integration window. Otherwise, the average reverts to the unconditional mean. For example, when several wavetrains with different inclinations are stacked along the wall-normal direction, the integral of their phase derivatives produces an average inclination angle which is similar to the global ensemble-averaged inclination.

2.2.4 Filter parameters

Table 2.2 gathers the information for the filters used in the paper. We use six wavelengths,

$$\Lambda_i/h = \frac{8\pi}{9} 2^{-i}, \quad i = 1, \dots, 6, \quad (2.19)$$

with the corresponding filtered fields denoted by $\bar{\mathbf{q}}_i = \bar{\mathbf{q}}(x, y, z, t; \Lambda_i)$ for the band-pass filters discussed in §2.3, and $\tilde{\mathbf{q}}_i = \tilde{\mathbf{q}}(x, y, z, t; \Lambda_i)$ for the low-pass ones in §2.4. They cover the energy spectrum with logarithmically equispaced bands of wavelengths, and span the range that can be expected to be relevant for the logarithmic layer. The long-wavelength limit of the widest filter is kept constant in outer units across Reynolds numbers, $\Lambda_1/h = 4\pi/9$. It follows from figure 2.4 that this filter corresponds to $y/h \in (0.5, 1.4)$,

NAME	Λ/h	y_0/h	y_1/h	S1000	L1000	S2000	F2000	L2000	S4000
Λ_1	1.4	1.0	0.51	○	○	○	○	○	○
Λ_2	0.7	0.7	0.25	○	○	○	○		○
Λ_3	0.35	0.35	0.13	○	○	○	○		○
Λ_4	0.175	0.175	0.063	○	○	○	○		○
Λ_5	0.087	0.087	0.032			○			○
Λ_6	0.044	0.044	0.016						○

TABLE 2.2: Properties of the filters used. Λ is the filter width defined in (2.11), and y_0 , y_1 are the integration limits in (2.17)–(2.18). A circle below a simulation indicates that that filter was computed for it.

which is well above the expected self-similar region, but it is retained here to compare it to the experiments in [68], who analysed the first streamwise mode of a minimal channel with $L_x = \pi/2$. It allows us to test whether the minimal domain used in that work affected the results. The short-wavelength limit, $\Lambda^+ \approx 180$, scales in wall units, resulting in an increasing number of filter bands as the Reynolds number increases. In the highest Reynolds number case S4000, this filter corresponds to Λ_6 . Since both our Reynolds numbers and our filter widths are approximately spaced by powers of two, there are matching filters for all the simulations, both in outer and in inner units, allowing us to compare scaling criteria.

To estimate how much energy is retained by our band-pass filtering and vertical integration operations, we compute the ratio between the energy of the filtered and unfiltered wall-normal velocity field within a given band of wall distances, as a function of the filter wavelength. It ranges from 10% to 20%, with the higher limit corresponding to the narrower filters. These values are similar to the one for the single most energetic mode in Jiménez [68]. In our case, we can study multiple sizes covering a broad range of scales, but the quantification of the total filtered energy is not straightforward because the spectral and wall-normal bands of the filters overlap. One way of getting the equivalent to a Parseval's theorem for localised basis functions would be to project the velocity onto an orthonormal wavelet basis instead of using continuous wavelets [112], but at the cost of limiting the spatial locations at which we could have information of a given scale, making it unsuitable to smoothly track the temporal evolution of the structures. Another approach is to rescale the kernels so that they recover the total energy when integrated over scale space [96], but this distorts the amplitudes, and only works for a particular spectrum. We have chosen to make our transfer functions unity at their nominal wavelength. The energy contained in the overlap of neighbouring filters is then about 2–5% (non-neighbouring bands are irrelevant because they do not overlap vertically). Accounting for overlaps, our filters

FIGURE 2.5: Snapshots of the amplitude, A_v^+ (left), and inclination angle, Ψ_v (right), of the band-pass-filtered wall-normal velocity of S4000. Top corresponds to $\bar{v}_1$ and bottom to $\bar{v}_6$. While the top snapshot has been carefully picked to contain a burst at this scale, the bottom one is randomly chosen, because all snapshots contain at least one burst. The inset amplifies a region of $\bar{v}_6$ with a burst. The amplification factor, 32, corresponds to the wavelength ratio between both filters. The contours are $A_v^+ = 0.5 \max \{A_v^+\}$.

approximately retain 20% of the total energy of the wall-normal velocity in $60 \nu/u_\tau < y < h$ at $Re_\tau = 2000$, with minor differences among Reynolds numbers.

2.3 BAND-PASS-FILTERED AMPLITUDE AND INCLINATION FIELDS

Figure 2.5 shows snapshots of the amplitude and inclination angle of the band-pass filtered wall-normal velocity, as defined in (2.17)–(2.18). Figure 2.5(a,b) uses the widest of our six filters, $\bar{v}_1$ ($\Lambda = 1.4h$), while figure 2.5(c,d) uses the narrowest one, $\bar{v}_6$ ($\Lambda = 0.044h$), which is 32 times narrower. To facilitate comparison, the insets in figures 2.5(c,d) are magnified by the ratio of the two filter widths, resulting in structures of similar size to those in figures 2.5(a,b). This supports our remark in §2.2.2 that the band-pass filter isolates structures of size proportional to Λ . The amplitude fields in figures 2.5(a,c) are smooth, with distinct intense regions which are candidates for Orr events at the moment of peak amplitude. They are highlighted as line contours in both the amplitude and inclination fields, and it is visually clear from figures 2.5(b,d) that the regions of high intensity are associated to vertical inclinations.

The joint probability density function (pdf) of A_v and Ψ_v is presented in figure 2.6(a–e), with the amplitude normalised by the modal value, A_{vM} , at

FIGURE 2.6: (a–e) Joint pdfs of the amplitude, $A_{\bar{v}}$, and inclination angle, Ψ_v , of the filtered wall-normal velocity. The amplitudes are normalised with their mode, given in (f). Contours contain 50, 90, 99% of the probability mass. (a) Joint pdf of $\bar{v}_3$ for S1000, S2000 and S4000. (b) Joint pdf of $\bar{v}_3, \bar{v}_4, \bar{v}_5$ for S1000, S2000 and S4000, respectively. (c) Averaged joint pdf for the filters that fall in the logarithmic region, $\bar{v}_2$ to $\bar{v}_6$, for the available cases, as in table 2.2. The solid lines are the averaged pdf. The dashed lines along each contour bracket the minimal area that contains all of the individual joint pdfs. (d) The solid lines are the average of all the joint pdfs of $\bar{v}_1$. The chain-dotted red line is as in figure 3 in [68]. (e) Joint pdf of $\bar{v}_3$ for F2000. The arrows are the Conditional Mean Velocity (CMV), normalised to unit length. The colours are the standard deviation of the CMVs normalised with their mean, ranging from 0.4 (dark blue) to 6.5 (red). (f) Modal values used to normalise the amplitudes. Symbols: ●, S1000; ○, L1000; ■, S2000; □, F2000; ★, L2000; ▲, S4000.

which its pdf is maximum. As in [71], the pdfs collapse better in this normalisation than with the mean. While figure 2.5 emphasises the self-similarity of the filtering operation, figure 2.6(a–c) tests the self-similarity of the filtered flow, which should hold both in inner and in outer units throughout the logarithmic region. Figures 2.6(a) and 2.6(b) each includes pdfs at three Reynolds numbers, but those in figure 2.6(a) have the same filter width in outer units ($\Lambda/h = 0.35$), and those in figure 2.6(b) have the same width in inner units ($\Lambda^+ \approx 350$). In both cases, the pdfs agree well. Figure 2.6(c) collects the pdfs for all the Reynolds numbers and all the filters whose vertical domain is contained within the logarithmic region (see table 2.2). The solid contours are the average of all the pdfs, and the dashed ones bracket the narrowest band that contains the corresponding contours of all the cases. It is clear from these figures that the similarity of the pdfs is satisfied extremely well.

Figure 2.6(d) displays the pdf of the widest filter $\bar{\nu}_1$, averaged over all the flows in table 2.2. This filter is too wide to collapse with the ones in the logarithmic layer, but all the pdfs used in figure 2.6(d) also agree well among themselves (not shown). The averaged pdf is compared in the figure with the results in [68], who analysed a single Fourier mode of wavelength comparable to Λ_1 . The good agreement validates the approximate equivalence between the new methodology and the monochromatic analysis in [68].

The collapse of the pdfs in figure 2.6 also supports that the behaviour of the ν -bursts is relatively independent of the numerical box. Figure 2.6(a–e) contains pdfs from a variety of numerical boxes, which agree well among themselves, and the pdf from [68] in figure 2.6(d) uses a very small box, $L_x/h = \pi/2$, $L_z/h = \pi/4$, which only represents well the structures below $y/h \approx 0.25$, but which is too small for the larger ones farther from the wall.

The pdfs in figure 2.6 have two distinct regions. Their core, which includes most of the probability mass, contains vertically oriented structures with amplitudes of the order of the modal value $A_{\bar{\nu}M}$. That the typical inclination of ν in channels is vertical had previously been shown using proper orthogonal decomposition by [117], and using auto-correlation functions by [153].

The upper part of the outermost isocontour of the pdfs contains large amplitudes and inclinations, and is approximately triangular. The extreme values and low probabilities in this region suggest that these points represent individual structures, while the triangular shape implies a definite statistical relation between the inclination angle and the intensity, as graphically suggested by figure 2.5. Weak regions are inclined either forward or backward, and strong ones are approximately vertical or slightly tilted backwards, as in the transient linearised bursts discussed in §2.1.1.

Moreover, the upper part of the pdf is traversed from left to right, as expected of shear-dominated structures. This was already shown to be the case for

monochromatic wavetrains in Jiménez [68], but is confirmed here for localised structures in large boxes. In the time-resolved cases, S1000 and F2000, we can define Conditional Mean Velocities (CMVs) in the $(\Psi_{\bar{v}}, A_{\bar{v}})$ parameter space as

$$\left\langle \left[\frac{D\Psi_{\bar{v}}}{Dt}, \frac{DA_{\bar{v}}}{Dt} \right] \right\rangle_{(\Psi_{\bar{v}}, A_{\bar{v}})}, \quad (2.20)$$

where $\langle \rangle_{(\Psi_{\bar{v}}, A_{\bar{v}})}$ denotes conditional averaging at $(\Psi_{\bar{v}}, A_{\bar{v}})$, and $D/Dt = \partial_t + C_\Lambda \partial_x$, is a semi-Lagrangian approximation to the total derivative that incorporates an advection velocity, C_Λ , estimated by linear regression of the time-dependent x location of the maximum of the x - t two-point two-time auto-correlation of $A_{\bar{v}}$, using at each instant a centred seven-point time stencil. In this approximation, the advection velocity is assumed to be unique for the whole (x, z) plane, independently of the number of bursts present at each moment, and depending only on the filter width. This is justified because the structures of the logarithmic layer were shown not to be dispersive by [103] and by [69], and because the uniform scale of the band-pass filtered variables would probably guarantee a uniform advection velocity even in they were. The rest of the temporal derivatives in (2.20) use five-point centred finite differences. Note that, although the advection velocity of individual structures is known to be approximately equal to the mean flow velocity, and is therefore a function of y , the velocity used here for the vertically integrated bursts is independent of the wall distance.

The result is presented in figure 2.6(e) for F2000, where the CMVs are plotted in the $(\Psi_{\bar{v}}, A_{\bar{v}})$ plane as arrows pointing toward the next most probable state. The arrows are coloured by the conditional standard deviation of the CMV, which is high in the core and lower edge of the pdf, and low in its upper edge, as in Jiménez [68]. This suggests again that the mean velocities in the upper part of the pdf are representative of individual coherent events, although it is unclear at this point whether the full periphery can be considered to represent a single burst, or whether it is formed by tangents of different burst trajectories at different locations.

While the coherent part of the pdf is traversed from left to right, its lower part is traversed from right to left, closing the cycle. This is inconsistent with linear models, but the standard deviation of the CMVs in this region is large, suggesting that it is populated by structures that cannot be characterised solely by their position in the $(\Psi_{\bar{v}}, A_{\bar{v}})$ plane, and whose evolution cannot be modelled by quasi-linear dynamics.

Finally, figure 2.6(f) shows the modal values used to normalise $A_{\bar{v}}$ in the rest of figure 2.6. Different box sizes yield the same modal value, which is

fairly constant when compared with the wall-normal velocity fluctuation averaged within the same bands,

$$(\mathcal{V}')^2 = \frac{1}{y_1 - y_0} \int_{y_0}^{y_1} (v')^2 dy. \quad (2.21)$$

This constancy is a consequence of the self-similar definition of the filter, whose transfer function has constant logarithmic width in λ_x , independently of y and Λ . Since figure 2.4 shows that the wavelengths of the spectrum of the wall-normal velocity also scale self-similarly with y , the energy selected by a truly self-similar filter should be approximately constant, but, because our band-pass filters are defined as low-pass in λ_z (see §2.2.2), the transfer function of the narrower filters spans a wider range of λ_z . The extra energy in these wavelengths explains the slightly larger modal values of the narrow filters in figure 2.6(f).

2.3.1 Wall-parallel velocities

We next compute the amplitude and inclination angles of the band-pass filtered wall-parallel velocities. Figure 2.7(a,b) shows the joint pdfs of (Ψ, A) for $\bar{u}$ and $\bar{w}$, respectively. Their structures are preferentially tilted towards the direction of the shear, in contrast to the wall-normal velocity structures, which are almost equally distributed between forward and backward inclinations. This is more pronounced for the spanwise velocity, which is rarely tilted backwards (10% probability), and subtler for the streamwise component, which is tilted backwards 40% of the time. As with the wall-normal velocity, there is a statistical correlation between the amplitude and inclination of the spanwise velocity, but the same is not true for the streamwise velocity, for which the two quantities are essentially unrelated. Scrambling $A_{\bar{u}}$ with respect to $\Psi_{\bar{u}}$ leaves their joint pdf unchanged (not shown).

The behaviour of the inclination angles and amplitudes of $\bar{u}$ and $\bar{w}$ in the upper part of their joint pdfs is qualitatively similar to the linearised trajectories in figure 2.1(b), although the inclination angles differ quantitatively. For example, the spanwise velocity shares with the transient-growth model an amplitude ‘peak’ towards positive inclinations, but the tilting angle of the linearised peak is always larger than that of the DNS by 0.1–0.2 radians. Using unconditional two-point correlations of the wall-normal and spanwise velocity components, [69] argued that the cross-stream velocities contributing to the correlation are part of a quasi-streamwise ‘roller’. Some of the features of those auto-correlation functions are shared by the intense core of the joint pdfs in figures 2.6 and 2.7. For example, the auto-correlation of the wall-normal velocity in the logarithmic layer is approximately vertical ($\Psi_{\bar{v}} \approx 0$), and the spanwise velocity is tilted forward by $\Psi_{\bar{w}} \approx \pi/6$, in agreement with the amplitude peaks of the joint pdfs of $\bar{v}$ and $\bar{w}$. Because the correlation

FIGURE 2.7: Joint pdfs of A and Ψ involving $\bar{u}_3$ and $\bar{w}_3$ for S1000. All amplitudes are normalised with their modal values. (a) $\Psi_{\bar{u}}$ and $A_{\bar{u}}/A_{\bar{u}M}$, contour levels are 90%, 50% and 10% of the probability mass. (b) $\Psi_{\bar{w}}$ and $A_{\bar{w}}/A_{\bar{w}M}$, contours as in (a). (c–f) Black contours are the joint pdf of $A_{\bar{v}}/A_{\bar{v}M}$ and $\Psi_{\bar{v}}$, contours containing 99.9%, 99%, 90% and 50% of the probability mass. Shaded contours are the mean of several quantities conditioned to the value of $A_{\bar{v}}$ and $\Psi_{\bar{v}}$. (c) Conditional average of $A_{\bar{u}}/A_{\bar{u}M}$. (d) $A_{\bar{w}}/A_{\bar{w}M}$. (e) $\Psi_{\bar{u}}$. (f) $\Psi_{\bar{w}}$.

NUMBER OF BURSTS	DATASET	Λ_1	Λ_2	Λ_3	Λ_4
In total	S1000	70	228	1301	9777
In total	F2000	863	5098	34802	274265
Per time-area in S_Λ/Λ_i^2	S1000	0.267	0.106	0.085	0.091
Per time-area in S_Λ/Λ_i^2	F2000	0.16	0.114	0.106	0.111

TABLE 2.3: Number of bursts found in different band-pass filtered time series. See table 2.1 for details of the simulations and section §2.3.2 for the details of the time-area normalisation.

function is dominated by strong events, these similarities are not surprising, but they confirm that the band-pass filter retains some of the structure of the intense events of the velocity.

To explore the relation between the different variables during the bursting cycle, we compute averages conditioned to $(\Psi_{\bar{v}}, A_{\bar{v}})$. The conditional mean amplitudes of $\bar{u}$ and $\bar{w}$ are presented in figures 2.7(c,d), and their conditional inclinations are shown in figure 2.7(e,f). If we assume that the CMVs in figure 2.6(e) reflect the temporal evolution of individual strong bursts, we can describe the burst by the evolution of the conditional mean values of the different quantities as we move from left to right along the upper edge of the pdf of $(\Psi_{\bar{v}}, A_{\bar{v}})$. The burst starts from intense $\bar{u}$, with weaker $\bar{v}$ and $\bar{w}$. At this stage $\bar{u}$ and $\bar{v}$ are tilted upstream, but $\bar{w}$ is too weak for its inclination to be defined. The ambient shear then tilts everything forward, and $\bar{v}$ and $\bar{w}$ are amplified while the amplitude of $\bar{u}$ decreases. Beyond the point where $\bar{v}$ is vertical, the amplitude of the wall-normal velocity decreases again, but $\bar{w}$ reaches its maximum value. The interpretation of these observations will be deferred to §2.3.2, after we have examined the conditional temporal behaviour of individual bursts.

2.3.2 Conditional mean evolutions

While in the previous section we inferred a hypothetical burst evolution from the statistical relations among the different quantities along the upper edge of the joint pdf of $(\Psi_{\bar{v}}, A_{\bar{v}})$, we now turn our attention to the direct study of their conditional temporal evolution by examining structures that pass through the upper ‘tip’ of the joint pdf,

$$A_{\bar{v}} > 2A_{\bar{v}M}, \quad |\Psi_{\bar{v}}| < 0.15, \quad (2.22)$$

at some stage of their lives. We know from figure 2.6(e) that the standard deviation of the CMVs passing through these very intense events is low, and that the statistics along the upper edge of the pdf are approximately the same as those of linearised Orr bursts. This suggests that they represent the

FIGURE 2.8: (a) Snapshot of the conditional bursts of $\bar{q}_2$ at $S_\Lambda T = 0$, the peaking time of $\bar{v}_2$. The contours, normalised with the maximum of each component are: shaded, $A_{\bar{v}} = (0.6, 0.8)$; solid $A_{\bar{u}} = (0.8, 0.9)$; dashed $A_{\bar{w}} = (0.8, 0.9)$. (b–e) Different aspects of the evolution of A and Ψ , conditioned to the presence of a burst. F2000. (b) Temporal evolution of the maximum amplitude, as a function of time. —, conditional $\bar{u}_2$; ---, conditional $\bar{u}_3$. Symbols are $\bullet$, $\bar{u}$; $\blacksquare$, $\bar{v}$; $\blacktriangle$, $\bar{w}$. (c) As (b), for the deviations of the inclination angle respect to their mean, at the position of the maxima of each component at every instant. (d) Transient amplification of a linearised Orr burst. Conditions are as in figure 2.1(b), but the energies are computed over the range of y corresponding to the Λ_2 filter. (e) The shaded contours are the joint pdf of $(\Psi_v, A_{\bar{v}})$ for Λ_2 . — $\bullet$ —, Evolution in (b–c), conditioned to (2.22); - - $\bullet$ - -, conditioned to the higher threshold (2.27); — $\blacksquare$ — evolution in (f–g), conditioned to (2.28). The five markers are $S_\Lambda T = [-1.5, -0.3, 0, 0.3, 1.5]$. (f, g) As in (b, c), but conditioned to a burst precursor.

evolution of individual bursts, at least as long as the amplitude threshold is chosen high enough to separate bursts from each other. The restriction in (2.22) to a small inclination angle is not strictly necessary, since it follows from figure 2.6 that intense wall-normal velocity structures are almost always vertical, but including it allows us to relax somewhat the requirement for large amplitudes, and to identify bursts that would be too weak otherwise. The analysis is most easily performed in a convective frame of reference, $(\xi = x - C_\Lambda t, z, t)$ where C_Λ is the convection velocity used in (2.20). We identify the position and time, (ξ_c, z_c, t_c) , of the maximum amplitude for each region as representative of the ‘peaking’ point of the burst, and use it as the centre for our conditionally averaged evolution $\langle \cdot \rangle_B$, defined as

$$\langle a \rangle_B(\Xi, Z, T) = \frac{1}{N} \sum_{c=1, \dots, N} a(\xi_c + \Xi, z_c + Z, t_c + T), \quad (2.23)$$

where N is the number of identified bursts for each filter (see table 2.3), and a is either $A_{\bar{u}, \bar{v}, \bar{w}}$ or $\Psi_{\bar{u}, \bar{v}, \bar{w}}$. To avoid ‘false positives’, we discard small bursts for which $\mathcal{S} < 2.25\Lambda^2$, where $\mathcal{S}$ is the temporally averaged area enclosed in the (x, z) plane by the threshold in (2.22). Independently of the chosen threshold, approximately 5% of the identified bursts are temporally clustered in quick succession, reminiscent of the ‘packets’ of vortical structures reported by [3], but most cases cannot be easily grouped into such packets.

Bursts defined in this way are spatio-temporal objects in (ξ, z, t) . To normalise their density per unit duration-area, we assume that two bursts cannot occupy the same volume in space-time, and that they scale spatially with Λ , and temporally with the inverse of the average shear across their integration band (2.17),

$$S_\Lambda = (U(y_1(\Lambda)) - U(y_0(\Lambda)))/(y_1(\Lambda) - y_0(\Lambda)). \quad (2.24)$$

Their expected duration-size would then be proportional to Λ^2/S_Λ , and the number of detected bursts should collapse when their total duration-area is normalised in these units. Table 2.3 shows that this is approximately true, with a density close to 10% except for the largest filter, which is too large for the scaling to hold. This density is of the same order as the volume fraction of the ‘Q’ structures studied in [101].

Figure 2.8(a) shows a snapshot of the conditional burst of $\bar{q}_2$ in F2000 at peaking time. Other filters collapse well with Λ , with average dimensions $(\ell_\Xi/\Lambda, \ell_Z/\Lambda) \approx (2, 1)$. The filtered spanwise velocity can be seen to be offset to one side of the streamwise and wall-normal components, while the latter fall on top of each other. If the average had been computed strictly as in (2.23), the conditional mean evolution would have two symmetrical lobes of $A_{\bar{v}}$ centred on the burst, but this symmetry is statistical, and does not imply the symmetry of individual events. The equations of motion and boundary conditions of the channel are invariant to reflections with respect to

(x, y) planes, so that any solution $(u(x, y, z), v(x, y, z), w(x, y, z))$ implies that $(u(x, y, -z), v(x, y, -z), -w(x, y, -z))$ is also a solution. Individual Orr events are equally likely to be chiral-left or chiral-right, and the reflection symmetry can be applied to each snapshot to obtain an average that is more representative of individual events [101, 161]. There is no unique way of deciding how to do this consistently for complete histories, and the criterion that we found to produce the least symmetric conditional mean was to always locate to the left of the burst the strongest maximum of $A_{\bar{w}}$ as it passes through the spatio-temporal neighbourhood of its peak,

$$|\Xi| \leq 0.2\Lambda, |Z| \leq \Lambda, S_\Lambda|T| \leq 0.2. \quad (2.25)$$

The conditional histories computed in this way only contain one region of strong w , as in figure 2.8(a), suggesting that most of the underlying events have only one strong spanwise velocity region, corresponding to a single streamwise roller. Double rollers are found more seldom. Visual inspection of a representative sample of cases shows that approximately 5% of the bursts are almost symmetric, with two w lobes.

Figure 2.8(b) shows the conditional evolution of the maximum amplitude of the three filtered velocity components as a function of time. Far from $S_\Lambda T = 0$, the three components tend to their unconditional mean, which is used to normalise the plot. The wall-normal velocity, which is used to condition the temporal evolution, is amplified the most, but the stream- and span-wise velocities are also amplified, peaking at $S_\Lambda T \approx \pm 0.3$, respectively before and after the peak of $\langle A_{\bar{v}} \rangle_B$, as also suggested by figure 2.7(d,e). The evolution of the deviation of the inclination angles from their unconditional means,

$$\langle \psi_a \rangle_B = \langle \Psi_a \rangle_B - \langle \Psi_a \rangle, \quad (2.26)$$

is presented in figure 2.8(c), which plots the inclination at the location of the maximum amplitude of each component. The evolution of the inclination angles of the streamwise and wall-normal velocity is from negative to positive, and tends to the unconditional mean beyond $|S_\Lambda T| \approx 1$. The changes in $\Psi_{\bar{w}}$ are less pronounced, with values closer to its unconditional mean across the whole evolution.

A reference to figure 2.1 shows that this behaviour is qualitatively similar to linearised equilateral Orr bursts, but there are some interesting differences. The inclination angles of $\bar{u}$ and $\bar{v}$ are proportional to each other in the upper part of the pdf of $(\Psi_{\bar{v}}, A_{\bar{v}})$, where the presumed quasi-linear cycle takes place, but the inclination angle of $\bar{w}$ is not. While $\Psi_{\bar{u}}$ evolves from negative to positive both in figures 2.1 and 2.7(e), the inclination $\Psi_{\bar{w}}$ is never negative in figures 2.7(f) and 2.8(c). Moreover the conditional amplification of $\bar{u}$, as given by the ratio between the highest and lowest conditional averages in figure 2.7(c), is almost three times higher than the one for $\bar{w}$ in figure

2.7(d), while the opposite seems to be true in figure 2.1(b), and the maximum amplitude of the two variables in figure 2.8(b) are approximately equal. The reason turns out to be the different spatial location of the various velocity components. We saw in figure 2.8(a) that, while the structures of $\bar{u}$ and $\bar{v}$ are spatially collocated in (x, z) , the $\bar{w}$ -eddy is offset to one side, resulting in a weaker footprint in the conditional joint pdfs. Even more important is the vertical offset. The amplitudes in figure 2.1(b) are integrated over the whole channel, but those in figure 2.7 are band-pass filtered. This includes the wall-parallel filtering by wavelength, which has no effect on the monochromatic wavetrain in figure 2.1, but also the restriction in (2.17)–(2.18) to a band of wall distances, which changes the balance among the different components as they drift vertically in and out of the filtering band. This effect is visually clear from figure 2.1(a), and figure 2.8(d) shows that it can be reproduced by restricting the amplitudes in figure 2.1(b) to the same band of wall distances as in figure 2.7. It was already mentioned when discussing figure 2.1 that the definition of the angles is biased by the different amplitude distribution of the three velocity components, and it is now clear that so is the definition of the band-pass amplitudes. Thus, while figures 2.7(c,d) and 2.8(b–d) might give the impression that the earlier part of the burst is dominated by the association of u and v (such as in a spanwise-oriented roller), while the later part is dominated by v and w (a streamwise roller), this behaviour is probably an artefact of how the flow is filtered. Moreover, it should be remembered that linear models cannot change the wall-parallel wavevector of a single mode, nor its wall-parallel orientation.

Figure 2.8(b) includes results for two different filters, which scale well when normalised with the mean shear (2.24), and imply lifetimes of the order of $S_\Lambda T \approx 1$, measured at one-half the maximum amplification. Using slightly different definitions of mean shear and lifetime, [103] and [68] report the somewhat longer value $S_\Lambda T \approx 5$. In fact, the evolution of the inclination angle in figure 2.8(c) suggests that the conditional evolution only represents individual bursts when they are relatively near the conditioning time, $S_\Lambda |T| \lesssim 0.5$, beyond which both the angle and the amplitude tend to their unconditional values.

This is better seen in figure 2.8(f) which displays the $(\Psi_{\bar{v}}, A_{\bar{v}})$ trajectory of figure 2.8(b,c) on top of the joint $(\Psi_{\bar{v}}, A_{\bar{v}})$ pdf. The trajectory follows the direction of the CMVs near the conditioning point, but soon drops into the high-probability region at the core of the pdf. Note that the trajectory crosses the lowest possible point consistent with (2.22). Most eddies are in the core of the joint pdf, and any conditional mean is dominated by the points closest to the core. In a similar way, once the effect of the condition is lost, trajectories

naturally fall towards the probability maximum. Thus, if condition (2.22) is substituted by the higher-amplitude one,

$$A_{\bar{v}} > 3A_{\bar{v}M}, \quad |\Psi_{\bar{v}}| < 0.15, \quad (2.27)$$

the conditional trajectory changes to the qualitatively similar, but higher-intensity, trajectory plotted as dashed in figure 2.8(f).

