

The
University
Of
Sheffield.

Numerical study of drag reduction through rotating actuators in a wall-bounded turbulent flow

Paolo Olivucci

A thesis submitted in partial fulfilment of the requirements for the degree of
Doctor of Philosophy

The University of Sheffield
Faculty of Engineering
Department of Mechanical Engineering

Submission Date
September 2020

Statement of Originality

I, the author, confirm that the Thesis is my own work. I am aware of the University's Guidance on the Use of Unfair Means (<https://www.sheffield.ac.uk/ssid/unfair-means>). This work has not been previously presented for an award at this, or any other, university.

Part of Chapter 3 has been published as the journal article [Olivucci et al. \(2019\)](#). Work developed during the preparation of this thesis has led to a correction being issued for the article [Wise et al. \(2018\)](#).

Acknowledgements

I am grateful to my supervisor Dr. Pierre Ricco for giving me the opportunity to pursue this degree, for his supervision and his feedback, especially on the importance of rigorous scientific writing and presentation. I also thank my second supervisor Dr. Franck Nicolleau for many insightful discussions.

I am indebted to Mrs. Julia Boler, the Gordon Franklin Scholarship and the Department of Mechanical Engineering at the University of Sheffield for providing financial support.

The EPSRC Grant EP/L000261/1 and the UK Turbulence Consortium provided computer time on the Archer supercomputer for the work of Chapter 3 and 5. I acknowledge Physical Review Fluids' referees for the useful comments to the published part of Chapter 3 and to Prof. Martin Skote for reading and commenting on a preliminary version of the manuscript.

Dr. Markus Rütten of the Deutsches Zentrum für Luft- und Raumfahrt (DLR) in Göttingen provided the technical specifications of his disc-fitted boundary-layer experimental setup. Prof. Rob Dwyer-Joyce (University of Sheffield) shared his precious expertise for the discussion of the bearing torque in Chapter 4.

I would like to thank Dr. Javier González, Andrei Paleyes and the EmuKit team at Amazon Research for their invaluable guidance in adapting their code for Chapter 5.

I want to express my personal gratitude to the examiners Prof. Ning Qin of Sheffield and Prof. Kunihiko Taira of UCLA for the thought-provoking and stimulating discussion that significantly improved the quality of the work; especially the precious suggestions on how to improve the dissertation cohesion, emphasise rigorous justification and for encouraging a strategic, broader perspective on my research.

Abstract

The reduction of frictional drag generated by turbulent fluid flow on solid surfaces is an important problem in fluid mechanics, with the long-term potential to impact the energy efficiency and the performance of many engineering systems. This thesis consists of three essays that elaborate on turbulent drag reduction through wall-mounted rotating discs, originally introduced by [Ricco and Hahn \(2013\)](#). The proposed methods are evaluated via direct numerical simulation of turbulent channel flow using the high-performance solver *Incompact3D*.

In the first, turbulent channel flow altered by the combination of flush-mounted spinning rings and vertical-velocity opposition control or hydrophobic surfaces is studied. The two types of distributed control are applied over the surface area that is not occupied by the spinning rings. Drag reduction is enhanced by the novel combined methods, with up to 27% reduction compared to 20% of the simple rotating rings. A idealized predictive model of the combined drag-reduction performance is presented. A spatially-dependent variant of the Fukagata-Iwamoto-Kasagi integral identity is developed to explain the influence of the highly non-uniform spatial structure of the flow on the skin-friction. The streamwise structures forming between discs have a drag-increasing contribution, while drag is highly reduced over the central region where the rings generate a triangular wave of spanwise velocity. The different outcome of combining the rings with the opposition control and the hydrophobic surface are clarified by the alteration to the elongated structures between rings.

The second part investigates the dynamics of rigid discs that are free to rotate under a turbulent channel flow. Simple yet realistic models are introduced for the actuator geometry and the frictional torques that are generated by fluid motion in the cavity underneath the disc and in the support bearing, and the related assumptions are discussed critically. Firstly, the dynamics of isolated, passively rotating discs of increasing diameter is simulated, finding that the discs oscillate around rest under the action of the shear-stress fluctuations at the wall. The root-mean-square of the turbulence-induced torque and the disc velocity show power-law dependence on the disc diameter, respectively positive and negative in the large-diameter range. A uncoupled model of the disc dynamics, where the disc do not move, is used as a small-velocity approximation of the disc-fluid system, leading to better understanding of the frequency-domain response of the disc to the turbulent excitation and the root-mean-square dependence on the disc diameter. Furthermore, half discs i.e. semicircular free-to-move actuators aligned to the streamwise direction and arranged in a rectangular array are considered. This configuration simulates the concept developed experimentally by [Koch and Kozulovic \(2013\)](#) of partially covered discs. The mean-shear generates a non-zero torque on the disc half and the discs rotate with a finite mean angular velocity, locally producing a slip velocity at the wall. Drag reduction is found to be directly correlated to the rotation rate of the discs, the fastest-rotating case yielding above 5% globally and 20% on the disc surface. Dependence of the rotation velocity on the disc model parameter is discussed and a series of simulations aimed at a more comprehensive investigation of the half-disc technique is laid out.

The third part of the work proposes a simple proportional feedback-control scheme for the array of discs, whereby the motor-powered discs are switched on or off to maintain the discs within a pre-determined velocity range. The actuator model developed previously is augmented with the inclusion of a model of a typical electric motor, used to power the discs. The optimal values of the control parameters i.e. the velocity range bounds and the motor torque are investigated by means of a Bayesian Optimization algorithm, capable of efficiently performing an autonomous search of the three-dimensional parameter space with a minimal number of simulations. The motorised steadily-rotating case is also simulated and its performance compared to the ideal model and to the feedback-controlled flow, concluding that it is highly likely that net power saving can be achieved even accounting for realistic losses in the disc housing. The power-budget dependence on the feedback-control parameters is analysed extensively and special consideration is given to its correlation to the disc mean velocity and to the characteristics of the transients. Although the on-off control can achieve marginal improvements on the reference steady activation of the discs, an optimised constant-velocity control is superior.

Contents

1	Wall-bounded turbulent flow and drag reduction	1
1.1	Introduction	1
1.2	Frictional drag reduction	2
1.2.1	Flow control for frictional drag reduction	2
1.2.2	Turbulent drag reduction by in-plane wall motions	4
1.3	Turbulent flow in an infinitely-extended channel	7
1.3.1	Statistical description of turbulence	9
1.3.2	Basic results in channel flow	11
1.3.3	Skin-friction and drag	13
2	Direct Numerical Simulation of turbulent channel flow	16
2.1	Direct Numerical Simulation of turbulent flow	16
2.1.1	The Incompact3D solver	17
2.1.2	High-performance computing	23
2.2	Reference channel flow	28
2.2.1	Numerical set-up	28
2.2.2	Computational performance	30
2.2.3	Verification and validation of computational simulations	30
2.2.4	Results	31
2.3	Statistical methods	35
2.3.1	Statistical estimation	36
2.3.2	Uncertainty quantification	37
2.3.3	Post-processing software	40
3	Open-loop turbulent drag reduction combining rotating actuators and other techniques	41
3.1	Introduction	41
3.2	Numerical and statistical procedures	44
3.2.1	Numerical simulations	44
3.2.2	Averaging operators	46
3.2.3	Flow symmetries	48
3.2.4	Flow decomposition	49
3.2.5	Governing equations for the three-component flow	49
3.2.6	Definition of turbulent drag reduction	50
3.2.7	Uncertainty analysis	50
3.3	Results	51
3.3.1	Ring flow	51

3.3.2	Turbulence statistics	55
3.3.3	Turbulent drag reduction	60
3.3.4	Ideal prediction of combined-forcing drag reduction	60
3.3.5	Spatial distribution of drag reduction	62
3.3.6	Decomposition through the Fukagata-Iwamoto-Kasagi identity	64
3.3.7	Details of the skin-friction budget terms	69
3.3.8	Detection height modification for case ROC	70
3.4	Summary and conclusions	71
3.5	Further work	73
3.5.1	Dependence of \mathcal{R} on the control parameters	73
3.5.2	Ring-like and wave-like boundary conditions	74
3.5.3	Disc edge sheltering	76
3.5.4	Higher-Reynolds flow	76
3.A	Identification of coherent vortical structures	79
4	Passively moving discs for frictional drag reduction	80
4.1	Introduction	80
4.2	Numerical and statistical procedures	81
4.2.1	Numerical solver	81
4.2.2	Reference channel flow and non-dimensional units	82
4.2.3	Averaging operators and performance quantities	83
4.3	Description and modelling of the actuators	84
4.3.1	Dynamics of the rotating discs	85
4.3.2	Torque from fluid shear stresses	85
4.3.3	Resisting torques from disc housing	85
4.3.4	Frictional torque of a ball bearing	87
4.3.5	Flow parameters and dimensional reference	88
4.4	Freely rotating discs	88
4.4.1	Description of the disc-fluid system	89
4.4.2	Numerical simulation of coupled discs-fluid dynamics	90
4.4.3	Uncoupled dynamics	96
4.4.4	Frequency-domain analysis	98
4.5	Freely rotating half discs	105
4.5.1	Review of the experimental results	106
4.5.2	Parameters of the channel flow with half discs	107
4.5.3	Numerical cases	108
4.5.4	Simulation results	110
4.5.5	Experiments and simulations	115
4.6	Summary and conclusions	117
4.7	Future research on half-discs	118
4.A	Details of the reduced-order model from Sec. 4.4.4	120
5	Feedback control of rotating discs	121
5.1	Introduction	121
5.2	Modelling and simulation of the actuators	121
5.2.1	Averaging operators	122
5.2.2	Model and dynamics of the rotating discs	123
5.2.3	Motor torque and active forcing	124
5.2.4	Dimensional reference	125

5.3	On-off control	125
5.3.1	Power-budget analysis	126
5.3.2	Transient simulations	129
5.3.3	Optimisation of the total spent power	131
5.4	Bayesian Optimization	131
5.4.1	Application to a closely related flow control problem	134
5.5	Results and data analysis	135
5.5.1	Steady forcing and optimisation	147
5.5.2	Power gain and power gain optimisation	150
5.6	Summary and conclusions	152
5.A	Bayesian Optimization	155
5.B	Exact model of the transients	158
6	Conclusions and future directions	160
6.1	General conclusions	160
6.2	Future outlook	161
	Bibliography	164

Chapter 1

Wall-bounded turbulent flow and drag reduction

1.1 Introduction

The interaction of the turbulent motion of a liquid or a gas with a solid surface is an ubiquitous natural phenomenon of great scientific and technological importance for many areas of engineering and the applied sciences. The unsteady and irregular character of turbulent motion distinguishes turbulent flow from the more orderly and predictable laminar flow, which is encountered for smaller scales and velocities than those prevalent in the majority of mechanical, aeronautical, civil, and chemical engineering systems. The resistance force generated by the fluid flow at the solid surface, called skin-friction drag, is naturally higher in turbulent flow than in laminar flow. One crucial technological objective has been the reduction of turbulent skin-friction drag, with the aim of diminishing the dissipation of energy and thus improve performance or limit fuel consumption of many engineering systems.

From the perspective of understanding (and potentially manipulating) the wall-bounded flow physics, it is convenient to concentrate on idealised problems, among which are parallel flows, namely plane channel and circular pipe, and turbulent boundary layers. This work concentrates on channel flow, that consist in the unidirectional motion of a fluid in the gap between two infinitely extended parallel flat plates and is convenient to study for its relative simplicity. The universal nature of the flow physics near the wall ensures the generality of the results obtained in channel flow, that are applicable to other canonical wall-bounded flows (pipe flow and boundary layer) and to more complex cases that can be locally approximated by one of the above. This also motivates the use of plane channel flow for the development and testing of control techniques that target the near-wall region.

Systematic research on the peculiar nature of turbulent flow and its implications for flow resistance started with the famous experiments of [Reynolds \(1883\)](#) who studied the appearance of irregular fluctuations in the flow through otherwise smooth and straight channels. Beforehand, flow resistance in pipes and channel had already been the subject of empirical investigation for its importance to practical problems in hydraulics. The study of turbulent flow has remained for a long time largely experimental, reflecting its complexity and dependence on the characteristics of each particular problem. Starting from the second half of the 20th century, the availability of increasing computational power and the progress in numerical methods has allowed fluid flows to be simulated from first principles ([Harlow, 2004](#)). Since around the 1980s, computers become powerful enough to tackle the resolution of the fine-scale fluctuations that characterise turbulence,

thus laying the foundation for many developments in the understanding of the physics and the design of methodologies to control turbulence and diminish drag. Numerical simulations provide a cheaper and more flexible alternative to experiments, with the additional advantage that a much more detailed measurement and manipulation of flow data is possible. Experiments are to this day unavoidable for flows whose complexity lays beyond current computers' capabilities and are necessary to validate numerical models.

In Chapter 1 and Chapter 2 we will lay out more precise definitions of the ideas outlined above, reviewing wall-bounded flow physics, its numerical simulation through high-performance computing systems and the methodologies devised to reduce turbulent friction.

The chapter is structured as follows. Section 1.2 contains a review of the existing turbulent drag reduction methods, focusing on in-plane wall motion techniques that form the basis of the research work developed in the following chapters. In Section 1.3 we define the plane channel flow problem, we define its mathematical modelling, introduce the physics of turbulent flow, its statistical description and some well-known results in wall turbulence. The basic definitions of skin-friction and drag reduction are also given.

1.2 Frictional drag reduction

1.2.1 Flow control for frictional drag reduction

Turbulent skin-friction drag reduction is a subject of great interest in fluid mechanics research, for its potential to limit fuel consumption and carbon emissions in the airline and marine transport industries, and to tackle the pressing problem of environmental pollution. Around 60% of the aerodynamic drag of a typical airliner at cruise conditions is due to frictional drag, with the other components being form drag and induced drag (Leschziner et al., 2011). In the case of a typical slow merchant ship, such as a tanker or a container ship, around 70% of the total hull resistance is frictional drag, the remainder being form drag and wave-making resistance (Larsson and Raven, 2010). For the aerospace industry, it has been estimated that a 1% reduction in skin-friction drag for a long-range commercial aircraft would decrease fuel consumption by approximately 0.45%. For a typical airline this is equivalent to a reduction in CO₂ emissions of 5.4 million equivalent tonnes per year (Leschziner et al., 2011), and an annual reduction in operating costs of 0.2% (Reneaux, 2004). In the following, we give an overview of the drag-reduction techniques that exist in the literature. Turbulent skin-friction is usually studied in channel flow, pipe flow or in boundary layers, as in these flow skin-friction is the only source of drag, allowing a better understanding of its mechanism and its response to control. More rigorous formulation of the ideas and the quantities discussed below is given starting from Section 1.3, where turbulent channel flow is introduced.

a. Classification of drag-reduction methods

Wall-based flow control methods for drag reduction can be broadly classified as being either active or passive. The graph in Fig. 1.1 provides a schematic classification of wall-forcing methods for drag reduction.

In passive methods, the flow is manipulated through macroscopic or microscopic modifications to the wall surface geometry, or via compliant surfaces that consist the wall being partially free to move under the action of the flow. Passive methods do not necessitate an external power input, although the interaction between the new geometry and the fluid may constitute an additional source of friction.

Active methods are those requiring an external energy input to manipulate the flow, which is often altered by devices that rely on an external power supply to operate. For such methods, the net power saving $\mathcal{P}_{net} = 100(P_0 - P_{tot})/P_0$ represents a critical performance metric, because the total power $P_{tot} = P_p + P_{in}$ accounts for the input power P_{in} required to drive the control devices, beside the flow pumping power P_p . Another performance metrics for active methods is the gain $\mathcal{G} = (P_0 - P_p)/P_{in}$ (Quadrio, 2011), which represents the efficiency of the control scheme to obtain pumping power saving per each unit of input power. It can be noted that, contrary to active methods, passive methods that achieve drag reduction always have infinite gain. An important results regarding the power efficiency of flow control methods is that of Fukagata et al. (2009) who found that “one cannot drive a given flux with less total power than that of the laminar flow”. The laminar case skin-friction has been demonstrated by Bewley (2009) to be the true minimum that can be reached (without violating the power principle cited above).

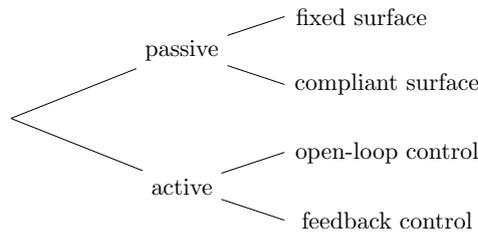


Figure 1.1: Classification of wall-based flow control techniques.

In this work, we research active open-loop methods (Chapter 3), active feedback-loop methods (Chapter 5), and passive compliant methods (Chapter 4).

b. *Passive methods*

Among the passive methods of drag reduction, the broadest class is represented by the permanent surface modification to the geometry of solid boundaries. Typically they come in the form of “roughness elements”, such as streamwise-aligned riblets that reduce the skin-friction of up to 10% (Garcia-Mayoral and Jiménez, 2011) for appropriate values of the spanwise spacing and the height. In other techniques, the wall geometry is altered in order to induce a secondary flow in the near wall region that mimics the effect of active methods (Ghebali et al., 2017). A lubrication-like mechanism forms the basis of the super-hydrophobic surfaces (Min and Kim, 2004; Daniello et al., 2009), where roughness elements and entrapped gas pockets are combined to generate a finite value of the velocity at the wall, called slip. Luchini (2015) has shown that the drag-reduction mechanisms of riblets and finite-slip surfaces are analogous and rooted in the modification of the near-wall high-velocity streaks that characterise wall turbulence. A more complete introduction to super-hydrophobic surfaces is given in Chapter 3, where a numerically model of a super-hydrophobic surface is used. In the framework of Fig. 1.1, super-hydrophobic surfaces can be classified as a type of compliant surface, where the wall surface is not fixed but realises a sort of natural feedback response to the fluid motion. Another example of a compliant surfaces are pressure-driven coating models with wall-normal velocity response (Fukagata et al., 2008), or the channel fitted with spanwise-oriented, free-to-move cylinders of Józsa et al. (2019) that generate a streamwise velocity response.

c. Active methods

In active frictional drag reduction methods, a distinction is made between methods that operate the devices at a pre-determined value of the control parameters (called “open-loop” methods) and those that respond to specific changes in some fluid variable following a prescribed feedback-control algorithm. Among the feedback-loop techniques, a large body of research has been dedicated to applications of suction and blowing actuators (Choi et al., 1994; Bewley et al., 2001) as a counter-measure to velocity fluctuations. Suction and blowing devices found applications either as a tool for fully-developed turbulence manipulation, or to delay the transition to turbulence in spatially developing flows by appropriately counter-acting flow instabilities. Various open-loop techniques has been researched such as volume-force methods (Du and Karniadakis, 2000) and the in-plane wall motions that are discussed thoroughly in Section 1.2.2. Out-of-plane motions, in the form of sinusoidal wall-deformation travelling waves, have also been successful for drag reduction (Nakanishi et al., 2012; Tomiyama and Fukagata, 2013). Open-loop drag reduction studies typically entail a systematic variation of the control parameters in order to locate the optimum combination that maximises drag reduction. Unconventional methods also exist, such as the wall-acceleration technique of Kühnen et al. (2018) that is based on the relaminarisation of the flow via a flow instability mechanism, unlike the rest of the techniques described in this section, where a turbulent state is preserved.

d. Drag reduction as a logarithmic-layer shift

Wall-based drag reduction techniques aim at altering the structure of near-wall turbulence, producing an attenuation of the wall-normal velocity bursts that promote the transport of high-momentum fluid in the near-wall region and lead to a flatter velocity profile and increased skin-friction. The consequence on the wall-normal velocity profile is a thickening of the viscous sub-layer (Choi, 1989) and an upwards translation of the logarithmic layer away from the wall (an example is found in Sec. 3.3.2). This effect can be assessed quantitatively using the friction relation defined in (1.26), where a positive (upwards) shift $\Delta B > 0$ (Gatti and Quadrio, 2016) generates a reduction of C_f .

1.2.2 Turbulent drag reduction by in-plane wall motions

Drag reduction by in-plane wall motion was originally introduced by Jung et al. (1992) in a numerically simulated turbulent channel flow at $Re_\tau^0 = 200$ (where Re_τ^0 is the friction Reynolds number of the uncontrolled flow) with an oscillating wall. Wall oscillation is defined by the following temporal law for the spanwise velocity w at the wall:

$$w = A \sin(\omega t), \quad (1.1)$$

where the maximum wall velocity A and the time period of the oscillation $T = 2\pi/\omega$ are the control parameters. The study demonstrated the possibility of decreasing the turbulent skin friction of up to 40% by oscillating the wall in the cross-flow direction. They also found that the flow-control performance is highly dependent on the oscillation period T , with $T^+ \sim 100$ being the best-performing value (the superscript “+” denoting scaling in viscous units of the uncontrolled flow). The optimal oscillation period of around $T^+ = 100$ was found by Quadrio and Ricco (2004) to produce maximum drag reduction because it matches a characteristic temporal scale of the near-wall turbulent structures and an optimal thickness of the generated Stokes layer. For any T^+ , the drag reduction increases monotonically with A^+ , but the growth slows at larger values of A^+ (Quadrio and Ricco, 2004). The oscillating technique has been applied successfully to pipe flow, where the oscillation is performed along the azimuthal direction (Orlandi and Fatica,

1997), and to boundary layers (Skote, 2013). Experimental confirmation of the oscillating-wall technique has also been established in boundary-layer (Choi and Graham, 1998) and pipe flow (Choi, 2002).

a. Spanwise waves

Quadrio et al. (2009) generalised the concept of wall oscillation through travelling waves of spanwise velocity at the wall:

$$w = A \sin(\kappa_x x - \omega t). \quad (1.2)$$

The control parameters are therefore the velocity amplitude A , the wavenumber $\kappa_x = 2\pi/\lambda_x$ and the forcing period $T = 2\pi/\omega$.

The dependence of the skin-friction drag reduction on the parameters at $Re_\tau^0 = 200$ and $A^+ = 12$ is shown in Fig. 1.2b. When $\omega = 0$ a stationary wave of spanwise velocity is formed (Viotti et al., 2009). The optimal wavelength of the stationary wave is in the interval $1000 < \lambda_x^+ < 1250$, where the stationary wave generates drag reduction of around 45%, which is larger than the best pure wall-oscillation case. Upstream-travelling waves ($\omega < 0$, not shown in Fig. 1.2b) always yield drag reduction, but the absolute maximum drag reduction of 48% is found for a downstream-travelling wave. A subset of downstream-wave parameter combinations (the yellow region in Fig. 1.2b) are instead found to give drag increases up to more than 20%. It is observed that for the drag-increasing cases, the phase speed $c = \omega/\kappa_x$ of the travelling waves is similar to the near-wall turbulent convection velocity.