Perhaps the most interesting feature of the conditional evolution is the peak of $\bar{u}$ preceding the burst, also represented by the high values of the conditional amplitude of $\bar{u}$ over the left part of the joint pdf of $(\Psi_{\bar{v}}, A_{\bar{v}})$ in figure 2.7(c, e). Note that a strong $A_{\bar{u}}$ should not be interpreted as a especially strong streak of the streamwise velocity, because the band-pass filter is designed to capture the energy of v , and the structures of u are between six and ten times longer than those of v [69]. The band-pass filtered amplitude of a uniform infinite streak is zero, and we should consider $A_{\bar{u}}$ as a measure of the ‘meandering’, or other forms of inhomogeneity of the streak. Streak inhomogeneities have often been identified as important features of the logarithmic layer cycle [48, 53].

The previous discussion shows that a burst of $\bar{v}$ is preceded by backwards-leaning perturbations of $\bar{u}$. The converse, that backwards-leaning perturbations of $\bar{u}$ act as precursors of the bursts of $\bar{v}$, is tested in figure 2.8(f,g) by the conditional evolution of bursts conditioned to

$$A_{\bar{u}} > 1.8A_{\bar{u}M}, \quad \Psi_{\bar{u}} < -0.4, \quad (2.28)$$

which is on the left edge of the joint $(\Psi_{\bar{v}}, A_{\bar{v}})$ pdf, but relatively far from its top. As shown in the figure, the conditional burst develops as predicted, with qualitatively similar temporal relations and delays among the different components as in figure 2.8(b,c). It is particularly striking that, even if figure 2.8(e,f) is conditioned on $\bar{u}$ rather than on $\bar{v}$, the amplification of the latter is even stronger than that of the former, suggesting that at least the left part of the joint pdf represents well organised bursts. A similar conclusion was reached by [68], who showed that burst could be linearly ‘predicted’ from conditions up to about half their lifetime before the peak amplification of v . The mean trajectory in $(\Psi_{\bar{v}}, A_{\bar{v}})$ space of the evolutions initialised within (2.28) is plotted with squares in figure 2.8(e). Attempts to use the later peak of $A_{\bar{v}}$ to ‘postdict’ the burst of $\bar{v}$ were not successful. This may be interpreted as that strong, forward-leaning perturbations of $\bar{w}$ are not generated solely by Orr bursts, but it is most probably just a reflection of the irreversibility of the Navier–Stokes equations, which are ill-posed when run backwards in time.

	Λ_1	Λ_2	Λ_3	Λ_4
Number of sweeps	637	2898	17079	147686
Number of ejections	797	3282	18159	160207
Number of pairs	410	1717	9995	82009
Sweeps per time-area in S_Λ/Λ^2	0.118	0.065	0.052	0.06
Ejections per time-area in S_Λ/Λ^2	0.147	0.073	0.055	0.065

TABLE 2.4: Number of bursts found in the low-pass-filtered temporal series of F2000. The details of the time-area normalisation are in §2.3.2.

2.4 LOW-PASS-FILTERED VELOCITY FIELDS

We mentioned in §2.2.2 that, while the band-pass-filtered field (2.10) is useful in isolating the amplitude and inclination of eddies of a given size, it is not easily interpreted as a velocity. As a consequence, the conditional results in the previous section give only limited information about the flow structure. In particular, the local sign of the velocities is missing, and so is the relation of individual features with the flow around them. Both are important. Sweeps (negative wall-normal velocity events) and ejections (positive wall-normal velocity events) are known to have different characteristics [101, 104, 178], and this asymmetry is strongest for the wall-attached structures responsible for most of the tangential Reynolds stress [39]. Similarly, bursts of the wall-normal velocity are known to be associated with streamwise-velocity streaks [39, 101], but the two velocity components have very different sizes. Figure 2.9(a) shows that the spectrum of u in the logarithmic region is much longer than that of v , and it is difficult to study the interaction of the two variables if they are band-passed to a single scale.

Both deficiencies are substantially remedied by the low-pass-filtered velocity (2.7), which is easily interpreted as a smoothed flow field, and retains the largest features of the flow, including the long streaks of u . Unfortunately, these low-passed fields are not approximate wave trains or packets, and we lose the information about the local inclination angle used in (2.22) as part of the condition to identify bursts. However, we saw when discussing that condition that the inclination angle was only intended to relax the intensity identification threshold, since figure 2.6 shows that high amplitudes are unlikely to be anything but vertical. As a consequence, we study in this section conditional flow histories conditioned only on the intensity of the events, and expect their inclination, if any, to emerge as a consequence of that conditioning. Since this requires both a large computational box to contain the largest flow structures, and temporal information, the rest of the section only uses the temporally resolved simulation F2000.

We create filtered fields using the low-pass filters described in §2.2.2 and, to obtain structures that are approximately equivalent to the band-passed ones in §2.3.2, integrate the resulting velocity over the same y -bands. The resulting two-dimensional fields recall the amplitudes studied in §2.3.2, but contain intense regions of negative as well as of positive velocity. Therefore, the process of isolating intense regions of the wall-normal velocity now consists of two independent thresholding operations, $\tilde{v}_i > 2\tilde{v}'_i$ for ejections, and $\tilde{v}_i < -2\tilde{v}'_i$ for sweeps, where $\tilde{v}'_i$ is the root-mean-square intensity of the filtered time series. This results in two sets of intense events, which are treated independently, and the rest of the section includes averages conditioned to one or to the other. As a consequence, the presence of a sweep in a flow conditioned to ejections should be considered a feature of the flow, not of the conditioning, and vice versa.

As in §2.3.2, the purpose of the threshold used to define the sweeps and ejections is mostly to separate individual events, and its value is not critical. We also discard bursts which are too small, and, in addition, merge into a single object bursts of the same sign whose centres, (ξ_c, z_c, t_c) , defined as the location of their maximum intensity, are too close to each other. The centre of the resulting object is taken to be the centre of the stronger of the two eddies being merged. We saw in §2.3.2 that the size of the bursts is $\ell_{\Xi}/\Lambda \sim \ell_Z/\Lambda \sim Tu_{\tau}/\Lambda \sim O(1)$, and, after some experimentation, choose as the merging criterion that

$$\sqrt{(\Delta t_c u_{\tau})^2 + (\Delta \xi_c)^2 + (\Delta z_c)^2} < 0.15\Lambda. \quad (2.29)$$

The number of centres determined in this way is given in table 2.4 for each filter size. Centring each burst on its defining extremum, the conditional structures are computed as in (2.23), using the four-dimensional filtered histories without integrating them in y , to retain the wall-normal burst structure,

$$\langle a \rangle_L(\Xi, y, Z, T) = \frac{1}{N} \sum_{c=1, \dots, N} a(\xi_c + \Xi, y, z_c + Z, t_c + T). \quad (2.30)$$

For each burst, we use a centred four-dimensional box spanning the channel half-height, $y = (0, h)$, the wall-parallel box with dimensions $\Xi/h = \pm 4\pi$, $Z/\Lambda = \pm 3$, and a time interval consistent with the burst lifetimes discussed in §2.3.2, $S_{\Lambda}T \approx \pm 5$. The conditional mean evolutions for the different filters are found to collapse spatially when normalised with the height of the most intense point of the conditional evolution, y_c , which can be interpreted as the distance from the wall of the centre of the conditional burst at the moment of its highest intensity. Figure 2.9(b) shows that y_c increases linearly with the filter width as $y_c^+ = \Lambda^+/2 + 100$, corresponding to bursts whose central height is proportional to $\Lambda/2$, with an offset of $100\nu/u_{\tau}$ representing

FIGURE 2.9: (a) Spectra of: —, the streamwise velocity; ---, the wall-normal velocity, integrated over $y/h \in [0.2, 0.4]$. Contours contain 50% and 10% of the integrated spectral mass. (b) Distance from the wall of the point of highest intensity of the conditional burst. ● ejections; ▲ sweeps. The diagonal line is $y_c/h = 0.5\Lambda/h + 0.05$.

the buffer layer. Because of this offset, the spatial structure of the bursts collapses better with y_c than with Λ , and this scaling will be used in the rest of the section.

Figure 2.10 provides three orthogonal sections of the three-dimensional structure of the bursts at $S_\Lambda T = 0$, and shows that the scaling with y_c is extremely good, only interrupted by the channel centreline. The largest filters produce ‘cropped’ bursts at $y = h$, but they agree well below that height. The dotted lines in each figure show the location of the other two sections. Individual ejections and sweeps are often accompanied by a single strong structure with opposite wall-normal velocity [101], sitting at $\bar{\Xi} \approx 0$, $Z/y_c \approx \pm 1.2$. Inspection of a representative sample of individual bursts shows that only about 5% of them have more than one companion, and are approximately symmetric. As in §2.3.2, the bursts in (2.30) are oriented to retain as much as possible the asymmetry of this arrangement in the conditional mean. This is done by placing on a negative Z the strongest ejection found in the neighbourhood

$$|\bar{\Xi}| \leq 0.2\Lambda, |Z| \leq \Lambda. \quad (2.31)$$

of the conditioning sweep (or vice versa), reflecting the velocity field as required. For all the filters and conditions tested, the companion appears in the conditional evolution as a single opposite-signed ‘partner’ with half the intensity of the main structure.

Small differences can be seen between the conditional sweeps and the ejections. For example, the ejections in the bottom panel of figure 2.10(b) have small upstream ($\bar{\Xi} < 0$) ‘tails’ near the wall, while the sweeps in the upper panel have a more rounded downstream ‘nose’ farther from the wall.

FIGURE 2.10: (a,b,d) Different sections of the conditional structure of $\tilde{v}$ at the peak of the ν -burst, $S_\lambda T = 0$, for the four filters. F2000. —, positive wall-normal velocity; ---, negative. The contours are (-0.15, 0.15, .45) of the corresponding extremum. In each figure the top panel is conditioned to the sweep, and the bottom one, to the ejection. The lines are: blue, $\tilde{v}_1$; orange, $\tilde{v}_2$; green, $\tilde{v}_3$; red, $\tilde{v}_4$. (c) Autocorrelation function of the wall normal velocity, $C_{vv}(\Delta x, y, y')$ at $y' = 0.05h = 100\nu/u_\tau$. The shaded contours are the unconditional autocorrelation function; —, autocorrelation only for positive ν events; ---, autocorrelation only for negative ν events. Contours are (0.15, 0.3, 0.6) of the maximum correlation in each case.

FIGURE 2.11: (y, Z) -section of the conditional mean structure of $\tilde{q}_2$ at the peak of the v -burst, $S_\lambda T = 0$. F2000. The shaded contours are streamwise velocity perturbations, $\tilde{u}_2^\pm = \pm(0.5, 0.9, 1.3)$. The arrows are the cross-plane velocities, with values in $(\tilde{v}_2^2 + \tilde{w}_2^2)^{1/2} = [0, u_\tau]$, represented by the arrow length.

Interestingly, both features are also found in the autocorrelation function of the wall-normal velocity, added to figure 2.10(b) as a shaded area, suggesting that different regions of the autocorrelation function of the wall-normal velocity are controlled by perturbations of different sign.

A similar effect was found by Sillero et al. [153] for the (x, z) sections of the autocorrelation of the spanwise velocity. He found them to be approximately square, and showed that, when the correlation is conditioned to only positive or negative values of w at the reference point, the correlation separates into two almost diagonal patterns, which form the square when added together. Here, sweeps and ejections contribute to the ‘nose’ and ‘tail’ of the unconditional correlation of v , respectively. This is confirmed in figure 2.10(c) which shows the autocorrelation functions of v at $y/h \approx 0.05$,

$$C_{vv}(\Delta x, y, y' = 0.05h) = \langle v(x + \Delta x, y)v(x, y' = 0.05h) \rangle, \quad (2.32)$$

where the average $\langle \cdot \rangle$ is performed in three different ways: unconditionally (shaded in the figure), and conditioned to either negative (dashed lines) or positive (solid) values of v at the reference point $v(x, 0.05h)$.

Figure 2.10(e) shows that the wall-parallel sections of the conditioning burst are roughly circular, whereas the conditional partners are bean-shaped, but we will show in §2.4.1 that the mismatch between the shapes of the bursts and of their partners is almost surely an artefact of the conditional average, not of the instantaneous flow fields.

Figure 2.11 shows a cross-flow section of the conditional $\tilde{q}_2$ field. The figure is conditioned on v , and the shaded contours are streamwise velocity perturbations. The sign of the primary u structure (near $Z = 0$) is opposite to the one of v , forming a classical $(Q_2, u < 0, v > 0)$ or $(Q_4, u > 0, v < 0)$ eddy. Both cases have a weaker partner of opposite polarity at $Z/y_c \approx -1.2$,

and the complete structure shares many features with the conditional pairs of uv -Qs in Lozano-Durán et al. [101] and [39], which were conditioned on their intensity, without regard to their temporal evolution. Those authors found that strong Qs can be classified as ‘attached’, with roots that extend very near the wall, or ‘detached’, which do not. The attached Qs are self-similar, as in the present case, form Q_2 – Q_4 pairs of size comparable to the present ones and, when conditionally averaged by centring them on the centre of gravity of the pair, they also contain a roller between streaks. They are responsible for most of the tangential Reynolds stress in the flow. Similar rollers have been identified as part of the self-sustaining cycle of the buffer layer, where they tend to be almost parallel to the wall [56, 73–75, 150], and are also believed to be important in the logarithmic layer [48, 53]. In our case, the roller remains inclined in the $(\bar{E}, y)$ -plane at approximately 15° during the evolution (not shown), sitting between the pair of streaks. This is similar to the inclinations previously found from the autocorrelation of u [see 153, for a review]. The similarity between all these structures suggests that bursts, Qs and quasi-streamwise vortices are different manifestations of the same phenomenon.

Up to now, we have described bursts at the moment of their maximum intensity, but figures 2.12(a) and 2.12(b) show the time evolution of a constant isosurface of $\tilde{v}_2$ for the conditional evolution of an ejection and a sweep, respectively. The burst is always located at $Z/y_c \approx 0$, and its partner structure at $Z/y_c \approx -1.2$, as in figure 2.10(d). The conditional evolution spans a total of $S_\Lambda T = 8$, clearly showing an Orr-like evolution of the inclination angle and amplitude of the burst in both cases. The partner structure is present during the full evolution, and is amplified by approximately the same amount as the conditioning burst, but with half its average intensity. Figure 2.12(c,d) shows the evolution of the streamwise velocity streaks during the same conditional event. They are considerably longer ($\sim 20y_c$) than the v -bursts ($\sim 3y_c$) and, unlike the latter, which have a single secondary lateral structure, they have a pair of comparable streaks of the opposite sign at each side. However, although there is some amplification of the streak during the burst, evidenced in the figure by the thickening and lengthening of their u isosurface, it is less clear than the amplification of v . In particular, the streaks are already present in the conditioning volume when the v -burst begins to form, and remain in it when the burst disappears.

Figure 2.13 shows longitudinal $Z = 0$ sections of the conditional evolutions in figure 2.12, at different moments during the burst. To be able to compare it with the band-pass filtered conditional evolution in §2.3.2, we also provide contours of the band-pass filtered streamwise velocity, obtained by filtering the conditional field (which already is an average of low-pass-filtered fields) with a high-pass filter chosen so that the scales retained by the combined

FIGURE 2.12: Mean evolution of $\tilde{v}_2$ and $\tilde{u}_2$ conditioned to the presence of a burst in F2000. The snapshots range from $S_\lambda T = -4(2)4$. The flow moves in the direction of the dashed arrow, an the mean advection between snapshots has been removed. (a, b) The blue surface corresponds to $\tilde{v}_2^+ < -0.15$ and the red to $\tilde{v}_2^+ > 0.15$. (c, d) Conditioned to an ejection. (c, d) As in (a, b), but for the evolution of $\tilde{u}_2$; the surfaces are, blue, $\tilde{u}_2^+ < -0.15$; red, $\tilde{u}_2^+ > 0.15$.

FIGURE 2.13: (a–c) $Z = 0$ section of the low-pass-filtered low-speed streak (dashed lines, $|\tilde{u}_2^+| = [0.15, 0.3]$) and its associated ejection (shaded contours, $|\tilde{v}_2^+| = [0.15, 0.3, 0.6]$). The solid lines are the shorter wavelengths of the filtered streamwise velocity, comparable to the band-pass-filter velocity in §2.3, as explained in the text. (a) $S_\Lambda T = -2$. (b) $S_\Lambda T = 0$. (c) $S_\Lambda T = 2$. (d–f) As in (a–c), but for the high-speed streak, and the wall-normal velocity sweep. F2000.

effect of the two filters is comparable to the band-pass filter in §2.3. Filtering the streamwise velocity in this way reveals an inner ‘core’ of the streak. While the longer ‘body’ of the streak is always tilted forward, the much-shorter core tilts both backwards and forwards in synchrony with the wall-normal velocity. This confirms the different nature of the short (band-pass filtered) and long (low-pass filtered) perturbations of the streamwise velocity discussed in §2.3. It is worth noting that the sweep is located towards the front of the high-speed streak, whereas the ejection tends to be closer to the back of the low-speed streak, recalling a similar arrangement of streaks and vortices in the buffer layer [see figure 9b in 70].

Before interpreting these results, it is important to understand that figures 2.12(c,d) and 2.13 do not represent instantaneous streaks, or even averaged ones, but the part of the streak conditioned to the presence of a strong burst of v . As such, the intensity of the structures in those figures reflects both the intensity of the fluctuations of u , and how is u correlated to the burst. For example, the isolines in figure 2.13 do not represent the intensity of u , and can not be used to estimate that intensity except probably near the inner core. Any attempt to derive from them whether the burst is a consequence of the presence of the streak, or the other way around, is bound to be speculative. In the same way, the fact that the streaks in figure 2.12(c,d) are not seen to strengthen during the burst does not mean that they do not do so (or vice versa).

With these restrictions in mind, the simplest interpretation of the positional bias in figure 2.13 is that the streaks are ‘wakes’ created by sweeps and ejections advected by the mean flow at the wall-distance at which they originate [70]. Sweeps, coming from above, move faster than the local mean speed and bring high-speed fluid down, leaving a high-speed wake upstream. The effect of the ejections is the opposite. However, it was argued by [10] that this explanation is unlikely, because the v -bursts do not live long enough. Thus, if we take their lifetime to be the same as for Qs, $T^+ \approx 2y_c^+$ and the velocity difference to be $O(1.5u_\tau)$ [103], the maximum length of their wake would be $O(3y_c)$. This is the length of the inner core in figure 2.13, but much shorter than the length of the streak. A more likely possibility is that each streak is created by several bursts, each of which contributes a small fraction to its length [10]. The fact that the conditional streaks weaken so little away from the peak of the burst in figure 2.12(c,d) also strongly suggests that the streaks are stable features of the flow, while the burst grows and eventually disappears within them. This is supported by the known streamwise distance between consecutive Qs, $\Delta x \approx 7y_c$ [39, 101], which is much shorter than the streak length. These dimensions will be confirmed for the bursts in the next section.

In fact, if we accept the argument above that very long wakes cannot be a consequence of short-lived bursts, the scaling of figure 2.13 with the filter width can be used to extract some causal information. We saw in discussing figure 2.10 that the size on the ν -burst is proportional to $y_c \propto \Lambda$. The inner core of the streak in 2.13 also scales with Λ , which is not surprising because it is created by a pseudo-band-pass filter of that size. The causal information is contained in the scaling of the rest of the streak, which is obtained with a low-pass filter, and therefore contains all the scales larger than Λ . Although not shown in the figure to avoid clutter, most of the streak dimensions scale with h , not with Λ , suggesting that most of the streak it is not caused by the burst, nor does directly causes it. The only regions that scale with Λ are the *upper edge* of the leftward tail of the low-speed streak in figure 2.13(a–c), and the *underside* of the rightward nose of the high-speed streak in figure 2.13(d–f). Both regions are approximately indicated by a dotted rectangle in figure 2.13(b,e).

We may now come back to the question of what causes the positional bias in figure 2.13, which is part of the wider question of what causes sweeps and ejections to be organised along streaks. It is generally accepted that the bursts create streaks by deforming the mean profile, but the discussion in the previous paragraphs implies that something else organises the bursts so that the short streak segments join into longer objects. The possibility that streak instability is responsible for the bursts has been mentioned often, although the detailed mechanism is unclear [see discussions in 46, 150], as well as the possibility that long streaks are compound objects [69]. However, most of these analyses deal with infinite uniform streaks, and cannot explain a preferential longitudinal distribution of the bursts within them. Although the question of which are the original perturbations that give rise to the formation of bursts, is beyond the scope of the present paper, an obvious suggestion of the above observations is that bursts are preferentially created at the ‘active’ end where the streak ‘collides’ with the ambient flow, either along the upper surface of the back of the slow streaks, as they are overrun by the faster flow behind, or along the bottom of the ‘nose’ of the faster flow, as it overruns the slower one.

A slightly different interpretation of the same data is that the streak interaction does not take place at the end of a streak, but in the front of a meander in which high-speed flow pushes into a low-speed one. The conditional data in figure 2.13 are not enough to distinguish between those two possibilities, even when inspected in other flow sections or in three-dimensional views, and all that can probably be said is that bursts tend to be created at pre-existing longitudinal inhomogeneities of the streaks.

FIGURE 2.14: Probability mass functions (pmfs) of the number of partners of each structure. (a) Sketch of the definitions of pair, partner and number of partners (node degree). (b) Pmf of the number of sweep-partners for ejections (open symbols) and ejection-partners for sweeps (closed symbols). (c) Pmf of the number of ejection-ejection partners (open symbols) and sweep-sweep partners (closed). (a–b) $\bullet$, Λ_1 ; $\blacksquare$, Λ_2 ; $\blacktriangle$, Λ_3 ; $\blacktriangledown$, Λ_4 . The solid line is the corresponding distribution for a Poisson process.

2.4.1 Space-time organisation of sweeps and ejections

It remains unclear from the previous discussion whether the secondary ν structure that appears to one side of the conditioning burst in figures 2.10 and 2.12(a,b) is an independent Orr burst, a true companion of the conditioning burst, or another unrelated structure. However, it seems unlikely that ν ‘monopoles’ of any given sign, strong enough to be reflected in the conditional average, appear isolated in the flow, because continuity requires that the mean wall-normal velocity over patches of dimensions comparable to the distance from the wall should vanish. Our baseline hypothesis, motivated in part by the similarity of the conditional averages in figure 2.10 to the conditional Qs in [101] and [103], is that bursts appear in side-by-side pairs of opposite sign, and that the secondary structure in figures 2.10 and 2.12 is just a poorly centred reflection of the opposite-signed partner of the burst being used to condition the evolution.

To verify this, we compute the distance between the centres of each detected sweep and its closest ejection, and vice versa, defined in the mean-convective frame of reference,

$$L_2 = [(\Delta t u_\tau)^2 + (\Delta z)^2 + (\Delta \xi)^2]^{1/2}. \quad (2.33)$$

[101] and [39] give similar statistics for unconditional Qs, although without the inclusion of the time difference. In their case, they require restrictions on the relative size of the eddies to avoid pairs with very different components, but this is made unnecessary in our case by only considering structures obtained with a particular filter. The definition of a pair is not straightforward,

and figure 2.14(a) shows a two-dimensional sketch of possible complications. Each sweep, such as s_1 , chooses as partner its closest ejection, such as e_1 , and may in turn be chosen as partner by other ejections. Partnership is not symmetric. A given sweep may be chosen as partner by several ejections, such as e_2 – e_4 in the case of s_3 , and each of those may or may not be the ejection chosen by the sweep as its closest partner. Couples which choose each other are defined as pairs, such as (e_1, s_1) and (s_3, e_4) in figure 2.14(a), and this relation is mutual and unique. Each structure can be part of a pair, or of none at all, but being paired does not preclude having been chosen as a partner by other structures. Examples of unpaired structures are e_2 , e_3 and s_2 , which have not been chosen by their preferred partner, and a paired structure with a large number of partners (or node degree, in graph notation) is s_3 . Approximately 60% of the sweeps are in pairs, and 14% of them are in exclusive pairs (see table 2.4).

There are no obvious statistical differences between the components of exclusive and non-exclusive pairs, but pairs are not random occurrences. The probability distribution of the degree of the sweeps and ejections is shown in figure 2.14(b), where it is compared with the expected density function for Poisson-distributed structures, represented by the solid line. It is clear that sweeps and ejections are clustered in the sense of having a higher number of partners of the opposite type than the Poisson prediction obtained by randomising the position of the structures. On the other hand, figure 2.14(c), shows the distribution of the degree of structures of the same type, which is considerable steeper than Poisson's. These two behaviours can be informally summarised as that bursts of opposite type 'attract' each other, while those of the same sign 'repel', and can probably be explained by the aforementioned effect of continuity, which requires that the net wall-normal mass flux over large areas should tend to vanish.

It is interesting that the probability of having $N > 3$ partners of the opposite type in figure 2.14(b) follows quite closely a geometric distribution $P(N) \sim q^N$ with $q = 0.4$, represented by the dashed line. This is intriguing, because it implies that the probability of being chosen as a partner is independent of how many partners a structure already has, probably implying that some structures have to choose a distant partner because they cannot find a closer one. This, in turn, suggests that bursts that are not part of a local pair are relatively isolated from prospective partners of the opposite sign. The same is not true for the same-sign partners in figure 2.14(c) which do not have an exponential tail or, if they do, it is one with a much smaller parameter than in the previous case ($q \approx 0.04$). The default model proposed by [101] from the analysis of the uv -Qs, is that Qs tend to cluster into streamwise trains of tight pairs. The previous discussion suggests that the organisation into trains is better developed than the one in pairs, or at least that many

FIGURE 2.15: (a–d) Probability density functions of the distance between components of sweep-ejection pairs. (e–h) Same for sweep-sweep pairs. (a,e) Pdf of the space-time distance L_2 , defined in (2.33). (b, f) Pdf of $|\Delta t u_\tau / y_c|$. (c, g) $|\Delta z / y_c|$. (d, h) $|\Delta \xi / y_c|$. In all cases, lines are: —, Λ_1 ; ---, Λ_2 ; - · -, Λ_3 ; ·····, Λ_4 ; —●—, pdf of the intra-pair distance for an equivalent Poisson process.

pairs are asymmetric, with one of the partners too weak to be identified as such by the analysis.

Figure 2.15 shows the probability distribution of the distance between the components of established pairs, compared with its Poisson counterpart. Figure 2.15(a–d) refers to sweep-ejection pairs, and figure 2.15(e–h) to pairs of the same type. More than 90% of the normalized distance between sweeps and ejections fall within $L_2 / y_c \in (0.5, 2)$, with its mode at $L_2 / y_c \approx 1.1$. Its pdf is significantly different from the Poisson distribution, shown in the figure as a line with solid circles. Figure 2.15(b–d) shows the contributions to L_2 of its different components. The main one is from the spanwise distance, which peaks at approximately $\Delta z / y_c \approx 1$, similar to the spanwise offset between the main and secondary eddies in figures 2.10 and 2.12. The streamwise component also plateaus for $\Delta \xi / y_c \lesssim 1$, but the temporal offset peaks for contemporaneous events ($\Delta t = 0$) and decays quickly. As with the total L_2 distance, the Poisson model fails for all components. Events tend to be spatially farther from each other than the Poisson model, but closer in the time of their peak development. The former is most probably due to the finite size of the objects being considered. Figure 2.10 shows that the (x, z) diameter of the sweeps and ejections is approximately y_c , and it is not surprising that there are few pairs in figure 2.15 spatially closer than that distance. On the

other hand, this restriction does not apply to time, and figure 2.15(b) shows that pairs preferentially evolve in synchrony, as in figure 2.12.

Figure 2.15(e–h) tells a slightly different story about homogeneous pairs. In the first place, although the average distances are approximately the same, the homogeneous pairs do not show the same probability deficit at short distances as in the heterogeneous case. Sweeps are closer to each other than their size would appear to allow. The reason can be understood from the individual pdfs in figure 2.15(f–h). While the exclusion in dissimilar pairs takes place in space, similar pairs are mainly spaced in time. The pdf for Δz peaks at the origin, in agreement with the model of a train of aligned pairs, and $\Delta \xi$ is also only weakly constrained at short distances, but these events do not happen simultaneously, and figure 2.15(f) reveals a clear deficit for $\Delta t u_\tau / y_c \lesssim 0.5$, or $S \Delta t \lesssim 2$. This is of the order of the lifetime of the burst, suggesting that, although consecutive pairs may align along streaks, they are created asynchronously, so that neighbouring sweeps are typically at different stages of their life. Note that this would be consistent with the observation in figure 2.13 that sweeps tend to form at the ‘active’ front of the streak, and suggests a model in which the trains of bursts propagate incrementally while they extend the streak length. Although it has to be emphasised that, as we have mentioned several times in this discussion, these bursts are only very seldom symmetrical hairpins, one cannot avoid noting the similarities between this propagation model and the incremental formation of hairpin packets analysed by [186] in the absence of turbulent background activity.

2.5 CONCLUSIONS

Using wavelets as band-pass filters, we have extended the identification in Jiménez [68] of Orr-like bursts in minimal channel flows, to computational boxes of arbitrary size. Wavelets allow us to reduce the flow to a field of local single-scale wave trains whose inclination and intensity can be computed, but lack the connection with the surrounding flow and the multi-scale character of turbulence. We have complemented this analysis with low-pass filters that retain the very large motions of the streamwise velocity. The evolution of both filtered velocities is tracked in time. Regardless of the filter used we find significant statistical evidence of Orr-like evolution of localised structures with strong wall-normal velocity.

The intense wall-normal velocity wavefronts are always vertical, and are consistently preceded by backwards-leaning perturbations. The length and time scales of these wave trains are self-similar with respect to the filter width, especially in the logarithmic layer. Increasing the Reynolds number in the range $Re_\tau = 950\text{--}4000$ simply increases the range of scales in the

logarithmic region, resulting in a wider hierarchy of self-similar eddies. When normalised with their filter-width, the events have roughly constant duration-area, approximately filling 10% of the time-area at a given height.

We also study the effect of the box-size, using short boxes of dimensions $2\pi h \times 2h \times \pi h$, long ones of $8\pi h \times 2h \times 3\pi h$, and comparisons with the single-mode analysis in minimal boxes by Jiménez [68]. We find good agreement between the three box sizes, suggesting that the dynamics of the Orr-like bursts are well captured in minimal domains. The evolution of the inclination angle and amplitude of the wavefronts of the three velocity components is reasonably well predicted by linearised optimal transient growth, corroborating that the evolution of these events is dominated by linear dynamics.

Because the statistical analysis of the wave trains shows that the most intense events of the wall-normal velocity are always vertical, Orr events can be identified solely from the amplitude, allowing low-pass filters to identify Orr-events while retaining the large-scale motions and the sign of the velocity perturbations. The result is that sweeps and ejections tend to form contemporaneous ‘pairs’, i.e. the peaking of the sweep and of the ejection happens at the same time. The components of the pairs are roughly aligned in the spanwise direction, forming streamwise-oriented tilted rollers, and the rollers tend to be aligned in the streamwise direction with respect to one another. However, rollers tend not to be contemporaneous, and the temporal offset among contiguous rollers is comparable to their lifespan. Sweeps and ejections are embedded in streamwise-velocity streaks of the opposite velocity sign, thus generating negative tangential Reynolds stresses that reinforce the streaks in which they reside. In fact, the evidence from the conditional flow histories in the space-time neighbourhood of strong bursts is that bursts tend to form along pre-existing streaks, especially near streamwise streak inhomogeneities.

The characteristics of the Orr events, such as their size, time scale, shape and spatio-temporal organisation, are consistent with the attached three-dimensional regions of intense tangential Reynolds stress studied in [39, 101, 103]. These works provide strong indications that the mean shear controls the dynamics of the attached momentum-carrying structures. The new evidence corroborates these findings, suggesting that the ‘tall’-attached momentum-transfer structures are Orr bursts, and can thus be reasonably well described by linear dynamics.

ACKNOWLEDGEMENTS

We are grateful to M. Lee for providing the data used in figure 2.4.

ESTIMATING LOGARITHMIC LAYER TURBULENCE FROM WALL OBSERVATIONS

The contents of this chapter are published in Physical Review Fluids (4, 114603), under the title “Logarithmic-layer Turbulence: a view from the wall”, with J. Jiménez as a coauthor; and in the Proceedings of the CTR Summer Program (2018, pp. 217–226), with the title “Reconstructing channel turbulence from wall observations”, and A. Lozano-Durán and J. Jiménez as coauthors.

This chapter deals with the possibility of observing the flow far from the wall in a turbulent channel, using only measurements at the wall. There are at least two reasons why this may be desirable. The first one is purely scientific, because the relation between the near-wall and outer flows has been one of the most controversial questions in wall-bounded turbulence, and it stands to reason that influence and observability are related. If A does not influence B, it is unlikely that the former can be reconstructed from observations of the latter, and, conversely, if A can be reconstructed from B, it is unlikely that the two are completely unrelated. Note that the inverse is not true, especially if, as here, we restrict ourselves to linear reconstructions. If A cannot be reconstructed from B, it cannot be concluded that the two are independent. In addition, reconstruction by itself does not give information on the direction of causality, if any. That B reconstructs A does not imply that B causes A, or the other way around. We focus on variables that can be reconstructed from wall measurements, and, in that regard, we will be able to propose criteria for which variables are ‘attached’ at what scales, in the sense that their influence extends to the wall, although these criteria might not coincide with Townsend’s dynamical definition of attached structures [170]. On the other hand, although it is probable that variables that cannot be reconstructed from the wall are ‘detached’, in the sense of having little or no influence over it, it follows from the previous discussion that this cannot be guaranteed.

A second reason for our interest in reconstruction is the possibility of active control of wall turbulence. This is a classic goal of turbulence research, whether to decrease or increase wall friction, reduce noise or other applications. Transportation of fluids through pipes, such as oil, gas or water, would benefit from drag reduction, and so would vehicles moving in fluids. In other engineering applications, such as combustion or heat transfer, or even aerodynamics, it may be desirable to increase the turbulent intensities close to

walls to promote mixing or to delay separation. There is a well-developed theory for the optimal control of linear systems, and, although turbulence is nonlinear, there is extensive evidence that a substantial fraction of the dynamics of shear-driven flows can be linearised [66]. This has led to heuristic and theoretical active-control schemes that manipulate friction in turbulent channels at moderate Reynolds numbers, at least in simulations [28, 45], but they are not free from problems and ambiguities.

The first problem has to do with the scales involved. Most of the dissipation in wall-bounded flows is contained within or below the logarithmic layer (see figure 3.1a), suggesting that effective control should act near the wall. As a consequence, most schemes have targeted the buffer-layer, which contains the strongest turbulent fluctuations, but technological considerations suggest that practical applications should rather centre on the logarithmic-layer eddies. Thus, for a flow of thickness h and Reynolds number $Re_\tau = u_\tau h/\nu$, where the ‘+’ superscript denotes wall units based on the friction velocity u_τ and on the kinematic viscosity ν , the sizes and passing times of the structures in the buffer layer decrease proportionally to Re_τ when expressed in outer units. On the contrary, when the distance, y , from the wall is a fixed fraction of h , sizes and times scale in outer units, or at most depend logarithmically on the Reynolds number. For example, in a water pipe with $h \approx 1$ m and bulk velocity $U \approx 1$ m/s ($Re_\tau \approx 10^5$), the wall-normal velocity eddies in the buffer-layer have length and time scales of $\Delta x \approx 1$ mm ($\Delta x^+ = 100$) and $\Delta t \approx 0.1$ ms, but those at a distance from the wall $y/h \approx 0.1$ are $\Delta x \approx 20$ cm and $\Delta t \approx 0.2$ s, where Δx is the streamwise length. Similar differences apply to the flow over aeroplane wings.

However, moving away from the wall is not without cost. Figure 3.1(a) presents the fraction of the total dissipation below a given distance, y , from the wall as a function of the Reynolds number [69]. An ‘ideal’ control strategy would probably result on the complete elimination of the turbulent energy dissipation over some part of the channel. Below $y^+ = 20$, most of the dissipation is due to the mean shear, and is independent of fluctuations. Centring on the remaining dissipation, in the example above, the velocity fluctuations in the logarithmic region, defined as $80\nu/u_\tau \leq y \leq 0.2h$, are responsible for about 40% of the energy loss, but the scales in its lower limit are those discussed above as being too small. If we want to restrict ourselves to fixed fractions of the flow thickness, to avoid wall scaling, the maximum fraction of the drag that could potentially be saved by completely removing all the dissipation between $y/h = 0.01$ and 0.2 would be approximately 25% (at $Re_\tau \approx 10^5$). Decreasing the lower limit of wall distances increases the potential gain, but at the cost of having to deal with smaller length and time scales. Raising the upper limit is less effective. These estimations are only upper bounds, because turbulent dissipation can probably not be completely

damped, but they show that the logarithmic layer is a ‘preferred region’ in which some control authority still remains, while the size of the eddies stays bounded from below as the Reynolds number increases.

A second problem is that most active-control schemes assume knowledge of the flow at wall distances of the order of the size of the structures to be controlled. Typically, the assumption is that the flow is fully known, as in Direct Numerical Simulation (DNS), or at least that a two-dimensional section of some variables is available, as in particle-image velocimetry. Unfortunately, most practical observations are limited to the wall, where the only accessible variables are the pressure and the two shears. This problem is especially acute when we are interested in the logarithmic layer, which is separated from the wall by the very active buffer layer. Figure 3.1(b) shows spectra of the three observable variables at the wall. All of them are dominated by a core whose position scales in wall units, but they also include a larger-scale component, which is long for the two shears and wide for the pressure, with wavelengths that scale proportionally to h , and which are associated with eddies farther from the wall [57].

A related question is which part of the outer flow is effective in generating pressure fluctuations at the wall, which are of interest in acoustics and structural loading. Wall pressure can be computed exactly if the flow is fully known, as in DNS. But lower-fidelity Large Eddy Simulations (LESs) computations only determine the outer flow approximately, and the question is which level of detail is required to predict wall pressure to a given accuracy.

This report addresses the second question first, and only then addresses the problem of estimating the state of the flow at interior points from measurements at the wall. In both cases, we try to use the fact that pressure satisfies a Poisson equation whose right-hand side is the object to be estimated. We will see that we eventually have to use statistical information to make our results physically plausible, but the use of two quantities linked by an equation at least guarantees that the observations are fully explained by the reconstruction. In a sense, equation-based give statistical ones a physical interpretation. The problem of determining the effectiveness of the right-hand side will turn out to reduce to a projection on an effective subspace, while that of optimal estimation can be understood as an extreme case of data assimilation [133].

The remainder of this chapter focuses on the properties of the estimations of logarithmic-layer turbulence using noiseless, although potentially incomplete, wall measurements from DNS. The widespread availability of data in the last years has allowed various research groups to test the performance of similar methods with considerable success [13, 14, 17, 60, 145]. These works differ from ours in that most of them focus on the use of physically motivated linear models [17, 60] to relate velocity measurements at one y to predict

FIGURE 3.1: (a) Fraction of the total energy dissipation below a given distance from the wall, as a function of the Reynolds number. The dashed lines are scaled in wall units, and the labels refer to the predominant contributor to the dissipation in each region. The solid lines are three distances from the wall, scaled in outer units [adapted from 69]. (b) Premultiplied spectra of [57] at $y^+ = 5$; ---, streamwise stress; ---, spanwise stress; —, pressure. The contours of each spectrum contain 50%, 70% and 90% of its spectral mass. λ_x and λ_z are respectively the streamwise and spanwise wavelengths.

the same velocity component at another wall-normal distance. Others [114, 145], also explore the performance of nonlinear models, with modest gains over their proposed linear method. In contrast, our methodology is linear regression, which has been successfully used for conditional flow reconstruction away from walls [4], or to synthesise buffer-layer velocities from experimental measurements farther away [13]. This Linear Stochastic Estimation (LSE), although apparently naive, is statistically optimal for linear systems, and we will show that it is adequate to reconstruct the Reynolds-stress structures in the logarithmic layer, using only variables that can be measured at the wall. A closely related investigation is [175], which analyses the correlation between wall pressure at the wall-normal velocity farther into the flow, but, to our knowledge the present work is the first to one to use the full set of wall observables to reconstruct the full internal flow field, including the relative importance of the different observables on the three velocity components, and the identification of the accessible scales. Finally, it should be noted that, although not specifically addressed in the present work, similar methods should apply to transitional flows, where scales are typically larger, the role of linearity is more pronounced, and are thus more amenable to active control.

The remainder of this chapter is organised as follows. The numerical datasets are presented in §3.1. Sections §3.2 and §3.3 respectively analyse the relation of the pressure with the second invariant, and the details of the re-

Name	Re_τ	$(L_x, L_z)/h$	$N_x \times N_y \times N_z$	$\Delta x^+, \Delta z^+, \Delta y_{\max}^+$	$T^{(h)}$	D_y	N_f	Ref.
S1000	932	$2\pi, \pi$	$512 \times 385 \times 512$	12, 6, 7.7	18	CH	1800	[103]
S2000	2003	$2\pi, \pi$	$1024 \times 633 \times 1024$	12, 6, 8.9	11	FD	800	[103]
F2000	2003	$8\pi, 3\pi$	$512 \times 512 \times 512$	24, 9, 8.4	14	FD	1200	[A, 174]
F5300	5300	$8\pi, 3\pi$	$1024 \times 1024 \times 768$	22, 12, 11.2	30	FD	3000	[A, 174]

TABLE 3.1: Parameters of the simulations. $L_{x,z}$ are the streamwise and spanwise period of the numerical box, and $N_{x,y,z}$ are the number of collocation points of the stored grid in each direction. Δ^+ are the grid resolution of the DNS at the channel centreline. $T^{(h)}$ is the time spanned by the approximately equispaced snapshots in eddy turnovers. D_y represents the discretisation used for the wall-normal direction: ‘CH’ refers to Chebyshev polynomials and ‘FD’ to compact finite differences. N_f is the number of fields used to compute statistics.

construction algorithm. Sections §3.4 and §3.5 discuss the properties of the reconstruction of Q and the velocity fields; and §3.6 concludes.

3.1 NUMERICAL EXPERIMENTS

We consider incompressible flow in channels whose half-height is h . The streamwise, wall-normal and spanwise directions are x, y, z , respectively, and u, v, w, p are the fluctuations of the corresponding velocity components and of the kinematic pressure. Capitals are used for averaged quantities, and primes for standard deviations. The computational domain is periodic in both wall-parallel directions, with periods L_x and L_z . The spatial discretisation is Fourier spectral in those directions, dealiased using the 2/3 rule, and compact finite differences or Chebyshev expansions in the wall-normal direction, depending on the Reynolds number. Table 3.1 collects basic information on the simulations, and refers to the original publications for further details.

Since our focus is on the logarithmic layer, and on scales that are large compared to the viscous length, the spatial resolution of the database is not critical except very near the wall, but the longest and widest scales, particularly those of the streamwise velocity, require large boxes and long integration times to be accurately represented. To complement the high-resolution, but relatively small, boxes of S1000 and S2000 while keeping the computations economical in terms of storage, we use the large-box simulations F2000 and F5300. They are computed as DNSs but stored at the resolution of large-eddy simulations, $\delta x^+ \approx 120$ and $\delta z^+ \approx 90$, which is still sufficient to study the energy- and stress-carrying eddies in the logarithmic layer. The pressure at the wall in these data sets is computed retaining the effect of the discarded scales [174].

3.2 RELATION BETWEEN PRESSURE AND THE SECOND INVARIANT OF THE VELOCITY GRADIENT TENSOR

The pressure satisfies the Poisson equation

$$\nabla^2 p = 2Q, \quad (3.1)$$

where $Q = (\omega_i \omega_i - 2S_{ij}S_{ij})/4 = \partial_i u_j \partial_j u_i / 2$ is the second invariant of the velocity gradient, and $S_{ij} = (\partial_i u_j + \partial_j u_i)/2$ is the rate-of-strain tensor. The pressure satisfies boundary conditions at both walls

$$\partial_y p = \nu \partial_{yy} v = \nu (\partial_x \omega_z - \partial_z \omega_x), \quad (3.2)$$

where the last equality follows from continuity and impenetrability.

Expanding in Fourier series along (x, z) , and denoting Fourier coefficients by carats, (3.1) becomes

$$(\partial_{yy} - k^2) \hat{p} = 2\hat{Q}, \quad (3.3)$$

where $k^2 = k_x^2 + k_z^2$ is the wavenumber magnitude. Following [84], $\hat{p}$ can be expressed as the sum of an “inertial” component, $\hat{p}_Q$, satisfying (3.3) with homogeneous Neumann boundary conditions, and a “Stokes” contribution,

$$\hat{p}_S = [-\partial_y \hat{p}(0) \cosh k(2h - y) + \partial_y \hat{p}(2h) \cosh ky] / k \sinh(2kh), \quad (3.4)$$

which satisfies Laplace’s equation with boundary conditions, (3.2). The inertial pressure can be obtained from $\hat{Q}$ by inverting (3.3) as

$$\hat{p}_Q(y) = \int_0^{2h} \hat{Q}(\zeta) G(y, \zeta) d\zeta, \quad (3.5)$$

where the Green’s function is

$$G(y, \zeta) = \begin{cases} -2 \cosh k(2h - \zeta) \cosh ky / k \sinh 2kh, & \text{if } y \leq \zeta, \\ -2 \cosh k\zeta \cosh k(2h - y) / k \sinh 2kh, & \text{otherwise.} \end{cases} \quad (3.6)$$

Note that, if the pressure and the two shears, $\partial_y u$ and $\partial_y w$, can be measured at both walls, the Stokes pressure follows from (3.2) and (3.4), and the inertial wall pressure, $\hat{p}_{Qw} = \hat{p}_w - \hat{p}_{Sw}$, is observable. Note also that, except for very large structures in which $kh \ll 1$, the Stokes pressure is essentially determined locally by the boundary condition at each wall, and is small everywhere. In what follows, we suppress the subindex ‘Q’ when referring to the inertial pressure, but the estimations of Q in Section 3.5 are corrected for the Stokes component. A preliminary analysis of these equations is presented by [71]. Defining wavelengths as $\lambda_i = 2\pi/k_i$, the location of the core of the pressure spectrum away from the wall is approximately equilateral, $\lambda_x \approx \lambda_z \approx 3y$, and the intensity scales in wall units, $p(y/h) \sim u_\tau^2$.

FIGURE 3.2: (a) Sketch of the pressure-generating subspace, with the orthogonal projection of Q on the Green's function. (b) Optimal reconstruction of Q based on the wall pressure. Solid line, Full $\hat{Q}$; dashed line, minimum-norm orthogonal reconstruction, $\hat{Q}_{\parallel}$; M950, $\lambda_x/h = 0.7$, $\lambda_z/h = 0.78$. (c) As in (a), emphasising the known probability distribution of Q . (d) As in (b), for the minimum-error reconstruction. Solid line, $\hat{Q}$; dash-dotted line, $\hat{Q}^{\dagger}$.

Particularising (3.5)–(3.6) at the lower wall,

$$\hat{p}_w = \int_0^{2h} \hat{Q}(\zeta) G(0, \zeta) d\zeta = \frac{-2}{k \sinh 2kh} \int_0^{2h} \hat{Q}(\zeta) \cosh k(2h - \zeta) d\zeta. \quad (3.7)$$

Near the wall, Figure 3.3(a) shows that the wavelengths of the spectrum scale well in wall units, suggesting that most of the pressure at the wall is due to flow structures in the buffer layer. But there is a substantial large-scale tail that lengthens with increasing Reynolds number, and therefore presumably originates farther into the flow. This is confirmed by Figure 3.3(b), which displays the contribution to the spectrum from different height intervals in the integral (3.7).

From the point of view of flow reconstruction, it is clear that the function $\hat{Q}(y)$ cannot be fully characterised by the single number $\hat{p}_w$, and that some kind of projection is required. There is a relatively well developed theory for linear optimal “state observers” in control theory [155] and in dynamical systems theory [37], whose key result is that an observer can only be built for observable quantities. In the present context, this means that $\hat{Q}$ should have a measurable effect at the wall to be reconstructible, but the exponential decay of the kernel in (3.7) suggests that this is not the case for many Q structures. Small eddies with $kh \gg 1$, whose Green's function decays fast with ky , are not observable and cannot be reconstructed far from the wall. We will see below that larger structures can.

If we represent discretised functions of y as vectors, such as $\mathbf{g} \equiv G(0, y)$ and $\mathbf{Q} \equiv \hat{Q}(y)$, (3.7) can be written as a matrix product,

$$\hat{p}_w = \mathbf{g}^* \mathbf{Q}, \quad (3.8)$$

where the asterisk denotes Hermitian transpose. In practice, (3.8) has to incorporate the integration weights, as explained by [117]. The situation is sketched in Figure 3.2(a). For a given $\widehat{p}_w$, (3.8) implies that $\mathbf{Q}$ lies in a hyperplane normal to $\mathbf{g}$, which is parallel to the null-space of the inner product. Any component in this subspace $\mathbf{Q}_\perp$, such that $\mathbf{g}^* \mathbf{Q}_\perp = 0$, does not contribute to the wall pressure. There are two optimisation problems to be considered. The first one concerns which is the smallest $\mathbf{Q} = \mathbf{q}$ that generates a given $\widehat{p}_w$, and is equivalent to minimising

$$\mathbf{q}^* \mathbf{q}, \quad \text{given} \quad \mathbf{g}^* \mathbf{q} = \mathbf{g}^* \mathbf{Q} = \widehat{p}_w. \quad (3.9)$$

The second one is to minimise the error between the predicted and observed $\mathbf{Q}$,

$$(\mathbf{Q}^* - \mathbf{q}^*)(\mathbf{Q} - \mathbf{q}), \quad \text{given} \quad \mathbf{g}^* \mathbf{q} = \mathbf{g}^* \mathbf{Q} = \widehat{p}_w. \quad (3.10)$$

For any given flow field, the answer to the second problem is trivially $\mathbf{q} = \mathbf{Q}$, with zero error, while the answer to the first one is the orthogonal projection of $\mathbf{Q}$ over $\mathbf{g}$,

$$\mathbf{Q}_\parallel = (\mathbf{g}^* \mathbf{Q}) \mathbf{g} / |\mathbf{g}|^2 = \widehat{p}_w \mathbf{g} / |\mathbf{g}|^2, \quad (3.11)$$

where $|\mathbf{g}|^2 = \mathbf{g}^* \mathbf{g}$. An example of $|\mathbf{Q}|$ and $|\mathbf{Q}_\parallel|$ for a particular Fourier mode and snapshot of channel M950 is given by the solid and dashed lines in Figure 3.2(b), illustrating that the optimum projection has the expected exponential form of the Green's function, but that most of $\mathbf{Q}$ is not contained in the projection.

Unfortunately, (3.11) is physically unrealistic because $\widehat{\mathbf{Q}}$ has to satisfy constraints that are not taken into account by the above procedure. For example, it follows from incompressibility and from the no-slip velocity condition that $\widehat{\mathbf{Q}} = 0$ at both walls, which is not true for the Green's function. In essence, (3.11) is a property of Poisson's equation, otherwise unrelated to the flow in a channel. On the other hand, it is the optimal solution of (3.8), and any other $\mathbf{Q}$, even if motivated by physical considerations, contains essentially arbitrary components which are unrelated to the pressure at the wall.

The simplest way to incorporate the properties of $\widehat{\mathbf{Q}}$ into (3.9)–(3.10) is to apply the optimisation to an ensemble average of flows, rather than to individual snapshots. Consider the reconstruction $\mathbf{q} \sim \boldsymbol{\phi}$, where $\boldsymbol{\phi}$ is a vector normalised as

$$\mathbf{g}^* \boldsymbol{\phi} = 1, \quad (3.12)$$

constant across snapshots. The condition that $\mathbf{q}$ reconstructs the correct pressure is

$$\mathbf{q} = \widehat{p}_w \boldsymbol{\phi}. \quad (3.13)$$

It is easy to check that this statistical modification has no effect on (3.9), whose solution is still the orthogonal projection, (3.11), but (3.10) contains Q explicitly, and the effect of incorporating its statistics is not trivial. The quantity to be minimised is

$$\langle (Q - \hat{p}_w \phi)^* (Q - \hat{p}_w \phi) \rangle + \mu(1 - \phi^* g) + \dots, \quad (3.14)$$

where $\langle \cdot \rangle$ is the ensemble average, the trailing dots represent terms that do not involve ϕ^* , and μ is the Lagrange multiplier that enforces (3.12). Differentiation with respect to ϕ^* leads, after some algebra, to $\mu = 0$ and to the reconstruction

$$Q_* = \hat{p}_w \langle \hat{p}_w^* Q \rangle / \langle |\hat{p}_w|^2 \rangle, \quad (3.15)$$

where the projector $\phi = \langle \hat{p}_w^* Q \rangle / \langle |\hat{p}_w|^2 \rangle$ plays the role of $g / |g|^2$ in (3.11). Some reflection reveals that ϕ is an approximation to the direction of the mean value of the distribution of Q in function space. The only statistical object that needs to be compiled in (3.15) is the two-point correlation matrix $\mathbf{C} = \langle Q Q^* \rangle \equiv \langle \hat{Q}(\zeta) \hat{Q}^*(\zeta') \rangle$, because (3.8) implies that all the second-order statistics of $\hat{p}_w$ can be derived by contracting g with $\mathbf{C}$. For example, $\langle \hat{p}_w Q^* \rangle = g^* \mathbf{C}$, and $\langle |\hat{p}_w|^2 \rangle = g^* \mathbf{C} g$.

Equation (3.15) is a variant of Linear Stochastic Estimation (LSE) [1], which uses the joint statistics of a series of observed events, $\hat{p}_w$, and unknowns $\hat{Q}(y)$ to be estimated, to compute the linear operator minimising the average error. In contraposition to (3.11), only statistical information is included in (3.15), which does not involve the Green's function g , and is essentially independent of the physics in (3.1).

LSE has often been used on turbulent flow, but usually to derive mean flow behaviour conditioned to particular events, or a local Taylor expansion for the velocities [1]. An early attempt to estimate the instantaneous buffer-layer vortices from wall shear was made by [91], but we know of few other cases in which it has been used as a real-time flow estimator. As it is used here, there is no implication that $\hat{Q}(y)$ can be locally expanded as a Taylor series of position. The assumption is that the perturbations of $\hat{Q}$ are linearly related to the perturbations of $\hat{p}_w$.

An example of this reconstruction for the case in Figure 3.2(b) is shown in Figure 3.2(d). It agrees better than $Q_{\parallel}$ with the known statistical behaviour of $\hat{Q}$, including vanishing at the wall, but it is still a small fraction of the original $\hat{Q}$. Most of the right-hand side of (3.3) is not observable from the wall.

FIGURE 3.3: (a) Premultiplied pressure spectra at $y^+ = 5$. Contours are $k_x k_z E_{pp}^+ = (0.25, 0.50, 0.75)$, and the dashed diagonal is $\lambda_x = \lambda_z$. Large-box channels at: solid, $h^+ = 550$; dashed, 934; dash-dotted, 2000. Various sources. (b) Shaded contours are the pressure spectrum at the wall. Contours contain 10%, 50%, and 90% of the spectral mass, M950. solid, contribution from (3.7) below $y^+ = 60$; dashed, $60 < y^+ < 200$; dotted, $y^+ > 200$. Contours contain 63%, 18%, and 4% of the total pressure, and 80%, 38%, and 31% of the pressure in their spectral region.

3.3 LINEAR STOCHASTIC ESTIMATION

Stochastic estimation is the approximation of an unknown, $\mathbf{u}$, using the statistical information from an observable, $\mathbf{E}$. To fix ideas, we consider $\mathbf{u}$ to be a vector with the streamwise velocity component at every grid point,

$$\mathbf{u}(y) = [u_i(y)], \quad (3.16)$$

where the i subindex represents the two grid indices $i_x = 1 \dots N_x$ and $i_z = 1 \dots N_z$. The wall-normal ($\mathbf{v}$) and spanwise ($\mathbf{w}$) velocity components are analogously expressed, and any procedure described for $\mathbf{u}$ is equivalently applied to them. Owing to the linearity of the method, optimising the individual velocity components is equivalent to optimising the velocity vector as a whole. Our observables are,

$$\mathbf{E} = [p_i(0), \partial_y u_i(0), \partial_y w_i(0)], \quad (3.17)$$

where indices are defined as above, but $\mathbf{E} = [E_{i(s)}]$ includes a second index, $s = 1 \dots 3$, representing the three observables. The best statistical estimation, $\mathbf{u}^\dagger(y)$, of $\mathbf{u}(y)$ under the L_2 -norm, is the conditional average of $\mathbf{u}(y)$ given $\mathbf{E}$, i.e. $\langle \mathbf{u}(y) | \mathbf{E} \rangle$, where $\langle \cdot \rangle$ stands for ensemble averaging. When expanded as a Taylor series in $\mathbf{E}$, and truncated at the linear term, it is called a linear stochastic estimator (LSE), and is numerically equivalent to the linear mean-square approximation of $\mathbf{u}$ in terms of $\mathbf{E}$ [1]. The estimator $\tilde{L}_{ij(s)}(y)$ in

$$u_i^\dagger(y) = \tilde{L}_{ij(s)}(y) E_{j(s)}, \quad (3.18)$$

where repeated indices imply summation, is obtained by minimising

$$\left\langle \left(u_i(y) - u_i^\dagger(y) \right)^2 \right\rangle = \left\langle \left(u_i(y) - \tilde{L}_{ij(s)}(y) E_{j(s)} \right)^2 \right\rangle, \quad (3.19)$$

which reduces to solving the linear system,

$$\langle E_{j(s)} E_{m(r)} \rangle \tilde{L}_{ij(s)}(y) = \langle u_i(y) E_{m(r)} \rangle, \quad (3.20)$$

where $\langle E_{j(s)} E_{m(r)} \rangle$ is the autocorrelation tensor of the observables, and $\langle u_i(y) E_{m(r)} \rangle$ is the cross-correlation between the unknowns and the observations. The presence of the latter in the right-hand side of (3.20) codifies the intuition expressed in the introduction that two quantities are only mutually observable if they are correlated. It also implies that reconstruction is a largely symmetric property. If A allows us to estimate B, B should allow us to estimate A. Note that the only information required to compute $\tilde{L}_{ij(s)}$ is contained in the two correlation tensors, and that $\mathbf{u}$ is only assumed to be linear with respect to $\mathbf{E}$. There is no assumption about the linearity of the velocity field in terms of the geometric coordinates.