Although drag reduction increases monotonically with A^+ , the same is not true for the power efficiency of the scheme. The largest net power saving at $A^+ = 12$ is 18% and an even larger saving of 26% is found at $A^+ = 6$. A minimal forcing amplitude has been found, i.e., the maximum velocity of the wall must be larger than a minimal value to affect the near-wall turbulence and cause drag reduction (Ricco and Quadrio, 2008).

The travelling-wave concept has been experimentally validated by Auteri et al. (2010), yielding up to 33% drag reduction in a pipe flow with independently-rotating wall slabs, and by Bird et al. (2018) who measured drag reduction rates of up to 21.5% for upstream-travelling waves in a boundary layer fitted with a actuated-membrane wall surface.

b. Stokes layer

A key finding for the oscillating-wall case was that the spatially-averaged spanwise-alternating flow follows an exact laminar Stokes layer solution (Quadrio and Sibilla, 2000). The visualizations by Ricco (2004) show that the streamwise-velocity streaks are dragged laterally by the spanwise shear of the Stokes layer, while the shorter, faster-travelling quasi-streamwise vortices located above the streaks (Adrian, 2007) and above the Stokes layer do not move laterally, but are only indirectly altered by the wall forcing as their energy is attenuated. This optimal Stokes layer directly affects the buffer-layer velocity streaks, as it is a few wall-units thick. This result was extended to the case of the streamwise-travelling waves by Quadrio and Ricco (2011), who derived a generalised Airy-function Stokes layer and revealed a striking correlation between its thickness and the drag reduction.

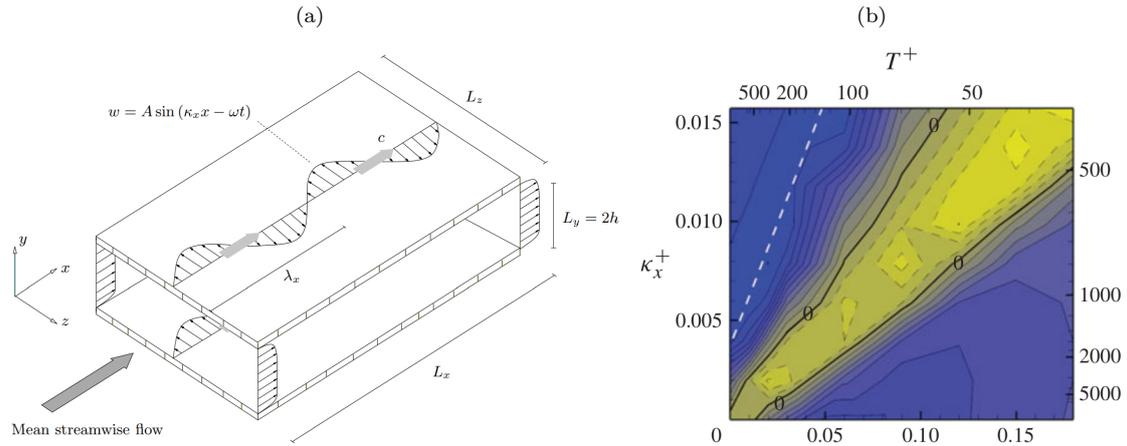


Figure 1.2: Streamwise-travelling waves of spanwise velocity: (a) sketch of a plane channel flow with spanwise-velocity waves at the wall from Quadrio et al. (2009), (b) parametric dependence of the drag-reduction of spanwise-velocity waves at $Re_\tau = 180$ from Hurst et al. (2014), yellow is drag increase and the dashed line marks the maximum ridge.

c. Reynolds number effect

An important question concerns the effect of the Reynolds number on the performance of in-plane wall motion techniques and is critical to assess their applicability to high-Reynolds-number scenarios. At the low friction Reynolds number value of $Re_\tau = 200$, the near-wall dynamics are responsible for the largest contribution to the skin-friction (Quadrio, 2011; de Giovanetti et al., 2016). However, at higher Reynolds numbers, the very large structures or “super-streaks” that characterise the upper layers become increasingly energetic and account for a growing contribution to the skin-friction on the wall (Agostini and Leschziner, 2014). It can be argued that this constitutes a fundamental limitation to the performance of near-wall control methods at high Reynolds numbers. Indeed, several numerical studies up to $Re_\tau = 2000$ (Hurst et al., 2014; Gatti and Quadrio, 2016; Yao et al., 2019) have ascertained that drag reduction rates of spanwise wall motions decrease at higher Reynolds numbers, possibly following a power-law scaling $\sim Re_\tau^{-\gamma}$. However, the

d. Discs and rings

Spinning circular actuators, or discs, mounted on the wall were first proposed by Keefe (1997, 1998) as a drag-reduction device combined with suction and blowing at the wall as shown by Fig. 1.3. Rotating discs arranged in a rectangular array on the wall have been introduced by (Ricco and Hahn, 2013) and simulated numerically in a channel flow at $Re_\tau = 180$. The disc belonging to adjacent rows rotate in opposite directions, thus realising, along their centreline, a spanwise-velocity pattern similar to a standing-wave forcing at the wall. In case of steadily-rotating discs, the forcing parameters are the disc-tip velocity W (analogous to the wave amplitude) and the disc diameter D (analogous to the wavelength). The maximum drag reduction of 23% is obtained for a diameter $D^+ = 900$ and a disc-edge velocity $W = 9^+$, while the maximum net power saved is 10% for $W = 6^+$ and $D^+ = 900$ (Ricco and Hahn, 2013; Wise and Ricco, 2014). Contrary to spanwise waves, the drag reduction rate of rotating discs does not increase monotonically with the forcing amplitude W^+ but has a maximum for any value of the forcing scale D^+ . The discs offer several practical advantages with respect to spanwise waves or oscillations, i.e. the

local character that avoids the impractical motion of the whole solid wall and the wider range of diameters leading to drag reductions of similar magnitude to the maximum (Ricco and Hahn, 2013; Wise et al., 2014).

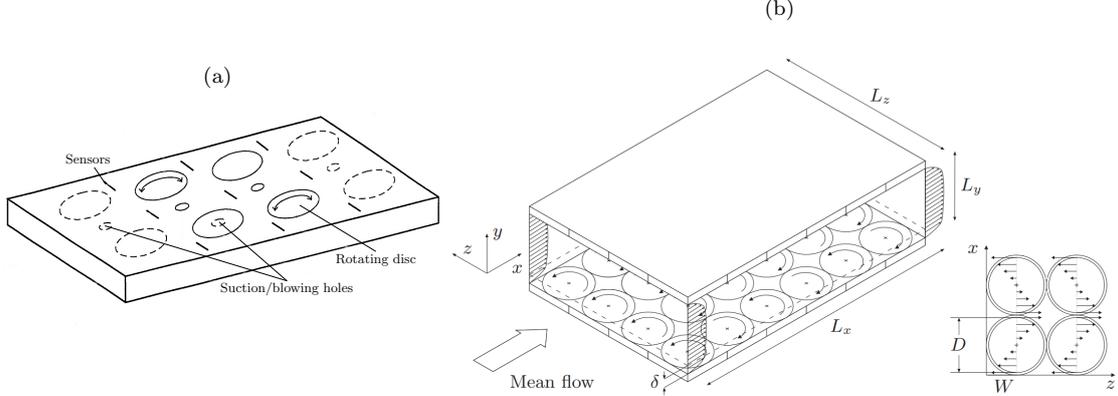


Figure 1.3: (a) rotating discs actuators combined with wall transpiration, original sketch from Keefe (1998), (b) sketch of a plane channel flow with array of rotating discs from Ricco and Hahn (2013).

Wise et al. (2014) showed that rotating annular actuators can deliver similar levels of drag reduction while requiring up to 20% less driving power. The same type of actuator, combined with opposition control and hydrophobic surfaces is the theme of Chapter 3.

1.3 Turbulent flow in an infinitely-extended channel

The physical problem we study is the motion of a fluid between two infinitely-extended, rigid, impermeable plates or walls, separated by a distance $2h^*$ as shown in Fig. 1.4. We adopt a right-handed Cartesian coordinate system x^*, y^*, z^* , such that the origin is located on one of the two walls, the second axis y^* is orthogonal to the wall and the first axis x^* is aligned with the direction of the pressure gradient \mathbf{F} that drives the flow. The positive direction of y^* is towards the other wall and defines the notions of “upper” and “lower” wall. x^* is called the streamwise axis and z^* is called the spanwise axis.

Adopting the Eulerian description of the flow (Batchelor, 1967), the fluid kinematics is described by its pointwise velocity field $\mathbf{u}(\mathbf{x}, t)$, where \mathbf{x} is the position vector in the Cartesian coordinate system and $t \geq 0$ is time. The scalar fields $u(\mathbf{x}, t)$, $v(\mathbf{x}, t)$ and $w(\mathbf{x}, t)$ denote respectively the streamwise, wall-normal and spanwise component of the velocity vector \mathbf{u} . The fluid is incompressible, isothermal, Newtonian and characterised by a constant viscosity μ^* and a constant density ρ^* . The fluid is subjected to a uniform pressure gradient $\mathbf{F} = \{P_x^*, 0, 0\}$. Since the channel is infinitely extended, we can measure the flow rate per unit area Q^* as:

$$Q^* = \frac{1}{2h^*} \int_0^{2h^*} u(\mathbf{x}, t) dy. \quad (1.3)$$

From the flow rate, a reference velocity can be derived as $U_b^* = Q^*/2h^*$, called the bulk velocity. Alternatively, a reference velocity can be defined as $U_p^* = 3U_b^*/2$, which is the centreline velocity produced in the plane Poiseuille flow at the flow rate Q^* . The reason of this choice will be made clear in the next paragraph.

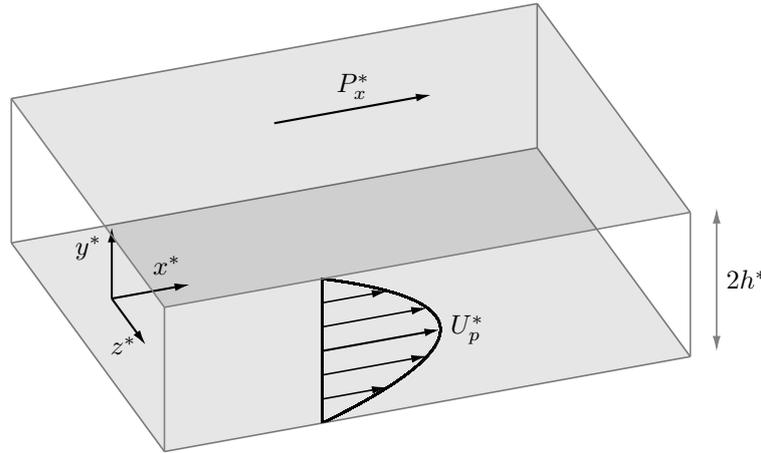


Figure 1.4: Plane channel flow: schematic of the fluid domain and the parabolic Poiseuille velocity profile.

The solution to the physical problem described above depends on a four constants or parameters, namely the properties of the fluid μ^* and ρ^* , the reference velocity U_p^* or the forcing P_x^* , and the geometrical parameter h^* . It is always convenient to reduce the problem to its non-dimensional form, which represents the most compact and general form of the problem by reducing the number of parameters to the minimum. Any flow variable e.g. the velocity \mathbf{u} can be expressed as a function of the dimensional parameters as $\mathbf{u}^* = f^*(\mu^*, \rho^*, h^*, U_p^*)$. A non-dimensional form of the above problem is obtained by normalising lengths with h^* , velocities with U_p^* and the density with ρ^* . By combining the fundamental units, we can reduce other quantities to their non-dimensional form: flow rates with $h^*U_p^*$, times with h^*/U_p^* , stresses with $\rho^*U_p^{*2}$ and pressure gradients with $\rho^*U_p^{*2}/h^*$. The non-dimensional form depends on a single parameter and reads $\mathbf{u}^*/U_p^* = g(\rho^*h^*U_p^*/\mu^*)$. The parameter is called the Reynolds number and indicated by the symbol Re_p , where the subscript p refers to the Poiseuille velocity U_p . Another possible normalisation is obtained by scaling velocities with U_b^* and leads to the definition of bulk Reynolds number $Re_b = \rho^*h^*U_b^*/\mu^*$. Dimensional variables are henceforth indicated with an asterisk and non-dimensional ones without, for example $\mathbf{u} = \mathbf{u}^*/U_p^*$, $x = x^*/h^*$ etc.

The laws of motion of the fluid can be derived (Batchelor, 1967) by considering the conservation of momentum and the mass on an infinitesimal volume of fluid, and using the incompressible Newtonian constitutive law. In accordance with all the assumptions described above, the dynamics of the fluid is described by the following differential problem, written in non-dimensional form:

$$\begin{aligned} \frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} &= -\nabla p + \frac{1}{Re_p} \Delta \mathbf{u} + \mathbf{F}. \\ \nabla \cdot \mathbf{u} &= 0 \\ \mathbf{u}(x, 0, z, t) &= \mathbf{0} \\ \mathbf{u}(x, y, z, 0) &= \mathbf{u}_0. \end{aligned} \tag{1.4}$$

where the first equation expresses point-wise the conservation of momentum and the second the point-wise conservation of mass. The two equations are known as Navier-Stokes equations (Navier, 1823).

The initial condition \mathbf{u}_0 is the plane Poiseuille flow, that is an exact laminar solution to (1.4) when no perturbations are present. The plane Poiseuille velocity field reads:

$$\mathbf{u}_p(x, y, z) = \{y(2 - y), 0, 0\}. \quad (1.5)$$

The pressure gradient needed to sustain Poiseuille flow is $P_x = 2Re_p$ and the flow rate is $Q = 4/3$.

a. Transition to turbulence

For sufficiently large values of the Reynolds number, wall-bounded laminar flows are unstable to disturbances. Instability means that when a disturbance of appropriate amplitude is introduced, the flow does not return to its laminar state but develops a more complex behaviour, characterised by spontaneous spatial and temporal fluctuations. This behaviour is markedly different from laminar flow where all the streamlines are parallel to the pressure gradient and do not change in time.

The plane Poiseuille flow (1.5) is known to be unstable to arbitrarily small perturbations when $Re_p > 5772$, but obtaining and sustaining a turbulent state is possible for approximately $Re_p > 1000$ when finite-amplitude disturbances are present (Schmid and Henningson, 2001). To obtain a turbulent state, the boundary condition is modified as $\mathbf{u}_0 = \mathbf{u}_p + \mathbf{u}_{noise}$ i.e. the plane Poiseuille flow \mathbf{u}_p is perturbed by a disturbance \mathbf{u}_{noise} that has to be chosen appropriately to generate the transition.

b. Fully developed turbulence

If the Reynolds number is beyond a critical threshold, the transition excites an increasing number of spatial and temporal modes, eventually reaching the so-called fully developed turbulent state, characterised by chaotic fluctuations and a continuous energy spectrum. Fully-developed turbulent flow is ubiquitous in engineering applications, where the Reynolds number is often very large and some form of random disturbance is almost always present. In fully developed turbulence the flow field appears as a collection of local, transient features called “turbulent structures”, that, when examined individually, are practically never identical to another. Such complexity leads to a statistical description of turbulent flow, as studying averages in space or time is more tractable than following individual structures (this operation can be justified as explained in Sec. 1.3.1). Flow statistics such as the structures’ length-scale, life-time, and behaviour vary in the wall normal direction. Three fundamental length-scales characterise wall-bounded turbulence: the viscous length-scale δ_ν^* , the Kolmogorov microscale η^* and the outer, or bulk, scale corresponding to h^* (Smits et al., 2011). η^* represent the scale at which most of the energy dissipation into heat takes place. The minimum spatial scale of motion correspond to the diameter of the finest structures, around 10η (Smits et al., 2011). On the other extreme, the very-large scale motions (VLSM) can span a streamwise size of around $40h$ (Lozano-Durán and Jiménez (2014)). The presence near the wall of elongated, roll-like structures (Adrian, 2007) has important implications for the shear force exerted on the wall because they intermittently blow high-velocity fluid from above to near the wall, generating steeper, on average, velocity gradients than in the laminar case. The origin of the emergence and self-sustained dynamics of turbulent structures from otherwise smooth and plane surfaces is the subject of active research, with the current understanding being that they are produced through the so-called near-wall regeneration cycle (Hamilton et al., 1995) that is capable of extracting energy from the mean shear through instability mechanisms immediately above the wall. These ideas lead to the idea that, if one wants to modify wall-bounded turbulence and potentially affect the drag, it is reasonable to focus on the manipulation the near-wall flow.

1.3.1 Statistical description of turbulence

In the fully-developed turbulent state, it is convenient, for all practical purposes, to describe the fluid motion in terms of statistical quantities, such as the mean and the variance of the velocity components, the pressure and other flow variables. The statistical approach to turbulence relies on some conjectures on the dynamics generated by the equations of motion (Ruelle and Takens, 1971; Frisch, 1995) that can be summarised as follows. A solution \mathbf{u} to problem (1.4) belongs to the functional space \mathbb{G} of the solutions to (1.4). Starting from an initial condition $\mathbf{u}_0 = \mathbf{u}(0) \in \mathbb{G}$, the dynamics $\mathbf{u}(t)$ evolves following a trajectory in \mathbb{G} . For example, the unperturbed laminar flow \mathbf{u}_p is a fixed point of the dynamics (1.4). The steady-state turbulent solutions are conjectured to be trajectories that evolve in a subset $\Omega_a \subset \mathbb{G}$, orbiting chaotically while remaining indefinitely in Ω_a . This means that turbulent solutions are highly sensitive to the initial condition, generating very different trajectories $\mathbf{u}(t) \in \Omega_a$ for very similar \mathbf{u}_0 . Ω_a is called the attractor set and it is conjectured that, given any initial condition that leads to fully developed turbulence, the dynamics will eventually reach Ω_a . Therefore the choice of a specific initial condition is not important as long as it can successfully trigger the transition to fully-developed turbulence. While the details of the fluctuations are unpredictable, being highly sensitive on the initial condition, the statistical properties of the turbulent solutions are instead predictable due to the uniqueness of the attractor set. The statistical study of the flow variables would be otherwise meaningless because, in general, each (random) initial condition would produce a unique dynamics.

Another important assumption in the statistical description of turbulent flows is ergodicity (Frisch, 1995). In a flow under stationary (constant in time) forcing, the following ergodic theorem can be formulated for the velocity field \mathbf{u} :

$$\lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \mathbf{u}(t; \xi) dt = \int_{\Omega_a} \mathbf{u}(\xi) dp(\mathbf{u}) \equiv E[\mathbf{u}], \quad (1.6)$$

where ξ is a generic random variable that represents the arbitrariness of the initial condition, and p is the probability distribution of \mathbf{u} on the attractor set Ω_a . The right-hand side is the expected value of \mathbf{u} . The left-hand integral is the time-average of \mathbf{u} , performed on a finite interval T . As the size T of the time interval goes to infinity, the average converges asymptotically to the expected value of \mathbf{u} and provides, for finite values of r , an estimate of the true mean $E[\mathbf{u}]$ with some level of uncertainty.

The formulation in time of (1.6) for stationary turbulence can be extended to other variables, depending on the symmetries of the problem. For example, in a spatially homogeneous flow the spatial position \mathbf{x} can replace time. It is expected that the same flow symmetries of the laminar flow are restored, this time in the statistical sense, in a fully-developed turbulent flow (Frisch, 1995). In this respect, studying the flow in a plane channel (or a pipe) is convenient in terms of the collection of statistics because three statistical symmetries are present:

$$E[\mathbf{u}(x, y, z, t)] = E[\mathbf{u}(x, y, z, t + C)], \quad (\text{stationary}) \quad (1.7)$$

$$E[\mathbf{u}(x, y, z, t)] = E[\mathbf{u}(x + A, y, z + B, t)], \quad (\text{homogeneous}) \quad (1.8)$$

$$E[\mathbf{u}(x, y, z, t)] = E[\{u, -v, w\}(x, 2 - y, z, t)], \quad (\text{top - bottom}) \quad (1.9)$$

where $A, B, C \in \mathbb{R}$ arbitrary constants. Each symmetry has an associated averaging operator:

$$\bar{\mathbf{u}} = \frac{1}{T} \int_0^T \mathbf{u}(x, y, z, t) dt, \quad (1.10)$$

$$\langle \mathbf{u} \rangle = \frac{1}{L_x L_z} \int_0^{L_x} \int_0^{L_z} \mathbf{u}(x, y, z, t) dx dz, \quad (1.11)$$

$$[\mathbf{u}]_{tb} = (\{u, v, w\}(x, y, z, t) + \{u, -v, w\}(x, 2 - y, z, t))/2 \quad (1.12)$$

Substituting any of the above operators in the left-hand side of (1.6), gives asymptotically the mean value of a flow variable. In practice, the averaging operators are always compounded to minimise statistical uncertainty. The mean velocity field is written as $\mathbf{U} = E[\mathbf{u}]$ and the mean pressure field as $\mathbf{P} = E[\mathbf{p}]$, other mean flow variables are usually denoted using the notation of one of the operators e.g. the mean Reynolds shear stresses are written $\langle uv \rangle$ or \overline{uv} . New symmetries and averaging operators will be introduced when they are needed, see for example the more complex symmetry belonging to the disc boundary conditions presented in Chapter 3.

Whenever the mean value of a flow variable is considered, an associated statistical decomposition (called ‘‘Reynolds decomposition’’) can be defined that separates the mean from the zero-mean fluctuating component. The most basic one is the decomposition of the flow field between the mean \mathbf{U} and the fluctuations \mathbf{u}_t :

$$\{\mathbf{u}, p\}(\mathbf{x}, t) = \{\mathbf{U}, P\}(y) + \{\mathbf{u}_t, p_t\}(\mathbf{x}, t). \quad (1.13)$$

Statistical decompositions for any other flow variable can be derived from the flow field decomposition, for example those for the kinetic energy or the Reynolds stresses shown in the next section.