Equation (3.20) can be manipulated to increase or decrease the number of observables used to estimate the velocity field. A reduced estimator may be desirable to separate the contributions of the different observables to the full model, or for practical considerations. For example, if only the pressure ($s = 0$) is used to estimate u , (3.20) becomes,

$$\langle E_{j(0)} E_{m(0)} \rangle \tilde{L}_{ij,p}(y) = \langle u_i(y) E_{m(0)} \rangle. \quad (3.21)$$

In the reduced system, $\langle E_{j(0)} E_{m(0)} \rangle$ is the autocorrelation of the pressure at the wall, and $\langle u_i(y) E_{m(0)} \rangle$ the cross-correlation of p at the wall with u at every wall distance. Note that $\tilde{L}_{ij,p}(y) \neq \tilde{L}_{ij(0)}$, as the later contains information about the shears through the inverse of $\langle E_{j(s)} E_{m(r)} \rangle$.

In our case, $\mathbf{E}$ is a vector of length $3N_x N_z$, resulting in a large correlation matrix that has to be inverted in order to solve for $\tilde{L}_{ij}$ in (3.20). This can be alleviated by exploiting the periodicity of the domain and projecting on the Fourier basis. For every pair of Fourier modes, $\widehat{\mathbf{E}}(k_z, k_x)$ and $\widehat{\mathbf{u}}(k_z, y, k_x)$, equation (3.20) becomes,

$$\langle \widehat{E}_s \widehat{E}_r^* \rangle(k_z, k_x) \widehat{L}_s(k_z, y, k_x) = \langle \widehat{u}(k_x, y, k_z) \widehat{E}_r^*(k_x, k_z) \rangle, \quad (3.22)$$

where the carat denotes Fourier transformation, k_i are the wavenumbers, and the asterisk is complex conjugation. The now more feasible calculation consists on solving $N_z N_x$ linear problems of dimension three. This procedure is also known as spectral linear stochastic estimation [SLSE, 168]. It is formally equivalent to LSE, but it behaves better with incomplete or noisy data because it avoids the spurious correlations between the orthogonal Fourier basis functions. The inverse Fourier transform of $\widehat{L}_s$, as obtained

from (3.22), can be used to reconstruct the velocity fields in physical space, and (3.18) becomes a discrete convolution,

$$u^\dagger(i_x, y, i_z) = \sum_{\forall j_x, j_z} \check{L}_s(i_x - j_x, y, i_z - j_z) E_s(j_x, j_z), \quad (3.23)$$

where the full operator $\check{L}_{ij(s)}$ has been expressed as a block-Toeplitz matrix whose blocks are shifted versions of the inverse transform $L_s(i_x, i_z)$ of the individual Fourier estimators. In a mean-square sense, the operator $\check{L}_{ij(s)}(y)$ is by construction the best possible data-driven linear operator for the reconstruction of the velocity components at a given wall distance. This can be seen from (3.23), or from its continuous equivalent

$$u^\dagger(x, y, z) = \iint \check{L}_s(x - x', y, z - z') E_s(x', z') dx' dz', \quad (3.24)$$

because it uses all the contemporaneous measurements at the wall to reconstruct the velocity at each point.

The accuracy of the reconstruction can be improved by enriching E with information from previous time steps. However, some experimentation with physically motivated time sampling resulted in a best-case reduction of 14% in the reconstruction error for S1000, compared with only using instantaneous data, at a cost between 8 and 16 times higher. As a consequence, these experiments were not pursued, and the results in this thesis only use contemporaneous observables.

3.4 ESTIMATIONS OF Q USING LSE

Wall-parallel two-dimensional spectra of Q are shown as line contours in Figure 3.4. They are approximately equilateral, except very near the wall, peaking at some multiple of the local Kolmogorov scale, $\lambda_i \approx 25\eta$. The same is true for the vorticities and for most of the velocity gradients [69], and it is known that Q , enstrophy, and the discriminant of the velocity gradient are approximately equivalent representations of vortices [27]. The shaded contours in the figure are the fraction of the spectrum of Q reconstructed by (3.15). Figure 3.4(a) shows that the reconstruction is essentially exact within the viscous layer ($y^+ = 9$), at least with regard to the magnitude of Q , while Figures 3.4(b,c) show that the reconstructible wavelengths are restricted to larger scales ($\lambda = 2\pi/k \approx 4y$), as y moves into the outer buffer and logarithmic layers. Modes where $|\widehat{Q}|^2 < 10^{-3}|\widehat{Q}|_{max}^2$, such as the largest scales, have been masked for clarity.

The evolution of the reconstruction with y is further analysed in Figure 3.5(a,b), which shows the spectra of the observed and reconstructed Q as a function of the wavelengths and of y . It is interesting that, whereas the spectrum of Q tracks the Kolmogorov scales, as it does for the vorticity, that of

FIGURE 3.4: Line contours are the premultiplied spectrum of $\widehat{Q}$ for M950, with contours containing 90%, 50%, and 10% of the spectral mass at each height. Shaded contours are the fraction of $|\widehat{Q}|^2$ contained in the reconstruction in (3.15), with the colour bar at right. Only wavenumbers for which $|\widehat{Q}|^2$ is above 10^{-3} of its maximum are included. (a) $y^+ = 9$. (b) $y^+ = 60$. The dashed line is $2\pi/kh = 0.25$. (c) $y^+ = 111$, and $2\pi/kh = 0.5$.

$Q^\dagger$ peaks at $\lambda \approx 4y$, which is characteristic of the energy-containing eddies. Álamo and Jiménez [11] studied wall-attached clusters of vortices in channels. They showed that they are distinct from individual vortices, that their size grows linearly with y , and that clusters are associated with strong wall-normal velocities. [101] and [38] later showed that vortex clusters are part of the sweep-ejection pairs responsible for a large fraction of the Reynolds stresses. (3.7) shows that the associated wavelengths are also responsible for the pressure fluctuations at the wall. All these observations suggest that the eddies of the reconstructed Q are identical to the vortex clusters, as further shown in Section 3.4.1.

Figure 3.5(c) displays profiles of the r.m.s. intensities of Q and $Q^\dagger$, and includes the two Reynolds numbers S1000 and S2000. Dimensional considerations imply that $Q' \sim u_\tau^3/y$ in the scale-free logarithmic layer. The figure shows that this is true: Q' scales in wall units in the buffer layer, and the collapse extends into the logarithmic layer. Interestingly, the same applies to Q'_* . The two quantities approximately agree below $y^+ \approx 10$, beyond which Q'_* decays, recovering the $1/y$ trend above $y^+ \approx 100$. Comparison with the spectra in Figure 3.4 shows that the near-wall peak corresponds to eddies at buffer-layer scales, while the outer logarithmic trend corresponds to larger structures. Both [11] and [101] found that vortex clusters are predominantly present near the wall, but that there is a population of self-similar Q structures that extends into the logarithmic layer. They cover a fraction of the wall-parallel area approximately independent of y . Álamo and Jiménez [11] justified this on the grounds of the percolation analysis used to determine the threshold that defines the clusters, but the proportionality in Figure 3.5(c) between the intensities of the observed and reconstructed Q , which use

FIGURE 3.5: One-dimensional spectra of Q (lines) and its reconstruction (shaded), normalised to unit spectral mass at each height. In both cases, the contours are (0.032, 0.1, 0.32). ---, $\lambda_i = 4y$; - · - · -, $\lambda_i = 25\eta$, where $\eta = \nu^{3/4}\varepsilon^{-1/4}$ is the Kolmogorov viscous length computed from the kinetic-energy dissipation ε . (a) As a function of the streamwise wavelength. (b) Spanwise. (c) R.m.s. fluctuation intensities for the true Q (without symbols), and for the reconstructed $Q^\dagger$ (symbols). —, M950; ---, W1900. The two diagonals are $Q'^+ = 5/y^+$, and $0.1/y^+$.

no threshold, suggests that it is a self-similarity property of the structures themselves.

3.4.1 Three-dimensional reconstructions of Q

More information on the geometry of Q is provided by the investigation of the three-dimensional regions where Q or Q'_* are particularly intense. We focus on points satisfying

$$Q(x, y, z) > \alpha Q'(y), \quad (3.25)$$

$$Q_*(x, y, z) > \alpha_* Q'_*(y), \quad (3.26)$$

where $\alpha = 1.25$ and $\alpha_* = 1.75$ are thresholds selected by a percolation analysis as done by Moisy and Jiménez [118], and Q and Q_* structures are defined as connected point sets satisfying (3.25)–(3.26), respectively. Connectivity is defined in terms of the six orthogonal neighbours in the Cartesian DNS mesh. We have confirmed that the conclusions below remain valid for values of α and α_* that are half and double of those discussed below.

The size of individual structures is defined by the streamwise (l_x), wall-normal ($l_y = y_{max} - y_{min}$), and spanwise (l_z) lengths of their grid-aligned bounding box, where y_{min} and y_{max} are the minimum and maximum distance of each object to the closest wall, respectively. Following [11], Q - and $Q^\dagger$ -structures are classified as wall-attached when $y_{min}^+ < 20$ from the bottom wall, and as wall-detached otherwise. We do not consider structures attached to the top wall, because the reconstruction only uses information from $y = 0$.

FIGURE 3.6: (a) Joint p.d.f. of the maximum and minimum wall distances of the identified Q and $Q^\dagger$ structures, $p(y_{min}, y_{max})$. Contours contain 98% and 50% of the probability. (b) Joint p.d.f. of the logarithms of the streamwise and spanwise sizes of Q and $Q^\dagger$ structures, $p(l_x/h, l_z/h)$. (c) P.d.f. of the intersected fraction of box-volume of Q boxes with $Q^\dagger$ boxes. Lines with and without squares are for wall-attached and wall-detached structures, respectively (the p.d.f. for detached structures is divided by two). Solid lines are the p.d.f. Dashed lines are computed for Q boxes with $Q^\dagger$ boxes at uncorrelated times. Note that the solid and dashed lines are indistinguishable for detached structures.

Figure 3.6(a) contains the joint probability density function (p.d.f.) of the minimum and maximum wall distances of the structures. Both Q and Q_* structures are organised into wall-attached and wall-detached families. The former are contained in the narrow vertical band for $y_{min}^+ < 20$, and the latter in the wider strip above the diagonal. Interestingly, the band for detached objects is narrower for Q_* structures, which tend to be smaller in the wall-normal direction. The streamwise and spanwise sizes of individual structures are quantified in Figure 3.6(b). The populations of Q and Q_* structures exhibit similar sizes, although the distribution of the reconstructed structures is biased towards smaller sizes, consistent with the lack of large detached structures reported above.

A better insight into the representation of true structures by the reconstructed field follows from the correlation between Q boxes (the bounding box of a Q structure) and $Q^\dagger$ -boxes. For each Q box of volume V_b , we select the spatially intersecting $Q^\dagger$ -box with highest overlap V_i , and the correlation fraction between Q and $Q^\dagger$ boxes is defined as $f_c = V_i/V_b$. To quantify the amount of random overlap between structures, the calculation is repeated for Q and $Q^\dagger$ boxes from different uncorrelated times.

Figure 3.6(c) shows that there is a non-trivial overlap between wall-attached Q and $Q^\dagger$ boxes, while there is no significant correlation between the detached Q and $Q^\dagger$ families. The result reinforces the observation that detached Q structures are not correctly captured by the reconstructed field, probably because their limited wall-normal extent does not produce the necessary footprint at the wall to allow for their reconstruction.

FIGURE 3.7: (a,b) Snapshots of S1000 at $y^+ = 10$. (a) v^+ . (b) $(v^+)^+$. (c-h) F5300 at $y/h = 0.1$. (c) u^+ . (d) v^+ . (e) $(u^+)^+$. (f) $(v^+)^+$. (g) u_G^+ , filtered with (3.27) and $\Delta_x \times \Delta_z = (4 \times 2)y$. (h) v_G^+ with $(2 \times 2)y$.

3.5 ESTIMATIONS OF THE VELOCITY FIELDS USING LSE

We use the methodology described in §3.3 to generate flow reconstructions for all the cases in table 3.1. First, we tested if reconstructing fields in-sample made a difference over reconstructing fields out-of-sample. For this purpose, we computed the operators of S1000 and F2000 using only half of the available velocity fields. The correlations can be symmetrised, effectively multiplying the number of independent samples [117]. Thus, the linear operator converges reasonably fast and, no difference can be found between the in-sample and out-of-sample snapshots. Considering the outcome of this experiment, we use all the flow fields available to compute the correlations in (3.22).

Figure 3.7(a,b) shows snapshots of the true and reconstructed wall-normal velocity very close to the wall ($y^+ \approx 10$) in S1000. The reconstructed field

captures the original in almost full detail, and the same is true for the other two velocity components, and for the Reynolds product $u^\dagger v^\dagger$ (not shown).

A similar level of accuracy only holds for $y^+ \lesssim 20$. Figures 3.7(c-h) compare the true u and v with their reconstructions in the logarithmic layer ($y/h = 0.1$, $y^+ \approx 500$). At this distance from the wall, most eddies with sizes smaller than $O(y)$ are missing from the reconstructions in figure 3.7(e,f). This agrees with the idea that ‘detached’ structures do not leave a footprint at the wall, and cannot be reconstructed from wall information. However, it is clear from the comparison of figures 3.7(e,f) and 3.7(c,d) that some large-scale information survives the reconstruction. This is best seen by comparing the reconstructed fields with those in figure 3.7(g,h), where the true velocities have been filtered with a low-pass Gaussian kernel,

$$G(x, y, z) = \left(\frac{2\pi}{\Delta_x \Delta_z} \right)^{1/2} \exp \left[-(\pi x / \Delta_x)^2 - (\pi z / \Delta_z)^2 \right], \quad (3.27)$$

where the filter widths are empirically adjusted to retain those large scales for which the error in the reconstruction is less than 50%. Figure 3.7 shows that, even far from the wall, the wall-normal velocity is organised into streaks. The reconstructed velocity fields include streaks of u , which account for most of the kinetic energy, and weaker v structures of similar size. The energy of the latter is only about 10% of the total v^2 -energy at that height, but their correlation with u is strong enough for the reconstructed velocity field at the height of figure 3.7(c-h) to contain approximately 52% of the tangential Reynolds stress. The filtered fields hold a similar amount of Reynolds stress, 45%, reinforcing the visual confirmation that the filter widths used are reasonable. The fair amount of Reynolds stress captured is not surprising if we consider that the spectral structure parameter, $-\phi_{uv}/(\phi_{uu}\phi_{vv})^{1/2}$, is known to be very close to unity for those large modes [70], and the long v structures seen in the figure are those responsible for ‘lifting’ and maintaining the streaks [87].

The relation between the two velocity components is even clearer in figure 3.8, which shows an instantaneous three-dimensional representation of the reconstructed velocity field below $y/h = 0.25$. The solid grey isosurfaces of figure 3.8(a) represent the low-speed streaks of $u^\dagger$, while the red translucent ones are those of u . The original field has been filtered with (3.27) to retain only reconstructible scales. It is clear that, except for some discrepancies in the smaller structures far from the wall, all of the large features of u , including short or ‘broken’ streaks are well reproduced. Most interesting are the isosurfaces of $-u^\dagger v^\dagger$, represented in solid grey in figure 3.8(b). As mentioned above, the reconstruction of v is poor at this wall distance, but it is clear from the figure that the v -structures that can be reconstructed are those associated with intense Reynolds stress. The isosurfaces of in the figure are the ‘ejections’ and ‘sweeps’ [Q_{2-4} , 178, 180], which are known to mostly occur within the streamwise velocity streaks, and this arrangement is well captured

FIGURE 3.8: Snapshot of the F5300 reconstruction. (a) The solid grey isosurfaces are $u^\dagger/u^\dagger'(y) < -1.5$, coloured with the distance from the wall, and the translucent red ones are $u_G/u'_G(y) < -1.5$. (a) The solid grey isosurfaces are $u^\dagger v^\dagger/(u^\dagger v^\dagger)'(y) < -1.5$, coloured with the distance from the wall, and the translucent red ones are $u_G v_G/(u_G v_G)'(y) < -1.5$. The real velocities u and v (but not $u^\dagger$ and $v^\dagger$) are filtered using (3.27) with $\Delta_x \times \Delta_z = (4 \times 2)y$ and $\Delta_x \times \Delta_z = (2 \times 2)y$ respectively.

by the reconstruction. Except for several ‘detached’ structures from the wall, most of the filtered Reynolds stress Q_{2-4} structures that extend from the wall are well captured, and they are known to carry half of the total Reynolds stress [103].

To quantify the performance of the reconstruction as a function of eddy size, we define the fractional spectral error

$$R_{ab}(k_x, y, k_z) = \frac{\text{Re}\langle (a - a^\dagger)(b - b^\dagger)^* \rangle(k_x, y, k_z)}{\text{Re}\langle ab^* \rangle(k_x, y, k_z)}, \quad (3.28)$$

where a and b stand for either u , v or w . The spectral error as defined above is related to other measures of correlation or ‘coherence’ across wall distances.

FIGURE 3.9: Spectral reconstruction error at different heights and Reynolds numbers, on top of the premultiplied velocity spectra and cospectrum (shaded) for F5300. The line contours are $R_{ab} = 0.5$: --- F2000; —, F5300. The shaded contours contain 90%, 50% and 10% of the total energy or of the tangential Reynolds stress. From left to right, $k_x k_z \phi_{uu}$, $k_x k_z \phi_{vv}$, $k_x k_z \phi_{wv}$ and $-k_x k_z \phi_{uv}$. From top to bottom $y^+ \approx 20$, $y^+ \approx 40$, $y/h = 0.1$, $y/h = 0.2$. The blue-shaded portion of the spectra marks the region reproduced with less than 50% error.

Other works using LSE evaluate its performance using the Linear Coherence Spectrum [13, 106, 168],

$$\gamma_{ab}^2(k_x, y, k_z) = \frac{\text{Re}\langle a^\dagger b^{\dagger*} \rangle(k_x, y, k_z)}{\text{Re}\langle ab^* \rangle(k_x, y, k_z)}, \quad (3.29)$$

which is related to R_{ab} by,

$$R_{ab} = 1 + \frac{C_{ab}^\dagger}{C_{ab}} - 2\gamma^2. \quad (3.30)$$

FIGURE 3.10: Scaling of the error isolines with the distance from the wall for F5300. From left to right, R_{uu} , R_{vv} , R_{ww} , R_{uv} . (a-d) Wavelengths scaled with the channel height. Contours are $R_{ab} = 0.5$ at $y/h = [0.012, 0.025, 0.05, 0.1, 0.15, 0.2, 0.3]$. The solid lines are within the logarithmic region, $150\nu/u_\tau < y \leq 0.2h$. (e-h) Wavelengths scaled with the distance from the wall, including only data within the logarithmic layer. Solid isolines, $R_{ab} = 0.5$; dashed, $R_{ab} = 0.25$; Shaded contours are the premultiplied spectra and cospectrum at $y/h = 0.1$.

It follows from (3.30) that the greater the coherence, the lower the reconstruction error. For the limiting cases where the reconstruction error R_{ab} is either unity or zero, both quantities are complementary, i.e. if the coherence is one, the error is zero and *vice versa*. One potential advantage of the spectral error is that the second factor in (3.30) penalises reconstructed energy without coherence, which may be significant for methods that are not L_2 -optimal.

The contour of $R_{ab} = 0.5$ is plotted at four wall distances in figure 3.9, compared with the energy spectrum of the three velocity components, and with the cospectrum of the tangential Reynolds stress. The regions shaded in blue and red indicate $R_{ab} < 0.5$ and $R_{ab} > 0.5$, respectively. In the first row, corresponding to $y^+ = 20$, the low-error region encloses most of the spectral plane and accounts for most of the energy, especially for v . As we saw in figure 3.7(a,b) the full flow field is well reconstructed at this distance from the wall. The accuracy decays quickly with y , especially for w . At $y^+ = 40$ (corresponding to the second row), the $R_{ww} = 0.5$ isoline only encloses wavelengths two to three times longer and wider than those at $y^+ = 20$. Figure 3.9 contains data at two Reynolds numbers. At the wall distances in the first two rows, which are in the buffer layer, the error isolines collapse well in wall units. In this range, the widest reconstructible spectral range is that of v , but, because the large scales of v contain much less energy than

those of u , the latter is the best reproduced variable in terms of energy. These relations persist up to $y^+ \approx 60$.

The trend to poorer reconstruction for larger y persists as we move into the logarithmic layer, as shown by the last two rows in figure 3.9. They are normalised in outer units and also collapse reasonably well for the two Reynolds numbers in the figure. The spectral region that can be well reconstructed is different for the four variables. In general, $w^\dagger$ has the narrowest spectral range, and $v^\dagger$ and $u^\dagger v^\dagger$ the widest one, although it is interesting that the Reynolds product, $u^\dagger v^\dagger$, can be reconstructed over a wider range than either $u^\dagger$ or $v^\dagger$. Note that the reconstruction efficiency only makes sense within the non-zero part of the spectra, since otherwise (3.28) is essentially undetermined.

The wavelengths of the error isolines above the buffer layer are approximately proportional to their distance from the wall, confirming that the eddies that are ‘attached’ in this sense are self-similar [170]. This is illustrated in figure 3.10, where the reconstruction boundaries are shown at different heights, scaled both with h and with y . The scaling with y is evident in the lower row of figures. The spanwise dimensions of the reconstructed $u^\dagger$, $v^\dagger$, and $u^\dagger v^\dagger$ are $\lambda_z \gtrsim 2y$, which is in fair agreement with the dimensions of the attached ‘quadrant’ (Q-)eddies isolated in the logarithmic region by thresholding uv [101]. The streamwise scales that can be reconstructed differ among variables. The longest one is $u^\dagger$, probably because the shorter features of u are less energetic, but $v^\dagger$ and $u^\dagger v^\dagger$ can be reconstructed down to $\lambda_x \approx 5y$, which is again in reasonable agreement with the dimension of Q-structures in [101]. In any case, while the relation between different variables appears robust, the absolute limits should not be taken too seriously. The 25% error isolines in the lower row of figure 3.10 show that the limits depend on the chosen error threshold.

As mentioned above, the spanwise velocity component is the worst reproduced at all heights. As seen in figures 3.9 and 3.10, its wavelength range is approximately 40% smaller than for u or v , although it is interesting to note that the instantaneous reconstructed fields satisfy continuity everywhere [4]. Even though the observable wavelengths of w are limited to larger (and wider) scales than those of v , the reconstructions of both velocity components contain similar kinetic energy, as these scales of w are more energetic than the large scales of v . We will see below that w is mostly associated with the shear at the wall, suggesting streamwise rollers or vortices. If this is the case, each roller would have a lower and an upper layer of w and, while the lower layer would leave a clear wall footprint, the upper one would be harder to reconstruct.

FIGURE 3.11: Reconstructed energy ratio with single-observable models for F2000 (top) and F5300 (bottom). The lines are: —, $a^\dagger$; ---, $a^\dagger_{u_y}$; ···, $a^\dagger_{u_y}$; — · —, $a^\dagger_p$. From left to right, uu , vv , ww , uv . The dashed vertical lines are $y^+ \approx 100$, approximately corresponding to the end of the buffer layer.

Figure 3.11 compares the overall performance of the reconstructions by means of the ratio

$$\beta_{ab}(y) = \frac{\sum_{\forall k_x, k_z} a^\dagger (b^\dagger)^*(k_x, y, k_z)}{\sum_{\forall k_x, k_z} ab^*(k_x, y, k_z)}, \quad (3.31)$$

between the energy of the reconstructed and original fields. The solid lines in the figure show that reconstruction recovers more than 50% of u^2 and uv below $y/h \approx 0.2$, but that the cross-flow velocities are only reconstructed to that accuracy within the buffer layer, $y^+ \lesssim 100$. All the reconstructions degrade above $y/h \approx 0.2-0.3$.

3.5.1 Incomplete wall data

Not all observables contribute equally to the reconstruction. Their relative contribution is different for each velocity component, and changes between the buffer and the logarithmic layer. In this section, we explore these differences by repeating the previous study using only one observable at a time. The restricted reconstructed velocity fields are denoted as $a^\dagger_p$, $a^\dagger_{u_y}$ and $a^\dagger_{u_y}$, where a stands for the variable being estimated, and the subindex is the observable.

Figure 3.11 compares the performance of the different restricted reconstructions with that of the full model. The reconstruction with the highest energy fraction is $u^\dagger$. Not surprisingly, the most important observable is $\partial_y u$

in this case, and $u_{u_y}^\dagger$ closely approximates the full model. The spanwise shear, and especially the pressure, have much less influence on this component. On the other hand, v is, on average, mostly reconstructed from the pressure, which seems reasonable, except very near the wall, where the spanwise shear dominates. This is probably associated with the presence of the near-wall streamwise vortices. The correlation between the wall pressure and v was studied in some detail by [175]. They find that the correlation is almost always maximum at zero spanwise offset, but that it has weak side-lobes at $\Delta z^+ \approx \pm 50$ when v is in the buffer region, consistent with the known structure of the streamwise vorticity in the sublayer [85]. The spanwise velocity is mostly controlled by the spanwise shear near the wall, but all observables have a similar influence on it farther away. As we have already discussed, the reconstruction of this variable is generally poor.

On the other hand, the tangential Reynolds stress, $u^\dagger v^\dagger$, is fairly well reconstructed up to $y/h \approx 0.3$. It is intriguing that $u^\dagger v^\dagger$, behaves better than $v^\dagger$, which is one of its factors, but we already mentioned when discussing figure 3.8 that the reason is that the large v -structures that can be reconstructed from the wall are precisely those associated with the large streaks of the streamwise velocity. The relevant observables for $u^\dagger v^\dagger$ are intermediate between those for $u^\dagger$ and $v^\dagger$. In the buffer layer, the two shears are dominant, while the pressure is not. Farther away, the three observables are comparable. Note that the good behaviour of this variable is encouraging from the point of view of observing dynamics from the wall.

The reconstructions for the two Reynolds numbers agree reasonably well in inner units close to the wall ($y^+ < 100$), and in outer units far from it. The obvious exception are the reconstructions based only on the pressure in figure 3.11, which contain considerably less energy at the higher Reynolds number. This trend is most pronounced for $u_p^\dagger$ and, consequently, for $u_p^\dagger v_p^\dagger$, for which the reconstructed energy is 50% smaller for F5300 than for F2000. The degradation is restricted to the reconstruction from the pressure, and does not affect either the full model or the reconstructions in terms of the shears. Inspection of the spectra of $u_p^\dagger$ and $v_p^\dagger$ at the two Reynolds numbers reveals that most of the energy missing from F5300 resides in relatively elongated scales, which we will see below to be dominated in the full model by the shear (see figure 3.12). A possible explanation for this behaviour follows from the decomposition of the pressure at the wall into an ‘inertial’ component,

$$\nabla^2 p_I = -\nabla \cdot (u \cdot \nabla u), \quad \partial_y p_I|_{y=0} = 0, \quad (3.32)$$

and a ‘Stokes’ one [84],

$$\nabla^2 p_s = 0, \quad \partial_y p_s|_{y=0} = \nu \partial_{yy} v, \quad (3.33)$$

Encinar et al. [42] showed that the latter can be computed from the wall shears using continuity,

$$\widehat{p}_s^+(0) = -i[k_x^+ \widehat{u}_y(0)^+ + k_z^+ \widehat{w}_y(0)^+] / |k|^+. \quad (3.34)$$

It can be shown that the ratio of between $\widehat{p}_s$ and $\widehat{p}_I$ is approximately independent of the Reynolds number when the wavenumbers are scaled in wall units, but that it decays as λ_x^+ increases along the diagonal $\lambda_z = \lambda_x$, which contains the cores of the velocity spectra (see figure 3.12), essentially because $k_x k_z \phi_{pp}(0)$ is located along this diagonal, but the two shears are not (figure 3.1b). The result is that, as Re_τ increases and we use our estimation to reconstruct $u_p^\dagger$ at a fixed y/h , and therefore at increasing $\lambda_x^+ \sim y^+ = Re_\tau y/h$, the importance of the Stokes component decreases. It follows from (3.34), that p_s is actually part of the shears, but there is no reason why the particular combinations of shears in (3.34) is in any way optimal to reconstruct the velocities. The pressure-only reconstructions are therefore contaminated by p_s , and should asymptote to a ‘pure’ reconstruction in terms of p_I as Re_τ increases. In this view, the ‘purest’ $u_p^\dagger$ is the one for F5300, and the apparently higher energy fraction in F2000 is accidental. The full model does not have this problem, because it uses information from the two shears to compensate the spurious effect of p_s . As a consequence, its performance in figure 3.11 does not change with Re_τ .

Figure 3.11 is contaminated below $y^+ \approx 100$ (indicated by a vertical red line) by the relatively low resolution of the data stored for F2000 and F5300, and we used the higher-resolution smaller box S1000 to estimate the effect of the contamination. The only difference found was a better performance of $v_p^\dagger$ and $w_p^\dagger$ (but not of $u_p^\dagger$), most probably because of a better estimation of the streamwise vortices of the buffer layer. This idea is supported by the high value of the cross-correlation between w_y and p at the wall, which is close to unity for the small scales, and the fact that w_y is the best predictor of the cross velocities in the buffer region.

Figure 3.12(a-d) displays the spectral distribution of the energy reconstructed by the different single observables. In most cases, the wall pressure and the two shears predict similar spectral regions, comparable to those of the the full reconstructions in figure 3.10, but the behaviour of $v^\dagger$ in the logarithmic layer is interesting, because different observables dominate in different spectral ranges.

The spectrum of v has a wide and a narrow component in the logarithmic layer [71], which appear in figure 3.12(b) as the two ‘horns’ on the upper right-hand corner of the spectrum. The spanwise wall shear contributes most to the estimation of the longer component, and the pressure dominates the shorter one, suggesting, respectively, streamwise rollers and compact sweeps and ejections. The reconstruction from the pressure is centred on equilateral wavenumbers, whereas the reconstruction from the streamwise

FIGURE 3.12: (a-d) Premultiplied spectra and cospectrum at $y/h = 0.1$ for F5300. Shaded, spectrum of the original fields; —, spectrum of the reconstruction using only the streamwise shear; ---, using only the pressure. The contour levels for all the spectra are 1.5%, 10% and 50% of the maximum of the corresponding true spectrum. The white contours in (b) are reconstructions at $y/h = 0.05$. The diagonal lines are: —, $\lambda_x = \lambda_z$; ---, $\lambda_x = 6\lambda_z$. (a) $k_x k_z \phi_{uu}$. (b) $k_x k_z \phi_{vv}$. (c) $k_x k_z \phi_{ww}$. (d) $-k_x k_z \phi_{uv}$. (e-f) Premultiplied spectra of $v_p^†$ (red) and $v_{uy}^†$ (black). (e) Modes along $\lambda_x = \lambda_z$. (f) Along $\lambda_x = 6\lambda_z$. (g) Streamwise sections at different $\lambda_z \in (0.1 - 0.84)$, from dark to light, of: —, $k_x k_z \phi_{vv}^+ = 0.05$; ---, $k_x k_z \phi_{uu}^+ = 0.3$. The straight lines are, —, $\lambda_x = \lambda_z$; ---, $\lambda_x = 6\lambda_z$; ·····, $y = 0.2\lambda_z$.

shear dominates at wavelengths elongated in the streamwise direction. Only $v_{uy}^†$ is shown in figure 3.12, but the effect on $v^†$ of the two shears is similar above $y^+ \approx 100$.

That this separation among v -structures is not restricted to a particular distance from the wall is shown in figure 3.12(e,f), which displays profiles of the energy spectrum of $v_p^†$, in red, and of $v_{uy}^†$, in black, at various points along the two diagonals in figure 3.12(b). Along the equilateral solid diagonal for which the pressure reconstruction is dominant, $\lambda_x = \lambda_z$, the spectrum reconstructed from the pressure is approximately three times more intense than that of $v_{uy}^†$ (figure 3.12e), and the opposite is true for the longer structures in figure 3.12(f), $\lambda_x = 6\lambda_z$. It is interesting that the peak of the reconstructed spectra is at $y/\lambda_z \approx 0.2$ for all the wavelengths for which this peak is in the logarithmic layer. When a similar plot is drawn for the true v , the peak is at $y/\lambda_z \approx 0.3 - 0.5$. Reconstruction from either observable only recovers struc-

FIGURE 3.13: Reconstruction of the wall-normal velocity at $y/h = 0.1$ in F5300. Top, using only the pressure, $v_p^\dagger$. Middle, using only the wall shear, $v_{u_y}^\dagger$. Bottom, true velocity field, filtered with (3.27) and $\Delta_x \times \Delta_z (2 \times 2)y$. The shading goes from $-u_\tau$ to u_τ .

tures closer to the wall than the average for the real flow field, presumably corresponding to the ‘roots’ of taller attached eddies.

Further insight into the difference among the two estimations is provided by figure 3.12(g), which contains streamwise sections of the premultiplied spectrum of u , in red, and of v , in green; represented against the wavenumber aspect ratio. The reconstruction peak at $y/\lambda_z = 0.2$ is marked by the horizontal dotted line, and intersects the lower edge of the v -spectrum. The two diagonals in figure 3.12(b) are represented by vertical lines in this plot, as

they cover wave numbers with constant aspect ratio. The solid vertical line is the equilateral diagonal, and crosses the centre of the v -spectrum. This is the spectral location at which $v^\dagger$ is best reconstructed by the wall pressure, and figure 3.12(g) suggests that it contains v -structures whose lower edge interacts with the impermeability of the wall to create a pressure peak, as larger modes are damped faster when they approach the wall. In contrast, the elongated diagonal $\lambda_x = 6\lambda_z$, represented in the figure by the dashed vertical line, crosses the peak reconstruction wall distance at the edge of the v -spectrum, but near the centre of the spectrum of u . This v -structures are too far from the wall to create an overpressure, but they can be detected from the wall shear because they are associated with the streaks of u , whose spectrum reaches the wall in this region. Finally, figure 3.13 displays two $v^\dagger$ snapshots reconstructed from the pressure (top), and from the wall shear (bottom). The difference in geometry is obvious, and the lateral spacing between streaky v -structures of similar sign in the bottom snapshot ($\Delta z/h \approx 0.6$, $\Delta z^+ \approx 3200$) approximately agrees with the separation between the streamwise velocity streaks at the distance from the wall of the figure (see the spectrum in figure 3.12a). It is interesting that, although it follows from figure 3.11(h) that the total uv -energy reconstructed from the streamwise shear is approximately twice larger than the one reconstructed from the pressure, the tangential Reynolds stress contained in intense structures [101] of the two snapshots in figure 3.13 are approximately equal (not shown), with similar intensity and area fraction.

3.5.2 Reconstruction in physical space

Although the spectral expression of the reconstruction operator discussed in the previous section is very useful in understanding the behaviour of the different flow scales, its physical-space counterpart defined in (3.23) and (3.24) gives a more intuitive representation of the operation that is actually being performed. It is probably also easier to generalise to wall-bounded flows besides the channel. We saw in §3.3 that the LSE operator is essentially a two-point correlation scaled with the cross-correlation tensor of the observables at the wall, and both correlations are likely to be reasonably independent of the type of flow in the near-wall region where the reconstruction is effective. For example, this is the case for boundary layers and channels below $y/h \approx 0.1$, although they differ farther from the wall [153].

Figures 3.14 and 3.15 show sections of the physical-space reconstruction operator $\tilde{L}_s^{(a)}$ for the three velocity components, $a = \{u, v, w\}$, and the three observables, $s = \{p, u_y, w_y\}$. Note that, although $\tilde{L}$ is formally a function of the two wall-parallel coordinate increments, $\Delta x = x - x'$ and $\Delta z = z - z'$, and of the reconstruction height y , their role is different. The operator is

FIGURE 3.14: Components of the operator $\tilde{L}_s$ for F5300. The rows are the operator for u , v and w , from top to bottom. The columns are the components corresponding to p , u_y and w_y , from left to right. Shaded contours are for $y/h = 0.1$, positive in red and negative in black. Lines are for $y/h = 0.05$, positive in green and negative in magenta. The contour levels are $[-0.5, -0.2, 0.2, 0.5]$ of the maximum absolute value at each height. The dashed lines indicate the vertical sections in figure 3.15.

FIGURE 3.15: As in figure 3.14. Streamwise sections: $z = 0$ for symmetric components, and at the location of the maximum absolute value at each height for antisymmetric ones. The dashed line is $y/h = 0.1$. The shading is scaled with the maximum absolute value at each height.

always two-dimensional in $(\Delta x, \Delta z)$, and acts as a convolution kernel on $E(x', z')$ at the wall. The wall-normal coordinate y is a parameter, and the operators for the reconstruction at different wall distances are, in principle, unrelated to each other. To emphasise this, the coordinates in these two figures have been inverted. The operator for one wall distance is given by the section of the three-dimensional function $\check{L}(\Delta x, y, \Delta z)$ at one y , as in figure 3.14, and $\Delta x < 0$ implies that the convolution kernel is acting on the wall upstream of the reconstruction point.

The wall-parallel reconstruction kernels in figure 3.14 show that the main reason for the lack of small scales in the reconstructions is that operators act as smoothing filters. All kernels have positive and negative parts, and their relative arrangement can partly be predicted from symmetry considerations. For example, $\check{L}_{u_y}^{(u)}$ in figure 3.14(b) is symmetric with respect to $\Delta z = 0$ because both u and u_y are symmetric with respect to $\Delta z \rightarrow -\Delta z$, but $\check{L}_{w_y}^{(u)}$ in figure 3.14(c) is antisymmetric because u is symmetric and w_y is antisymmetric. This has to be taken into account in interpreting figure 3.14, because the shape of some of the kernels suggests that the relevant variable is not the observable but some of its derivatives. For example, the positive-negative alternation along the Δx axis, centred at $\Delta x = 0$, of $\check{L}_p^{(u)}$ in figure 3.14(a) probably implies that the relevant observable is not p but $\partial_x p$, because there is no statistical symmetry to $\Delta x \rightarrow -\Delta x$, and thus $\check{L}_{\partial_x p}^{(u)}$ should be maximum at $\Delta x = 0$. The similar alternation in Δz , which would suggest $\partial_{zz} p$, may be due to the spurious symmetry of the correlation along the Δz axis, and could equally be interpreted as a statistical distortion of $\partial_z p$ in which 50% of the events fall to either side of the reconstruction point. For a discussion of symmetry artefacts of LSE in the context of the local reconstruction of the buffer layer, see [161].

The evolution of the reconstruction kernels with the distance from the wall of the target point is displayed in figure 3.15. Symmetric kernels are plotted at the symmetry plane, $\Delta z = 0$, but antisymmetric ones, for which this plane is identically zero, are plotted at the Δz plane for which the kernel is most intense. A general property of all the kernels is that they become wider as y increases, and that the growth is almost linear, suggesting self-similarity with respect to y . In fact, the wall-parallel plot in figure 3.14 is scaled with y , and includes two wall distances, whose kernels collapse well in this representation. Another property of most kernels in figure 3.15 is that they tilt forward, implying that the flow is reconstructed from upstream wall variables. Reference [175] found that the correlation of v with the pressure at the wall is antisymmetric in Δx , with both upstream and downstream lobes, but figure 3.15(d) shows that the corresponding kernel, $L_p^{(v)}$, is one of the few that is not tilted in Δx . This approximately applies to all the kernels based on the wall pressure, but those based on the shears are tilted by an angle that

depends on the particular case, but which is of the order of 10^0 – 20^0 to the wall, similar to other structural angles measured for u and w in the logarithmic layer of shear flows [153]. The disparity between the tilting of the shear and pressure kernels can be explained if we consider the processes generating the signals at the wall. If an attached eddy of size $O(l)$ is centred at a distance to the wall of $y_l \sim l$, the time it takes for a change of its dynamics to affect the shear at the wall should be of the order of $t_l \sim y_l/u_\tau \sim l/u_\tau$, equivalent to its local turnover time. Coincidentally, this time is related to the shear time by the Corrsin shear parameter, which is of the order of 10 for the logarithmic region of channel flows [66]. This implies that the information carried by the eddy is advected upstream during the time it takes for it to reach the wall. In contrast to that, the pressure of incompressible fluids instantly transmits sudden changes globally, generating a signal at the wall with no offset.

3.5.3 Off-design behaviour

Although the reconstruction operators are *Re*-dependant, they are obtained from two-point correlation tensors which are known to collapse in inner units close to the wall, and in outer units farther away [153]. This scaling extends to the linear estimators, allowing operators developed from DNS to be used at a different Reynolds number than the one at which they were obtained. Figure 3.16(a,b) presents a snapshot of the true v and of the reconstructed $v^\dagger$ in the buffer layer of S2000, and figure 3.16(c) presents a similar reconstruction of u in the logarithmic layer. Both reconstructions are performed using an operator computed from S1000. For the buffer-layer reconstruction ($y^+ \approx 10$), the observables, velocities and operators have been scaled in inner units, and, for the logarithmic region ($y/h \approx 0.1$) they have been scaled in outer units. Although there is some degradation of accuracy with respect to the reconstructions at the nominal Reynolds number in figure 3.7, the result is still reasonable. The worst errors are in scales that are not contained in the rescaled operators. For example, for computational boxes with similar dimensions in outer units, operators obtained at lower Reynolds numbers lack some of the longest scales when scaled in wall units. Similarly, some of the smallest scales are missing from those operators when scaled in outer units. The reconstructions in figure 3.16(b,c) have been correspondingly mildly filtered to avoid aliasing artefacts. The reconstruction operators are models of the structure of the flow, and the success of the rescaling supports the intuition that the attached turbulent structures near and far from the wall approximately scale in inner and outer units, respectively.

FIGURE 3.16: Snapshots of different reconstructions. (a) S2000, v^+ at $y^+ \approx 10$. (b) S2000, $(v_{S1000}^\dagger)^+$ at $y^+ \approx 10$. (c) S2000 at $y/h \approx 0.1$. Shaded, u^+ ; —, $(u_{S1000}^\dagger)^+ = 1$; ---, $(u_{S1000}^\dagger)^+ = -1$. S2000 is filtered with G and $\Delta_x \times \Delta_z = (0.4 \times 0.2)$. (d-f) S1000, $(u^\dagger)^+$ at $y/h \approx 0.1$, with information limited to discrete sensors (black crosses). (d) Fully dealiased. (e) No dealiasing. (f) Sensors averaged over 25% of the wall surface.

3.5.4 Practical application to limited input data

Lastly, we explore some of the issues pertaining to the practical application of this methodology. Figure 3.16(d-f) shows $u^\dagger$ at $y/h = 0.1$ for S1000, reconstructed from wall observations which are assumed to be only known at a coarse grid of sensors spaced by $\Delta_x/h = 0.2$ and $\Delta_z/h = 0.1$, represented in the figure by black crosses. Sampling has the effect of limiting the spatial resolution of the observables to wavelengths longer than the Nyquist limit ($2\Delta_x, 2\Delta_z$), and of aliasing the information of the shorter wavelengths into longer ones that fall within this limit [25]. If the information can be properly dealiased by filtering the short wavelengths before sampling, the only effect is the loss of resolution, as in figure 3.16(d). Because of the linearity of the reconstruction and the homogeneity in the wall-parallel planes, this is equivalent to *a posteriori* filtering of a reconstruction performed at full resolution. Unfortunately, this is usually impossible, because the full measurements are unavailable. Figure 3.16(e) shows the worst-case scenario in which ‘point’ sensors are used without dealiasing. The reconstruction is heavily distorted, mostly in the form of higher intensities, and the effect is more severe than what could be expected from simple aliasing of the wall observables. The latter contain a lot of small scales which should have no effect on reconstruc-

tions far from the wall, but which are aliased into larger scales and amplified by the linear operator into large spurious features of the reconstruction. A more realistic scenario is presented in 3.16(f), in which the sensors are modelled as finite rectangles with size $\Delta_x/h = 0.1$ and $\Delta_z/h = 0.05$, and assumed to provide values averaged over their surface. This *a priori* filtering does not fully eliminate the aliasing error, but reduces it significantly by reducing the short-wavelength content of the wall observables. In a real experimental application, where the temporal evolution of the sensor signal is available, time filtering can also be used to limit the aliasing error. The lifetimes [103] and convection velocities [8] of the turbulent structures in the logarithmic layer are known, and are typically significantly longer and faster than those of the scales responsible for the aliasing of the wall signal. Time filtering of the signals at those scales mimics streamwise averaging through Taylor's approximation, and supplements the benign effect of non-zero sensor size.

3.6 CONCLUSIONS

Because the pressure and the second invariant, Q , of the velocity gradient tensor are linked by a Poisson equation, the latter can be expressed as a projection of the former on a Green's function, and the optimal reconstruction can be posed as an orthogonal projection problem. However, the result of this projection is physically unsatisfactory, and a better strategy is to take into account the statistics of Q , in the form of its two-point autocorrelation function. The resulting procedure is equivalent to linear stochastic estimation, and provides reconstructed Q s which agree very well with the observed ones in the viscous layer, reasonably well below $y^+ \approx 100$, and well for large-scale structures farther from the wall. Correspondingly, only large enough structures far from the wall create a footprint on the wall pressure. The reconstructed Q structures share many characteristics with the attached vortex clusters described by [11].

We have also shown that the three velocity components in turbulent channels can be reconstructed reasonably well below $y/h \approx 0.2$ with linear stochastic estimation, using only the observed pressure and the streamwise and spanwise shear at the wall. Only eddies attached in the sense of having sizes comparable to y [171], are reconstructed, but these are found to contain more than approximately 50% of the kinetic energy of u , and of the tangential Reynolds stress, uv , in the logarithmic layer. The optimum reconstruction uses the three wall observables, but not all them are equally important, and the relative balance among them depends on the variable being reconstructed and on the distance from the wall. We have proposed that this can be used to sharpen the definition of which velocity components are attached to the wall in which spectral ranges, as those that can be

reconstructed from wall information, regardless of the particular reconstruction method. The results differ from the generally accepted ones in some cases. For example, the wall-normal velocity, which is usually considered a detached variable because it is inhibited near the wall by the impermeability condition, is linked to the wall by the pressure, for equilateral scales, and by the wall shears, for longer scales in which v is associated with the u -streaks. Interestingly, although the two sets of attached v -structures defined in this way, are very different, with a different balance of u and v components, both are active in the sense of generating a comparable amount of tangential Reynolds stress.

Although most of the chapter deals with mode-by-mode reconstructions in the Fourier representation of the flow, we have shown that the expressions of the operators in physical space gives interesting information on the flow physics. Thus, while the wall pressure predominately reconstructs velocities immediately above the observation point, the effect of the shear is transmitted along lines inclined by 10° – 20° from the wall, comparable to previous observations of the structures of u and w in wall-bounded flows.

We have also shown that the reconstruction operators obtained at one Reynolds number can be adapted to a different one by proper scaling, making them potentially useful for control strategies in which only wall observations are available, or, in turbulence physics, for the study of the flow dynamics from the wall. They can be adapted to experimental setups in which the instantaneous measurements are only available on a coarser grid of sensors, but only after careful dealiasing of the observations. Although not explicitly shown, it can be inferred from figure 3.11 that the information provided by the two shears is equivalent in the logarithmic region, and a practical grid of sensors may obviate the spanwise shear with almost no penalty in the accuracy of the reconstructions. Note that this does not apply to the flow in the buffer layer, where the spanwise shear is the most relevant observable for the cross flow.

Considering the results presented previously, this chapter explores the possibility of using the techniques introduced in Chapter 3 to detect the structures identified in Chapter 2. The approach we propose consists on band-pass filtering the reconstructed velocity fields using (2.10), and detect structures in the amplitude and inclination angles computed from those fields. Several considerations are required for proper identification, such as the difference between the reconstructed and real velocity perturbations.

The remainder of this chapter is organised as follows. We review the methodology, with considerations on its practical application in §4.1. We present the results in §4.2 and §4.3 closes with the discussion and conclusions.

4.1 THE FILTERED-RECONSTRUCTION ESTIMATOR

In Chapter 2 we detected the Orr-like structures in two ways, using a band-pass filter and using a low-pass filter. Here we explore the properties of the band-pass filtered reconstructed fields, and we leave aside, for the moment, the analysis of low-pass filtered fields. This is in part because we saw in 3.5 that the reconstructed field can be considered a low-pass filtered version of the original, and the direct detection of vortex clusters and Qs was already hinted in Chapter 3. Another reason is that is more practical to use the band-pass filter for the detection because it can be done more economically, as the quantities implied are integrated across the wall-normal direction and thus lower dimensional. Let us recall the equations of the amplitude and inclination angles of the band-pass filtered fields,

$$A_a^2 = \frac{1}{y_1 - y_0} \int_{y_0}^{y_1} |\bar{a}|^2 dy, \quad (4.1)$$

$$\Psi_a = \arctan \frac{\Lambda \int_{y_0}^{y_1} \text{Im}(\bar{a}^* \partial_y \bar{a}) dy}{2\pi \int_{y_0}^{y_1} |\bar{a}|^2 dy}; \quad (4.2)$$

where $\bar{a}$ stands for a band-pass filtered velocity component. The reconstruction of $\hat{a}$ for a given (k_x, k_z) mode is,

$$\hat{a}^\dagger(y) = \hat{L}_j(y) \hat{E}_j, \quad (4.3)$$

which band-pass filtered yields,

$$\hat{\bar{a}}^\dagger(\Lambda) = \hat{G}(\Lambda) \hat{L}_j(y) \hat{E}_j. \quad (4.4)$$

FIGURE 4.1: Pseudo-spectra of the wall-normal velocity component (left) and the reconstructed wall-normal velocity (right) at $Re_\tau \approx 5000$. The shaded contours are 50% and 90% of the maximum at each height, indicated by the chained lines. (left) The solid line is $y = \Lambda$, and the dashed line is $y = 4\Lambda/11$. (right) The red lines are analogous to the black lines in (a). The black lines in (b) are 2.4 times closer to the wall than in (a).

Note that we can define the operator $\bar{L}_j$ as the convolution of the band-pass filter G with the components of the operator L , which in turn computes the filtered-reconstructed velocity $\bar{a}^\dagger$ from wall measurements. Consider also that equations (4.1–4.2) only require the values of the reconstructed velocities in the range (y_0, y_1) and thus the operator $\bar{L}_j$ only needs to be evaluated in that range, further reducing the computational cost.

Because the spectrum of the reconstructed velocity field $v^\dagger$ is different from that of the real velocity (see §3.5), the integration bands (y_0, y_1) need to be reassessed for this case. Figure 4.1 shows the pseudo-spectra $\mathcal{E}_{\bar{v}\bar{v}}$ of the wall-normal velocity (equivalent to figure 2.4b) and of the reconstructed velocity, $\mathcal{E}_{\bar{v}\bar{v}}^\dagger$, at $Re_\tau \approx 5000$. The energy in $\mathcal{E}_{\bar{v}\bar{v}}^\dagger$ is concentrated in a thinner band than $\mathcal{E}_{\bar{v}\bar{v}}$, and displays a comparable if not larger self-similar range. This is not surprising, as we saw in §3.5 that the reconstructed part of the flow is ‘attached’ to the wall and remarkably self-similar in the logarithmic region, and so is the filter. However, the band is displaced respect to the original by a factor of 2.4, which can be interpreted in one of two ways. A possible interpretation of this displacement would be that the structures detected are larger than the most energetic v in $\mathcal{E}_{\bar{v}\bar{v}}$. The second interpretation, more plausible, is that only the bottom part of the structures, or the ‘feet’, leaves a (pressure) footprint at the wall and thus can be reconstructed from it. We recall that the pressure is associated to the equilateral part of the reconstruction and thus retained by the filter, unlike the elongated structures of v that can be detected from the shear. This hypothesis is backed by the behaviour of the large structures of $\bar{v}^\dagger$, represented in the range $\Lambda > 2h$, where the effect can be clearly seen as a displacement of the spectra towards the wall and not a increase in the length scale of the observed perturbations.

FIGURE 4.2: Components of the operator $\check{\tilde{L}}_{ij}$ for Λ_2 for F5300, at the center of the integration interval. The rows are the operator for u , v and w , from top to bottom. The columns are the components corresponding to p , u_y and w_y , from left to right. The contour levels are $[-0.5, -0.2, 0.2, 0.5]$ of the maximum absolute value at each height. (a-i) Real part. (j-r) Imaginary part.

FIGURE 4.3: Joint probability density function of $A_{\bar{v}^\dagger}$ and $\Psi_{\bar{v}^\dagger}$. The amplitudes are normalised with their modal values. The lines are: $\text{---}\bullet\text{---}$, Λ_1 ; $\text{---}\blacksquare\text{---}$, Λ_2 ; $\text{---}\blacktriangle\text{---}$, Λ_3 ; $\text{---}\blacklozenge\text{---}$, Λ_4 . Contours are 90%, 50% and 10% of the probability mass.

The real and imaginary part of the operators are shown in figure 4.2. Most of the signal of $\bar{v}^\dagger$ comes from the imaginary part of the pressure component, in line with the findings of [175]. An alternating pattern of structures can be seen in the spanwise direction of the pressure operators of $\bar{u}^\dagger$, reminiscent of the streak cores. The pressure operators of $\bar{v}^\dagger$ in figure 4.2(d), and $\bar{u}^\dagger$ in figure 4.2(a) are in phase opposition, which is not surprising considering that the strong events are usually Q_2 or Q_4 events. However, it is surprising that within the spectral band of the filter, the streamwise-shear operators are in phase, although their imaginary parts are not collocated. We saw in §3.5.1 that for the equilateral modes that this band-pass filter retains, the signal that comes from the shear stress is minimal, still but is surprising that is associated with Q_1 and Q_3 events.

4.2 RESULTS

Figure 4.3 shows the joint probability density function of amplitude and inclination angles of the reconstructed wall-normal velocity of F2000. Most features of the joint pdf resemble those of the real wall-normal velocity, for example, the triangular shape that relates the amplitude and the inclination angle or the higher intensities of the smallest filters. It has lost part of the probability mass in the areas dominated by linear dynamics, i.e. the upper edges of the triangle, and the different filters do not agree as well between themselves as they do in the joint pdf of the real velocities. The source of this disagreement is the bottom right tail of the pdf that drifts based on the filter width. It was argued in [68] that this bottom right corner was affected by the

FIGURE 4.4: The solid lines are the joint probability density function of $A_{\bar{v}}^\dagger$ and $\Psi_{\bar{v}}^\dagger$ of $\bar{v}_2^\dagger$. Contours are 90%, 50% and 10% of the probability mass. The shadings are the mean value of $A_{\bar{u}}^\dagger$ (left) and $\Psi_{\bar{u}}^\dagger$ (right) conditioned to the values of the pdf.

(eddy) viscosity, but is hard to imagine how that difference affects only the reconstructed velocities and not the real ones.

We showed in §2.3 that the Orr burst was promptly preceded by backwards-leaning streamwise velocity perturbations, with strong intensity. Figure 4.4 explores the relation between $\bar{u}^\dagger$ and $\bar{v}^\dagger$, and it shows that it is similar to what we found in the real velocities. The top left part of the joint pdf is characterised by negative values of the inclination angle of u , but the strong intensity has shifted towards the amplitude peak of the joint pdf. Despite these differences, the distribution of the inclination angles of u corroborates that the detected events are the ‘feet’ of the bursts detected in §2.3.

Following the procedures developed in Chapter 2, we computed conditional mean evolutions based both on the intensity and inclination angles of the wall-normal velocity and of the streamwise velocity. While the former criterion aims to directly detect the burst, the latter was used in §2.3 as a predictor of bursts. In order to compute these conditional evolutions, first we need to compute a moving frame, $\xi = x + C_\Lambda t$ that follows the mean advection of the structures in the streamwise direction. Compared to the convection velocities of $\bar{v}$, the convection velocities of $\bar{v}^\dagger$ are roughly constant across the different filter widths and integration bands, and further away from the mean velocity profile. These features are reminiscent of the ‘global’ modes described in [8], although the length scales of those modes are large ($\lambda_x > 2h$) and independent with y , suggesting that the nature of our structures is different. Nevertheless, if the convection velocities are different in the reconstructions than in the real fields, it is safe to assume that the detected ‘wavepackets’ are dispersive in nature. This means that the reconstructed ‘feet’ of the structures can only be shadowed by their ‘heads’ for

FIGURE 4.5: Conditional mean evolutions of the reconstructed velocities. (Left) Conditioned to the presence of a burst ($A_{\bar{v}} > 2A_{\bar{v}M}$, $|\Psi_{\bar{v}}| < 0.15$). (Right) Conditioned to the presence of a precursor ($A_{\bar{u}} > 1.8A_{\bar{u}M}$, $\Psi_{\bar{u}} < -0.4$). The symbols are $\bullet$, $\bar{u}^{\dagger}$; $\blacksquare$, $\bar{v}^{\dagger}$; $\blacktriangle$, $\bar{w}^{\dagger}$. The lines are — , λ_1 ; -- -- , λ_2 .

short times, and cannot remain coherent during the lifetime of the Orr burst. This is shown in figure which shows the conditional evolution of both the burst and its precursor. The precursor fails to precede the burst, and instead acts as a detector, roughly equivalent to the burst conditioned to $\bar{v}^{\dagger}$. In both cases, while the negative inclination angle of $\bar{u}^{\dagger}$ is clearly recovered (even though in the conditional burst nothing is imposed on $\bar{u}^{\dagger}$), the inclination angle of $\bar{v}^{\dagger}$ is close to zero during the whole evolution, and its ‘tilting’ from negative to positive cannot be objectively observed. This could be already inferred from figure 4.3, where the upper edges of the joint pdf of $A_{\bar{v}^{\dagger}}$ and $\Psi_{\bar{v}^{\dagger}}$ are missing, and figure 4.4, where the footprint of $\bar{u}$ is shifted towards the peak of $\bar{v}$.