It is also possible to write conservation laws for the mean or the fluctuating quantities. These equations are useful to understand the mean dynamics of the flow (which is usually the main object of interest) and its relation with the fluctuating component. The most important of such equations are the so-called Reynolds-averaged equations for the mean flow, that read:

$$\frac{\partial \mathbf{U}}{\partial t} + (\mathbf{U} \cdot \nabla) \mathbf{U} = -\nabla P + \nabla \cdot \mathbf{R} + \frac{1}{Re_p} \Delta \mathbf{U} + \mathbf{F}(t), \quad (1.14)$$

$$\nabla \cdot \mathbf{U} = 0, \quad (1.15)$$

$$\mathbf{U}(x, 0, y, t) = 0. \quad (1.16)$$

where the matrix \mathbf{R} contains the Reynolds shear-stresses (RSS) e.g. $R_{12} = E[u_t v_t]$. A noteworthy property of (1.14) is the fact that the conservation equation for the mean flow (a first-order statistical moment) contains the RSS term (a second-order statistical moment), implying that it is not possible to solve for the mean flow without knowing the fluctuating dynamics, or, in more technical terms, the equations are not closed. The closure problem is a general rule, as the equations for any n -th order moment of a flow variable always contain a moment of order $n + 1$ or higher.

1.3.2 Basic results in channel flow

In this section we briefly recall a few important results that can be derived analytically from the mean-flow equations and that are used to validate and discuss the numerical results in Section 2.2.4 and throughout this work.

a. *Linear total shear-stress profile*

First, we observe that from the problem symmetries $W = 0$. Therefore, the mean-flow equations (1.14) are reduced to:

$$0 = \frac{d}{dy} \left(\frac{1}{Re} \frac{dU}{dy} - \langle uv \rangle \right) + P_x, \quad (1.17)$$

$$0 = -\frac{d\langle v_t^2 \rangle}{dy} - \frac{\partial P}{\partial y}, \quad (1.18)$$

$$\frac{dV}{dy} = 0 \quad (1.19)$$

From continuity (1.19) and the BCs it follows that $V = 0$ and thus the mean velocity field is $\mathbf{U} = \{U, 0, 0\}$. The first equation is the balance of the mean forces along x , meaning that the total mean shear stress τ , i.e. the sum of the mean viscous stress $Re_p^{-1}dU/dy$ and the turbulent stresses $-\langle uv \rangle$, is balanced out by the mean pressure gradient P_x . From equations (1.17)-(1.18), a few exact results can be derived.

From continuity it can be shown that the mean total shear-stress $\tau(x, y, z) = \tau(y)$ and the mean pressure $P(x, 0, z) = P(x)$, which implies that both terms in (1.17)-(1.18) are constant. Knowing that τ is anti-symmetric about $y = 1$ we obtain a linear profile upon integration of (1.17)-(1.18):

$$\tau(y) = \tau_w(1 - y) \quad (1.20)$$

where $\tau_w = \langle \bar{\tau}(0) \rangle$ is the mean wall shear-stress. Such a conclusion is not valid for the boundary-layer flow where the transport of momentum takes place also along the wall normal direction and thus $V \neq 0$.

Starting from the mean wall-shear stress τ_w^* we can define the friction velocity $u_\tau^* = \sqrt{\tau_w^*/\rho^*}$ and the viscous length-scale δ_ν^* , such that $u^*(y^* = \delta_\nu^*) = u_\tau^*$. These quantities also define a new viscosity normalisation, called the friction Reynolds number $Re_\tau = u_\tau^* h^*/\nu^*$. Since $Re_\tau^{-1} = \delta_\nu^*/h^*$, the friction Reynolds number gives a measure of the relative size of the near-wall length-scale to the global scale of the flow. Re_τ has a sub-linear power-law dependence on the Reynolds number as $Re_\tau \propto Re_p^{0.88}$. Quantities scaled with u_τ^* and δ_ν^* , τ_w^* or combinations thereof are called viscous or inner units and are denoted by the superscript ‘+’.

b. *The mean velocity profile*

In a fully-developed turbulent channel flow, the mean flow is characterised by two length-scales in the wall-normal direction, the outer scale h^* and the near-wall scale δ_ν^* . Using dimensional analysis it can be proven that the flow at a distance y^* from the wall is completely specified by the ratio to the outer scale y and the ratio to the near-wall scale y^+ . These observations can be used to derive asymptotically exact results on several mean-flow variables. The wall-normal velocity profile $U(y)$ can be partially formulated as a piecewise function of the wall-normal coordinate (Townsend, 1976; Pope, 2000), based on the asymptotic considerations on magnitude of y and y^+ . The most important exact results are shown in Table 1.1.

The profile can be primarily subdivided into two regions, the near-wall layer and the outer layer. In the inner layer, y is very small and the effects of viscosity are not negligible, thus U depends only on y^+ . In the outer layer, y^+ is very large and the effects of viscosity are negligible. The near-wall layer can be further subdivided into the viscous sub-layer, where y^+ is very small and U follows a linear profile, and a second region, where y is still very small but y^+ is much larger than unity. In such conditions, the velocity gradient depends on y^{-1} and the velocity

profile has logarithmic dependence specified by the constants k and B . The logarithmic region is also called the equilibrium layer because in that region, there is no spatial transport of kinetic energy but a local equilibrium between production and dissipation. Furthermore, the ratio of the Reynolds shear-stress to the kinetic energy is constant. Between the viscous sub-layer and the logarithmic layer, a region called buffer layer exist where the viscous and Reynolds stresses have a similar magnitude and the velocity follows neither a logarithmic profile nor a linear one. The buffer layer is noteworthy as it contains the maximum of the velocity fluctuations, the maximum the turbulent kinetic energy production and is the maximally anisotropic region of the flow (see Section 2.2).

The uppermost region follows the velocity defect law (Coles, 1956), which does not scale in viscous units and is specified by the flow-dependent constant B_1 . For a more complete treatment of all the wall-normal subdivisions of the flow, see for example Table 7.1 and Figure 7.8 in Chapter 7 of Pope (2000).

As the wall-normal regions are defined asymptotically by y^+ , by definition they hold exactly only for $Re_\tau \rightarrow \infty$ i.e. when the inner-outer scale separation is infinitely large. Experiments and numerical simulations (including our channel-flow simulation of Section 2.2.4) establishes that at low values of Re_τ the velocity profile deviates from a perfect logarithmic law. Many studies exist on the log-law asymptotic dependence on Re_τ (Panton, 2007). In order to observe a true logarithmic profile, higher values of Re_τ must be reached e.g. Lee and Moser (2015) at $Re_\tau = 5200$ and Yamamoto and Tsuji (2018) at $Re_\tau = 8000$. Asymptotic convergence is observed as Re_τ increases, the near-wall layer shrinks and the logarithmic layer extends, extending closer to the wall. Tab. 1.1 also reports the approximate ranges where the exact asymptotic forms agrees well with the experiments (Pope, 2000).

Table 1.1: Exact results for the wall-normal velocity profile in channel flow, adapted from Pope (2000). U_c^+ is the centreline velocity, k , B and B_1 are constants.

name	asymptotic range	empirical range	U^+
viscous sub-layer	$y^+ \ll 1$	$0 < y^+ < 5$	y^+ (1.21)
logarithmic law	$y^+ \gg 1, y \ll 1$	$y^+ > 30, y < 0.1$	$\frac{1}{k} \log(y^+) + B$ (1.22)
velocity-defect law	$y^+ \gg 1, y \ll 1$	$0.1 < y < 0.3$	$U_c^+ + \frac{1}{k} \log(y) - B_1$ (1.23)

c. Universality of the velocity profile

The velocity profile in the near-wall layer, scaled in wall units as $U^+(y^+)$ is called “law of the wall” and applies to all the wall-bounded canonical flows (channel, pipe and boundary layer) and to any flow that can be locally approximated by one of the above. This property of $U^+(y^+)$ (and other statistics e.g. the Reynolds stresses) is called universality. In boundary-layer flows, this also implies self-similarity along the spatially-developing direction. The constants of the logarithmic law are universal and determined empirically as $k \approx 0.41$ and $B \approx 5.2$. The discrepancies observed in experiments and simulations across different flow types can be corrected based on the influence of the pressure gradient proposed by Luchini (2017) that retrieves the law of the wall. The intercept constant B depends on the type of boundary condition and changes when the boundary conditions are modified such as in presence of wall surface roughness or flow control for drag reduction.

1.3.3 Skin-friction and drag

Using the wall-shear stress and the bulk velocity, the skin-friction coefficient C_f can be defined as:

$$C_f = \frac{2\tau_w^*}{\rho^* U_b^{*2}}, \quad (1.24)$$

where τ_w^* is the mean wall shear-stress and U_b^* is the bulk velocity, defined in Section 1.3. C_f is the ratio of the resisting force to the flow bulk kinetic energy and can be thus understood as a measure of the (in)efficiency of the flow. The physical effect quantified by C_f is called “frictional drag”, “skin-friction drag” and the sub-field of fluid mechanics concerned with the reduction of C_f is called “frictional drag reduction”, “skin-friction reduction”.

For turbulent flow in plane channels, several experimental measurements of C_f exists, such as Dean’s correlation in square ducts (Dean, 1978) that reads:

$$C_f = 0.073 Re_b^{-1/4}, \quad C_f = 0.0336 Re_\tau^{-0.273}. \quad (1.25)$$

An useful relation between the skin-friction and Re_τ can be derived (Gatti and Quadrio, 2016) using the definition of the friction coefficient, and substituting the velocity profiles (1.22)-(1.23), which reads:

$$\sqrt{\frac{2}{C_f}} = \frac{1}{k} \log Re_\tau + B + B_1 - \frac{1}{k}. \quad (1.26)$$

The above relation applies well to the skin-friction dependence if Re_τ is high enough for the velocity to agree well with the logarithmic profile. Another important definition is that of mean pumping power or mean energy flux (per unit wall area), that is expressed in non-dimensional terms as:

$$P_p = 2\tau_w Q = C_f U_b^3. \quad (1.27)$$

a. Drag reduction

This study is primarily concerned with decreasing C_f through various flow-control methods, usually consisting in modifications to the boundary conditions (see Sec. 1.2 for a review). The performance of such techniques is measured by a quantity called “frictional drag reduction”, defined as the reduction of the skin-friction coefficient relative to the uncontrolled case and usually expressed in percent:

$$\mathcal{R}(\%) = 100 \cdot \frac{C_f^0 - C_f}{C_f^0}, \quad (1.28)$$

where C_f^0 is the skin-friction coefficient of the uncontrolled flow and C_f of the controlled flow. In this work and in the literature, frictional drag reduction is simply referred to as “drag reduction”, even though in practical engineering cases the total hydrodynamic drag includes other phenomena other than skin-friction, primarily form drag due to the curved geometry of real-world solid bodies.

b. Forcing paradigm

When a comparison is carried out between the controlled and the uncontrolled channel or pipe flow, a choice has to be made regarding how to impose the forcing. Typically, the two possibilities are to drive the controlled flow maintaining either the same pressure gradient or flow rate of the uncontrolled flow (Hasegawa et al., 2014). The first is known in the literature as constant-flow-rate (CFR) forcing and the second as constant-pressure-gradient (CPG) forcing. While the

difference does not matter for a single flow case, for which the flow rate is uniquely determined by the pressure gradient and vice versa, it does when a comparison is made between the controlled and uncontrolled case. At CFR, Re_b is by definition constant, therefore if $\mathcal{R} > 0$ then Re_τ decreases and thus the mean pumping power P_p decrease. At CPG Re_τ is by definition constant, therefore if $\mathcal{R} > 0$ then Re_b decreases and thus P_p increases. It is useful to introduce the definitions of drag reduction when the channel flow is driven by a constant-flow-rate (CFR) and a constant-pressure-gradient (CPG), respectively:

$$\mathcal{R}_{CFR}(\%) = 100 \cdot \frac{\tau_{w,0} - \tau_w}{\tau_{w,0}} \quad (1.29)$$

$$\mathcal{R}_{CPG}(\%) = 100 \cdot \frac{U_b^2 - U_{b,0}^2}{U_b^2}. \quad (1.30)$$

In this work, as it is done in the majority of the numerical literature, CFR forcing is used.

The forcing paradigm also bears consequences on the numerical simulation of the flow, owing to the finite size of the fluid domain. For CFR forcing, in order to keep U_b constant at all times, the gradient P_x must be adjusted instantaneously to counteract the fluctuation of the skin-friction that arise due to the incomplete statistical convergence of spatial fluctuations. Using CFR results in a shorter transition time to the fully turbulent state compared to the CPG approach (Oliver et al., 2014) and is convenient for drag-reduction studies.

Chapter 2

Direct Numerical Simulation of turbulent channel flow

In this Chapter, the numerical techniques and computational procedures adopted for the simulation of turbulent flow in a plane channel on state-of-the-art high-performance computers are described and motivated. A simulation of a reference plane-channel flow case is carried out and its outcome validated. Section 2.2 gives an overview of *Incompact3D*, the computer code that is used to efficiently simulate turbulent channel flow on high-performance computers. Section 2.1 applies *Incompact3D* to the numerical solution of plane channel flow, motivates its numerical set-up and compares the results with sources from the literature. In Section 2.3, a short account of the methods used to extract statistics from the simulation data is given.

2.1 Direct Numerical Simulation of turbulent flow

The goal of this section is to provide an overview of the procedures adopted to perform the numerical simulations that are used to investigate the model channel-flow problem and generally all the research problems discussed in the later chapters.

All the flow cases presented in this work are carried out through Direct Numerical Simulation (DNS). The term ‘direct’ means that the Navier-Stokes equations are discretised and integrated in their natural form, resolving numerically the entire range of spatio-temporal scales that characterise the turbulent fluctuations. DNS is distinguished from other techniques, such as Large Eddy (LES) or Reynolds-averaged Navier-Stokes (RANS) simulations, that use models of varying degrees of sophistication to avoid the computation of a certain range of scales and thus achieve a computational advantage. For example, LES uses coarser grids and sub-grid models to avoid the computation of the smallest scales (Meneveau and Katz, 2000), while RANS use turbulence models (Wilcox, 2006) to solve directly for the mean flow, avoiding the computation of fluctuations altogether and allowing much coarser grids and larger time-steps than a DNS.

Historically, Direct Numerical Simulations (Moin and Mahesh, 1998) were introduced in the work of Orszag and Patterson (1972) for homogeneous isotropic turbulence with tri-periodic boundary conditions at low Reynolds number. The first true channel-flow DNS is the seminal work of Kim et al. (1987) that, similarly to earlier work, used a spectral method for spatial derivatives. Finite-difference methods started to be employed later because they offered smaller aliasing errors than spectral methods, albeit with reduced spectral resolution. A major advancement in finite-difference techniques came with the introduction of compact schemes (Lele, 1992)

that made possible for the large grid size requirements of standard finite-difference schemes to be reduced.

This section is organised as follow: Section 2.1.1 explains the numerical methods of the solver Incompact3D and Section 2.1.2 provides an introduction to parallel computing and the parallelisation technique used by Incompact3D.

2.1.1 The Incompact3D solver

In order to simulate the flow physics using finite digital computers, the continuum-mechanical problem (1.4) is approximated through a numerical method. A numerical method consists of *i*) finite temporal and spatial domains, *ii*) a discrete version of the PDEs and the boundary conditions (Quarteroni et al., 2010). The type and discretisation depends on both the physics of the problem and the numerical techniques adopted.

The simulations in the present work are carried out using the open-source Incompact3D solver. Incompact3D is a finite-difference based, highly parallel code that is capable of DNS and LES of incompressible flows¹ on Cartesian grids. A detailed description of the solver is outlined in the papers Laizet and Lamballais (2009) and Laizet and Li (2011). Here we provide a brief review of the code as applied to the plane channel flow problem (1.4).

Incompact3D discrete spatial domain consists in an orthogonal, partially stretched Cartesian grid. The domain $\mathbb{R} \times [0, 2] \times \mathbb{R}$ of (1.4) is approximated by a finite domain called computational grid or mesh and defined as:

$$\begin{aligned} & \{x_i\}_{i=1}^{n_x} \times \{y_j\}_{j=1}^{n_y} \times \{z_k\}_{k=1}^{n_z} \\ \text{with } & x_1 = 0, y_1 = 0, z_1 = 0, x_{n_x+1} = L_x, y_{n_y} = 2, z_{n_z+1} = L_z, \end{aligned} \quad (2.1)$$

where the indexes i, j, k are integers. Along x and z the grid is uniformly spaced i.e. $\Delta x = x_{i+1} - x_i$ and $\Delta z = z_{k+1} - z_k$ are constants. Along the wall-normal direction the grid is stretched, meaning that the grid spacing depends on its location as $\Delta y_j = y_{j+1} - y_j$. In general, a non-uniformly spaced sequence of points can be defined by a mapping $y_j = h(s_j)$ where $\{s_j\}_{j=1}^{n_y}$ has a constant spacing. More precisely, y_j is non-uniformly stretched i.e. $h(s_j)$ is non-linear. The use of a stretched grid along the wall-normal direction of a turbulent channel or pipe flow can lead to a significant reduction of the grid size with respect to a uniformly-spaced grid. Stretching is possible because the wall-normal spatial scales are about ten times smaller in the near-wall region than at the centreline. Incompact3D is capable of working with computational grids that are non-uniform along one direction. As Incompact3D solves the pressure in spectral space (see Section a.), the functional form of h is constructed (Cain et al., 1984) to reduce the computational cost of calculating the the FFT on a non-uniform grid (which would generally entail a very costly convolution product across all wavenumbers). h is adjusted by varying a single parameter, called β , that controls the stretching along y and determines the size of the near-wall and the centreline spacing.

The grid size defines the discrete domain in spectral space i.e. the sequence of wavenumbers that can be ideally resolved on the grid. For instance, the spectral domain along the x direction is

$$\{\kappa_l\}_{l=n_x/2}^{n_x/2-1} \quad \text{with} \quad \kappa_l = \frac{2\pi l}{\Delta x}. \quad (2.2)$$

Therefore, the $2\pi/L_x$ is largest wavenumber and $\kappa_l = \pi l/\Delta x$ is the smallest Nyquist wavenumber.

The discrete temporal domain consists in a finite sequence of time points $\{t_n\}_{n=1}^N$, uniformly spaced by Δt .

The choice of the computational grid parameters $L_x, L_y, L_z, \Delta x, \Delta y_j, \Delta x, \Delta t$ depends on both the flow physics and the numerical methods adopted as discussed in greater detail in the next paragraphs.

a. Description of the solver

Algorithm 1 is a pseudo-code representation of Incompact3D's time-advancement loop. The Incompact3D Fortran subroutine names corresponding to each algorithmic step are also reported as comments.

Before entering the time-advancement loop (Step 0), the algorithm reads in an initial condition, which can consist of a laminar velocity field perturbed by noise (used to produce a turbulent state) or a previously-generated turbulent flow field (when e.g. resuming a simulation or applying new forcing/BCs). The random perturbation \mathbf{u}_{noise} is defined as:

$$\mathbf{u}_{noise} = sA \exp\left(-\frac{(y-h)^2}{5}\right) \quad (2.3)$$

where s is a uniformly distributed random vector in $[0, 1]$ and A is the noise intensity. The above expression modulates a white-noise perturbation onto a Gaussian profile. The components of the velocity-gradient tensor are also initialized with the same perturbation.

In Step 1, the forcing \mathbf{F}^{n+1} is applied to the previous-step velocity field. For the reference channel flow problem (1.4), different forcing strategies are possible: the two most common in drag-reduction studies are the constant-flow-rate and the constant-pressure-gradient forcing. In the latter case, the pressure gradient is constant at every time-step and equal to the Poiseuille pressure gradient. When a constant flow rate is required, a uniform streamwise pressure gradient is applied at every time-step, which in non-dimensional form reads:

$$\mathbf{G}^{n+1} = \left\{ \frac{Q^n - Q}{\Delta t}, 0, 0 \right\}, \quad (2.4)$$

where Δt is the time-step size, Q^n is the average flow rate of the previous-step flow field and $Q = 4/3$ is the target Couette flow rate.

Step 2 calculates the spatial derivatives of the momentum equation. The terms \mathbf{H}^n in Step 2 contains all the convective terms in their skew-symmetric form:

$$H_i^n = -\frac{1}{2} \left(\frac{\partial}{\partial x_i} (u_{ij}^n u_{ij}^n) + u_j^n \frac{\partial u_i^n}{\partial x_j} \right), \quad (2.5)$$

where the repeated-index notation is used. The skew-symmetric form is used because of its energy-preserving properties in presence of aliasing errors (Kravchenko and Moin, 1997). The term \mathbf{F}^n groups the all the compact-finite-difference discrete form of the spatial terms and therefore a big-o error term is also shown.

b. The fractional-step method

Incompact3D uses as a variant of the fractional-step method as the time-advancement algorithm of the Navier-Stokes Equations. The fractional-step method (Chorin, 1968) addresses the problem of enforcing the incompressible continuity condition which is an instantaneous, kinematic

¹Several variants of the code for different types of flows exist and are available at the main code repository <https://github.com/xcompact3d>, including compressible and free-surface flows.

constraint and therefore cannot be naively decoupled from the momentum equation. To achieve this, the fractional-step method splits a time step into two parts: the first evolves the velocity field without preserving incompressibility, the second restores incompressibility by applying a pressure correction.

The four steps that are marked with a star in Algorithm 1 correspond the fractional-step time-advancement, where the actual physics is computed. The remaining steps are preliminary, intermediate, or I/O operations. Following Algorithm 1, the structure of the fractional-step method reads:

- 3★: integrates the associated compressible problem i.e. the time- and space-discrete approximation of the momentum equation together with the Dirichlet boundary conditions of (1.4). There is no pressure gradient and the resulting $\nabla \mathbf{u}^{**} \neq 0$
- 4★: corrects the Dirichlet boundary conditions in order to prevent the build-up of a $\mathcal{O}(\Delta t^\alpha)$ error at the solid boundaries, where α is the order of the time-advancement scheme (in the present case, $\alpha = 2$).
- 6★: solves the Poisson problem for the numerical pressure π . With respect to the numerical solution, this is the trickiest step: more background and discussion on the numerical treatment of pressure is given in Section e..
- 8★: subtracts the correction $-\Delta t \nabla p^{n+1}$ from \mathbf{u}^{**} to obtain the incompressible solution \mathbf{u}^{n+1} . This step is usually called “projection step” because mathematically, it projects \mathbf{u}^{**} into a divergence-free functional space by subtracting its component.