4.3 CONCLUSIONS

We explored the possibility of using band-pass filters, together with LSE, to identify Orr-like bursts using information limited at the wall. The detection operation can be done economically because the required quantities, amplitude and inclination angles, are wall-normal averages, and thus the resolution of the reconstruction in y is not critical, and the sampling space of the operator is limited to the integration intervals. We found that Orr-like structures can be detected, although important differences exist when they are compared with the Orr-like structures of Chapter 2. First, the observed length scales are larger at the same wall-normal distance than the most energetic structures of ν studied in §2.3. Following the discussion in Chapter 3, this disparity of scales can be interpreted as the ‘feet’ of the structures been the only part that can be reconstructed from the wall. The statistical proper-

ties of the ‘feet’-structures are similar to those of the reconstructed velocities, both of the wall-normal and the streamwise velocity components. The similarities between $\bar{v}$ and $\bar{v}^\dagger$ allow the detection of Orr bursts from the wall. However, the differences among them are enough to make the predictor ‘faulty’ at detecting them before peaking time. A possible explanation for these differences is that the wall-normal velocity perturbations have to be intense enough to leave a pressure footprint at the wall, as they are undistinguishable for the turbulent background otherwise. It is then reasonable that they can only be detected then they are the most intense, and an the predictor fails at its intended task. For this reason, the interpretation of the effect of the time axis in figure 4.2 should be carefully considered, as it is probably only indicative of the variables reverting to the unconditional mean when separating from the condition centre, and barely related to their time evolution.

DISCUSSION AND CONCLUSIONS

In this thesis we have investigated the dynamics of the logarithmic layer of turbulent channel flows. For this purpose, we analysed several databases of time-resolved channel flows at high Reynolds numbers, in an attempt to identify and detect coherent structures which are simple enough to be described by linear dynamics. We have also explored the possibility of observing turbulence, with emphasis in turbulent structures, exclusively from wall-measurements, which would be the only observables in any practical detection system.

This chapter concludes this dissertation with a summary of the original results presented in this work.

5.1 COMPLEX ‘VELOCITY’ FIELDS

We introduced in §2.2 a relatively novel methodology that can be used to analyse turbulent flow fields. Similar techniques, such as wavelet transforms [43, 112] or Hilbert transforms [160] predate this work, and a inverse citation look-up on the citations thereof, reveals thousands of publications using wavelets. Nevertheless, to the best of our knowledge, this is the first time complex wavelets have been used to isolate and track structures in high-Reynolds-numbers turbulent channel flows.

The complex filters can be interpreted as a complex-wave approximation to the real signal locally, possessing both an amplitude, and a phase. The filter can be approximately¹ interpreted as a real-valued band-pass filter acting over the *analytic signal* of the velocity. The complex analytic signal $S_u(x)$ of a real-valued signal $u(x)$ can be defined in several ways, for example,

$$S_u(x) = u(x) + i\mathcal{H}[u(x)], \quad (5.1)$$

where $\mathcal{H}$ denotes the Hilbert transform, in our case, along the streamwise direction. In turn the Hilbert transform can be defined in Fourier space as,

$$\mathcal{H}[\hat{u}(k)] = \begin{cases} i\hat{u}(k) & \text{for } k < 0, \\ 0 & \text{for } k = 0, \\ -i\hat{u}(k) & \text{for } k > 0. \end{cases} \quad (5.2)$$

¹ Only for fluctuating (meanless) signals, and this yields a different normalisation compared to our approach.

It follows from equations (5.1) and (5.2) that the positive frequencies of S_u are zero, and the negative frequencies with $k < 0$ are twice the amplitude of the original signal. A direct definition of S_u would be,

$$S_u(k) = \begin{cases} 2\hat{u}(k) & \text{for } k < 0, \\ \hat{u}(0) & \text{for } k = 0, \\ 0 & \text{for } k > 0. \end{cases} \quad (5.3)$$

Because $u(x)$ and $\mathcal{H}[\hat{u}(k)]$ are orthogonal and have the same power spectrum, the modulo of S_u represents the envelope of u . In turn, the derivative of the phase of S_u is a local wavenumber. Under this framework, the real-valued filter simply isolates a band of frequencies of a signal that contains the local envelope and local wavenumber of the original velocity field at each point. An alternative approach is a physical interpretation of the functional (2.14), where the resulting complex signal is a ‘local’ Fourier coefficient acting inside a gaussian window. The value of the functional itself could be used as a figure of merit for the local approximation. Non-coincidentally, both the definition of analytic signal and the wavelet we used for our analysis can be attributed to Gabor [51].

The amounts of works that link linear theories to shear flows is huge. From linear stability, to resolvent operators [110], transient growth models [11, 44] or more complicated frameworks [30, 58], it has become apparent from the results in §2.3 that only portions of the turbulent flows evolve as dictated by linear dynamics. Here, ‘local’ implies locality in time, space and scale space, and each of them are or can be considered in the definition of the filter. Because the linear models are formulated in Fourier space, the complex wave is adequate to study the local behaviour of models, as an alternative to using minimal domains. This is particularly useful in the multi-scale logarithmic region, where a box can only be minimal for one height.

5.2 ORR BURSTS IN THE LOGARITHMIC REGION OF CHANNEL FLOWS

In §2.3 we used the complex filters to detect Orr-like bursts. The wall-normal velocity bursts, characterised by the relation between their amplitude and their inclination angle, are found to populate a wide range of heights ($60\nu/u_\tau < y < 0.7h$) of channel flows. Their evolution is qualitatively similar to the linear Orr bursts that can be computed using optimal transient growth algorithms across the equilateral $\lambda_x \sim \lambda_z$ wavelengths.

Their spatio-temporal scales are self-similar with to the distance to the wall, and higher Reynolds numbers simply contain a wider scale separation between the largest and smallest bursts. We did not find any difference between the bursts in small ($2\pi \times \pi$) or large ($8\pi \times 3\pi$) boxes, and the half-height structures detected by Jiménez [68] agree well with their long-box

counterpart. This suggests that the dynamics of the Orr bursts are well captured in minimal domains for the logarithmic region. The statistical analysis revealed that only a small fraction of the channel flow time series is occupied by self-similar Orr bursts ($\sim 10\%$), and the local approximation to the velocities as streamwise wavepackets fails elsewhere. This filling fraction is roughly constant with the length scale of the burst, with more bursts occupying the same spatio-temporal volume the smaller they are.

Analysing the wall-parallel velocity components showed that the spanwise and the wall-normal component are part of a roller, as a significant spanwise velocity signal is always located at a small spanwise offset of ν , and evolves in agreement with the linear theory. The signal of u precedes the burst, and can be used to predict it, for short times. This strong footprint in relatively short wavelengths of u is indicative of inhomogeneities in the well-known streamwise velocity streaks. Streak inhomogeneities are believed to be part of the self-sustaining cycle [48, 56] of the logarithmic region, although it cannot be confirmed solely from the evidence found in this work.

We enriched our analysis using low-pass filter fields. The key idea here is that low-pass filtered fields are real valued, and thus easier to interpret as coarse-grained velocity fields. Most of the characteristics that are directly comparable between the band-pass and the low-pass methodologies, such as density of bursts, observed tilting angles, and relation between the velocity components are in agreement with each other. This allowed us to compute a conditional evolution that is reminiscent of the quadrant structures studied by various authors [38, 101, 103, 104, 178]. The spatial and temporal scales, as well as the general shape of the conditional evolutions are qualitatively and quantitatively similar to the attached Q_{2-4} pairs, as they also generate tangential Reynolds stress and come in pairs of opposite wall-normal velocity sign, strongly suggesting that they are essentially the same objects. If the dynamics of the largest attached quadrant objects are controlled by the Orr mechanism, as evidenced by the data, this would be against the idea of structures being created at the level of the buffer layer and ‘growing’ into the logarithmic layer. It does however support the growing evidence that the shear generates these structures, and that the role of the wall is providing such shear.

The similarities between the quadrant pairs and the Orr-bursts do not end their local structure. We studied the spatio-temporal organisation of Orr bursts of the same scale, revealing a strong statistical evidence for intense events of opposite sign to be in pairs, and same-signed events to be away from each other. The pairs are preferentially organised along the spanwise direction, with a small offset in the streamwise direction, and a better-than-Poisson chance to be contemporaneous. As proposed by Lozano-Durán and Jiménez [103], the data suggests that the pair, connected by the spanwise

component, can be interpreted as a single ‘roller’ that sits in the edge of a streamwise streak. The organisation of these rollers turns out to be preferentially along the streamwise direction, with a preference for small spanwise offset, in agreement with the idea that several Q s ‘collaborate’ to generate or sustain a single streamwise streak [69]. The events tend not to be simultaneous, having waiting times between closest events of the order of their lives.

If the spanwise and wall-normal velocities found in the conditional evolutions have similar sizes, and are self-similar when normalised with the filter width, the streamwise velocity is not. The conditional bursts are embedded in large streaks which do not collapse when normalised with the filter width. However, the ‘core’ of the streak, which can be isolated from the streak with a high-pass filter, does scale as the wall-normal velocity. The core can be identified without the aid of the high-pass filter by scaling the streaks with the filter width and observing the only portion that collapses under this normalisation. The portion thereof turns out to be the ‘leading edge’ of the streaks, in the lower-front of the high-speed streaks, and in the upper-back of the low-velocity streaks. This suggests that these edges are the dynamically active parts of the streaks, although further research is required in this matter.

5.3 LOGARITHMIC LAYER TURBULENCE: THE VIEW FROM THE WALL

Motivated by the desire of carrying over structural models to real applications, the second part of this thesis explores the possibility of observing the flow field far away from walls, exclusively from the information available in its surface. We explored the linear relation between the pressure at the wall and the second invariant Q of the velocity gradient tensor far away from it. The information contained in the equation results insufficient for a physically meaningful estimation of Q , and a statistical approximation is preferred. A linear approximation is proposed, as it conserves the linear properties of the relation. The approximation recovers a small portion of the Q -field, which is still enough to generate the correct pressure signature at the wall. The reconstructed Q -field turns out to be composed by ‘sponges’ that shadow the location of the vortex clusters identified in wall-bounded flows [10, 83].

Owing to the satisfactory performance of the Q reconstructions, we pursued the reconstructions of the velocity fields, even if they lack a linear relation linking them to the wall. Nevertheless, the linear reconstruction turns out to produce reasonable results, and has several advantages over other methods. For example, is L_2 -optimal, meaning that no ‘fake’ perturbations are reconstructed; the reconstructed velocities satisfy continuity everywhere, as linear relations are preserved by the reconstruction, and the effects of

noise in the input signals can be predicted *a priori*, as linearity implies that the noise does not interact with itself in the reconstruction process.

The properties of Fourier space are exploited to obtain an operator that uses full contemporaneous information at the wall to reconstruct each coordinate of the wall-parallel planes of the flow field. The reconstructions are excellent in the buffer layer, but degrade with the distance from the wall. The accuracy within the buffer layer suggests that applications at low Re or aimed at the buffer layer will not benefit from more complex methods. In the logarithmic layer, the reconstructed flow fields are remarkably self similar, which suggests that the eddies reconstructed are those attached in the sense of Townsend. Moreover, it suggests an alternative definition of attached eddies, being those that can be observed from the wall, independently of the particular methodology employed. The main implication of this alternative definition is that it can be quantified, as the flow can be decomposed in attached and detached portions. The cross-velocity components are poorly reconstructed, as the scales that can be observed are larger and wider than the typical eddies of v and w . That is not the case for the streamwise velocity or the attached tangential Reynolds stresses which can be recovered to great extent.

The reconstructions use the three observables at the wall, but not all of them are equally important in practice. Unsurprisingly, the streamwise component is mostly reconstructed from the knowledge of the streamwise shear, and the pressure is almost superfluous. We found that the wall-normal component is observed both by the streamwise shear and the pressure, and that the typical eddies that can be detected from both are quite different. The pressure can be used to observe equilateral objects, reminiscent of the Orr bursts studied herein. Their wall-normal location suggests that they are the bottom part of the bursts, which leave a pressure footprint at the wall due to its impermeability. The streamwise shear is associated to long (and weak) features of v , which nevertheless carry a comparable fraction of the total reconstructed Reynolds stress, making both observables useful in this regard.

The reconstructions are satisfactory below $y = 0.3h$, and the simplicity of the method makes it suitable for practical applications. The implications of having realistic measurements of the observable rather than noiseless high-resolution DNS data can be studied *a priori*, making this another key advantage of sticking to a linear model, at least for practical applications.

5.4 FUTURE WORK

5.4.1 Local wave approximation

We developed a method that approximates the local flow field as a wave-packet. In our case, it was convenient to approximate the flow field as a

streamwise wave, in order to compare the flow structure with linear models. The key advantage of the method is that the wavepackets are *local*, and thus can isolate portions of the flow. There are many other potential applications of this method. An immediate one could be estimating the convection velocities of different structures of the flow. Relating the phase velocity of waves with the convection velocities of structures is far from new [8, 181], but these methods so far have been *global*. A different approach consists in tracking different objects, e.g. quadrant structures, in time and directly computing the propagation velocity of said objects [103]. If the former methods proved that the convection velocities depend on the wavelengths (sizes) of the eddies, the latter showed that, evidently, high- and low- speed regions propagate at different speeds. A hybrid methodology like the one we propose could be used to isolate different regions locally in space and scale and study their convection velocities.

5.4.2 *What are the events seeding the Orr bursts?*

We were able to unequivocally relate streamwise perturbations to the formation of the Orr bursts. It is still unclear where do those perturbations come from and if they represent any known structure of the flow field. A popular model known as ‘self-sustaining cycle’ or ‘self-sustaining process’ is believed to maintain turbulence in the buffer layer [56, 73, 74, 177]. This model (roughly) consists in a turbulent streaky background that breaks due to meandering (sinous) instabilities and generates quasi-streamwise vortices which in turn generate new streaks through the lift-up of the mean profile. A similar scenario is believed to happen in the logarithmic region [48, 53] although the observations of sinous instabilities are only clear in minimal domains, and a more disorganised scenario can be observed in large boxes [40]. The discovery of the ‘active’ edge of the streaks hints at a good place to look for streak instabilities in large-box DNSs, which could shed some light in the aspect and dynamics of this ‘burst precursors’.

5.4.3 *Multi-scale interactions of Qs-Orr bursts*

In this work we isolated the different Orr bursts from each other through the filtering process and the integration bands. However, it is possible that different structures of the self-similar hierarchy interact with each other. These interactions, if present, are removed by our filters, and thus the current methodology does not allow any conclusions about their mutual dynamics or lack thereof. Cardesa et al. [26] showed that structures in HIT interact with each other when the scale separation is approximately two, and that this interaction is mostly gone by scale separation of order of four. Because in

the channel flow the dynamics of the Orr bursts are controlled by the mean shear, their interactions, if any, should be different in nature than in HIT, but a dedicated study in this matter should help in our understanding of wall-bounded turbulence.

5.4.4 Towards a quantitative definition of wall-attached

We proposed that the flow can be decomposed into attached and detached components, i.e

$$\mathbf{u} = \mathbf{u}_{\text{attached}} + \mathbf{u}_{\text{detached}} = f(\mathbf{E}^\star) + \mathbf{u}_{\text{detached}},$$

where the attached component is dynamically related to the wall. In Chapter 3, the relation used is Linear Stochastic Estimation (LSE) $f(\mathbf{E}^\star) = \mathbf{L}\mathbf{E}^\star$, and the observables are limited to contemporaneous measurements of the three stresses at the wall, $\mathbf{E}^\star = \mathbf{E} = [p, \partial_y u, \partial_y w](x, z)$. While these limitations are reasonable for practical applications, they probably are unjustified for theoretical approaches. However, it may be possible to define both $f(\bullet)$ and $\mathbf{E}^\star$ in a way that $\mathbf{u}_{\text{attached}}$ is optimal, and thus a objective definition of ‘attached’. The definition of $\mathbf{E}^\star$ is secondary to $f(\bullet)$, as one can define concepts as *attached under contemporaneous pressure observations* or *attached under time-resolved shear-stress observations*. If the maximum observability is desired then the vector

$$\mathbf{E}_{\text{sup}}^\star = \left[\partial_x^i \partial_y^j \partial_z^k \partial_t^m \mathbf{u}(x, y, z, t), \partial_x^i \partial_y^j \partial_z^k \partial_t^m p(x, y, z, t) \right]_{y=0},$$

$$\forall x, z, t; \quad \forall i, j, k, m \in \mathbb{N};$$

contains all the dynamical information of the flow available at the wall. Because full streamwise, spanwise and temporal information is contained in $\mathbf{E}_{\text{sup}}^\star$, the vector

$$\mathbf{E}_{\text{sup}}^\star = \left[\partial_y^j \mathbf{u}(x, y, z, t), \partial_y^j p(x, y, z, t) \right]_{y=0}, \quad \forall x, z, t; \quad \forall j \in \mathbb{N};$$

contains the same information, as the missing derivatives can be obtained from it if necessary. Note that $\mathbf{E}_{\text{sup}}^\star$ is ‘useless’ in practical applications but may be useful for a extremal definition of attached. Physical considerations may make desirable to limit $\mathbf{E}_{\text{sup}}^\star$ to measurements of the stresses at the wall, as it is hard to imagine the physical relevance of high-order wall-normal derivatives of the stresses. The function $f(\bullet)$ is more complicated to define, although some approaches may be fruitful. One option could be to use non-linear optimisation [120, 131], which could be closer to the upper bound of what is observable. Another alternative is to use state- (phase-) space reconstruction methods [37, 162], whose downside is their prohibitive cost when the number of degrees of freedom of the problem increases.

APPENDIX

SECOND-ORDER CONSISTENT DATABASES FOR CHANNEL FLOWS

Part of the contents of this appendix are published online in a technical report, under the name ‘A second-order consistent, low-storage method for time-resolved channel flow simulations up to $Re_\tau = 5300$ ’, with Adrián García-Gutierrez and Javier Jiménez as coauthors, and Alberto Vela-Martín as corresponding author.

Wall bounded flows are fundamental building blocks of many industrial applications or scientific research, and developing accurate models for their prediction and control is a crucial challenge for the next decades. One fourth of the energy in advanced economies is used in transportation, and about 20% of that amount is dissipated in wall-bounded turbulent flows. Therefore, 5% of the total energy, and a disproportionate amount of the resulting CO₂ emissions, are spent that way [67]. Wall turbulence is the main contributor to aerodynamic friction, and is responsible for most of the pressure drop in internal flows such as pipelines. Its importance cannot be overemphasized for the energetic future of this century. Among wall bounded flows, channels are the simplest to simulate numerically and represent a valuable tool for the accurate study of these phenomena. Their direct numerical simulation has been a basic component of their study [7, 57, 85, 93, 119] for some time.

While the behaviour of turbulence in the viscous layer near the wall has been reasonably well understood [56, 74, 150] in the last decades, the dynamics of the flow farther from the wall is less clear. One of the most vexing remaining problems is the logarithmic layer that separates the near-wall and outer regions. It is in this layer that most of the velocity drop of a boundary layer takes place at high Reynolds numbers, and where the transition between the very different length-scales of turbulence in the two regions takes place. The logarithmic layer has been studied experimentally for a long time, but numerical data have only been available in the last decade because it only exists at relatively high Reynolds numbers [65]. Although different limits have been proposed it is safe to assume that the logarithmic layer extends from above $yu_\tau/\nu \approx 150$ to $y/h = 0.15$, where y is the distance from the wall, u_τ is the friction velocity, ν is the kinematic viscosity, and h is the boundary-layer thickness or the channel half-width. Therefore, a logarithmic layer does not exist if the friction Reynolds number is $Re_\tau = u_\tau h/\nu \lesssim 1000$, and it is only appreciably wide if Re_τ is substantially higher. The first direct numerical

simulation (DNS) of a channel flow at $Re_\tau = 180$ had no logarithmic layer in the sense just mentioned [85]. Since then, the Re_τ has increased steadily, and simulations with a short logarithmic layer have become available. The current state of the art is $Re_\tau \sim 4000\text{--}8000$ [18, 93, 103, 184], where the logarithmic layer extends over a scale ratio of 5 to 8 in the wall-normal direction.

Another open question in wall-bounded turbulence is the origin and dynamics of the large scales in the outer region, which extends above $y \sim 0.1h$. These scales play a relevant role in wall-bounded flows as they contain most of the total kinetic energy and Reynolds stresses of the flow, especially at high Reynolds numbers [5, 9, 62, 86, 97], but their study has been hindered by the difficulty to obtain quality data. In the case of channels, the flow between two infinite parallel walls is modelled by imposing periodic boundary conditions in the two wall-parallel dimensions, where the period represents the size of the numerical box. It was soon realized that the size of the box is an important parameter that conditions the size of the largest ‘well-resolved’ structures [9, 103]. The computational box has to be large enough for the structures in the outer region to be unconstrained, in the same way that the grid has to be fine enough to capture the smallest eddies and to reproduce the correct dissipation. A *de-facto* standard for large-box simulations is $(\text{length} \times \text{span}) = (8\pi \times 3\pi)h$ [65]. Experiments, and simulations in boxes up to $(60\pi \times 6\pi)h$ [103], have shown that this is a reasonable size to properly reproduce the dynamics of the largest structures that arise in channels.

Another reason for considering the large scales has recently become apparent. The buffer region near the wall has traditionally been considered the prime target for turbulence control, but the length scales and frequencies involved are demanding. In a wing at transonic speeds, where $u_\tau \approx 10 \text{ ms}^{-1}$ and $\nu/u_\tau \approx 10^{-6} \text{ m}$, the buffer-layer structures have wall-parallel dimensions $(500 \times 100)\nu/u_\tau \approx (0.5 \times 0.1) \text{ mm}$, and approximately 10^{10} sensors and actuators would be required to control the flow over a typical 300 m^2 wing. Their passing frequency is 20 KHz, which generates 2×10^{14} floating point data per second. Supposing, for simplicity, 5 operations per datum to obtain the desired control, one petaflop of computing power is required for this task. An alternative is to target the structures in the logarithmic layer, whose size is about $(0.3 \times 0.1)h$. For $h \approx 2 \text{ cm}$, the number of logarithmic-layer structures over a typical wing is a ‘more manageable’ 10^5 , with a passing frequency of 30 Hz. However, the possibility of controlling these larger scales depends on a theoretical understanding of their dynamics that is still missing.

The increasing availability of computational resources has provided DNSs of turbulent channel flow whose small and large scales are well resolved, but which focus mainly on producing high-quality snapshots and average statistical properties. In general, these simulations do not offer *a posteriori*

access to the dynamics of the flow and their practical utility is limited. A much more valuable insight into the dynamics of these flows is possible by also storing temporally resolved data. While all DNSs are by definition temporally resolved while they are being computed, the possibility of storing and post-processing temporally resolved time series, instead of a few independent snapshots, has become possible only recently. Since the early efforts in the 90s [141], the use of temporal series at higher Reynolds numbers has become a relevant tool of turbulent research, especially in wall-bounded flows [48, 103]. These temporal series, which are freely available to the community, marked an inflection point in turbulence research, because they allowed several groups to test hypotheses on the same data, and to refine them interactively. The limiting resources to produce these databases are storage space and post-processing time, as analysing temporally and spatially resolved time series of high-Reynolds numbers channels in large boxes is extremely challenging, specially when dynamics of the large scales are targeted. Part of the reason for the high cost of generating and storing useful temporal series of the large scales is that the simulation time has to be long enough to capture their relatively slowly temporal dynamics. In the logarithmic layer, the lifetime of structures centred at a distance y from the wall is $u_\tau T \approx 6y$ [48], or about an eddy-turnover time (ETT, $T_{eto} = h/u_\tau$) for structures at $y/h = 0.15$. Reasonable statistics require that the simulation time is $u_\tau t/h \gg 1$.

In the previous paragraphs, we have stressed the requirements of a turbulent database directed towards the study of the large scales of a turbulent channel flow: a high Reynolds number, an adequate box size, and storing temporally resolved and sufficiently long series. Reviewing the characteristics of the currently available databases, it is clear that none of them meets all these desirable properties, although some combinations of two of them can be found. For example, recent non-time-resolved simulations at high Reynolds numbers in large boxes are found in Lee and Moser [93]. Time-resolved series at moderate Reynolds numbers, $Re_\tau = 1000$ – 2000 , have been available for some time [40, 48, 103], but in small box sizes. A recent addition to the available databases is [126], which includes an $Re_\tau = 1000$ time-resolved channel in a large box, although only for a relatively short temporal period ($u_\tau t/h \approx 1.4$). While large boxes require larger computational resources than smaller ones at the same Reynolds number, both share the limiting factor of storing, sharing, and post-processing the computed data. For example, assuming a maximum data storage of 200 TB, and a target $Re_\tau = 4000$, the choice is between saving 500 snapshots of a large $(8\pi \times 3\pi)h$ box, or 6500 snapshots of a simulation in a smaller $(2\pi \times \pi)h$ box. The latter database would provide six ETT with a reasonable time interval between snapshots [103], while the former would only contain 0.5 ETT.

In this chapter we present a novel method to produce affordable, high quality, time-resolved databases of turbulent channel flow aimed at the study of the intermediate and large scales. This method, which can be easily extended to other flows, relies on two aspects. First, a dynamically meaningful reduction of the amount of information stored, which covers only large scales while retaining the effect of the small scales, and, second, the speed up of the computations by means of a controlled reduction of the numerical resolution, which only affects the small scales but not the scales of interest stored in the database. Further speed-up is achieved by implementing a new parallel GPU solver, which exploits the advantages of powerful heterogeneous architectures. This paper covers the computation and generation of a database for a state-of-the-art Reynolds number, $Re_\tau \sim 5000$, in a $(8\pi \times 3\pi)h$ box, providing time-resolved data for the unprecedented long time of $36T_{eto}$.

The storage limitations can be alleviated if one mostly wishes to study the dynamics of the large and intermediate motions, which, as mentioned above, are the ones that require large boxes. Although DNS simulations have to be properly resolved to reproduce the physics of turbulence, large and small scales are relatively independent, and do not need to be post-processed together. Recent simulations have shown that damping [59], or even removing [39, 116] the small scales near the wall, has negligible effects on the logarithmic-layer structures, and it has been known for some time that large-eddy simulations, which have no small scales, reproduce the large structures correctly [151]. Supported on these evidences, we store only accurate data obtained by *a posteriori* filtering of a DNS. This approach reduces the storage requirements by reducing both the spatial and temporal resolution of the database, which are set to capture the dynamics of the intermediate and large scales. The dynamics of the small scales are only preserved in the form of on-the-fly post-processed statistics. The novelty of our approach lies in retaining also their dynamical effect on the filtered data by storing also the filtered Reynold stresses. This approach allows a complete reconstruction of the dynamics of the three velocity components above the database filter. With this new approach, storing a time-resolved channel simulation in a large box becomes feasible, and opens new possibilities for studying the temporal evolution of accurately computed large-scale motions in wall-bounded flows. Discarding the dynamics of the small scales is justified on the availability of time-resolved simulations in boxes which, although smaller, are large enough to allow the study of these scales [48, 103, 126].

Since our aim is to capture the dynamics of the intermediate and large scales only, we require that the effect of the dissipative scales on these scales is well reproduced, but not that their dynamics is faithfully resolved to the level of high-quality DNSs. Considering that a relevant fraction of the total computational resources in DNSs is used to resolve the dissipative scales,

we obtain a substantial speedup in our simulations by adequately relaxing the standard resolution requirements while ensuring that the dynamics of the large scales are not affected. This approach has been widely used in homogeneous isotropic turbulence, for which simulations at high Reynolds numbers are produced by reducing the resolution of the small scales [61]. In this simulations, the average locality of scale interactions ensures that the effect of the insufficient resolution is confined to the small scales, but does not affect the inertial scales. This approach not only reduces the number of grid points, but also relaxes the stability conditions of the temporal integrator, providing a significant speedup in terms of computational time per ETT. We reduce the numerical resolution in the wall-parallel directions to reduce the computational requirements, and modify the distribution of points close to the wall to alleviate the stability requirements of the temporal integrator. We compensate this reduction of the numerical resolution by implementing spectral methods in the wall-parallel directions, and high-order, high-resolution compact finite differences in the wall normal direction. These modifications of the standard DNS requirements allow us to run for unprecedented long times, and do not affect the truncated data stored in the database. We present results that indicate that a reduction of the numerical resolution with respect to standard high-quality DNSs is acceptable for simulations that target the large scales of wall-bounded flows.