It can be shown that the pressure in a fractional-step method converges as $\mathcal{O}(\Delta t/Re_p)$.

The next paragraphs examine more closely the numerical properties of the temporal and spatial discretisation and the Poisson problem for the pressure.

Algorithm 1 Incompact3D's time-advancement loop.

```

0:  $\mathbf{u}^{n*} = \mathbf{U}_p + \text{noise}$  [read  $\mathbf{u}^{n*}$ ]      ▷ initial condition init() [or restart simulation with restart()]

for  $n \leq N$  do

1:  $\frac{\mathbf{u}^n - \mathbf{u}^{n*}}{\Delta t} = \mathbf{G}^{n+1}$                 ▷ apply pressure gradient channel()

2:  $\mathbf{F}^n + \mathcal{O}(\Delta x^6) = -\mathbf{H}^n + \frac{1}{Re} \Delta \mathbf{u}^n$         ▷ spatial derivatives convdiff()

3★:  $\frac{\mathbf{u}^* - \mathbf{u}^n}{\Delta t} = \frac{3}{2}\mathbf{F}^n - \frac{1}{2}\mathbf{F}^{n-1} + \text{B.C.}$         ▷ time advancement intt()

4★:  $\frac{\mathbf{u}^{**} - \mathbf{u}^*}{\Delta t} = \nabla p^n$                 ▷ pressure correction for BC pre_correc()

5:  $\nabla \cdot \mathbf{u}^{**}$                                 ▷ calculate divergence divergence()

6★:  $\Delta p^{n+1} = \frac{1}{\Delta t} \nabla \cdot \mathbf{u}^{**}$         ▷ solve poisson decomp_2d_poisson_stg()

7:  $\nabla p^{n+1}$                                     ▷ calculate pressure gradient gradp()

8★:  $\frac{\mathbf{u}^{n+1} - \mathbf{u}^{**}}{\Delta t} = -\nabla p^{n+1}$         ▷ projection step corgp()

9: optional output to disc.                    ▷ VISU_INSTA() and others

end for

```

c. Temporal discretisation

The time-advancement is fully explicit and uses a 2nd-order Adams-Bashforth scheme, which is a two-step, linear time advancement scheme. Implicit methods are generally uncompetitive for DNS because of the small time-step sizes that are needed to resolve the small scales and sustain a turbulent state (Moin and Mahesh, 1998). Given a general first-order system $u'(t) = f(t)$, its discrete AB2 time advancement reads:

$$u^{n+1} = u^n + \Delta t \left(\frac{3}{2}f^n - \frac{1}{2}f^{n-1} \right) \quad (2.6)$$

which applied to the semi-discretised momentum equations gives (3★) step in Algorithm 1. The choice of the time-step size Δt must ensure that *i* the time-advancement scheme is absolutely stable (Quarteroni et al., 2010), *ii* the smallest physical time-scales are adequately resolved and *iii* the computational cost is optimised. In order to meet the condition *iii*, the time-step must be as large as possible while ensuring that conditions *i*) and *ii*) are still satisfied. such as the non-linear Runge-Kutta scheme which Runge-Kutta schemes allow larger time-step sizes. but are more computationally intensive, therefore at the small time-step sizes needed for a fully-resolved DNS, AB2 smaller computational costs.

For the simpler case of periodic BC in all directions, the limits for absolute stability of the compact 6-th order scheme 2.8 coupled with the AB2 time advancement are given in Lele (1992) for the case of a purely advection and diffusion equations, respectively:

$$\Delta t \leq \frac{\Delta x}{U_c} \frac{\sigma_i}{1.989}; \quad \Delta t \leq \frac{\Delta x^2}{\nu} \frac{\sigma_r}{6.857} \quad (2.7)$$

where σ_i is the CFL number and σ_r is the Péclet number. The CFL (convective) condition dominates. These formulas are used to make first guesses on Δt , based on the Δy_{min} and considering $\sigma_i = 0.1$, but the final value of Δt is obtained heuristically by running a number of short simulations and seeking the smallest possible stable time-step size (see Table 2.2 for the values used in the reference channel case (1.4)).

d. Spatial derivatives

The discretisation of the spatial derivatives is obtained through compact finite differences schemes (Laizet and Li, 2011; Lele, 1992). In the following, we discuss the motivation for the use of such schemes and illustrate their definition and numerical properties (Lele, 1992). In general, four characteristics of a spatial discretisation scheme have to be considered for turbulent flow simulation:

- (a) *formal order of accuracy*: this corresponds to the order of the truncation error of the finite approximation.
- (b) *spectral resolution*: this requirement is a critical for DNS, where all the physically relevant small scales have to be resolved by the computational grid. The main reason of the widespread use of spectral methods in DNS codes is that they allow the maximum resolving efficiency (down to the Nyquist frequency of the grid) without numerical dissipation. Given equal spatial resolution, a finite difference scheme always has lower resolution in the high wavenumber range, due to its inherently local nature that acts as a low-pass filter. However, for the same reason FDs also suffer of less severe aliasing errors than spectral methods.
- (c) *versatility* to adapt to different/complex boundary conditions. Spectral methods present constraint on the geometry, mostly periodic or semi-periodic. Finite differences instead allow the use of more complex geometry and boundary conditions.
- (d) *computational cost*: due to the non-linear nature of the convective terms, finite differences are computationally cheaper than (pseudo-)spectral methods. The spectral approach is very efficient for linear problems such as the Poisson equation, and is therefore used in the treatment of pressure.

Compact finite differences offer many advantages with respect to the above points. Compact FD are capable of higher spectral resolution than normal FD schemes while also ensuring very low dispersive errors. At the same time, they maintain the flexibility of the FD schemes to adapt to complex boundary conditions. Incompact3D uses the following central compact scheme for the first derivatives:

$$\alpha f'_{i+1} + f'_i + \alpha f'_{i-1} = a \frac{f_{i+1} - f_{i-1}}{2\Delta x} + b \frac{f_{i+2} - f_{i-2}}{4\Delta x}, \quad (2.8)$$

where $\alpha = 1/3$, $a = 14/9$ and $b = 1/9$. The highest-order term of the truncation error for the above approximation is $4f^{(7)}\Delta x^6/7!$, making it formally sixth-order accurate while using only five stencil points. The implicit relation 2.8 constitutes a tri-diagonal system $\mathbf{A}\mathbf{f}' = \mathbf{b}$ that can be efficiently inverted using $\mathcal{O}(N)$ algorithms. On a non-periodic boundary (e.g. a solid wall), the wall-normal derivative can't be evaluated through a centred scheme, therefore the following one-sided, third-order accurate scheme is used:

$$f'_1 + 2f'_2 = \frac{-5f_1 + 4f_2 + f_3}{2\Delta x}. \quad (2.9)$$

On the first point above the boundary f_{i-2} is not available and a three-point, fourth-order scheme is used. Similar relations to (2.8) and (2.9) can be derived for second derivatives that preserve the same sixth and third order of accuracy respectively.

The spectral resolution of several finite-difference schemes can be assessed (Lele, 1992) through the normalised wavenumber:

$$w_l = \frac{2\pi l \Delta x}{L_x} = \frac{2\pi l}{n_x}, \quad -n_x/2 < l < n_x/2 - 1. \quad (2.10)$$

Using w_l , the spectral resolution of a first-derivative FD scheme can be evaluated by considering the Fourier series expansion of a periodic function f and its derivative in spectral space. The exact Fourier coefficients of the first derivative $i w_l \hat{f}_l$ are compared to those given by the Fourier series of its FD spectral approximation, that can be written as $i w'_l \hat{f}_l$. The graphs in Fig. 2.1 show $w'(w)$ for first-derivative schemes, including the exact spectral differentiation (a straight line) and the FD approximation (2.8).

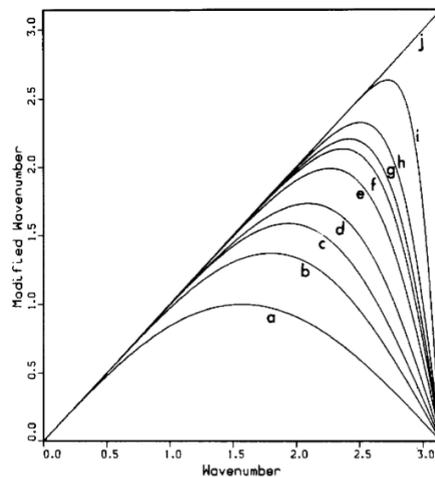


Figure 2.1: Plot of the modified wavenumber of several first-derivative finite difference approximations from Lele (1992). The curve (a) is the standard second-order central difference, the curve (c) a standard sixth-order central difference and the curve (e) is the sixth-order compact scheme (2.8).

e. Pressure treatment

The pressure p is stored on a mesh that is staggered by one half grid spacing with respect to the velocity mesh in order to avoid the odd-even decoupling phenomenon. The odd-even decoupling is a classical issue in computational fluid mechanics, that arises when central differences are used on a collocated grid (i.e. both p and \mathbf{u} on the same grid). When solving the Poisson equation on a collocated grid, velocity and pressure are decoupled and spurious oscillations (the “checker-board effect”) may arise, leading to instabilities. Computing the pressure gradient on a staggered grid greatly reduces this effect. Solving the pressure on a staggered grid entails two interpolation operations, one before solving the Poisson equation and one after computing the pressure gradient to perform the velocity correction on the collocated grid. Interpolation is through a that also requires solving to a tri-diagonal system.

The Poisson equation in Step 6★ is solved in spectral space, which is computationally convenient as it is a linear problem. The actual solution of the Poisson equation only entails one

division per Fourier mode to be solved in spectral space:

$$\widehat{p}_{lmn}^{n+1} = \frac{\widehat{D}_{lmn}}{F_{lmn}} \quad (2.11)$$

where \widehat{D} is the Fourier transform of $\nabla \cdot \mathbf{u}^{**}$ and F are coefficients that depend on the spatial discretisation and the boundary conditions and the indexes l, m, n are the discrete points of the staggered grid.

The entire procedure entails: 1) interpolating $\nabla \mathbf{u}^n$ (Step 5) on the staggered grid, 2) calculating its 3D Fourier transform, 3) solve (2.11) and 4) calculate inverse FT of π 5) interpolate π on the collocated grid. The spectral approach based on 3D Fourier transforms is perfectly adapted to periodic boundary conditions, however it has to be modified in case of the Neumann BC $\nabla \pi_y = 0$ that is used on solid walls². In that case, a Cosine transform is used along the wall-normal direction. The use of this technique introduces an additional error, reducing the spatial accuracy on the boundary to $\mathcal{O}(\Delta y^2)$.

The Discrete Fourier Transforms are computed using a Fast Fourier Transform (FFT) algorithm, with various FFT implementations available in Incompact3D (see Section e.). FFTs are algorithms that reduce the computational complexity of computing a discrete Fourier transform to $\mathcal{O}(n \log n)$ floating-point (additions and multiplications) operations (Ferguson Jr., 1982), where $N = n_x \times n_y \times n_z$ is the total number of grid points. The FFTs remain the single most computationally intensive step of the whole algorithm.

f. Immersed Boundary Method

The Immersed Boundary Method (IBM) is a numerical technique used to model a solid body that is located inside the fluid domain (Parnaudeau et al., 2004). The term ‘‘immersed boundary’’ means that the solid surface is not modelled through a conventional boundary condition but through the specification of a local velocity field \mathbf{u}_0 that results in no-slip and/or no-penetration conditions on the desired surface within the fluid. In the IBM terminology, \mathbf{u}_0 is called ‘‘target velocity’’. The target velocity is realised by modifying the discrete integration of the Navier-Stokes equations (Step 3 in Algorithm 1) as:

$$\frac{\mathbf{u}^* - \mathbf{u}^n}{\Delta t} = \frac{3}{2} \mathbf{F}^n - \frac{1}{2} \mathbf{F}^{n-1} - \nabla p^n + \tilde{\mathbf{f}}^* \quad (2.12)$$

The ad-hoc forcing term $\tilde{\mathbf{f}}^*$ is specified by:

$$\tilde{\mathbf{f}}^* = e \left(-\frac{3}{2} \mathbf{F}^n + \frac{1}{2} \mathbf{F}^{n-1} + \nabla p^n + \frac{\mathbf{u}_0^* - \mathbf{u}^n}{\Delta t} \right) \quad (2.13)$$

where the mask function $e(\mathbf{x})$ is $e = 1$ at the grid points in the interior of the solid region and $e = 0$ outside. The simplest choice of a target velocity is to set $\mathbf{u}_0 = \mathbf{0}$. The result of this procedure is to obtain that in the solid region $\mathbf{u} = \mathbf{0}$ to same order of accuracy of the time advancement.

The choice of the simple forcing $\mathbf{u}_0 = \mathbf{0}$ generates discontinuities in the velocity gradients when high-order finite differences are used (Parnaudeau et al., 2004).

2.1.2 High-performance computing

The term high-performance computing is used to describe the hardware and software technologies developed to solve large-size computational problems as rapidly as possible. The computing

paradigm also bears important consequences for energy efficiency of the computation (Zwart, 2020). The size of a computational problem can be defined by the number of elementary arithmetical operations that are required for its solution. This leads to the definition of Floating Point Operations per Second (Flop/s) as a measure of computational performance for scientific applications. The amount of volatile memory and the amount of data read or written to storage memory are also critical factors to determine the performance of a machine in executing a given program. A problem such as the turbulent channel flow outlined above entails a number of elementary operations of the order of 10^{15} . In the following paragraphs, we outline some basic ideas of digital computer performance and discuss the realisation of high-performance computing systems.

The execution time ET of a program (ORNL, 1995) is the product of the number of instructions to execute n_i , the clock cycles per instruction CPI, and the time per clock cycle t_c :

$$ET = n_i \times \text{CPI} \times t_c \quad (2.14)$$

The last two factors depend entirely on the design of the CPU. t_c is lower in smaller circuits, that are made possible by advancements in material science and manufacturing techniques and also depend on the chip organisation. The CPI count is a matter of engineering at the micro-architecture level, depending on the instruction set architecture of the CPU and the adoption of techniques such as instruction-level parallelism and instruction pipelining (ORNL, 1995). Since scientific calculations are generally performed on general-purpose, commercial CPUs, the critical factor to computational performance is typically the reduction of n_i , which can be achieved through more efficient algorithms, better compilers and through parallelisation. For well-known computational problems, highly optimised numerical algorithms are already available that reduce the number of operations required to solve a problem of a given size (for example, the FFT algorithms discussed in Sec. 2.1.1).

However, very large-size problems such as channel flow DNS would still take tens of thousands of hours to be solved on a single processor, even using state-of-the-art chips capable of very low CPI and t_c . Therefore, a significant reduction of ET for large-size problems can often only be achieved by parallelisation, that consists in splitting the program into many independent parts which are executed on many processors concurrently, resulting in a lower n_i per each individual processor. The speed-up obtained when executing a program requiring N operations on P processors is measured as:

$$S_P = \frac{ET_{\text{serial}}}{ET_P}. \quad (2.15)$$

where the serial execution time ET_{serial} corresponds to (2.14) with $n_i = N$ and the parallel execution time E_P is:

$$ET_P = ET_{\text{comp}} + ET_{\text{comm}} + ET_{\text{sync}} + ET_{\text{I/O}} \quad (2.16)$$

where the computation time ET_{comp} is (2.14) with $n_i = N/P$, plus the time needed to execute the parts of the program that cannot be parallelised. The other two terms account for the time spent in exchanging data between the different processors (ET_{comm}) and the time spent synchronising the processors when required by the algorithm (ET_{sync}). Moreover, the time $ET_{\text{I/O}}$ spent performing I/O operations is to be accounted, and depends on the data transfer speed of the hardware. In large-size problems $ET_{\text{I/O}}$ can represent a large fraction of the total computation time. Under ideal circumstances (no communication overheads and fully parallel program) the speed-up equals the number of processors, but in real-world applications $S_P < P$. In order to better understand these performance concepts, we briefly consider the practical design of parallel computers.

a. Distributed-memory parallelism

The most convenient way of building a highly scalable parallel computer is the so-called distributed-memory architecture, schematized in Figure 2.2a. A distributed-memory machine consists of many (typically identical) independent computers, called nodes, that are interconnected in a network through which makes possible the exchange of data. Each node sees the others uniquely through input/output signals sent over the network following an appropriate communication protocol. This architecture can be scaled up simply by building a larger network with more nodes.

The major performance bottleneck of this architecture is due to the latency occurring in the exchange of data from one node's memory to another, corresponding to ET_{comm} in (2.16). To reduce latency, several design aspects have to be considered, the major ones being the network topology, the interconnect technology, and the software protocols used to organise the exchange of data between the nodes, called message passing standards. Specialist, high-bandwidth ($> 10\text{Gbit/s}$) interconnect technologies are necessary for the largest supercomputers with tens of thousands of processors. Examples of network topologies are the simple ring, fully connected, butterfly, or the dragonfly topology used in the Archer supercomputer. In most modern topologies the nodes are connected hierarchically to reduce the total number of connections.

In the next section we are going to examine more closely the message passing standards, which is the most relevant aspect for developers and users of parallel scientific applications such as Incompact3D.

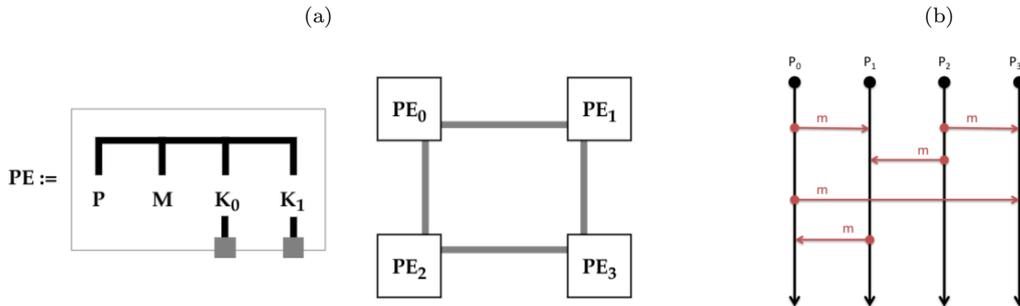


Figure 2.2: Distributed-memory architecture. (a) schematic of a 4-node distributed-memory machine (from ORNL, 1995): P = processor, M = memory, K = network interfaces, PE = node, black line = local interconnect (BUS), grey line = network interconnect. The interconnect topology is of the ring type. (b) schematic of a message-passing computation on 4 processes from EPCC, 2016-2020.

b. Message Passing Interface

The execution of computations on a distributed-memory machine requires a message passing protocol to manage data transfer to and from remote memory locations. Figure 2.2b shows a program executed in parallel on four processors. During the computation, the machines exchange data (symbolised by the horizontal arrows) at different points in time according to what required by the code.

The Message Passing Interface (MPI) is a specification that defines a communication model for distributed-memory computers. MPI is the *de-facto* industry standard in parallel computing, allowing the development of portable applications. Several libraries exist that implement the MPI standard for a variety of programming languages including Fortran and C, such as OpenMPI or MPICH. MPI defines a set of processes, that typically correspond to the physical cores being used to execute the program. The workload is distributed among these processes, that can only

operate on local variables in their memory space. Processes can not access each other’s memory spaces, but communicate data by passing messages over the network between a send and a receive buffer. Regardless of any form of memory sharing such as multi-core or multi-socket systems, MPI considers each core as a fully independent process, communicating via message passing. Messages can be passed between two individual processes, in what are called point-to-point communications (such as those in Fig. 2.2b), or can take place between several processes simultaneously, which is the case of collective communications. Moreover, data from several processes can be reduced to a single result by performing a global mathematical operation. MPI subroutines are called at the appropriate points in the program source code, meaning that each communication must be explicitly defined by the programmer.

An example of an MPI subroutine is the global sum reduction that is used to implement the coupled boundary conditions in Chapter 4. What follows is an example of its syntax with minimal context:

```
USE MPI           !loads the MPI library
MPI_INIT()
...
MPI_ALLREDUCE(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, COMM, IERROR)
...
```

The subroutine `MPI_INIT` initialises MPI and creates the communicator `COMM`, which defines a communication environment shared among all the processes. `COMM` groups the N processes together in a global space and orders them from 0 to $N - 1$, a number known as the process “rank”. `SENDBUF` and `RECVBUF` contain respectively the local variables that have to be sent and their destination in the remote machine’s memory. In this case, `MPI_ALLREDUCE` is a collective, all-to-all communication, therefore a send and a receive buffer exist on all the processes in `COMM`. `OP` is an object that specifies the type of global operation to be performed e.g. `MPI_SUM`. `COUNT` specifies how many elements of type `DATATYPE` are in the send buffer. The function `MPI_ALLREDUCE` is a synchronous operation, meaning that messages are passed and the sender process stays in stand-by until the operation has completed. Synchronous communications affect the term ET_{sync} in (2.16).

c. HPC systems

The parallel simulations were carried out using two specialised systems of different computational capability. The “Archer” system is the UK’s national (Tier-1) supercomputer located at the University of Edinburgh, while the Sheffield Advanced Research Computer “ShARC” is the local (Tier-3) HPC cluster at the University of Sheffield. The main specifications of the two architectures are listed in Tab 2.1.

ShARC was used to perform preliminary or moderate-size simulations, typically requiring resources for about 1000 core-hours. Archer was used for the majority of the computationally demanding cases, with the median case requiring around 6,000 core-hours and largest cases more than 10,000 core-hours. Access to the ShARC compute nodes is unlimited for UoS researchers, while access to Archer is granted through fixed-budget calls issued every six months by the UK Turbulence Consortium (<https://www.ukturbulence.co.uk/>). Both systems adopt the PBS job scheduler (<https://www.openpbs.org/>) to manage submission queues.

Table 2.1: Specifications of the two HPC systems used to perform the simulations presented in this work.

	Archer	ShARC
Machine	Cray XC30	Dell PowerEdge C6320
CPU model	Intel Xeon E5-2697 v2	Intel Xeon E5-2630 v3
Clock frequency	2.7 GHz	2.40 GHz
No. nodes	4920	98
CPUs per node	2	2
Cores per CPU	12	8
Total no. cores	118,080	1,568
Memory per node	64 GB	64 GB 1866MHz DDR4
NUMA domains per node	2	2
Interconnect	Cray Aries	Intel Omni-Path
Storage technology	SAS 12.0Gbit/s HDDs	SATA 6.0Gbit/s HDDs
Benchmark performance	1,642.54 TFlop/s ¹	44.88 TFlop/s ²
Operating System	Cray Linux	CentOS

The Archer system is currently being replaced by its successor Archer2, which is due to enter service in late 2020 and will have 748,544 total cores and a peak performance of around 28,000 TFlops/s.

d. Data management

The DNS simulations produced a significant amount of data, Data in excess of 20TB were produced over the course of the development of this work. The data have been stored on two facilities: the EPSRC Research Data Facility at the University of Edinburgh (<http://www.rdf.ac.uk/>) and the local Research storage at the University of Sheffield (<https://www.sheffield.ac.uk/it-services/research-storage>). The RDF facility provides a flexible amount of storage space while Sheffield’s Research storage allocates 10TB. Both facilities provide fast data access to the HPC clusters (Archer and ShARC respectively) and they provide some form of data redundancy. The main advantage of storing data at Sheffield’s local cluster consists in the reduced data processing latencies and remote data transfer times, as well as increased flexibility in data management.

Data are prepared for long-term storage by collecting them into tarball archives, using LZMA compression to reduce their size, and splitting them into multiple smaller files. Remote data transfer were performed through the `rsync` utility using the Secure Shell (SSH).

e. *Incompact3D*’s parallel strategy

Incompact3D exhibits very good (linear) scalability on massively parallel systems (Lai² and Li, 2011), achieved by using the 2DECOMP&FFT library (Lai² and Li (2010); <http://www.2decomp.org>). 2DECOMP&FFT implements a 2D domain decomposition strategy for spatial derivatives on Cartesian meshes, parallel Fast Fourier Transform and parallel I/O subroutines and is based on MPI.

To compute spatial derivatives, the 3D mesh is subdivided in $p_{row} \times p_{col}$ blocks as specified by the user, each of which is assigned to a separate MPI process. The blocks are organized in

¹LINPACK benchmark retrieved at <https://www.top500.org/system/178188>

²Double-precision matrix multiplication benchmark retrieved at <https://rse.shef.ac.uk/blog/iceberg-vs-sharc/>

a $p_{row} \times p_{col}$ 2D grid, so that each processor is assigned a block in a 2D array which spans the whole third direction of the domain. This strategy makes it possible to compute the compact finite difference along the third direction of each block, a task that would not be possible with a 3D decomposition because all the points along the direction of the derivative are necessary in an implicit formulation. For example, Fig. 2.3a shows the decomposition of a cubic domain in a 4×4 grid along y and z , leaving it uninterrupted along x so that each processor can compute the x -gradients autonomously. Each block is assigned to a different processor and once the finite difference has been calculated in parallel, the whole subdivision process is repeated using $p_{row} \times p_{col}$ blocks but spanning the y direction. This operation involves a transposing the domain across different processors, which is a very communication-intensive process, much more so than a halo-cell 3D decomposition. This methodology can require a large number of transpose operations, 24 being necessary to compute the convective and diffusive terms and 16 to compute the divergence. The transpositions are always performed sequentially i.e. from the x -pencils to the y -pencils to z -pencils in order to reduce the number of global communication operations. Despite the large number of communication operations, the high scalability properties of the software makes this strategy convenient on thousands of compute cores (Laizet and Li, 2011). The choice of $p_{row} \times p_{col}$ is also critical for computational performance and the best practice outlined in Laizet and Li (2011) recommends that $\max(p_{row}, p_{col}) = \min(n_x, n_y, n_z)$ and $p_{row} = p_{col}$.

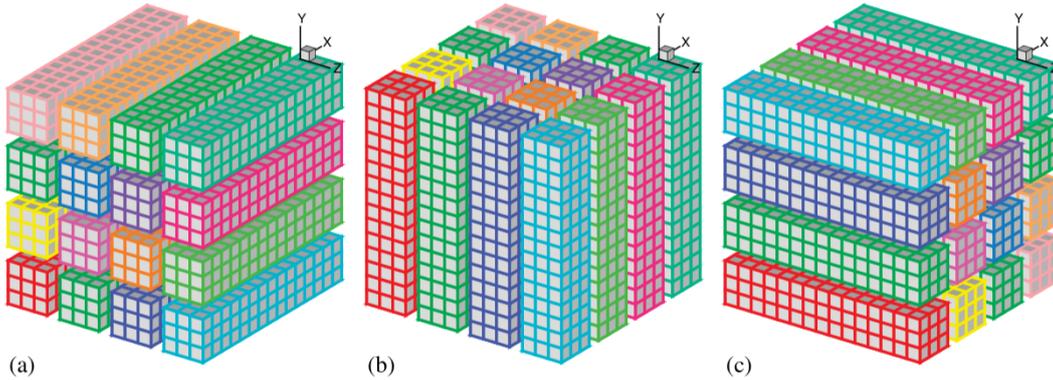


Figure 2.3: . 2D decomposition of a computational domain into 16 parallel sub-domains, from Laizet and Li (2011): (a) x -pencils, (b) y -pencils, (c) z -pencils. Each colour identifies an individual MPI process.

2DECOMP&FFT includes parallel ports of several Fortran FFT engines through a 1D decomposition strategy of the domain. To perform a 3D FFT, 2 domain transpositions are used, and the same holds for the inverse FFT. This brings the total number of communication operations for a periodic channel flow to 61. The FFT implementation can be specified by the user at compile time: in this work, we used MKL-FFTW3 on ShARC and Cray-FFTW on Archer, which are the versions of the well-known FFTW3 library (Frigo and Johnson, 2005) optimised for running on each machine respectively. The native machine-optimised compilers are used where possible, namely Intel compilers on ShARC and the Cray compilers on Archer. Performance improvements can be sizeable: compiling the code on ShARC using Intel compilers and MKL-FFTW3 implementation on ShARC reduces ET of 10% with respect to using `gfortran` compiler and the generic FFT – the average execution time per time-step is respectively 0.076sec and 0.085sec when run on 64 computational cores with a 8×8 block decomposition.

Incompact3D also makes use of the 2DECOMP&FFT parallel I/O subroutines that make possible for multiple processes to write and read to and from a single file (EPCC, 2016-2020).

2.2 Reference channel flow

This section illustrates the numerical solution of the flow problem (1.4). The numerical parameters used as input to Incompact3D to realise the flow are found and motivated in 2.2.1. A number of flow statistics is presented in Section 2.2.4. Relevant aspects of the flow physics are highlighted in the text as the data are presented.

2.2.1 Numerical set-up

In a Direct Numerical Simulation, all the spatial and temporal scales relevant to achieve transition to turbulence and sustain the fully-developed turbulent state must be resolved. We recall that wall turbulence has three length-scales, i.e. the viscous length-scale δ_ν^* , the Kolmogorov scale $\eta = (\nu^3/\varepsilon^*)^{1/4}$ (which is around $\eta = 1.5\delta_\nu^*$ at the wall (Pope, 2000)) and the outer scale h^* (Smits et al., 2011). Ideally, the computational grid must be designed in order to resolve the viscous sub-layer (a few δ_ν^* thick), the smallest-scale and the largest-scale turbulent structures. The longest time scale is of the order of the largest-eddy turnover time h/u_τ , while the shortest time-scales correspond to a few viscous time units δ_ν/u_τ (Quadrio and Luchini, 2003). While it is important to ensure the resolution of the physical time-scales, the time-step size is most often dictated by numerical stability reasons when explicit methods are used.

Table 2.2 below lists the simulation parameters used to simulate the channel flow using Incompact3D.

The simulations are performed at a nominal $Re_\tau = 180$. As the value of Re_τ is not known in advance, the value of Re_p needed to obtain Re_τ can be estimated through the empirical correlation $Re_\tau = 0.123Re_p^{0.875}$ (Pope, 2000).

Since the time-advancement scheme is explicit, the CFL stability condition must be satisfied. The approximate CFL formula (2.7) yields $\Delta t = 0.0012$, however we find by trial-and-error that the larger $\Delta t = 0.0025$ is still stable and is therefore used for these simulations.

The best practice (Pope, 2000) for DNS meshes suggests that the first near-wall point should be located below δ_ν and the grid spacing at the centre of the channel should be roughly equal to Δz . This allows to resolve the near-wall viscous layer and save grid points in the nearly-isotropic region at the centreline. The stretching parameter β and the grid count n_y must be tuned to yield a minimum spacing $\Delta y_{min}^+ < 1$ at the wall $\Delta y_{max}^+ < 10$ at the centre of the channel. The choice of $\beta = 0.25$ and $n_y = 129$ is satisfactory.

The chosen box dimensions L_x and L_z are identical to Moser et al. (1999) and are large enough to correctly reproduce the one-point statistics (Lozano-Durán and Jiménez, 2014), especially the wall-shear stress (de Giovanetti et al., 2016). The domain is smaller than the largest naturally occurring scale and thus expected to be unable to resolve VLSMs.

Once the box dimensions are fixed, the grid size n_x and n_z are chosen, according to good practice, such that $\Delta x^+ \leq 10$ and $\Delta z^+ \leq 5$. n_x and n_z must ideally be multiples of the number of processors in a compute node, so that each node is fully exploited and the workload is distributed evenly among all processors, thus avoiding communication bottlenecks. Moreover, choosing a grid count which are small multiples of a power of 2 is generally good practice to maximise the efficiency of the FFT algorithms.

The total integration time of the simulation T_{int} is the time interval over which, after reaching steady state, the statistics are calculated. A value of 40 eddy turnovers is suggested in the literature to produce well-converged statistics, here we pick $T_{int} = 43u_\tau/h = 1000U_p/h$. This means that a total of 400,000 time steps (given that $\Delta t = 0.0025$) are required to integrate the problem after discarding the initial transients.

A series of preliminary runs with a refined resolution (and time-step) have been carried out and achieving an almost neglectable change in statistics, which indicates consistency. For an

easier comparison with the data available in our group’s database, the numerical set-up in terms of time-integration scheme, time-step and spatial resolution or stretching has been only slightly altered from the reference channel simulation of [Khosh Aghdam \(2016\)](#).

Table 2.2: Simulation parameters for the direct numerical simulations of the reference channel flow. Δy_{max}^+ is the spatial resolution in the centre of the channel, in wall units.

Re_p	nom. Re_τ	L_x	L_z	$n_x \times n_y \times n_z$	Δx^+	Δz^+	Δy_{min}^+	Δy_{max}^+	Δt^+	$T_{int} u_\tau / h$
4200	180	4π	$4\pi/3$	$256 \times 129 \times 128$	8.8	5.9	0.9	7.8	0.04	43

The initial-condition noise intensity A in (2.3) is set at 12% of U_p . The channel is also subjected to a rotation around the z axis for the initial 5000 time-steps, realised by introducing ([Kristoffersen and Andersson, 1993](#)) in the momentum equation the following forcing term:

$$- Ro \mathbf{k} \times \mathbf{u} \quad (2.17)$$

where \mathbf{k} is the span-wise unit vector and $Ro = 2\omega^* h^* / U_p^*$ is the Rossby number, a non-dimensional parameter that specifies the ratio of the mean-shear vorticity to the background imposed vorticity ω^* . In this simulation, the rotation intensity is chosen by setting $Ro = 1/9$. The rotation is necessary to obtain the turbulent transitions from a noisy initial condition, otherwise the initial noise would dissipate and the flow relaminarise.

2.2.2 Computational performance

The numerical problem defined by the parameters of Table 2.2 is specified on $N = 4, 227, 072$ discrete grid points and entails the execution of more than 400,000 time-steps. In Incompact3D, its solution requires around 2×10^9 floating-point operations per time step or 10^{15} in total. The total number of parallel communication operations is of the order of 10^6 .

The simulation is run on the ShARC system (Tab. 2.1) using 64 individual cores on 4 compute nodes. The domain is decomposed into 8×8 blocks, as an equal number of columns and rows is generally recommended to maximise performance (see Sec. e.).

The problem takes a total of 8 hours and 26 minutes to compute, equivalent to an average 0.07 seconds per time-step and an average computational performance of around 32 GFlop/s. The volatile memory usage amounts to 6GB. The storage memory requirement is 67GB, needed to write to disc 2000 double-precision arrays (4 scalar arrays every 1000 time-steps), each of size $256 \times 129 \times 128 = 4227072$.

2.2.3 Verification and validation of computational simulations

In general, performing numerical simulations of a physical system entails a verification step and a validation step. Borrowing [Roy \(2005\)](#) definitions, “verification deals with mathematics and addresses the correctness of the numerical solution to a given model. Validation, on the other hand, deals with physics and addresses the appropriateness of the model in reproducing experimental data.”

a. Code verification

The first step entails code verification, that consists in satisfying the requirements of convergence and consistency of the numerical algorithm and the correctness of the code with respect to the underlying continuous mathematical model. Code verification is a crucial requirement that

code developers address carefully. Generally, code verification must be carried out once and for all integrated into the code design phase. For the Incompact3d code, the convergence and consistency of the numerical schemes, parallel scaling and performance are illustrated in [Laizet and Lamballais \(2009\)](#) and [Laizet and Li \(2011\)](#).

b. Solution verification

A second verification step deals with verifying that the code can correctly reproduce a known exact solution. This might not be the case even when the code is verified due to its failure to accurately resolve some property of the solution or the presence of numerical instabilities. Therefore solution verification must be carried out for each specific case as it depends on the physics to be simulated. Solution verification may entails adjusting numerical parameters such as the spatial discretisation, the time-step size, the degree of the interpolation schemes, implicit solver iterations and many others. In many practical cases, such as turbulent flow, an exact benchmark is not available therefore techniques exist to assess the numerical error in terms of grid refinement, such as the Richardson extrapolation. More detail on grid refinement in Incompact3D, including data for turbulent channel flow, can be found in [Laizet and Lamballais \(2009\)](#) and [Khosh Aghdam \(2016\)](#). Section 2.2 showcases with validation by benchmarking against literature fixed-wall data.

c. Solution validation

Solution validation refers to comparison of the exact solution against experimental data in order to evaluate the ability of the mathematical and numerical models and their code implementation to correctly simulate all the relevant physics. Explicit validation against experimental data is, for the reference flows presented in this work, avoided because the numerical data of [Moser et al. \(1999\)](#) are shown in the original and later publications to be in excellent agreement with channel-flow experimental data. Since no experimental results exist for the novel drag-reduction methods developed in this work, the appropriateness of the numerical time and spatial discretisation to resolve the physics have to be generally compared to closely related methods such as the spanwise waves.

2.2.4 Results

a. Convergence to the statistically steady state

The flow develops from the perturbed laminar flow to the statistically stationary, fully developed turbulent state. The noise level and the rotation parameters are chosen, ideally, in order to reach the stationary turbulent state in a reasonably small amount of time-steps and therefore save computer time to collect the steady-state statistics. The stationarity of the flow is assessed by visually monitoring several flow variables time-series³ and verifying the flow satisfies exact steady-state equilibrium.

Figure 2.4a shows the time evolution of three quantities, the skin-friction coefficient, the total kinetic energy of the flow and the distance of the wall-normal mean velocity $V(y)$ from its exact value of zero. The latter is calculated as the L_2 norm of the wall-normal velocity profile as $\|\Delta V\|_2 = \|\langle v \rangle_{xz}(y)\|_2$. All three variables are normalised with the steady-state mean value.

Monitoring $V(y)$ is important to verify that the effects of the initial channel rotation have vanished. It can be observed that after around $t = 200$ or 80000 time-steps the flow has reached

³Several approximate yet rigorous stationarity tests for time series also exist e.g. [Dickey and Fuller \(1979\)](#), that may provide a useful tool for assessing such transient problems quantitatively.

the stationary turbulent state. Statistics are therefore accumulated and steady-state integration time started from $t = 200$.

The shear-stress profiles obtained by time and space averaging on the stationary interval are shown in Figure 2.4b, matching the exact linear functional form expected from (1.20). The graph is useful to visualize the separation between the viscous sub-layer and the upper inertia-dominated layers. The so-called buffer layer occurring between the two is located in proximity of the point y_b where the viscous stresses τ_{lam} and Reynolds stresses τ_{RSS} are equal, which is also where the peak production of turbulent kinetic energy occurs.

A number of other steady-state statistics are discussed in detail in the following sections and compared to exact and numerical results from the literature.

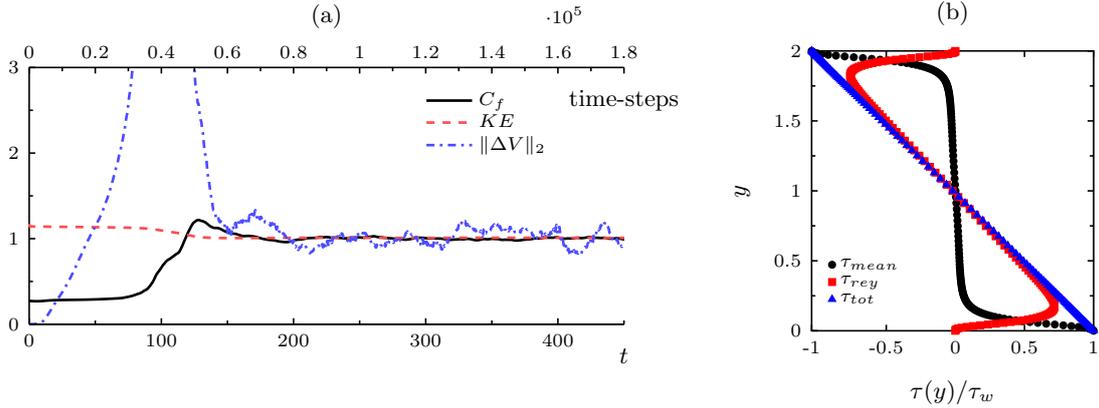


Figure 2.4: Convergence to a statistically stationary turbulent state. (a) evolution of several variables during the initial transition to turbulence. All variables are normalised with their final steady-state mean value. (b) viscous ($Re_p^{-1}dU(y)/dy$), turbulent ($-\langle \overline{uv} \rangle(y)$), and total ($\tau(y)$) mean shear-stress profiles.

b. Global parameters at steady state

Table 2.3 reports a list of the main integral parameters or functionals defined using the non-dimensional units used in the simulation. Since the equations are scaled in outer units (U_{Pois} , h), the definitions below can be seen as a the mapping between outer and inner units. The inner units, wall-units or plus-units are defined based on τ_w and represent the characteristic scale of dynamical variables in the near-wall region.

Table 2.3: Global functionals as computed from the simulation data.

Name	Symbol	Definition	Mean value
Wall-shear stress	τ_w	$\langle \partial_y U \rangle / Re_p$	1.810×10^{-3}
Skin-friction coefficient	C_f	$2\tau_w / U_b^2$	8.123×10^{-3}
Friction velocity	u_τ	$\sqrt{\tau_w Re_p}$	4.248×10^{-2}
Friction Reynolds number	Re_τ	$u_\tau Re_p$	1.785×10^2
Viscous time	t_ν	Re_p / Re_τ^2	1.319×10^{-1}
Bulk velocity	U_b	$(1/4) \int_0^2 U dy$	6.667×10^{-1}
Bulk Reynolds number	Re_b	$Re_\tau U_b / u_\tau$	5.600×10^3

c. *Local statistics*

In Fig. 2.5 the mean streamwise velocity and the six components of the Reynolds-stress tensor are plotted. The mean-velocity profile and the components of the velocity-correlation tensor are found to be in excellent agreement with the DNS data of Moser et al. (1999) which are shown therein to be in accordance with experimental evidence. A linear viscous sub-layer can be observed below $y^+ \approx 5$ and a region with logarithmic dependence on y^+ is obtained between around $30 < y^+ < 100$.

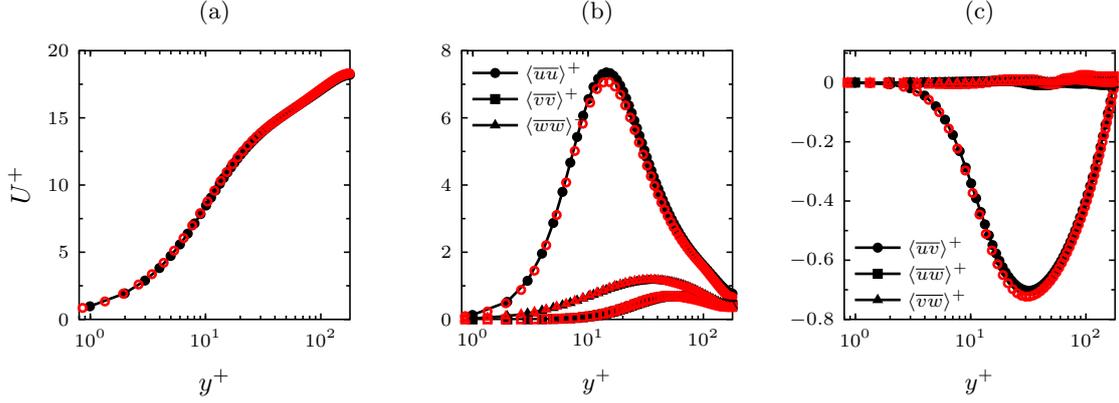


Figure 2.5: Reference channel results: velocity and Reynolds-stress statistics along y . The simulation data are plotted against Moser et al. (1999) (empty red symbols).

d. *Kinetic Energy budget*

The study of the turbulent kinetic energy $k = \langle \mathbf{u} \cdot \mathbf{u} \rangle / 2$ has an important role in turbulent flow research. When turbulence is statistically inhomogeneous, such as over walls, studying the TKE becomes important because its production, transport and dissipation become also inhomogeneous and their spatial distribution provide insight on the physics occurring under to the influence of the wall and the boundary conditions.