In this work we also validate the suitability of GPUs for large-scale direct numerical simulations of turbulent flows. Although GPUs have been part of many supercomputers for some time, their use for solving fluid mechanics problems is not general. Some previous codes have made use of GPUs to gain a considerable speed-up [77, 82], but they generally use low-resolution spatial discretisation schemes that make them less attractive for high-Reynolds number turbulence. Our results prove that the use of GPUs can also be advantageous for the computation of communication-intensive simulations such as those needed to compute turbulence with spectral and high-order compact schemes. For implementations that spend about 40% of the time transferring data, asynchronous overlapping between communications and computation can speed up simulations by a factor of two with respect to a standard CPU code running on the same platform. For large simulations spanning months, this advantage is crucial. As network bandwidth and clusters architecture improve, and the fraction of time spent communicating decreases, the advantage of using GPUs should be even clearer. When the next step in supercomputing is taken, and new exascale machines become available, efficient devices will be required to keep energy consumption down to a reasonable level. GPUs have proved to be efficient and fast and are present in ever more super-computer centres, as they deliver good performance while keeping energy consumption low. The

FIGURE A.1: Sketch of channel flow between two parallel walls

code used in this investigation takes advantage of the use of GPUs on a large scale, and represents a step into using future heterogeneous CPU/GPU exascale architectures for DNS simulations. The code publicly is available in https://github.com/albertovelam/CHANNEL_GPU.

This paper is organized as follows. The algorithm and code employed are described in §A.1, with details on the running performance and of the test simulations performed. Section A.2 describes the storage approach, the simulations performed for the database, the main statistics of those simulations, and their validation. Finally, conclusions are offered in §A.3

A.1 METHODS

A.1.1 Simulation algorithm

We consider the incompressible turbulent flow between two parallel planes separated a distance of $2h$. Both the spanwise, z , and the streamwise, x , directions are periodic, with period L_z and L_x , respectively. A constant mass flux is imposed in the streamwise direction by a time-variable mean pressure gradient. Due to incompressibility and impermeability at the wall, the plane-averaged velocity in the wall-normal direction, y , is zero at all distances from the wall. The mean value of the spanwise velocity component is allowed to drift, since no pressure gradient is applied along that direction, but is zero in the long-term average. A sketch of the flow is shown in figure ???. Magnitudes expressed in ‘wall’ units, which constructed from the friction velocity u_τ and the kinematic viscosity ν , are denoted by a ‘+’ superindex, so that $Re_\tau = h^+$. Upper case symbols are used for ensemble-average quantities, such as the mean velocity profile U , while primes are reserved for root-mean-square intensities, as in u' .

The evolution equations for the streamwise, wall-normal and spanwise velocities, $\mathbf{u} = (u, v, w)$, respectively, are the Navier–Stokes equations for an incompressible fluid,

$$\partial_t \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p + \nu \nabla^2 \mathbf{u}, \quad (\text{A.1a})$$

$$\nabla \cdot \mathbf{u} = 0, \quad (\text{A.1b})$$

where p is the kinematic pressure. For convenience, (A.1) are reduced to two equations for the vorticity in the wall-normal direction, ω_y , and for the Laplacian of the wall-normal velocity, Δv , thus removing the computation of the pressure [85].

The equations are then projected on a Fourier basis with N_x and N_z modes in each periodic direction. In real space, the mesh is uniformly spaced, with N_x and N_z collocation nodes chosen to provide the desired numerical resolution. Non-linear terms are computed with a pseudo-spectral method and are fully dealiased by zero padding [25]. In the wall-normal direction, the equations are discretised on a non-uniform mesh of N_y points, adjusted to keep, at all wall distances, an approximately constant resolution in terms of the local Kolmogorov scale [47]. Derivatives in that direction are computed directly on the mesh using seven-point compact finite differences with spectral-like resolution [94]. A semi-implicit third-order low-storage Runge–Kutta is used for temporal integration [158].

A.1.2 *Enforcing boundary conditions at high Reynolds numbers*

The use of high-order, high-resolution numerical approximations for the derivatives in the wall-normal direction is essential in turbulent channel DNSs. However these methods are prone to numerical errors if they are not developed together with the mesh. The most numerically sensitive operation in the vorticity–Laplacian formulation of the Navier–Stokes equations is the solution of the bilaplacian operator [85], which requires the solution of multiple homogeneous Helmholtz equations to impose the boundary conditions. For high Reynolds numbers and small wavelengths, these solutions include very thin boundary layers near the wall, and we found important numerical errors in the homogeneous solutions when compact finite differences of up to twelfth order of consistency are used for the discretisation of the second derivative.

Figure A.2 compares the worst-case solution for the Helmholtz problem used in the simulation (the one with the highest wavenumbers). The grid spacing of the original mesh, represented in the figure by circles, are optimized considering only physical arguments, as in [57]. This mesh worked well with low-order schemes, but the wide stencils used here led to large errors in the first few grid points, which, in extreme cases, were enough to make the simulation unstable. A refined mesh, plotted in squares, solved the problem of code instability by slightly displacing the first 20 grid points in order to minimize the error of the most dangerous solutions without affecting too much the individual grid spacing, thus preserving the physical properties of the mesh.

FIGURE A.2: Solution of the Helmholtz equation for the most critical wavenumber at $Re_\tau \approx 5000$. —, Analytic solution; ●, standard mesh; ■, optimised mesh.

A.1.3 Code

We implement a novel hybrid CUDA-MPI code to run the simulations on many distributed GPUs. While all numerical algorithms use exclusively the GPU using custom CUDA kernels, the CPU is used for node-node communications. Appropriate care is taken to optimize the CUDA kernels as much as possible. A plane-plane decomposition is used to partition the computational domain among GPUs. A configuration in yz planes is used to perform operations in the wall-normal direction, where compact finite differences require the use of banded-matrix solvers and global access to all points in that direction. For the computation of the non-linear convolution in the wall-parallel planes, the domain is reorganised in xz planes in order to apply direct and inverse Fourier transforms, using the highly optimized CUFFT library. This decomposition allows to implement the dealiasing by zero-padding locally in each GPU, reducing the amount of communications required to compute the convolution. All the transpose operations are performed using custom CUDA kernels to take advantage of the high memory bandwidth of the GPUs. Several strategies are used for the MPI communication part, from all-to-all to send/receive. Particular care is taken to overlap as much as possible communications and computations. This is done using asynchronous execution of the computation and the communications by using three GPU streams: one for computation and two for device-to-host and host-to-device communications; and a CPU stream for MPI communications.

The code shows good scaling up to 1024 GPUS, proving the suitability of GPUs for DNSs at high Reynolds number using high-resolution spectral methods. We also report a strong dependence of the scaling and performance on the balance between computation power, and network bandwidth.

A.2 A SECOND-ORDER CONSISTENT, TIME-RESOLVED DATABASE

A.2.1 Methodology

A time-resolved turbulent channel flow simulation at $Re_\tau \sim 5000$ spanning $36 Tu_\tau/h$ (eddy-turnover time) in a standard large box ($L_x = 8\pi h$ and $L_z = 3\pi h$) would produce 55 PB of data if the evolution of all the flow scales were retained [103, 126]. Information contained in that disproportionately large amount of data is used to represent the dynamics of viscous scales, which can be studied in much more affordable simulations in smaller boxes requiring much less computational effort.

We store only scales above the viscous limit, reducing in three orders of magnitude the storage required. This reduction comes from the lower spatial and temporal resolution necessary to accurately capture the dynamics of the large scales. The inertial range can be estimated to start at around 50η (η is the Kolmogorov length-scale [88]), which translates roughly into $100\nu/u_\tau$ (wall units) [65]. If an eddy of size l_e advected at velocity u_e (with respect to the mean velocity profile) is desirable to be time-tracked, the characteristic time between snapshots (Δt_e) needed to accurately represent its motion is of order of l_e/u_e . As we consider the eddies of size $l_e^+ \sim 100$ as the smallest we want to retain, we can expect these eddies to be advected with respect to the mean profile at approximately the root mean square of the velocity perturbations. The most restrictive condition arises from the stream-wise velocity perturbations, which for the logarithmic layer are $u'^+ \sim 2$. Thus, a reasonable estimate is $\Delta t_e^+ \approx 100/2 = 50$, which implies storing around $Re_\tau/\Delta t_e^+ \sim 100$ snapshots per turnover, $tu_\tau/h = 1$.

The velocity fields are filtered to retain only scales above $l_x^+ \sim 120$ in x and $l_z^+ \sim 60$ in z and then stored with the aforementioned temporal resolution. The difference between the characteristic lengths in x and z accounts for the elongated nature of the velocity perturbations in channel flows [57]. At that resolution the number of points in x and z needed in a $8\pi \times 3\pi$ box at $Re_\tau \sim 5000$ are $M_x = L_x Re_\tau / l_x^+ \sim 1000$ and $M_z = L_z Re_\tau / l_z^+ \sim 800$. The effect of the removed small scales into the large scales is not negligible, and impossible to recover from the filtered velocity, which would considerably limit what could be studied and computed from the database. To partially overcome this limitations, we also store the filtered Reynolds-stress tensor $u_i u_j$, which are the only second-order non-linear quantities that directly appear in the equations of motion. From the six filtered components of this tensor, the effect of small scale dynamics on the filtered velocities can be recovered, preserving dynamical information of the evolution of large scales, such as their temporal derivatives. Lastly, we also stored the filtered enstrophy, as is often used in turbulent research, and cannot be recovered from the filtered stress

tensor. Thus, the ten variables of the database are $\bar{\mathbf{u}}$, $\overline{\mathbf{u} \otimes \mathbf{u}}$, $\overline{\boldsymbol{\omega} \cdot \boldsymbol{\omega}}$, where $(\bar{\bullet})$ denotes Fourier cut-off filtering. Although the properties of this filter are far from optimal in real space [176], it commutes with homogeneous linear operators, allowing ‘exact’ *a posteriori* computations with the stored quantities. Some of the ‘derived’ variables that can be trivially computed from the database are the vorticities, $\bar{\boldsymbol{\omega}} = \nabla \times \bar{\mathbf{u}}$; or the kinetic energy, $\bar{E} = \frac{1}{2} \overline{\mathbf{u} \cdot \mathbf{u}}$. The computation of the pressure is less evident but equally possible. Taking the divergence of the momentum equation yields,

$$\nabla \cdot (\mathbf{u} \cdot \nabla \mathbf{u}) = -\Delta p, \quad (\text{A.2})$$

where continuity $\nabla \cdot \mathbf{u} = 0$ makes every other term banish. This is the well-known Poisson equation for the pressure [84], and is linear in p and $\mathbf{u} \otimes \mathbf{u}$,

$$\Delta \bar{p} = -\nabla \cdot [\nabla \cdot (\overline{\mathbf{u} \otimes \mathbf{u}})], \quad (\text{A.3})$$

which requires only the cross-derivatives of the filtered Reynolds stress tensor, which are part of the database. An identity worth noting is,

$$-\nabla \cdot [\nabla \cdot (\overline{\mathbf{u} \otimes \mathbf{u}})] = \overline{\mathbf{S} \cdot \mathbf{S}} - \overline{\boldsymbol{\omega} \cdot \boldsymbol{\omega}}, \quad (\text{A.4})$$

which gives a recipe to compute the norm of the rate-of-strain tensor $\mathbf{S}$ from the filtered enstrophy and the right hand side of the filtered pressure equation.

Finally, possessing a closed second-order database can be useful for *a priori* testing of LES models at high Reynolds numbers. The subgrid stress tensor can be computed from the filtered velocities as,

$$\boldsymbol{\tau} = \overline{\mathbf{u} \otimes \mathbf{u}} - \bar{\mathbf{u}} \otimes \bar{\mathbf{u}}, \quad (\text{A.5})$$

procuring a exact representation of what the model needs to account for, in a time-resolved fashion. Similarly, wall models for LES can be tested as the correct stresses for the filtered velocities are stored at every wall-normal location.

As stated before, our interest resides only on the larger scales of motion, making the resolution of DNS not a key factor in this simulation as long as the dynamics of the stored scales are unaffected. We will show that is possible to degrade the standard DNS resolution ($\Delta x^+ \approx 12$ and $\Delta z^+ \approx 9$) to ($\Delta x^+ \approx 24$ and $\Delta z^+ \approx 15$) producing a relative error of less than 2% in the retained scales, allowing in turn for running approximately four times more eddy turnovers with the same computational time.

A.2.2 Simulations

Table A.1 summarizes the characteristics of the simulations performed, as well as of those used for validation. The new simulations ran in PizDaint,

Name	Re_τ	$(L_x, L_z)/h$	$N_x \times N_y \times N_z$	$M_x \times M_y \times M_z$	$T^{(h)}$	N_{snap}
F2000	2000	$8\pi, 3\pi$	$2048 \times 512 \times 2048$	$512 \times 512 \times 512$	14	1300
F5300	5303	$8\pi, 3\pi$	$6144 \times 1024 \times 4096$	$1024 \times 1024 \times 1024$	36	3900
L2000	2003	$8\pi, 3\pi$	$4096 \times 633 \times 3072$	—	10	232
L5200	5200	$8\pi, 3\pi$	$10240 \times 1536 \times 7680$	—	7.8	—

TABLE A.1: Simulations performed (F2000 and F5300) and simulations used to validate them (L2000, from [57] and L5000, from [93]). N_d is the simulation resolution while M_d is the storing resolution. $T^{(h)} = Tu_\tau/h$ is the simulation time in eddy turnovers. N_{snap} is the number of snapshots available in the database.

a GPU-based supercomputer part of the PRACE Tier-0 network, for a total of 65000 node hours on 128 GPUs for F2000 and 1400000 node hours on 512 GPUs for F5300. The total size of the filtered databases is 5.6 TB and 100 TB respectively. To ensure equilibrium before collecting the data, the first 5 eddy turnovers run immediately after interpolating the initial condition were discarded. The initial conditions were interpolated from similar Reynolds-number simulations in equilibrium.

A.2.3 One-point statistics and energy spectra

Figure A.3 shows the main one-point statistics of the simulations in table A.1. The mean profiles agree with the ones of the full resolution simulations across all the channel, collapsing under the defect velocity form. When plotted in terms of the absolute velocity, they agree initially but depart later, by a constant shift in the value of the intercept constant. The logarithmic region can be expressed as

$$U^+ = \frac{1}{\kappa} \log(y^+) + C^+, \quad (\text{A.6})$$

where κ is the Kármán constant and C^+ is the intercept parameter. The values of both constants for the different simulations are computed by fitting (A.6) in the region $150\nu/u_\tau < y < 0.16h$ for $Re_\tau \approx 2000$ simulations and $350\nu/u_\tau < y < 0.16h$ for $Re_\tau \approx 5000$ simulations. The difference in the value of the intercept constant is similar to the effect of transitional roughness, where a roughness function accounts for the departure of the mean profile due to the wall roughness. The profile of F2000 is shifted respect to L2000 by $\Delta C^+ \approx -0.12$ and F5000 is shifted respect L5000 by $\Delta C^+ \approx -0.45$. The root mean square (rms) of perturbation velocities are plotted against their references in figures A.3(c, d), and their relative errors, $(u'_i - (u'_i)_{\text{ref}})/(u'_i)_{\text{ref}}$ are represented in figure A.3(e, f). They agree well except near the wall, where they depart from the reference simulations. The value of u' is equal to the reference

FIGURE A.3: Mean profiles and perturbation velocities rms for the simulations in table A.1. Symbols are triangles for the reference simulations and circles for the new ones. Filled symbols are $Re_\tau \approx 2000$ and open ones $Re_\tau \approx 5000$. (a) Mean velocity profile in velocity defect form for all the simulations. (b) Mean velocity profile for all the simulations. (c) Perturbation velocities rms for simulations at $Re_\tau \approx 2000$. (d) Perturbation velocities rms for simulations at $Re_\tau \approx 5000$. (e) Relative error in the rms of the velocity perturbations for simulations at $Re_\tau \approx 2000$ (f) Relative error in the rms of the velocity perturbations for simulations at $Re_\tau \approx 5000$. The vertical line in (e,f) is $y^+ = 150$.

simulations near the wall, whereas v' and w' have more energy than the references, and their relative error is maximum at this height. Moving away from the wall, the cross components converge to the reference value whereas u' lowers in intensity until the production peak of the buffer layer, $y^+ \approx 14$, where the relative error of the streamwise component is maximum. Farther away from the wall, at $x_2^+ > 100$, the deviation of either component is less than 2% (or within the statistical error for the F2000 case). From the larger error in the cross-velocity intensities near the wall is possible to trace the lower resolution of the grid as responsible for the difference in the intercept constant of the mean profile. Close to the wall, very small quasi-streamwise vortices populate the flow field, whose size is close to the resolution of the mesh in the reference simulations. Thus, these vortices are not well resolved in the new simulations, with lower resolution, as reflected by the error in the rms of the cross velocity components. This lack of resolution affects the friction at the wall, and the intercept constant changes slightly as if it were under the effect of very fine roughness. However, these phenomena do not affect the flow above the buffer layer, whose mean profiles collapse under the velocity-defect form with the reference simulations, and so do the intensity of the fluctuation velocities above $y^+ \approx 100$.

The same level of agreement is found in the spectra of the velocity components, which collapse with the reference simulations for the resolved scales. Figure A.4, gathers the spectra of both simulations and references at $y^+ \approx 150$. If we consider this wall-normal location the closest to the wall within the logarithmic region, the lower resolution of the database should be most critical here. The spectra of the new simulations agree extremely well with the references, over all the scales retained by the filter. As expected, the wall-normal velocity component is the worst resolved with $\approx 13\%$ of energy below the filter cutoff for F5300. The stream and spanwise velocities retain more than 98% and 93% of their energy respectively. As with the one point statistics, we have no indication of the lower resolution of the simulation affecting the turbulence scales stored in the database above the buffer layer.

A.2.4 Temporal statistics

Figure A.5(a) shows the temporal evolution of the instantaneous friction velocity u_τ in the bottom wall, across approximately 25 ETT. The largest variations of the value of u_τ have a slow characteristic time, while the changes that it experiences over times smaller than one ETT are smaller. The largest differences in the value of u_τ within time windows of one ETT are approximately 5 times smaller than those within 10 ETT. This implies that structures with lifetimes longer than one ETT are present in the flow, and contribute to the shear stress at the wall. Further insight can be gained from the premulti-

FIGURE A.4: Spectra of the velocity components at $x_2^+ = 150$ for the simulations in table A.1. (a-c) Spectra of u_i for L2000 (red shaded patches) and F2000 (black line). (d-f) Spectra of u_i for L5200 (red shaded patches) and F5300 (black line). The blue square marks the database resolution in all cases. Contours contain 90, 50 and 10% of the spectral mass of the reference simulations, and the same contour level is used for the new ones.

FIGURE A.5: (a) Time evolution of the friction velocity at the bottom wall. (b) Premultiplied temporal spectrum of the signal in (a). (c) Premultiplied (z, t) spectra for the first few streamwise Fourier modes of u at $y/h = 0.4$. From right to left, $\lambda_x = \infty$ (blue), $\lambda_x/h = 8\pi$ (red), $\lambda_x/h = 4\pi$ (yellow), $\lambda_x/h = 2\pi$ (purple), $\lambda_x/h = \pi$ (green). The contour is 40% of the maximum for each streamwise wavelength. F5000.

plied temporal spectra of this signal, shown in figure A.5(b), computed by smoothly windowing the temporal signal to prevent its lack of periodicity from harming the spectrum [179]. Approximately, 34% of the spectral mass, and thus of the skin friction fluctuations, is contained in periods longer than one ETT. Figure A.5(c) shows the z - t spectrum of several streamwise Fourier modes at $y/h = 0.4$, which is the distance from the wall where the streamwise velocity structures are longest [65]. The two longest streamwise modes have characteristic times longer than one eddy turnover, independently of their spanwise wavelength. Most of the energy of the infinitely long streamwise streaks, represented by the $k_x = 0$ wavenumber, is contained in lifetimes longer than 6 ETT. These structures have a typical width of $\lambda_z/h \approx 2$ –3, wider than the typical size of the outer-layer streaks [153]. The structures represented by the modes with $\lambda_x/h = 8\pi$ can be thought of as streamwise velocity streaks with length of approximately $12h$ and width $1.5h$, and have much shorter lifetimes, peaking at $Tu_\tau/h = 1$ –2. The temporal span of the database is sufficiently large to cover the dynamics of the large scales of the flow.

A.3 CONCLUSIONS

We have presented a new data base of turbulent channel flows containing time-resolved simulations at $Re_\tau = 2000$ and $Re_\tau = 5000$, for long times

$Tu_\tau/h \sim 30$, in large boxes of size $(8\pi \times 3\pi)h$. With a traditional resolution in time and space, storage size becomes the limiting factor. In the present case, it would require approximately 60 PB. To alleviate these requirements, and since the database is mainly intended to study the relatively large scales of the logarithmic layer, we only store structures filtered above the viscous limit, even if the simulation runs at full resolution. On the other hand, the data include the long histories required to capture the large-scale dynamics. To capture the influence of the discarded scales on the stored ones, the database also includes the instantaneous filtered Reynolds-stress tensor and the filtered enstrophy. This is still more economical than storing the full flow fields (~ 150 TB), and allows the reconstruction of other quantities of interest, such as the time derivatives of the filtered velocities, the filtered pressure, and the filtered norm of the rate-of-strain tensor. In this sense, the database is ‘second order consistent’, because any observable that depends on quadratic velocity products can be computed a posteriori. In particular, this property, together with relatively high Reynolds numbers, makes the database potentially useful for a-priori testing and development of new LES models.

In order to achieve the required running times, we employ a novel code which runs on many distributed GPUs and takes advantage of their highly efficient and powerful architectures. We show that this code scales up to 1024 GPUs, proving its suitability for the next generation of heterogeneous exascale computers.

To further reduce the cost of the simulations we slightly degrade the resolution of the computational grid, but we show that the first- and second-order, one- and two-point statistics of the retained scales are not affected by this resolution change, and that we generate ‘healthy’ logarithmic and outer regions. We show that this approach has potential applications in other DNSs that target the large scales.

From a preliminary analysis of the data, we show that time scales longer than one eddy-turnover time are present in the channel flow, especially for the streamwise velocity component. Some of these structures, responsible for about one third of the variation in the total turbulent skin friction, have characteristic times of $Tu_\tau/h \approx 6$. The same is true of the longest structures of u , but the long simulations make the study of their dynamics possible. A more detailed study of the temporal statistics will be conducted in future research.

ACKNOWLEDGMENTS

The new simulations ran in PizDaint under the PRACE Tier-0 project TREC. We are grateful to M. Fatica and E. Phillips, from Nvidia, for invaluable help in optimizing the performance of the MPI-CUDA code.

BIBLIOGRAPHY

- [1] **R. J. ADRIAN** 1994 Stochastic estimation of conditional structure: a review. *Appl. Sci. Res.* (53), pp. 291–303.
- [2] **R. J. ADRIAN** 2007 Hairpin vortex organization in wall turbulence. *Phys. Fluids* (19.4, 041301), p. 041301.
- [3] **R. J. ADRIAN, C. D. MEINHART AND C. D. TOMKINS** 2000 Vortex organization in the outer region of the turbulent boundary layer. *J. Fluid Mech.* (422), pp. 1–54.
- [4] **R. J. ADRIAN AND P. MOIN** 1988 Stochastic estimation of organized turbulent structure: homogeneous shear flow. *J. Fluid Mech.* (190), pp. 531–559.
- [5] **J. C. DEL ÁLAMO AND J. JIMÉNEZ** 2001 Direct numerical simulation of the very large anisotropic scales in a turbulent channel. *CTR Annual Research Briefs*. Stanford University, 2001, pp. 329–341.
- [6] **J. C. DEL ÁLAMO AND J. JIMÉNEZ** 2002 Characteristics of scalar dispersion in turbulent-channel flow. *Proc. Summer Prog.* Stanford University, 2002, pp. 203–214.
- [7] **J. C. DEL ÁLAMO AND J. JIMÉNEZ** 2003 Spectra of the very large anisotropic scales in turbulent channels. *Phys. Fluids A* (15), pp. L41–L44.
- [8] **J. C. DEL ÁLAMO AND J. JIMÉNEZ** 2009 Estimation of turbulent convection velocities and corrections to Taylor’s approximation. *J. Fluid Mech.* (640), pp. 5–26.
- [9] **J. C. DEL ÁLAMO, J. JIMÉNEZ, P. ZANDONADE AND R. D. MOSER** 2004 Scaling of the energy spectra in turbulent channels. *J. Fluid Mech.* (500), pp. 135–144.
- [10] **J. C. DEL ÁLAMO, J. JIMÉNEZ, P. ZANDONADE AND R. D. MOSER** 2006 Self-similar vortex clusters in the logarithmic region. *J. Fluid Mech.* (561), pp. 329–358.
- [11] **J. C. DEL ÁLAMO AND J. JIMÉNEZ** 2006 Linear energy amplification in turbulent channels. *J. Fluid Mech.* (559), pp. 205–213.
- [12] **M. AVILA, F. MELLIBOVSKY, N. ROLAND AND B. HOF** 2013 Streamwise-localized solutions at the onset of turbulence in pipe flow. *Phys. rev. lett.* (110.22), p. 224502.

- [13] W. J. BAARS, N. HUTCHINS AND I. MARUSIC 2016 Spectral stochastic estimation of high-Reynolds-number wall-bounded turbulence for a refined inner-outer interaction model. *Phys. Rev. Fluids* (1), p. 054406.
- [14] W. J. BAARS, N. HUTCHINS AND I. MARUSIC 2017 Reynolds number trend of hierarchies and scale interactions in turbulent boundary layers. *Philos. T. Roy. Soc. A* (375.2089), p. 20160077.
- [15] G. K. BATCHELOR The theory of homogeneous turbulence. Cambridge U. Press, 1953.
- [16] G. K. BATCHELOR AND A. A. TOWNSEND 1949 The nature of turbulent motion at large wave numbers. *Proc. R. Soc. London A* (199), pp. 513–520.
- [17] S. BENEDDINE, R. YEGAVIAN, D. SIPP AND B. LECLAIRE 2017 Unsteady flow dynamics reconstruction from mean flow and point sensors: an experimental study. *J. Fluid Mech.* (824), pp. 174–201.
- [18] M. BERNARDINI, S. PIROZZOLI AND P. ORLANDI 2014 Velocity statistics in turbulent channel flow up to $Re_\tau = 4000$. *J. Fluid Mech.* (742), pp. 171–191.
- [19] L. BIFERALE 2003 Shell models of energy cascade in turbulence. *Ann. Rev. Fluid Mech.* (35.1), pp. 441–468.
- [20] R. F. BLACKWELDER AND R. E. KAPLAN 1976 On the wall structure of the turbulent boundary layer. *J. Fluid Mech.* (76), pp. 89–112.
- [21] D. G. BOGARD AND W. G. TIEDERMAN 1986 Burst detection with single-point velocity measurements. *J. Fluid Mech.* (162), pp. 389–413.
- [22] J. BOUSSINESQ 1877 Essai sur la théorie des eaux courantes. *Mem. de L'Académie des Sciences* (23).
- [23] G. L. BROWN AND A. ROSHKO 1974 On density effects and large structure in turbulent mixing layers. *J. Fluid Mech.* (64), pp. 775–816.
- [24] K. M. BUTLER AND B. F. FARRELL 1992 Three-dimensional optimal perturbations in viscous shear flow. *Phys. Fluids A: Fluid Dynamics* (4.8), pp. 1637–1650.
- [25] C. CANUTO, M. Y. HUSSAINI, A. QUARTERONI AND T. A. ZANG Spectral Methods in Fluid Dynamics. Springer-Verlag, Heidelberg, 1988.
- [26] J. I. CARDESA, A. VELA-MARTÍN AND J. JIMÉNEZ 2017 The turbulent cascade in five dimensions. *Science* (357.6353), pp. 782–784.
- [27] P. CHAKRABORTY, S. BALACHANDAR AND R. J. ADRIAN 2005 Relationships between local vortex identification schemes. *J. Fluid Mech.* (). Aceptado.
- [28] H. CHOI, P. MOIN AND J. KIM 1994 Active turbulence control and drag reduction in wall-bounded flows. *J. Fluid Mech.* (262), pp. 75–110.

- [29] F. CHRISTIANSEN, P. CVITANOVIC AND V. PUTKARADZE 1997 Spatiotemporal chaos in terms of unstable recurrent patterns. *Nonlinearity* (10.1), p. 55.
- [30] N. C. CONSTANTINOU, B. F. FARRELL AND P. J. IOANNOU 2014 Emergence and equilibration of jets in beta-plane turbulence: applications of stochastic structural stability theory. *J. Atmos. Sci.* (71.5), pp. 1818–1842.
- [31] P. CVITANOVIĆ AND B. ECKHARDT 1989 Periodic-orbit quantization of chaotic systems. *Phys. rev. lett.* (63.8), p. 823.
- [32] H. DARCY Recherches expérimentales relatives au mouvement de l'eau dans les tuyaux. Mallet-Bachelier, Paris, 1857.
- [33] P. A. DAVIDSON Turbulence: an introduction for scientists and engineers. Oxford University Press, 2015.
- [34] P. A. DAVIDSON, Y. KANEDA, K. MOFFATT AND K. R. SREENIVASAN A voyage through turbulence. Cambridge University Press, 2011.
- [35] K. DEGUCHI AND P. HALL 2014 Canonical exact coherent structures embedded in high Reynolds number flows. *Philos. T. Roy. Soc. A* (372), p. 20130352.
- [36] D. J. C. DENNIS AND T. B. NICKELS 2011 Experimental measurement of large-scale three-dimensional structures in a turbulent boundary layer. Part 1. Vortex packets. *J. Fluid Mech.* (673), pp. 180–217.
- [37] E. R. DEYLE AND G. SUGIHARA 2011 Generalized theorems for nonlinear state space reconstruction. *PLoS One* (6.3), e18295.
- [38] S. DONG, A. LOZANO-DURÁN, A. SEKIMOTO AND J. JIMÉNEZ 2017 Coherent structures in statistically stationary homogeneous shear turbulence. *J. Fluid Mech.* (816), pp. 167–208.
- [39] S. DONG, A. LOZANO-DURÁN, A. SEKIMOTO AND J. JIMÉNEZ 2017 Coherent structures in statistically stationary homogeneous shear turbulence. *J. Fluid Mech.* (816), pp. 167–208.
- [40] UNIVERSIDAD POLITÉCNICA DE MADRID E.T.S.I.A Fluid Dynamics Lab. 2009.
- [41] J-P. ECKMANN AND D. RUELLE 1985 Ergodic theory of chaos and strange attractors. *The theory of chaotic attractors*. Springer, 1985, pp. 273–312.
- [42] M. P. ENCINAR, A. LOZANO-DURÁN AND J. JIMÉNEZ 2018 Reconstructing channel turbulence from wall observations. *Procs. CTR Summer School*. Stanford Univ. 2018, pp. 217–226.
- [43] M. FARGE 1992 Wavelet transforms and their applications to turbulence. *Annu. rev. fluid mech.* (24.1), pp. 395–458.

- [44] B. F. FARRELL AND P. J. IOANNOU 1993 Optimal excitation of three-dimensional perturbations in viscous constant shear flow. *Phys. Fluids A: Fluid Dynamics* (5.6), pp. 1390–1400.
- [45] B. F. FARRELL AND P. J. IOANNOU 1996 Turbulence suppression by active control. *Phys. Fluids* (8), pp. 1257–1268.
- [46] B. F. FARRELL AND P. J. IOANNOU 2012 Dynamics of streamwise rolls and streaks in turbulent wall-bounded shear flow. *J. Fluid Mech.* (708), pp. 149–196.
- [47] O. FLORES AND J. JIMÉNEZ 2006 Effect of wall-boundary disturbances on turbulent channel flows. *J. Fluid Mech.* (566), pp. 357–376.
- [48] O. FLORES AND J. JIMÉNEZ 2010 Hierarchy of minimal flow units in the logarithmic layer. *Phys. Fluids* (22.7, 071704), p. 071704.
- [49] U. FRISCH Turbulence. Cambridge U. Press, 1995.
- [50] U. FRISCH, P.L. SULEM AND M. NELKIN 1978 A simple dynamical model of intermittent fully developed turbulence. *J. Fluid Mech.* (87), pp. 719–736.
- [51] D. GABOR 1946 Theory of communication. Part 1: The analysis of information. *J. I. El. Telecom. Eng.* (93.26), pp. 429–441.
- [52] J. F. GIBSON, J. HALCROW AND CVITANOVIĆ P. 2009 Equilibrium and travelling-wave solutions of plane Couette flow. *J. Fluid Mech.* (638), pp. 243–266.
- [53] M. DE GIOVANETTI, H. J. SUNG AND Y. HWANG 2017 Streak instability in turbulent channel flow: the seeding mechanism of large-scale motions. *J. Fluid Mech.* (832), pp. 483–513.
- [54] G. HÄGEN 1839 Über den Bewegung des Wassers in engen cylindrischen Röhren. *Poggendorfs Ann. Physik Chemie* (16).
- [55] G. HÄGEN 1854 Über den Einfluss der Temperatur auf die Bewegung des Wassers in Röhren. *Math. Abh. Akad. Wiss Berlin* (17), pp. 17–98.
- [56] J. M. HAMILTON, J. KIM AND F. WALEFFE 1995 Regeneration mechanisms of near-wall turbulence structures. *J. Fluid Mech.* (287), pp. 317–348.
- [57] S. HOYAS AND J. JIMÉNEZ 2006 Scaling of the velocity fluctuations in turbulent channels up to $Re_\tau = 2003$. *Phys. Fluids* (18.1, 011702), p. 011702.
- [58] Y. HWANG AND C. COSSU 2010 Linear non-normal energy amplification of harmonic and stochastic forcing in the turbulent channel flow. *J. Fluid Mech.* (664), pp. 51–73.
- [59] Y. HWANG AND C. COSSU 2011 Self-sustained processes in the logarithmic layer of turbulent channel flows. *Phys. Fluids* (23), p. 061702.

- [60] S. J. ILLINGWORTH, J. P. MONTY AND I. MARUSIC 2018 Estimating large-scale structures in wall turbulence using linear models. *J. Fluid Mech.* (842), pp. 146–162.
- [61] T. ISHIHARA, T. GOTOH AND Y. KANEDA 2009 Study of high-Reynolds number isotropic turbulence by direct numerical simulation. *Annu. Rev. Fluid Mech.* (41), pp. 165–180.
- [62] J. JIMÉNEZ 1998 The largest scales of turbulence. *CTR Ann. Res. Briefs*. Stanford Univ. 1998, pp. 137–154.
- [63] J. JIMÉNEZ 2000 Turbulence. *Perspectives in Fluid Dynamics*. Ed. by G. K. Batchelor, H. K. Moffat and M. G. Worster. Cambridge U. Press, 2000, pp. 231–288.
- [64] J. JIMÉNEZ 2004 Turbulent flows over rough walls. *Annu. Rev. Fluid Mech.* (36), pp. 173–196.
- [65] J. JIMÉNEZ 2012 Cascades in wall-bounded turbulence. *Ann. Rev. Fluid Mech.* (44), pp. 27–45.
- [66] J. JIMÉNEZ 2013 How linear is wall-bounded turbulence? *Phys. Fluids* (25.11), p. 110814.
- [67] J. JIMÉNEZ 2013 Near-wall turbulence. *Phys. Fluids* (25.10), p. 101302.
- [68] J. JIMÉNEZ 2015 Direct detection of linearized bursts in turbulence. *Phys. Fluids* (27.6), p. 065102.
- [69] J. JIMÉNEZ 2018 Coherent structures in wall-bounded turbulence. *J. Fluid Mech.* (842).
- [70] J. JIMÉNEZ, J. C. DEL ÁLAMO AND O. FLORES 2004 The large-scale dynamics of near-wall turbulence. *J. Fluid Mech.* (505), pp. 179–199.
- [71] J. JIMÉNEZ AND S. HOYAS 2008 Turbulent fluctuations above the buffer layer of wall-bounded flows. *J. Fluid Mech.* (611), pp. 215–236.
- [72] J. JIMÉNEZ, S. HOYAS, M. P. SIMENS AND Y. MIZUNO 2010 Turbulent boundary layers and channels at moderate Reynolds numbers. *J. Fluid Mech.* (657), pp. 335–360.
- [73] J. JIMÉNEZ, G. KAWAHARA, M. P. SIMENS, M. NAGATA AND M. SHIBA 2005 Characterization of near-wall turbulence in terms of equilibrium and bursting solutions. *Phys. Fluids* (15.015105).
- [74] J. JIMÉNEZ AND P. MOIN 1991 The minimal flow unit in near-wall turbulence. *J. Fluid Mech.* (225), pp. 221–240.
- [75] J. JIMÉNEZ AND A. PINELLI 1999 The autonomous cycle of near-wall turbulence. *J. Fluid Mech.* (389), pp. 335–359.
- [76] J. JIMÉNEZ, A. A. WRAY, P. G. SAFFMAN AND R. RO GALLO 1993 The structure of intense vorticity in isotropic turbulence. *J. Fluid Mech.* (255), pp. 65–90.

- [77] K. I. KARANTASIS, E. D. POLYCHRONOPOULOS AND J. A. EKATERINARIS 2014 High order accurate simulation of compressible flows on GPU clusters over Software Distributed Shared Memory. *Comput. Fluids* (93.0), pp. 18–29.
- [78] T. VON KÁRMÁN 1930 Mechanische Ähnlichkeit und Turbulenz. *Proceedings Third Int. Congr. Applied Mechanics, Stockholm*. 1930, pp. 85–105.
- [79] T. VON KÁRMÁN AND L. HOWARTH 1938 On the statistical theory of isotropic turbulence. *Proc. Royal Soc. London (A 164)*, pp. 192–215.
- [80] G. KAWAHARA AND S. KIDA 2001 Periodic motion embedded in plane Couette turbulence: regeneration cycle and burst. *J. Fluid Mech.* (449), pp. 291–300.
- [81] S. KAWAI AND J. LARSSON 2012 Wall-modeling in large eddy simulation: Length scales, grid resolution, and accuracy. *Phys. Fluids* (24.1), p. 015105.
- [82] A. KHAJEH-SAEED AND J. B. PEROT 2013 Direct numerical simulation of turbulence using GPU accelerated supercomputers. *J. Comput. Phys.* (235.0), pp. 241–257.
- [83] H. T. KIM, S. J. KLINE AND W. C. REYNOLDS 1971 The production of turbulence near a smooth wall in a turbulent boundary layer. *J. Fluid Mech.* (47), pp. 133–160.
- [84] J. KIM 1989 On the structure of pressure fluctuations in simulated turbulent channel flow. *J. Fluid Mech.* (205), pp. 421–451.
- [85] J. KIM, P. MOIN AND R. D MOSER 1987 Turbulence statistics in fully developed channel flow at low Reynolds number. *J. Fluid Mech* (177), pp. 133–166.
- [86] K. KIM AND R. J. ADRIAN 1999 Very large-scale motion in the outer layer. *Phys. Fluids A* (11), pp. 417–422.
- [87] S. J. KLINE, W. C. REYNOLDS, F. A. SCHRAUB AND P. W. RUNSTADLER 1967 The structure of turbulent boundary layers. *J. Fluid Mech.* (30), pp. 741–773.
- [88] A. N. KOLMOGOROV 1941 Dissipation of energy in isotropic turbulence. *Dokl. Akad. Nauk. SSSR* (32), pp. 19–21.
- [89] A. N. KOLMOGOROV 1941 Local structure of turbulence in an incompressible fluid at very high Reynolds numbers. *Dokl. Akad. Nauk. SSSR* (30), pp. 9–13.
- [90] A. N. KOLMOGOROV 1962 A refinement of previous hypotheses concerning the local structure of turbulence in a viscous incompressible fluid at high Reynolds number. *J. Fluid Mech.* (13.1), pp. 82–85.

- [91] A. KRAVCHENKO, H. CHOI AND P. MOIN 1993 On the relation between near-wall streamwise vortices to wall skin friction in turbulent boundary layers. *Phys. Fluids A* (5), pp. 3007–3009.
- [92] L. D. LANDAU AND E. M. LIFSHITZ Fluid mechanics. Pergamon Press, 1944.
- [93] M. LEE AND R. D. MOSER 2015 Direct numerical simulation of turbulent channel flow up to $Re_\tau \sim 5200$. *J. Fluid Mech.* (774), pp. 395–415.
- [94] S. K. LELE 1992 Compact finite difference schemes with spectral-like resolution. *J. Comput. Phys.* (103.1), pp. 16–42.
- [95] M. LESIEUR Turbulence in fluids. Third Edition. Kluwer Academic Dordrecht, 1997.
- [96] T. LEUNG, N. SWAMINATHAN AND P. A. DAVIDSON 2012 Geometry and interaction of structures in homogeneous isotropic turbulence. *J. Fluid Mech.* (710), pp. 453–481.
- [97] Z. LIU, R. J. ADRIAN AND T. J. HANRATTY 2001 Large-scale modes of turbulent channel flow: transport and structure. *J. Fluid Mech.* (448), pp. 53–80.
- [98] E. N. LORENZ 1963 Deterministic nonperiodic flow. *J. Atmos. Sci.* (20), pp. 130–141.
- [99] A. LOZANO-DURÁN. Private Communication. 2019.
- [100] A. LOZANO-DURÁN AND H. J. BAE 2019 Characteristic scales of Townsend’s wall-attached eddies. *J. Fluid Mech.* (868), pp. 698–725.
- [101] A. LOZANO-DURÁN, O. FLORES AND J. JIMÉNEZ 2012 The three-dimensional structure of momentum transfer in turbulent channels. *J. Fluid Mech.* (694), pp. 100–130.
- [102] A. LOZANO-DURÁN AND J. JIMÉNEZ 2014 Effect of the computational domain on direct simulations of turbulent channels up to $Re_\tau = 4200$. *Phys. Fluids* (26.1), p. 011702.
- [103] A. LOZANO-DURÁN AND J. JIMÉNEZ 2014 Time-resolved evolution of coherent structures in turbulent channels: characterization of eddies and cascades. *J. Fluid Mech.* (759), pp. 432–471.
- [104] S. S. LU AND W. W. WILLMARTH 1973 Measurements of the structure of the Reynolds stress in a turbulent boundary layer. *J. Fluid Mech* (60), pp. 481–511.
- [105] J. L. LUMLEY 1967 The structure of inhomogeneous turbulent flows. *Atmospheric turbulence and radio wave propagation* ().

- [106] A. MADHUSUDANAN, S. J. ILLINGWORTH AND I. MARUSIC 2019 Coherent large-scale structures from the linearized Navier–Stokes equations. *J. Fluid Mech.* (873), pp. 89–109.
- [107] W. V. R. MALKUS 1956 Outline of a theory of turbulent shear flow. *J. Fluid Mech.* (1.05), pp. 521–539.
- [108] I. MARUSIC 2001 On the role of large-scale structures in wall turbulence. *Phys. Fluids* (13), pp. 735–743.
- [109] I. MARUSIC, J. P. MONTY, M. HULTMARK AND A. J. SMITS 2013 On the logarithmic region in wall turbulence. *J. Fluid Mech.* (716).
- [110] B. J. McKEON AND A. S. SHARMA 2010 A critical-layer framework for turbulent pipe flow. *J. Fluid Mech.* (658), pp. 336–382.
- [111] C. D. MEINHART AND R. J. ADRIAN 1995 On the existence of uniform momentum zones in a turbulent boundary layer. *Phys. Fluids* (7), pp. 694–696.
- [112] C. MENEVEAU 1991 Analysis of turbulence in the orthonormal wavelet representation. *J. Fluid Mech.* (232).
- [113] C. MENEVEAU AND K. R. SREENIVASAN 1991 The multifractal nature of energy dissipation. *J. Fluid Mech.* (224), pp. 429–484.
- [114] M. MILANO AND P. KOUMOUTSAKOS 2002 Neural network modeling for near wall turbulent flow. *J. Comput. Phys.* (182.1), pp. 1–26.
- [115] C. B. MILLIKAN 1938 A critical discussion of turbulent flows in channels and circular tubes. *Proceedings 5th Int. Congr. Applied Mechanics, New York*. Ed. by J. P. Den Hartog and H. Peters. Wiley, 1938, pp. 386–392.
- [116] Y. MIZUNO AND J. JIMÉNEZ 2013 Wall turbulence without walls. *J. Fluid Mech.* (723), pp. 429–455.
- [117] P. MOIN AND R. D. MOSER 1989 Characteristic-eddy decomposition of turbulence in a channel. *J. Fluid Mech.* (200), pp. 471–509.
- [118] F. MOISY AND J. JIMÉNEZ 2004 Geometry and clustering of intense structures in isotropic turbulence. *J. Fluid Mech.* (513), pp. 111–122.
- [119] R. D. MOSER, J. KIM AND N. N. MANSOUR 1999 Direct numerical simulation of turbulent channel flow up to $Re_\tau = 590$. *Phys. Fluids A* (11), pp. 943–945.
- [120] S. MOTOKI, G. KAWAHARA AND M. SHIMIZU 2018 Maximal heat transfer between two parallel plates. *J. Fluid Mech.* (851).
- [121] J. NIKURADSE 1932 Turbulente strömung in Nichtkreis förmigen Rohren. *Forshg. Arb. (Ing.–Wes)*, p. 356.
- [122] M. OBERLACK 2001 A unified approach for symmetries in plane parallel turbulent shear flows. *J. Fluid Mech.* (427), pp. 299–328.

- [123] L. ONSAGER 1945 The distribution of energy in turbulence. *Phys. Rev.* (68), p. 286.
- [124] W. M. ORR 1907 The stability or instability of the steady motions of a perfect liquid and of a viscous liquid. Part I: A perfect liquid. *Proc. R. Ir. Acad.* Vol. A27. 1907, pp. 9–68.
- [125] S. A. ORSZAG AND G. S. PATTERSON 1972 Numerical simulation of turbulence: statistical models and turbulence. *Lecture Notes in Physics* (12), pp. 127–147.
- [126] E. PERLMAN, R. BURNS, Y. LI AND C. MENEVEAU 2007 Data exploration of turbulence simulations using a database cluster. *Proc. SC07*. <http://turbulence.pha.jhu.edu>. ACM, New York, 2007, pp. 23.1–23.11.
- [127] A. E. PERRY AND M. S. CHONG 1982 On the mechanism of wall turbulence. *J. Fluid Mech.* (119), pp. 173–217.
- [128] A. E. PERRY, S. HENBEST AND M. S. CHONG 1986 A theoretical and experimental study of wall turbulence. *J. Fluid Mech.* (119), pp. 163–199.
- [129] S. B. POPE Turbulent flows. Cambridge U. Press, 2000.
- [130] L. PRANDTL 1925 Bericht Über die Entstehung der Turbulenz. *Z. Angew. Math. Mech.* (5), pp. 136–139.
- [131] C. CT. PRINGLE AND R. R. KERSWELL 2010 Using nonlinear transient growth to construct the minimal seed for shear flow turbulence. *Phys. rev. lett.* (105.15), p. 154502.
- [132] G. PUJALS, M. GARCÍA-VILLALBA, C. COSSU AND S. DEPARDON 2009 A note on optimal transient growth in turbulent channel flows. *Phys. Fluids* (21.1), p. 015109.
- [133] F. RAWLINS, S. P. BALLARD, K. J. BOVIS, A. M. CLAYTON, D. LI, G. W. INVERARITY, A. C. LORENC AND T. J. PAYNE 2007 The Met Office global four-dimensional variational data assimilation scheme. *Q. J. R. Meteorol. Soc.* (133), pp. 347–362.
- [134] L. RAYLEIGH 1892 VIII. On the question of the stability of the flow of fluids. *Lond. Edinb. Dubl. Phil. Mag.* (34.206), pp. 59–70.
- [135] H. REICHARDT 1951 Vollständige Darstellung der Turbulenten Geschwindigkeitsverteilung in glatten Leitungen. *ZAMM* (31), pp. 208–219.
- [136] O. REYNOLDS 1883 An experimental investigation of the circumstances which determine whether the motion of water should be direct or sinuous, and the law of resistance in parallel channels. *Philos. Trans. R. Soc. London Ser. A* (174), pp. 935–982.

- [137] O. REYNOLDS 1894 On the Dynamical Theory of Incompressible Viscous Fluids and the Determination of the Criterion. [Abstract]. English. *P. R. Soc. London* (56), pp. 40–45.
- [138] W. C. REYNOLDS AND W. G. TIEDERMAN 1967 Stability of turbulent channel flow, with application to Malkus' theory. *J. Fluid Mech.* (27), pp. 253–272.
- [139] L. F. RICHARDSON 1920 The Supply of Energy from and to Atmospheric Eddies. English. *P. R. Soc. Lond. A-Conta.* (97.686), pp. 354–373.
- [140] L. F. RICHARDSON Weather prediction by numerical process. Cambridge U. Press, 1922.
- [141] S. K. ROBINSON 1991 Coherent motions in the turbulent boundary layer. *Ann. Rev. Fluid Mech.* (23), pp. 601–639.
- [142] R. ROGALLO 1981 Numerical experiments in homogeneous turbulence. ().
- [143] C. W. ROWLEY 2005 Model reduction for fluids, using balanced proper orthogonal decomposition. *Int. J. Bifurcat. Chaos* (15.03), pp. 997–1013.
- [144] D. RUELLE AND F. TAKENS 1971 On the nature of turbulence. *Commun. math. phys.* (20.3), pp. 167–192.
- [145] K. SASAKI, R. VINUESA, A. V. G. CAVALIERI, P. SCHLATTER AND D. S. HENNINGSON 2019 Transfer functions for flow predictions in wall-bounded turbulence. *J. Fluid Mech.* (864), pp. 708–745.
- [146] P. SCHLATTER AND R. ÖRLÜ 2010 Assessment of direct numerical simulation data of turbulent boundary layers. *J. Fluid Mech.* (659), pp. 116–126.
- [147] P. J. SCHMID 2010 Dynamic mode decomposition of numerical and experimental data. *J. Fluid Mech.* (656), pp. 5–28.
- [148] P. J. SCHMID AND D. S. HENNINGSON Stability and transition in shear flows. Springer, New York, 2001.
- [149] T. M. SCHNEIDER, J. F. GIBSON AND J. BURKE 2010 Snakes and ladders: localized solutions of plane Couette flow. *Phys. rev. lett.* (104.10), p. 104501.
- [150] W. SCHOPPA AND F. HUSSAIN 2002 Coherent structure generation in near-wall turbulence. *J. Fluid Mech.* (453), pp. 57–108.
- [151] G. SCOVAZZI, J. JIMÉNEZ AND P. MOIN 2001 LES of the very large scales in a $Re_\tau = 920$ channel. *Proc. Div. Fluid Dyn.* Am. Phys. Soc., 2001, KF–5.

- [152] J. A. SILLERO, J. JIMÉNEZ AND R. D. MOSER 2013 One-point statistics for turbulent wall-bounded flows at Reynolds numbers up to $\delta^+ \approx 2000$. *Phys. Fluids* (25.10, 105102), p. 105102.
- [153] J. SILLERO, J. JIMÉNEZ AND R. D. MOSER 2014 Two-point statistics for turbulent boundary layers and channels at Reynolds numbers up to $\delta^+ \sim 2000$. *Phys. Fluids* (26.10), pp. 105–109.
- [154] A. SOMMERFELD 1907 Ein Beitrag zur hydrodynamischen Erklärung der turbulenten Flüssigkeitsbewegungen. *Proc. 4th Internat. Cong. Math. Rome*. Vol. 3. 1907, pp. 116–124.
- [155] E. D. SONTAG Mathematical control theory. Springer, 1998.
- [156] P. R. SPALART 1987 Direct simulation of a turbulent boundary layer up to $Re_\theta = 1410$. *J. Fluid Mech.* (187), pp. 61–98.
- [157] P. R. SPALART 2009 Detached-eddy simulation. *Annu. Rev. Fluid Mech.* (41), pp. 181–202.
- [158] P. R. SPALART, R. D. MOSER AND M. M. ROGERS 1991 Spectral methods for the Navier-Stokes equations with one infinite and two periodic directions. *J. Comp. Phys.* (96), pp. 297–324.
- [159] H. B. SQUIRE 1933 On the stability for three-dimensional disturbances of viscous fluid flow between parallel walls. *Proc. R. Soc. Lond. A* (142.847), pp. 621–628.
- [160] K. R. SREENIVASAN 1985 On the fine-scale intermittency of turbulence. *J. Fluid Mech.* (151), pp. 81–103.
- [161] D. D. STRETCH 1990 Automated pattern eduction from turbulent flow diagnostics. *CTR Ann. Res. Briefs*. Stanford Univ. 1990, pp. 145–157.
- [162] F. TAKENS 1981 Detecting strange attractors in turbulence. *Dynamical systems and turbulence, Warwick 1980*. Springer, 1981, pp. 366–381.
- [163] G. I. TAYLOR 1921 Diffusion by continuous movements. *Proc. London Math. Soc.* (20), pp. 196–212.
- [164] G. I. TAYLOR 1935 Statistical theory of turbulence. Parts I to IV. *Proc. Royal Soc. London (A)* 151, pp. 421–478.
- [165] G. I. TAYLOR 1938 The spectrum of turbulence. *Proc. R. Soc. Lond.* (164), pp. 476–490.
- [166] H. TENNEKES AND J. L. LUMLEY A first course on turbulence. MIT Press, 1972.
- [167] T. THEODORSEN 1952 Mechanism of turbulence. *Proc. 2nd Midwestern Conf. of Fluid Mechanics, Ohio State University*. ()

- [168] C. TINNEY, F. COIFFET, J. DELVILLE, A. HALL, P. JORDAN AND M. GLAUSER 2006 On spectral linear stochastic estimation. *Exp. Fluids* (41.5), pp. 763–775.
- [169] C. D. TOMKINS AND R. J. ADRIAN 2003 Spanwise structure and scale growth in turbulent boundary layers. *J. Fluid Mech.* (490), pp. 37–74.
- [170] A. A. TOWNSEND 1961 Equilibrium layers and wall turbulence. *J. Fluid Mech.* (11), pp. 97–120.
- [171] A. A. TOWNSEND The structure of turbulent shear flows. Second. Cambridge U. Press, 1976.
- [172] A. TSINOBER An informal introduction to turbulence. Vol. 63. Springer Science & Business Media, 2001.
- [173] A. TSINOBER Essence of Turbulence as a Physical Phenomenon: With Emphasis on Issues of Paradigmatic Nature. 3rd edition. Springer, 2019.
- [174] A. VELA-MARTÍN, M. P. ENCINAR, A. GARCÍA-GUTIÉRREZ AND J. JIMÉNEZ 2018 A second-order consistent, low-storage method for time-resolved channel flow simulations. 2018.
- [175] C. SANMIGUEL VILA AND O. FLORES 2018 Wall-based identification of coherent structures in wall-bounded turbulence. *J. Phys. Conf. Ser.* (1001), p. 012007.
- [176] B. VREMAN, B. GEURTS AND H. KUERTEN 1994 Realizability conditions for the turbulent stress tensor in large-eddy simulation. *J. Fluid Mech.* (278), pp. 351–362.
- [177] F. WALEFFE 1997 On a self-sustaining process in shear flows. *Phys. Fluids* (9.4), pp. 883–900.
- [178] J. M. WALLACE, H. ECKELMANN AND R. S. BRODKEY 1972 The wall region in turbulent shear flow. *J. Fluid Mech.* (54), pp. 39–48.
- [179] P. WELCH 1967 The use of fast Fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms. *IEEE Trans. Audio Electroacoust.* (15), pp. 70–73.
- [180] W. W. WILLMARTH AND S. S. LU 1972 Structure of the Reynolds stress near the wall. *J. Fluid Mech.* (55), pp. 65–92.
- [181] J. WILLS 1964 On convection velocities in turbulent shear flows. *J. Fluid Mech.* (20), pp. 417–432.
- [182] M. WOSNIK, L. CASTILLO AND W. K. GEORGE 2000 A theory for turbulent pipe and channel flows. *J. Fluid Mech.* (421), pp. 115–145.
- [183] X. WU AND P. MOIN 2009 Direct numerical simulation of turbulence in a nominally zero-pressure-gradient flat-plate boundary layer. *J. Fluid Mech.* (630), pp. 5–41.

- [184] Y. YAMAMOTO AND Y. TSUJI 2018 Numerical evidence of logarithmic regions in channel flow at $Re_\tau = 8000$. *Phys. Rev. Fluids* (3), p. 012602.
- [185] M. YOKOKAWA, K. ITAKURA, T. ISHIHARA AND Y. KANEDA 2002 16.4 Tflops Direct Numerical Simulation of Turbulence by a Fourier Spectral Method on the Earth Simulator. *IEEE/ACM SC2002 Conf.* 2002.
- [186] J. ZHOU, R. J. ADRIAN, S. BALACHANDAR AND T. M. KENDALL 1999 Mechanisms for generating coherent packets of hairpin vortices in channel flow. *J. Fluid Mech.* (387), pp. 353–396.