The evolution equation for the TKE can be derived from the Reynolds-averaged momentum equations (Pope, 2000) and reads:

$$\frac{Dk}{Dt} + \nabla \cdot \mathbf{T} = \mathcal{P} - \varepsilon, \quad (2.18)$$

where $\mathcal{P} = \langle uv \rangle \partial_y U$ is the power extracted locally by the Reynolds stresses from the mean flow (called the *production* term) and $\varepsilon = 2 \langle \partial_i u_j \partial_j u_i \rangle / Re$ is the viscous *dissipation* of TKE in the small-scale fluctuations. The term \mathbf{T} is a vector representing the non-equilibrium *turbulent transport* of TKE:

$$T_i = \frac{1}{2} \langle u_i u_j u_j \rangle + \langle u_i p \rangle - 2 \langle u_j s_{ij} \rangle / Re_p \quad (2.19)$$

the first term is the advection of TKE by the fluctuating velocity, the second is the pressure transport and the third one is the diffusion. The term-by-term budget is reported in Fig. 2.6a, showing good overlap with the validation data. The plot blows up on the near-wall portion, as all the terms decay in the centre of the channel, with \mathcal{P} and ε being slowest to decrease. The ratio of the production to the dissipation term is plotted in Fig. 2.6b. The region where the curve

flattens to a nearly constant $\mathcal{P}/\varepsilon \approx 0.9$ corresponds to the logarithmic layer, between about $30 < y^+ < 100$, where almost all k is produced and dissipated locally and spatial transport of k is reduced. The imperfect energy equilibrium in the logarithmic layer is due to the low Reynolds number of the flow, whereas at asymptotically high Reynolds number $\mathcal{P}/\varepsilon = 1$ is expected. The peak TKE production is located in the buffer layer at $y_b^+ \approx 10$ where the viscous and Reynolds shear stresses are equal (Fig. 2.4b). The maximum production can be calculated exactly as $\mathcal{P}_{max} = \frac{1}{4} (\tau_{y_b}/\tau_w)^2$.

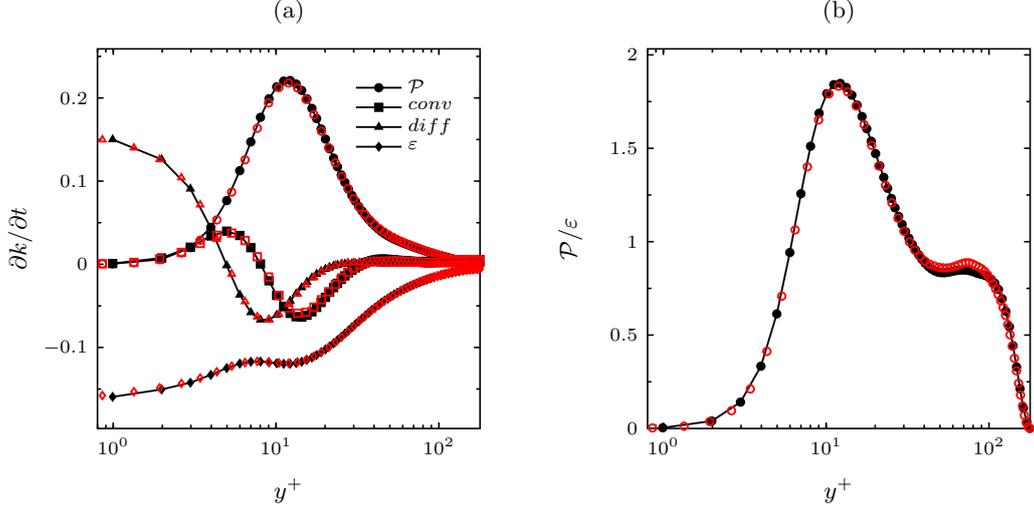


Figure 2.6: Turbulent kinetic energy and related statistics. Current data are compared to the data of Moser et al. (1999) (empty red symbols). (a) TKE budgets from the TKE equation (2.18). (b) Production to dissipation ratio.

e. Two-point statistics

In this paragraph, we briefly present some two-point statistics, that are statistics of flow variables calculated at two spatially separated locations, as opposed to the “one-point” statistics discussed above. Two-point statistics are used to reveal the spatial structure of the velocity fluctuations that characterise turbulent flows. The most basic two-point statistics is the two-point covariance matrix of the velocity fluctuations:

$$C_{ij}(\delta x, \delta z; y) = \overline{\langle u_i(x, y, z) u_j(x + \delta x, y, z + \delta z) \rangle}, \quad (2.20)$$

that goes back to the foundational work of Kolmogorov (1941). C_{ij} allows the study of the spatial structure of the velocity fluctuations in terms of spatial separation δ_x and δ_z . This tool can also be extended to spatio-temporal correlations which are important in near-wall turbulence for the discussion of topics such as the Taylor approximation (Zhao and He, 2009).

Fig. 2.7a shows that the one-dimensional streamwise correlation shown in are in good agreement with the data of Moser et al. (1999) at different wall-normal heights y .

The Wiener-Khinchin theorem (Chorin and Hald, 2013) states the spatial information carried by a two-point auto-covariance C_{ii} corresponds, after a Fourier transform, to the power spectral density. For example, the following Fourier transform pair holds for the streamwise kinetic energy $\overline{\langle u^2 \rangle}/2$:

$$E_{uu}(\kappa) = \mathcal{F}_{\delta x} \{ C_{uu}(\delta x) \}(\kappa), \quad (2.21)$$

where E_{uu} is the streamwise-velocity PSD, κ_x is the streamwise wavenumber and $\mathcal{F}_{\delta x}$ is the forward Fourier transform operator. In order to visualise the scale-wise distribution of the kinetic energy, the PSD is plotted on semi-logarithmic axes and pre-multiplied by the wavenumber to preserve area proportionality. Figure 2.7b shows the pre-multiplied streamwise TKE spectra E_{uu} at two different wall-normal locations, proving good agreement with the literature data and demonstrating the good spectral resolution of the compact FD schemes. Due to aliasing errors, spectral codes often report a cusp at high wavenumbers (Vreman and Kuerten, 2014), which is artificially cut off in the reference data of Moser et al. (1999) but is not present in the Incompact3D spectra. Pre-multiplication is also useful to visualise the spectra of other energy-related statistics, for example the PSD of the dissipation ε is $E_\varepsilon(\kappa) \propto \kappa^2 E_{ii}(\kappa)$ (Smits et al., 2011) because multiplication by the wavenumber in Fourier space correspond to differentiation.

In Figure 2.7c the entire wall-normal domain is represented as a contour plot. It can be noticed that the spectrum is incomplete at low wavenumbers due to the streamwise extension of the computational domain, that is not long enough to resolve the largest wavelengths. The cut-off coincides with the smallest wavenumber resolved within the computational box i.e. $\kappa_x = 4\pi/L_x$. The spectral peak (\times) lies at around $y^+ = 11$ and $\kappa_x^+ = 0.0056$, i.e. a wavelength $\lambda_x^+ = 1100$, in agreement with the literature data on the streamwise length of the near-wall high-velocity streaks.

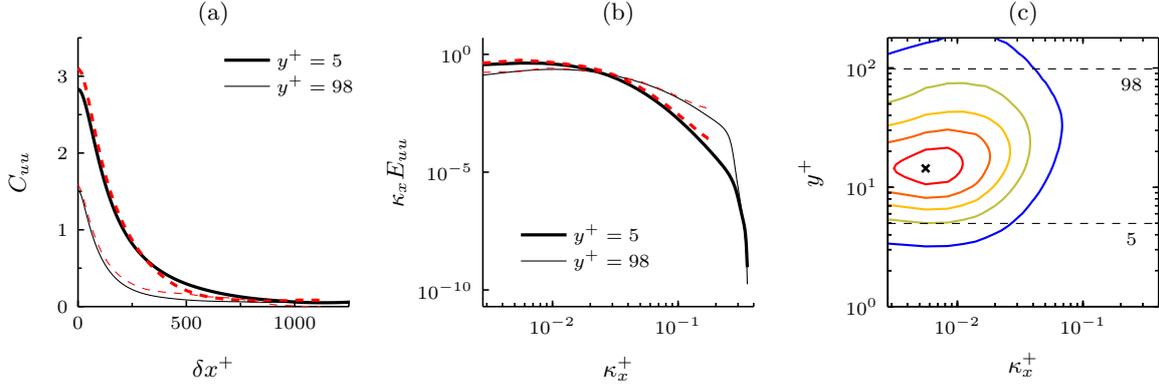


Figure 2.7: Reference channel: two-point statistics. The present simulation data (—) are compared to Moser et al. (1999) (---). (a) Longitudinal auto-covariance functions C_{uu}^+ of the streamwise turbulent fluctuations. (b) pre-multiplied power spectral densities E_{uu}^+ of the streamwise turbulent fluctuations at the same wall-normal locations as (a), normalised with their maximum. (c) contour map of $\kappa_x E_{uu}(\kappa_x, y)$.

The numerical spectral estimation is performed using a Fortran program that first computes the auto-correlation function using the naive estimator (von Storch and Zwiers, 2001) and then takes its Fourier transform using 2DECOMP&FFT's Glassman algorithm. The program can compute both 1D and 2D correlations and spectra: an example of the 2D spectrum can be seen in Fig. 4.16 applied to the wall-shear stress. A better alternative to this approach, especially for shorter 1D time-series, is provided by using the Welch estimator for 1D time-series, as done in Section 4.4.

Other statistics of the plane reference channel flow are shown in the next chapters and applied to the research problems. For example, fine-scale vortical structures are shown in Section 3.3.2 and their identification techniques are reviewed in Appendix 3.A. Another useful statistical tool are the integral identities for the skin-friction budget, that are also explained in Chapter 2.

2.3 Statistical methods

A large portion of this work relies on estimating the moments of unknown probability distributions from the finite amount of data that is produced by the numerical simulations. In this section, we review some basic concepts and good practices in statistical estimation that have been used throughout the work to produce reliable statistics from numerical data. The basic concepts are recalled in Section 2.3.1, criteria for sampling are described in b. and b. presents the uncertainty quantification models.

2.3.1 Statistical estimation

An estimator is a statistics on some property of a random variable, i.e. a method used to guess that property from a set of N observations or data, called the sample. In DNS of turbulent flows, the objects of study are random functions of time and space, therefore the sample is typically a time-series, a discrete spatial array or simply a collection of repeated experiments depending on which statistical symmetries are present. Typically, the property we are interested in is some moment μ of a probability distribution of one or more flow variables.

An estimator M for a quantity μ must converge to its true expected value as the sample size tends to infinity, according to the Law of large numbers. This is equivalent to requiring that the mean squared error $MSE = E[(M - \mu)^2]$ goes to zero. The MSE of the estimator M is given by its bias B and its variance, according to the decomposition $MSE(M; \mu) = B^2(M) + \text{Var}(M)$, where the bias is defined as $B(M) = E[M] - \mu$. For a random variable X with exact mean μ and variance σ^2 , the estimators M for the mean and for the variance s are given by the familiar formulas:

$$\begin{aligned} M_N &= \frac{1}{N} \sum_{n=1}^N w_i X_n \\ s_N &= \frac{1}{N-1} \sum_{n=1}^N w_i (X_n - M_N)^2 \end{aligned} \tag{2.22}$$

where N is the sample size and $w_i = 1$ for uniformly spaced samples. Non-unitary weights are used to account for sampling non-uniformity, such as in the case of the wall-normal direction where $w_i = \Delta y_i$. The estimators (2.22) are unbiased, with $B = 0$ for any N . The square root of the variance estimator is called “root mean square” and is used extensively in the text. The remainder of the MSE is given by $\text{Var}(M)$ i.e. the uncertainty, which depends on the distribution of X and the sample size. The uncertainty is discussed at length in Section 2.3.2.

Throughout this work, a number of other estimators are used for covariance (e.g. in Section 2.2), probability distributions (e.g. in Section 4.4), two-point correlations (e.g. in Section 4.4) and power spectral densities (e.g. in Section 4.4). Whenever a new statistics is introduced, an appropriate description of the estimator used will be provided.

a. Numerical estimation

The numerical aspect of statistical estimation is also of great importance (Chan et al., 1983) for performance and stability. In DNS the computational cost of moving large files in and out of memory can be significant and is therefore desirable that it be reduced. Welford’s algorithm computes the mean and variance estimators (2.22) accessing each data point only one time and ensuring good numerical properties (Chan et al., 1983): General one-pass algorithms for arbitrary higher-order moments also exist, see Pebay (2008).

b. Optimal time sampling

Run-time data storage affects $ET_{I/O}$ slowing down the simulation and leads to high disc space utilisation, especially if large double-precision arrays are stored very frequently. It is therefore important to reduce data writing to disk to the essential.

In this work, most of the data processing has been postponed after the simulation and only the three velocity components and the pressure are stored to disc. Spatial gradients and other costly spatial operations have been executed in post-processing through the same high-order compact finite difference schemes of the solver.

Another aspect that has to be carefully considered is the sampling rate, which ideally has to be reduced as much as possible, bringing the write-out period \mathcal{T}_{out} to the minimum without compromising the quality of the statistics. Since a turbulent time-series is always correlated in time, sampling at a very small T_{out} is not convenient because it does not contribute to the information content of the sample. The minimum T_{out} can be understood as either: a) the smallest interval of time for which the shortest-correlated flow variable is uncorrelated, b) the smallest interval of time for which the longest-correlated flow variable is uncorrelated. The first is the best choice if one wants to minimize uncertainty for the shortest-correlated variables, and has been used here for time-series obtained through high-sampling-rate probing of certain variables such as the wall gradient or the unsteady disc velocity in Chapter 4. Choice b) is more reasonable for whole-field snapshot sampling of a stationary flow, because it ensures a similar quality to the statistics of flow variables while minimising $ET_{I/O}$ and disc utilisation. These scales can be determined in a preliminary study upon reaching the statistically steady state. We refer to Zhao and He (2009) who carried out DNS investigations of the space-time correlations in channel flows at $Re_\tau = 180$. The integral time-scale grows closer to the wall and reaches values around $30 t^+$, which is taken as the sampling interval. The same value of $30t^+$ is calculated from the reference channel simulations of Section 2.2. The time-step expressed in wall units is $\Delta t^+ = 0.04$, therefore saving to disk every ~ 760 time-steps in order to work with well uncorrelated data, however for practical reasons the sampling is performed every 1000 time-steps which ensures even more independence.

2.3.2 Uncertainty quantification

When performing a DNS of a turbulent flow, it is highly beneficial to have an estimate of the uncertainty on the statistics of interest. Uncertainties in DNS data stem from the finite numerical accuracy and from sampling errors (Roy, 2005). Numerical errors are inherent to the spatial and temporal discretisation, while sampling uncertainty depends on the sampling frequency for the statistics and the total integration time. The total steady-state integration time T_{int} of the simulation shall be long enough to collect good-quality statistics but also as short as possible in face of limited computing resources. As evident in the definitions (2.22), the choice of the time-averaging interval and the number of samples affect the accuracy of the estimator. It is of interest to give bounds on the uncertainty of a statistics, in the form of confidence intervals.

We discuss here the sampling uncertainty quantification procedure for the mean of the spatially-averaged wall shear-stress $\langle \tau_w \rangle(t)$ as an example. The confidence interval for an estimator M of some statistics of X is defined by its upper and lower bounds M_L and M_U , such that:

$$Pr[M_L < \mu < M_U] = \tilde{P}, \quad (2.23)$$

where μ is the true value and \tilde{P} is a given probability. Unless otherwise stated, we will assume $\tilde{P} = 0.95$.

The most basic approach to confidence intervals is based on the Central Limit Theorem (CLT), which requires the sample to be identically distributed and independent (i.i.d.) Under

the i.i.d. assumption, then the 95% confidence interval on the mean is $(M - c_L SE_N, M + c_U SE_N)$. SE_N is the Standard Error for the sample size N which is defined as:

$$SE_N = \frac{s_N}{\sqrt{N}} \quad (2.24)$$

where N is the number of samples and s_N is the estimator of the standard deviation of τ_w calculated as in (2.22). The constants c_L and c_U depend on the distribution of the underlying variable X and on the statistics that we are trying to calculate. For example, if X is normally distributed and M is the estimator for the mean, then c_L and c_U are determined from Student's t -distribution. If the quantity being estimated is the variance, the χ^2 distribution is used. If X is not normally distributed, then more general methods have to be used such as bootstrapping.

More references on the handling of uncertainty in DNS are provided in [Oliver et al. \(2014\)](#).

a. Sampling uncertainty in DNS time-series

The theory explained until here relies on the assumption of an i.i.d. variable. The assumption of identical distribution can be justified for a stationary turbulent channel flow both in time and across different realizations of the same flow. However, any variable measured in a turbulent flow displays some degree of (spatio-)temporal correlation, invalidating the assumption of independence. Down-sampling or batching a turbulent time-series are two ways in which an uncorrelated sample can be obtained.

Down-sampling means sampling at a large sampling period or discarding a certain number of data points in order for the newly obtained time-series to be less correlated. The most obvious disadvantage of this approach is the loss of information, especially if the correlation time of the original time series is of the same order of the total sampling time.

Batching methods instead are conceived to avoid such losses while still yielding an uncorrelated sample. The idea of this class of methods is to process a correlated time-series in batches and then use the partial batch-wise statistics to assemble an uncorrelated sample – see e.g. [Russo and Luchini \(2017\)](#) and references therein.

The effect of correlation on the accuracy of a statistics μ can be quantified by defining an equivalent sample size n' such that:

$$\text{Var}(M) = \frac{\text{Var}(X)}{n'}. \quad (2.25)$$

When X is an uncorrelated process, (white noise) trivially $n' = n$. For red-shifted processes (such as the velocity fluctuations) it turns out that $n' < n$, implying a larger statistical uncertainty than an uncorrelated noise with equal σ_X . Therefore, neglecting the correlation would result in underestimating the uncertainty. An equivalent definition is that of ‘decorrelation time’ τ_D which is defined asymptotically:

$$\tau_D = \lim_{n \rightarrow \infty} \frac{n}{n'} \quad (2.26)$$

Given a finite sample of a discrete process with auto-correlation function C , the decorrelation time can be estimated as:

$$\tau_D = 1 + 2 \sum_{k=1}^{n-1} \left(1 - \frac{k}{n}\right) C(k). \quad (2.27)$$

Therefore, the problem of estimating n' (or τ_D) consists in finding a good approximation of the true auto-correlation function $C(k)$. Direct estimation of C from the data is not recommended because suffers from both finite-time bias and variance errors ([Thiébaux and Zwiers, 1984](#)). In

this work we follow [Oliver et al. \(2014\)](#) and we use an alternative approach, namely fitting an auto-regressive (AR) model to the data and using its exact correlation function as the estimate of C . An auto-regressive process of order p , denoted as $AR(p)$, is a discrete-time process with a linear dependence on the previous p steps:

$$X_t = \alpha_0 + \sum_{i=1}^p \alpha_i X_{t-i} + W_t. \quad (2.28)$$

where W_t is a Gaussian white noise term of variance σ_W ([von Storch and Zwiers, 2001](#)). An $AR(p)$ model is therefore specified by the choice of the parameters $\{\sigma_W, \alpha_1, \alpha_2, \dots, \alpha_p\}$. It can be proven that an $AR(p)$ approximation to any weakly stationary process exists and is unique ([von Storch and Zwiers, 2001](#)). It must be pointed out that time-series of turbulent flow variables are not an AR process because, in general, they are not Gaussian and have a non-linear inter-temporal dependence. The AR -based uncertainty quantification procedure follows the steps outlined in [Algorithm 2](#). The process order is selected by maximising the Akaike information criterion (AIC). The correlation function is calculated using Wilson's algorithm ([Tunnicliffe Wilson, 1979](#)), implemented in own code. The procedure is implemented in Python, using the statistical functions of the library SciPy that includes some AR modelling classes.

Algorithm 2 $AR(p)$ based estimation of the uncertainty on the mean in a correlated time-series ([von Storch and Zwiers, 2001](#)).

- 1: select the order p
 - 2: fit the parameters to the simulation data using Maximum Likelihood
 - 3: calculate the correlation function C using Wilson's algorithm
 - 4: substitute the exact correlation function into [\(2.27\)](#)
-

Figure 2.8 shows the validation of the AR-based uncertainty quantification method through its application to a synthetic $AR(2)$ process. Among the methods compared in [Fig. 2.8](#), the AR fitting appears to be the best for large sample sizes, of the same order of those available from the DNS. For smaller samples, of size closer to the decorrelation time, its accuracy is worse than using the direct correlation estimator.

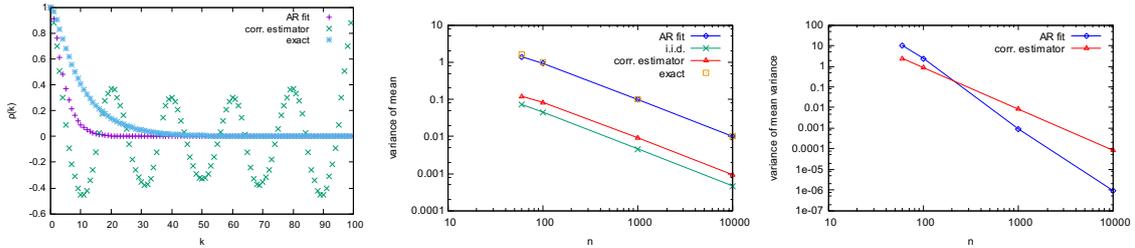


Figure 2.8: Estimation of a simulated $AR(2)$ process with zero exact mean, exact variance $\sigma_X^2 = 4.63$ and exact $\tau_D = 21.62$. a) Estimated and exact auto-correlation function of the same process. The estimates are made from a single realization of a 100-element time-series. The direct estimator is truncated at $n/2$. b) Convergence of the variance of the sample mean. c) convergence of the variance of the sample variance. Comparison between three different methods and the exact case. The values are the average of 1000 Monte-Carlo realizations of each sample size.

b. Application to τ_w

We now provide an estimate of the mean wall-gradient from the $Re_\tau = 180$ reference channel, based on the theory above. The data consist of a time series of the spatially averaged gradient $d\langle u \rangle(t_i)/dy|_{y=0}$ sampled at each time-step for 100000 time-steps, adding up to a total integration time of $250h^*/U_p^*$. Our model does not account for spatial correlation, which is more complex to assess and requires spatial correlation models.

Four levels of down-sampling are applied to the time series, from 1/1000 (the typical sampling rate for e.g. Chapter 3 simulations) to the 1/1 full resolution. The 95% confidence intervals on the mean value $dU/dy|_{y=0}$ are calculated for each sample according to both the uncorrelated model and the AR-based correlated model.

Table 2.4 the convergence the $\bar{\tau}_w$, showing that correlation can significantly affect the level of uncertainty on the wall-shear stress statistics. It also shows the sampling interval, sample size N , the AR order and the decorrelation lag τ_D . It can be noted that for large sampling period, the uncertainties on the correlated and the iid estimates are much closer than for higher sampling rates, which is explained by the increasing sample correlation. Moreover, the iid-based uncertainty decreases following an exact $N^{-1/2}$ dependence, while the AR-based one is much more stable.

Table 2.4: Statistical estimation of the mean wall gradient $dU/dy(y=0)$ from a correlated time-series.

sampling Δt	N	p	τ_D	mean	95% CI (AR)	95% CI (iid)
2.5	100	2	4.538	7.613	0.067	0.031
0.25	1000	6	64.90	7.613	0.082	0.010
0.025	10000	32	807.3	7.612	0.092	0.003
0.0025	100000	67	10980	7.612	0.109	0.001

2.3.3 Post-processing software

The statistics were computed using different software, selected based criteria such as ease of use and the availability of existing code and libraries for the type of post-processing task. Post-processing Fortran code has been used mostly when dealing with reading and processing the large raw flow field data into statistics. We have developed (mostly serial) fully autonomous in-house post-processing framework leveraging Incompact3D's I/O routines and made use of the existing post-processing code available since the 2018 Incompact3D update.

Nevertheless, we aimed towards a more extensive use of Python for scripting and post-processing DNS data. This tool has a number of advantages, such as *i*) the availability of a large number of open-source numerical and statistical libraries, *ii*) the relative simplicity of modifying and tailoring the code to a specific case and *iii*) the integration of the entire process from raw data to visualisation into a single framework.

Data visualisation has also been carried out through a diverse set of software. The Python library Matplotlib (<https://matplotlib.org/>), and the open-source software Gnuplot have been employed for quickly visualising time-series and flow fields. ParaView (<https://www.paraview.org/>) is a powerful open-source tool that has been essential for quickly inspecting 3D spatial data, observe and plan higher quality visualisations. The latex package PFGplots (<http://pgfplots.sourceforge.net/>) has been used for producing publication-quality vector graphs. Flow animations have been produced with Matplotlib's `animation` class.

Chapter 3

Open-loop turbulent drag reduction combining rotating actuators and other techniques

3.1 Introduction

Flush-mounted spinning discs have been proved to be a promising technique for turbulent drag reduction, yielding up to around 23% drag reduction at low Reynolds numbers (Ricco and Hahn, 2013; Wise and Ricco, 2014; Wise et al., 2014). Wise et al. (2014) showed that annular actuators, or “rings”, with an optimal inner-to-outer radius ratio equal to 0.6 deliver a slightly larger drag reduction than the full discs, while requiring up to 20% less driving power. The rings used by Wise et al. (2014) covered 50% of the wall area, while the other half of the surface was stationary. Although the rings generate a maximum drag reduction of about 20%, thus smaller than the uniform spanwise-wall oscillations, the net power saved, computed by accounting for the frictional resistance of the fluid, is comparable for the two methods, about 10%. Advantages of the rings with respect to the in-plane motions are the localized actuation, the double optimal spanwise forcing scale, and the wide range of diameters leading to drag reductions that are comparable with the maximum (Ricco and Hahn, 2013; Wise et al., 2014).

In the present work, we study the spinning rings in detail and we utilize the stationary wall surface confined between the rings and within the circles defining the inner-ring edges to enhance the drag-reduction performance by enforcing two types of distributed control. The word “distributed” is used herein to indicate that the control scale matches that of the near-wall turbulence, acting locally on the turbulent structures. This implies that the control performance is expected to scale proportionally to the actuated area. The two distributed control methods considered in this work are opposition control and hydrophobic surfaces.

Opposition control (OC) is an active control technique (Choi et al., 1994; Hammond et al., 1998) that works by detecting the instantaneous wall-normal velocity on a wall-parallel plane located at a height of y_d and, at the same time, by forcing the wall-normal wall velocity by the same amount of that detected in the y_d -plane, but with opposite sign:

$$v(x, 0, z, t) = -v(x, y_d, z, t). \quad (3.1)$$

This may result in drag reduction as a consequence of the attenuation of the sweeping motions of high-momentum fluid toward the wall and of the ejections from the near-wall region towards

the bulk of the flow. The detection plane acts as a virtual wall, which means that any convective transport of momentum between the plane and the physical wall is greatly hampered. The optimal distance of the detection plane is $y_d^+ \approx 15$ (Choi et al., 1994). Opposition Control is challenging to realise practically because it entails detecting and actively controlling small spatial and temporal scales.

The practical realisation of opposition control in wall turbulence (Kasagi et al., 2009) necessarily entails compromising with respect to the ideal numerical case. In the latter, sensing and actuation are both instantaneous and spatially co-located, whereas in a real-world design both control elements are necessarily lagged, sparse and occur at separate locations. Furthermore, the sensing of flow quantities must ideally occur at the wall rather than at a positive wall-normal height. This can be achieved by means of either wall-shear stress or pressure sensors, typically thermal or mechanical devices, whose measurements can be correlated to the wall-normal velocity sweeps. Many types of actuators have been devised in the experimental literature, such as synthetic jets driven by pistons (Rebbeck and Choi, 2001) or membranes (Rathnasingham and Breuer, 2003), flap devices (Jacobson and Reynolds, 1998) or plasma actuators (Moreau, 2007).

In academic wind-tunnel experiments (Rathnasingham and Breuer, 2003; Rebbeck and Choi, 2006), the actuation was provided by blowing/suction orifices or slits, and the sensing by hot-wire anemometers collocated upstream of the actuator (at a distance of $76\delta_\nu$ in Rebbeck and Choi (2006)). In real-world applications, an added layer of complexity is present due to fabrication, installation and concurrent operation of a significant number of micro-devices (given that the control performance is proportional to the controlled wall area).

The experimental evidence with suction/blowing devices demonstrated the viability of the approach to reduce turbulence intensity, such as in the feedforward-controlled Rathnasingham and Breuer (2003) who obtained respectively 30% turbulence intensity reduction and estimated 7% drag reduction.

This flow has been further examined by Chung and Talha (2011), who considered several detection planes and wall-forcing amplitudes. The dependence of drag reduction on the wall-normal location of the detection plane is linked to the relative position of this plane with respect to the near-wall eddies. If the plane is placed underneath the quasi-streamwise vortices, the OC actuation has a beneficial effect due to the disruption of the near-wall ejections and sweeps. If instead the plane is too distant from the wall, the OC actuation promotes the momentum transport and the drag increases. In our study the location of the detection plane has been fixed to $y^+ = 14.5$, scaled in terms of the wall units of the reference fixed-wall channel.

The drag reduction effect produced by the hydrophobic surfaces (HS) has been explained in terms of the shift of the virtual origin of the mean velocity profile (Min and Kim, 2004). The slip length that models the hydrophobic surface allows the streamwise velocity to take a finite value at the wall, which means that the point where the extrapolated velocity vanishes (the virtual origin) is shifted below the wall surface. This attenuates the intense sweeps and ejections typical of the no-slip wall layer. A spanwise wall slip instead promotes the growth of streamwise rolls leading to drag increase due to an enhanced lift-up effect. Both constant (Min and Kim, 2004) and shear-dependent (Aghdam and Ricco, 2016) slip-length wall boundary conditions have been proposed to describe the effect of different types of hydrophobic surfaces, modeling lotus-leaf-type surfaces, where air pockets are trapped in small cavities (Ling et al., 2017; Seo et al., 2018; Reholon and Ghaemi, 2018), and pitcher-plant-type surfaces, where oil is imbibed in the porous surface (Wong et al., 2011) or in transverse microcavities (Ge et al., 2018).

In the present work we follow the approach of Min and Kim (2004) where the no-slip boundary condition on the streamwise velocity is replaced by the following constant-slip-length condition:

$$u(x, 0, z, t) = l_s \left. \frac{\partial u}{\partial y} \right|_{y=0}. \quad (3.2)$$

l_s is chosen from the literature and fixed at $l_s^+ = 3.6$ which is the optimal value at $Re_\tau = 180$ found through direct numerical simulations by [Min and Kim \(2004\)](#).

A superhydrophobic surface ([Rothstein, 2010](#)) is made of an hydrophobic material and has a roughness texture that can span a size range from nanometers to micrometers. Once a superhydrophobic is submerged in water, gas bubbles become entrapped in the grooves of the surface roughness, effectively preventing most of the surface from direct contact with the liquid. This non-wetting state represent the ideal working state of superhydrophobic surfaces ([Cassie and Baxter, 1944](#)), because the solid-liquid interface is partially replaced by gas-liquid interfaces and thus a "lubrication" effect is achieved due to the dynamic viscosity of air being much lower than that of water. The resulting finite slip at the wall leads to reduced skin friction compared to conventional no-slip (smooth) walls.

Conversely, in the fully wetted state, where no gas bubbles are present and the liquid fills roughness elements, drag increase is typically produced. The fully wetted state can be reached when the unsteady shear stresses produced by a turbulent flow destabilise the bubbles, a phenomenon called bubble depletion. Bubble depletion represents one of the major technological issues for the wide-scale application of superhydrophobic surfaces to engineering problems and is the subject of active research aimed at increasing the wetted state robustness ([Seo et al., 2015](#)).

Experimental works in laminar flow ([Ou and Rothstein, 2005](#)) demonstrated slip lengths of $l_s = 25\mu\text{m}$ and drag reductions of more than 25%. A variety of experiments on turbulent flows over superhydrophobic surfaces have been carried out, consisting of either structured textures with regular arrays ([Daniello et al., 2009](#)), or randomly distributed textures ([Bidkar et al., 2014](#)). Drag reductions of up to 50% and slip lengths $l_s \approx 100\mu\text{m}$ were found by [Daniello et al. \(2009\)](#) thorough near-wall PIV measurement.

Riblets ([Garcia-Mayoral and Jiménez, 2011](#)) are one of the most well-studied passive drag-reduction method, consisting in covering the wall surface with a series of parallel grooves that run along the streamwise direction for the entire length of the channel. Riblets are largely inspired by the grooved surface of shark skin scales ([Dean and Bhushan, 2010](#); [Lloyd et al., 2021](#)).

Riblets are a geometric surface modification and thus require zero driving power, even though wall-wave methods lead to a theoretically higher overall power saving due to their superior performance in reducing friction. On the other hand, rotating discs entail the installation of actuators whose operation requires spending energy to overcome frictional losses (for more on the power budget of rotating discs, see [Section 4.3](#)).

Furthermore, wall waves, discs and other active and passive near-wall turbulence control methods can be analysed using the concept of equivalent roughness (or equivalent slip), meaning that their action is confined to the inner layer and they similarly induce a shift to the logarithmic velocity profile scaled in native wall units [Luchini \(2015\)](#). This theory can be leveraged to extrapolate the drag reduction performance of near-wall control methods to high Reynolds number regimes ([Gatti and Quadrio, 2016](#)). While the theory provides an unifying framework for bulk flow statistics far from the wall, it does not imply that their mechanisms of action on the near-wall turbulence are related in any obvious way. In fact, while the streak-modifying mechanism of riblets is based in the spanwise vortices from the wall ([Garcia-Mayoral and Jiménez, 2011](#)), for wall-waves and related methods it is rooted in a complex transient growth mechanism ([Blesbois et al., 2013](#)).

a. Objectives

The objectives of the present work are

- to further the understanding of physical mechanisms behind the drag-reduction effects induced by the spinning discs. In particular, we are interested in studying the streamwise-

elongated structures appearing between spanwise-adjacent discs that have been reported by Ricco and Hahn (2013), Wise and Ricco (2014), and Wise et al. (2014) to contribute significantly to the alteration of the near-wall momentum exchange;

- to enhance the drag-reduction performance given by the spinning rings by enforcing two different types of distributed control, i.e., opposition control based on wall-normal velocity actuation and hydrophobic surfaces. We aim to maintain the localized rotation that Ricco and Hahn (2013), Wise and Ricco (2014), and Wise et al. (2014) have shown to be effective in operating at large-flow scales and, at the same time, we want to evince whether the distributed control methods can weaken the elongated structures between rings that contribute detrimentally to drag reduction Wise et al. (2018). The original publications by Keefe (1997, 1998) do suggest to utilize the space between and within the rotating actuators to enforce wall-normal transpiration to fully specify a three-dimensional flow field. Keefe (1998) states: “Suction and blowing holes can be interspersed among, or made coaxial with, the disks for creating general three-dimensional velocity perturbations in the near-surface region.” An adaptation of the original sketch by Keefe (1998) is shown in Figure 1.3a, where the flow is actuated by rotating discs and wall transpiration;
- to make a case for the combined scheme as a strategy to improve the robustness of the individual methods. At the high Reynolds numbers typical of engineering boundary layers, super-hydrophobic surfaces notoriously suffer a progressive increase of drag, which originates from the depletion of the surface bubble layer by high shear (Aljallis et al., 2013; Bidkar et al., 2014). Opposition-control schemes have the well-known drawback that the ideal continuous coverage of the wall is hard and real-time detection of is not practical. Therefore localised pressure or spanwise shear-stress sensing is used and actuation systems have to be put in place that are often impractical. Furthermore, at large Reynolds numbers, increasing frequencies and smaller spatial scales imply that sophisticated devices are needed to be capable of a very fast response at miniature scale (Kasagi et al., 2009). This opens up the possibility that rotating actuators and wall-normal suction and blowing can prove mutually beneficial to each other, as employing steady forcing on (marginally) larger scales and using an open-loop, sensing-free control method can prove more practical than the localised, near-instantaneous actuation. Another option is to operate the combined scheme adaptively, activating discs or suction/blowing in order to improve performance or increase the method’s robustness to disturbances.
- for the first time, to develop and use an extended form of the Fukagata-Iwamoto-Kasagi identity (Fukagata et al., 2002) that can be employed to study active control strategies that are non-uniform in both the streamwise and spanwise directions.
- to advance a simple idealized model for predicting the upper bound of the drag reduction that can be achieved by combining the spinning-ring method and the distributed control actuations. This approach can be easily generalized to other combinations of drag-reduction techniques.

Section 3.2 describes the solver of the Navier-Stokes equations, the numerical procedures, the statistical tools, and the decomposition of the flow variables. Section 3.3 discusses the flow visualizations, the turbulence statistics, and the turbulent drag reduction. A summary of the results is presented in §3.4.

3.2 Numerical and statistical procedures

3.2.1 Numerical simulations

The combined effect of the spinning rings and the distributed forcing on wall turbulence is tested numerically in a channel with periodic boundary conditions along the streamwise direction x^* and the spanwise direction z^* . Dimensional quantities are henceforth denoted by the superscript $*$. The wall-normal coordinate is y^* and the parallel walls are separated by a distance $2h^*$. The numerical domain of the channel has dimensions L_x^* and L_z^* along x^* and z^* , respectively, and the velocity components are u^* , v^* , and w^* along x^* , y^* , and z^* , respectively. A ring is characterized by the outer diameter D^* and the inner diameter D_i^* . The shape factor $a = D_i^*/D^*$ is fixed for all the cases at the optimal value of 0.6 found by [Wise et al. \(2014\)](#). The portion of the wall surface that is not covered by the rings is subjected to the distributed actuation. For $a = 0.6$ this part of the surface amounts to 50% of the total wall surface. The ring motion is defined by the constant velocity W^* at the outer edge. Each ring has two $0.05D^*$ -wide gaps along the internal and external edges, where the velocity decays from the edge values to zero, following a linear profile. This greatly reduces the numerical oscillations arising from the sharp velocity discontinuity that would occur if the gaps were not present ([Ricco and Hahn, 2013](#)) and also models the absence of rigid wall between the rings and the stationary surface that would unavoidably occur in a laboratory set-up. Extensive resolution checks have been carried out to ensure the robustness and accuracy of the drag reduction values and of the turbulent statistics. The rings are fitted on both walls and form spanwise rows spanning the entire channel width, as shown in [Fig. 3.1](#). Rings belonging to the same row share the same sense of rotation, which is opposite to that of the next row. This arrangement creates triangular waves of spanwise velocity along the streamwise lines passing through the centers of the rings, thereby resembling the standing wall waves that have been shown to be very effective in reducing the turbulent drag ([Quadrio et al., 2009](#)). The periodicity of the computational box ensures that the flow field over a periodic array of rings is simulated correctly. Along the spanwise direction, the periodic BC ensures that the ring neighboring the domain limit faces the flow field produced by its analogous on the opposite side. The same is established along the streamwise direction, provided that the device array has an even number of rows along x .

The flow is simulated using the code Incompact3D ([Laizet and Li, 2011](#); [Laizet and Lamballais, 2009](#)), presented extensively in [Section 2.1.1](#). The simulations are run on the Cray XC30 ‘‘Archer’’ supercomputer of the National Supercomputing Service. The simulations are run using 1024 parallel computational cores and adopting a 32×32 block decomposition.

Quantities scaled in outer units, i.e., by the channel half-height h^* and the centerline velocity U_p^* of the laminar parabolic Poiseuille flow at the same flow rate, are not indicated by any symbol. The diameter and the tip outer velocity of the rings are kept constant to $D = 3.39$ and $W = 0.38$ to focus on the combined effect of the spinning rings and the distributed control. This pair of W and D leads to the near-optimum drag reduction of 19.5% for the full discs ([Ricco and Hahn, 2013](#)) which increases to 20% for the rings with $a = 0.6$. Scaled in the wall units of the reference channel, these parameters are $W^+ = 9$ and $D^+ = 604$. The fluid flows at the fixed volumetric flow rate per unit spanwise depth of $2h^*U_p^*/3$ in all the cases. This is maintained by applying a numerical correction to the mean velocity profile at every time step. The Poiseuille Reynolds number is $Re_p = U_p^*h^*/\nu^* = 4200$, where ν^* is the kinematic viscosity of the fluid.

The computations are initiated from the laminar Poiseuille channel flow between solid stationary walls, disturbed by random velocity fluctuations. This initial flow field evolves to the fully-developed uncontrolled turbulent flow, identified by the main statistics displaying convergence and the total mean stress profile being linear ([Orlandi, 2012](#)). This reference fully-developed

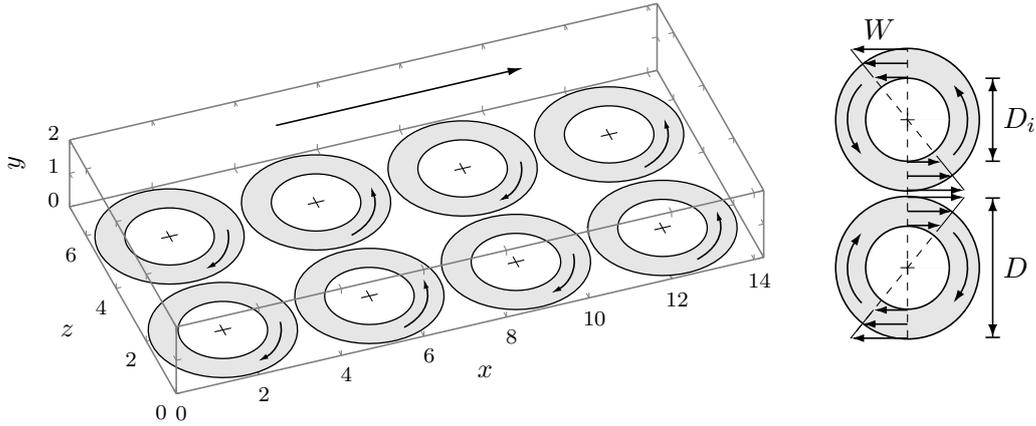


Figure 3.1: Schematic of the arrangement of the ring-fitted channel. Only the rings on the bottom wall are shown. The rings are shown in grey, while distributed actuation is applied in the white regions. The sense of rotation of the rings is highlighted by the curved arrows, while the longer arrow indicates the mean-flow direction. The sketch on the right illustrates the triangular-wave spanwise velocity pattern created along the centerline of the rings by two streamwise-adjacent counter-rotating rings.

turbulent flow is used as the initial flow for the computations of the controlled cases. The initial transient flow response to the modified boundary conditions is discarded to acquire meaningful statistics of the controlled flow. The flow statistics are collected by averaging well-delayed velocity and pressure flow fields at a temporal distance of about $t_d^+ = 20$ to reduce the statistical correlation between the flow fields. The total averaging time is $t_a = 950$, starting from the end of the transient flow.

Table 3.1 summarizes the main numerical parameters of the simulations. The wall no-slip Dirichlet boundary condition is used to impose the ring motion and the wall-transpiration boundary condition (3.1) is enforced for the OC simulations. When OC is applied to the entire wall surface through condition (3.1), the suction/blowing mean flow rate at the walls is zero, but in the combined ROC case a mean negative wall velocity $V_{oc}(y = 0)$ of the order of 10^{-4} exists if (3.1) is applied over the portion of the wall not covered by the rings. Above this controlled-wall region in the plane $y = y_d$ the wall-normal velocity has a positive mean value, which is mainly due to the wall-normal ejections associated with the streamwise-elongated structures existing between rings. It is clear from the continuity equation that, in order to maintain a constant mass flow rate at each streamwise location and at each time step, it is therefore not sufficient to impose (3.1) only, but it is also necessary to impose a blowing wall velocity $-V_{oc}(y = 0)$ over the surface not occupied by the rings to obtain $V = 0$ at the wall and at any wall-normal location. The Robin boundary condition at the wall (3.2) is applied to simulate the HS. We take the slip length to be $l_s = l_s^*/h^* = 0.02$, yielding a drag reduction of about 30% (Min and Kim, 2004) at low Reynolds numbers. This slip length corresponds to a realistic physical value $l_s^* = 100\mu\text{m}$ for a channel height of about 1 cm (Lauga and Stone, 2003; Choi and Kim, 2006; Bidkar et al., 2014). The numerical routines for the implementation of the combined control methods were written by Dr S. Khosh Aghdam for his Ph.D. research work (Khosh Aghdam, 2016), although his original routines did not enforce the zero-mass-flow-rate condition at the walls and the wall velocity adjustment along the inner circle of the rings. All the data presented herein are the output of novel simulations.

In table 3.1 and henceforth, the letter “R” indicates ring forcing, while “OC” and “HS” refer to the cases where the distributed control is applied on the entire wall surface. “R0” denotes cases

Table 3.1: Simulation parameters. The ring outer diameter D^+ and tip velocity W^+ are scaled in the wall units of the reference channel flow. The number of grid points are along x , y , and z , respectively. The Reynolds number is $Re_p = 4200$ and the inner-outer diameter ratio is $a = 0.6$ for all the cases.

Case	$L_x \times L_z$	grid points	D^+	W^+
Reference flow	$4\pi \times 1.33\pi$	$256 \times 129 \times 128$	-	-
Opposition control (OC)	$4\pi \times 1.33\pi$	$256 \times 129 \times 128$	-	-
Hydrophobic surface (HS)	$4\pi \times 1.33\pi$	$256 \times 129 \times 128$	-	-
Rings (R)	$4.53\pi \times 2.26\pi$	$256 \times 129 \times 256$	605	9
Rings + oppos. control (ROC)	$4.53\pi \times 2.26\pi$	$256 \times 129 \times 256$	605	9
Rings + hydroph. surface (RHS)	$4.53\pi \times 2.26\pi$	$256 \times 129 \times 256$	605	9
Hydroph. surface + fixed rings (HSR0)	$4.53\pi \times 2.26\pi$	$256 \times 129 \times 256$	605	0
Oppos. control + fixed rings (OCR0)	$4.53\pi \times 2.26\pi$	$256 \times 129 \times 256$	605	0

where the rings do not move and the distributed control is applied over the portion of the surface that is not covered by the rings. The letters ‘‘ROC’’ and ‘‘RHS’’ denote the combined cases of rings with opposition control and rings with hydrophobic wall, respectively. These notations are also used as subscripts to indicate that a quantity refers to that specific case.

3.2.2 Averaging operators

In this section the averaging operators are defined. The time average of a flow variable $f(x, y, z, t)$ over a time interval $[0, T]$ for which f is statistically stationary is defined as:

Time average

$$\bar{f}(x, y, z) = \frac{1}{T} \int_0^T f(x, y, z, t) dt. \quad (3.3)$$

The flow is statistically periodic along x with period $2D$ and along z with period D , that is, the minimal geometrical flow unit that repeats itself along the streamwise and spanwise directions consists of two streamwise-adjacent rings spinning in opposite directions. A spatial ensemble-averaging operator is therefore defined as:

Spatial ensemble average

$$[f]_e(x_e, y, z_e, t) = \frac{2}{N_x N_z} \sum_{n_x=0}^{N_x/2-1} \sum_{n_z=0}^{N_z-1} f(x + 2n_x D, y, z + n_z D, t), \quad (3.4)$$

where $0 \leq x_e \leq 2D$ and $0 \leq z_e \leq D$ are the ensemble spatial coordinates, while N_x and N_z are the number of rings in the numerical domain along x and z , respectively. The two streamwise-adjacent rings share a statistical mirror symmetry, which allows us to symmetry-average the ensemble-averaged flow quantities as follows:

Two-disc symmetry average

$$[u]_s(x_s, y, z_s, t) = \{[u]_e(x_e, y, z_e, t) + [u]_e(x_e + D, y, -z_e, t)\} / 2, \quad (3.5)$$

$$[v]_s(x_s, y, z_s, t) = \{[v]_e(x_e, y, z_e, t) + [v]_e(x_e + D, y, -z_e, t)\} / 2, \quad (3.6)$$

$$[w]_s(x_s, y, z_s, t) = \{[w]_e(x_e, y, z_e, t) - [w]_e(x_e + D, y, -z_e, t)\} / 2, \quad (3.7)$$

$$[p]_s(x_s, y, z_s, t) = \{[p]_e(x_e, y, z_e, t) + [p]_e(x_e + D, y, -z_e, t)\} / 2, \quad (3.8)$$

where $-D/2 \leq x_s, z_s \leq D/2$ and $\{x_e, z_e\} = \{0, 0\}$ is at center of one ring. We choose to maintain the spatial ensemble average (3.4) and the two-disc symmetry average (3.5) separate for clarity of notation and because flow visualizations in Fig. 3.3(a) utilize the x_e, z_e coordinates.

Spatial averaging along x_s and z_s , i.e., averaging over wall-parallel planes and over a square of size D^2 confining the minimal symmetry-averaged flow unit generated by one spinning ring, is defined as:

Spatial average

$$\langle f \rangle(y, t) = \frac{1}{D^2} \int_{-D/2}^{D/2} \int_{-D/2}^{D/2} [f]_s(x_s, y, z_s, t) dx_s dz_s. \quad (3.9)$$

The statistical sample is doubled by averaging a quantity across the two channel halves. A capital letter indicates a global average, $F(y) = [f]_g$, defined as:

Global average

$$[f]_g = \langle \bar{f} \rangle. \quad (3.10)$$

For example, the global-averaged streamwise velocity is $U(y) = [u]_g$.

3.2.3 Flow symmetries

In this section, we examine closely the derivation of the two-disc symmetry transformation (3.8), that is applied in the following to obtain higher-quality statistics of the single-ring flow unit.

We seek out a transformation mapping the flow field $\mathbf{u}(\mathbf{x})$ on a disc into another of opposite rotation. The transformation necessarily acts on x and z since the flow has no symmetry along y .

The statistical symmetry of the flow descends from the symmetry of the ring boundary conditions (inducing the swirling component of the flow) and that of the mean pressure gradient (inducing the mean shear). We first examine the symmetry of the swirling flow and then the symmetry of the mean flow.

For the swirling flow, it is necessary to find a transformation that turns a clockwise-rotating ring flow into a anti-clockwise one i.e. performs a sign change of the wall-normal component of the angular velocity Ω_y . The angular velocity is defined as the cross product $\boldsymbol{\Omega} = (\mathbf{x} \times \mathbf{u})/|\mathbf{x}|^2$. We recall that under a linear transformation \mathbf{R} , the cross product of two vectors \mathbf{a} and \mathbf{b} has the following property:

$$\mathbf{R}\mathbf{a} \times \mathbf{R}\mathbf{b} = (\det \mathbf{R})(\mathbf{R}^{-1})^T \mathbf{a} \times \mathbf{b}. \quad (3.11)$$

which applied to the angular velocity definition gives:

$$\mathbf{R}\mathbf{x} \times \mathbf{R}\mathbf{u} = |\mathbf{R}\mathbf{x}|^2 (\det \mathbf{R}) \mathbf{R}\boldsymbol{\Omega}. \quad (3.12)$$

This suggests that mapping Ω_y into $-\Omega_y$ requires $\det \mathbf{R} = -1$ and $|\mathbf{R}\mathbf{x}|^2 = |\mathbf{x}|^2$. These properties define an improper orthogonal transformation, i.e. a reflection. A suitable choice is the z -reflection \mathbf{R}_z , defined by:

$$\mathbf{R}_z = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -1 \end{pmatrix}. \quad (3.13)$$

The analogous x -reflection can also be applied to \mathbf{u}_d^p with the same result.

The mean flow $\mathbf{U} = (U, 0, 0)$ is homogeneous for any given y and invariant under z -reflections:

$$\mathbf{R}_z \mathbf{U} = \mathbf{U}, \quad (3.14)$$

whereas it is not invariant under an x -reflection because its sign changes. Any interaction between \mathbf{U} and \mathbf{u}_d^p is therefore expected to be invariant under \mathbf{R}_z but not under \mathbf{R}_x . This establishes \mathbf{R}_z

as the symmetry of both \mathbf{u}_d^r and the total flow \mathbf{u}_d . The above arguments can be validated by considering the relation:

$$\langle [C_f]_s \rangle = \langle C_f \rangle_x, \quad (3.15)$$

where the spatial average on the left-hand side is performed over two attached symmetry-averaged flow units of opposite rotation. This equation must be satisfied if the transformation (3.13) represents the correct symmetry of the flow. Figure 3.2 shows good agreement between the two approaches. The discrepancies are to be ascribed to a small residual spanwise velocity gradient W .

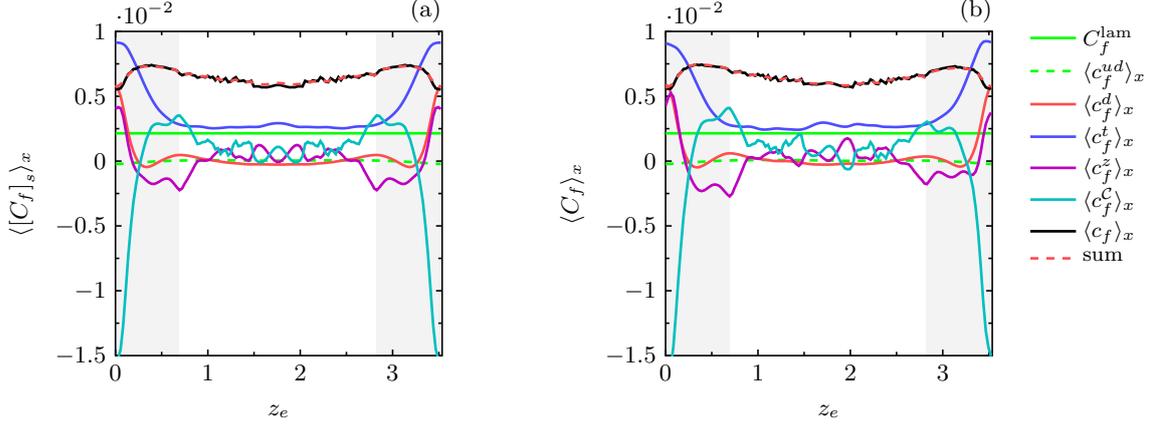


Figure 3.2: Symmetry transformation applied to the skin-friction data (Equation. (3.15)). Spanwise skin-friction budgets using (a) spatial averaging on the symmetry-averaged flow and (b) direct spatial averaging. The shaded areas represent the spanwise extension of the ring on a cross-flow section passing through its centre.

3.2.4 Flow decomposition

A useful analytical tool is the three-component (“triple”) decomposition first introduced by Reynolds and Hussain (1972) for time-periodic flows, who split the flow field into the uniform, steady component U , the time-periodic, zero-mean component \tilde{u} and the random turbulent fluctuations u' . Each component is derived from the full flow by sequentially performing partial averages and then removing the averaged component, realising a hierarchy of increasingly detailed time resolution. A set of RANS-like governing equations can be derived, one for each component, and used to study their reciprocal influence. The idea of a triple decomposition can be adapted to spatially periodic flow such as the flow over the array of rings, following Ricco and Hahn (2013). The key difference is that spatial averaging replaces the original time averaging:

$$\mathbf{u}(x, y, z, t) = \mathbf{U}(y) + \mathbf{u}_d(x_s(x), y, z_s(z)) + \mathbf{u}_t(x, y, z, t), \quad (3.16)$$

$$\mathbf{u}_d(x_s, y, z_s) = \{u_d, v_d, w_d\} = [\bar{\mathbf{u}}]_s - \mathbf{U}(y), \quad (3.17)$$

$$\frac{\partial p}{\partial x}(x, y, z, t) = \frac{dP}{dx} + \frac{\partial p_d}{\partial x}(x_s(x), y, z_s(z)) + \frac{\partial p_t}{\partial x}(x, y, z, t), \quad (3.18)$$

where the mean pressure gradient dP/dx is constant in fully-developed conditions and $\mathbf{U}(y) = \{U(y), 0, 0\}$. For clarity, we express the symmetry-average coordinates x_s, z_s as functions of the physical coordinates x, z . The flow field $\{\mathbf{u}_d, p_d\}$ is named the ring flow, while $\{\mathbf{u}_t, p_t\}$ indicates the fluctuating flow field. The governing equations for these quantities are obtained in Section

3.2.5 by applying the averaging operators to the full Navier-Stokes equations. A relation of central importance is found by substituting the decomposition (3.16) in the definition of the x - y Reynolds stresses:

$$[uv]_g = \langle u_d v_d \rangle + [u_t v_t]_g, \quad (3.19)$$

which appears in the mean-flow equation (3.22).

3.2.5 Governing equations for the three-component flow

We herein summarize the derivation of the governing equations obtained upon the three-dimensional flow decomposition (3.16)-(3.18), following the procedure of Reynolds and Hussain (1972). Substituting (3.16)-(3.18) in the Navier-Stokes equations and applying the time- and symmetry-averaging operators in §3.2.2 leads to the equation for the i -th component of the average flow $U + u_d$:

$$U^j \frac{\partial}{\partial x_s^j} (U^i + u_d^i) + U^j \frac{\partial u_d^i}{\partial x_s^j} = -\frac{\partial p_d}{\partial x_s^i} + \frac{1}{Re_p} \frac{\partial^2}{\partial x_s^j \partial x_s^j} (U^i + u_d^i) - \frac{\partial}{\partial x_s^j} \left(u_d^i u_d^j + [\overline{u_t^i u_t^j}]_s \right) + \frac{dP}{dx_s^i}, \quad (3.20)$$

where the Einstein repeated-index summation convention is used. The streamwise momentum equation (3.20) reads:

$$\begin{aligned} U \frac{\partial u_d}{\partial x_s} + v_d \frac{\partial U}{\partial y} = & -\frac{\partial p_d}{\partial x_s} + \frac{1}{Re_p} \frac{\partial^2}{\partial x_s^j \partial x_s^j} (U + u_d) - \frac{\partial}{\partial y} (u_d v_d + [\overline{u_t v_t}]_s) \\ & - \frac{\partial}{\partial z_s} (u_d w_d + [\overline{u_t w_t}]_s) - \frac{\partial}{\partial x_s} \left([\overline{u_d^2}]_s + [\overline{u_t^2}]_s \right) + \frac{dP}{dx_s}. \end{aligned} \quad (3.21)$$

The spatial average of (3.21) leads to the equation for the globally-averaged flow:

$$\frac{1}{Re_p} \frac{\partial^2 U}{\partial y^2} - \frac{\partial}{\partial y} (\langle u_d v_d \rangle + [u_t v_t]_g) + \frac{dP}{dx_s} = 0. \quad (3.22)$$

The governing equations for u_d are derived by subtracting (3.22) from (3.20).

3.2.6 Definition of turbulent drag reduction

The separation of scales typical of wall-bounded turbulence is measured by the friction Reynolds number $Re_\tau = h^* / \delta_v^*$, where $\delta_v^* = \nu^* / u_\tau^*$ is the near-wall viscous length scale, $u_\tau^* = \sqrt{\tau_w^* / \rho^*}$ is the wall-friction velocity, $\tau_w^* = \nu^* \rho^* dU^* / dy^* |_{y=0}$ is the global-averaged wall-shear stress, and ρ^* is the density of the fluid. The skin-friction coefficient is $C_f = 2\tau_w^* / \rho^* U_b^{*2}$, where $U_b = \int_0^1 U(y) dy = 2/3$ is the bulk velocity. The drag reduction \mathcal{R} is defined as the percentage decrease of the skin-friction coefficient:

$$\mathcal{R}(\%) = 100(\%) \cdot \frac{C_f^0 - C_f}{C_f^0}, \quad (3.23)$$

where C_f^0 is the skin-friction coefficient of the uncontrolled flow. We also define a spatially dependent turbulent drag reduction as follows:

$$\mathcal{R}_{xz}(x_s, z_s)(\%) = 100(\%) \cdot \frac{C_f^0 - c_f(x_s, z_s)}{C_f^0}, \quad (3.24)$$

where

$$c_f(x_s, z_s) = \frac{2}{Re_p U_b^2} \frac{\partial (U + u_d)}{\partial y} \Big|_{y=0}. \quad (3.25)$$

It follows that $\langle \mathcal{R}_{xz} \rangle = \mathcal{R}$ because $\langle c_f \rangle = C_f$, which is a direct consequence of $\langle u_d \rangle = 0$, given by the definitions (3.17) of u_d and (3.10) of U .

3.2.7 Uncertainty analysis

The sampling uncertainty is estimated via normal confidence intervals, corrected for the presence of temporal correlation by fitting an auto-regressive model to the time series of the sampled variable (von Storch and Zwiers, 2001; Oliver et al., 2014). The uncertainty is first computed for the wall-shear stress τ_w^0 in the reference conditions, i.e., ϵ_0 , and for the reduced wall-shear stress τ_w , i.e., ϵ . The uncertainty ϵ_R of the drag reduction \mathcal{R} is then found as follows:

$$\epsilon_R = 1 - \frac{\mathcal{R}}{100} - \frac{\tau_w - \epsilon}{\tau_w^0 + \epsilon_0} \quad (3.26)$$

Table 3.2 presents the uncertainty values.

Table 3.2: Confidence intervals (95%) of the wall-shear stresses τ_w^0 and τ_w , i.e., ϵ and ϵ_0 , respectively, and of the drag reduction \mathcal{R} , i.e., ϵ_R .

	Reference flow	R	ROC	RHS
$100\epsilon_0/\tau_w^0, 100\epsilon/\tau_w$	$\pm 0.7\%$	$\pm 0.3\%$	$\pm 0.5\%$	$\pm 0.5\%$
ϵ_R	–	$\pm 0.9\%$	$\pm 0.9\%$	$\pm 0.9\%$

3.3 Results

3.3.1 Ring flow

The main features of the non-uniform ring flow \mathbf{u}_d are discussed in this section. Some visualizations of \mathbf{u}_d for the case without distributed control (case R) are shown in Figs. 3.3 and 3.4. The three-dimensional contour plot of Fig. 3.3a clearly shows that the ring flow \mathbf{u}_d mainly consists of the near-wall swirling motion above the rings and streamwise-elongated structures at the sides of the rings. The isosurfaces represent the magnitude of the velocity vector $|\mathbf{u}_d| = 0.1$, while the color shows v_d . The structures are mainly characterized by wall-normal and upstream ejections. Over the rotating annular actuators, the wall-normal ring-flow velocity is directed towards the wall. These features were also observed in the full-disc case (Ricco and Hahn, 2013; Wise and Ricco, 2014; Wise et al., 2014, 2018), although they have not been described exhaustively. The swirling motion on top of the rings is due to the diffusive shearing action of the rings, while the elongated structures are caused by the interaction of the radial flows of two spanwise-adjacent rings and by the radial flows encountering the stationary surface at the outer ring edge. The radial flow is produced by the wall-normal pressure gradient that draws fluid towards the wall on top of the rotating parts and then toward the outer edge of the rings. This is a consequence of the centrifugal acceleration of the fluid near the wall and it is analogous to the three-dimensional swirling laminar flow over a rotating infinite plane (Von Kármán, 1921).

The streamwise component u_d exhibits a near-wall region of negative values centered at $z_s = D/2$ and confined below $y = 0.2$. Above this region, it acquires weakly positive values directly above the central region of the ring and a narrower, intensely negative region between spanwise-adjacent rings, as shown in the left graphs of Fig. 3.3b in cross-sectional planes perpendicular to the mean flow and in the right graphs of Fig. 3.3b in planes parallel to the streamwise direction. These u_d -regions balance one another because of continuity, i.e., the net flow rate of u_d across y - z planes, is null.

On x - y planes between rings (right column in Fig. 3.3b), the wall-normal ring-flow component v_d shows large positive values corresponding to ejections towards the channel core. These ejections persist in the streamwise direction and largely correspond to the low-speed regions of

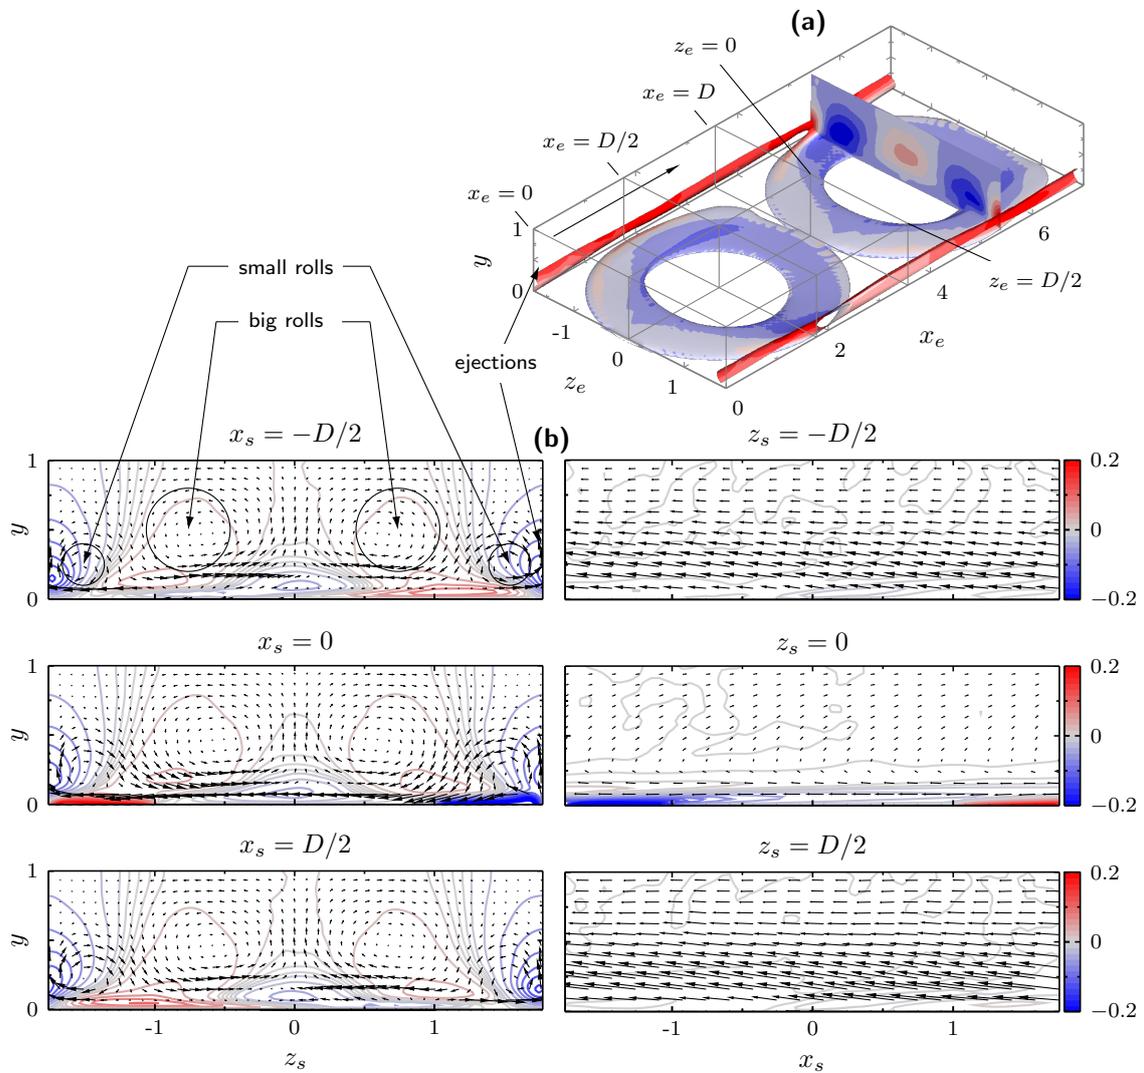


Figure 3.3: Rings-only flow (case R in Tab. 3.1). (a) The isosurface represents the magnitude $|\mathbf{u}_d| = 0.1$ and the colors show the wall-normal velocity v_d on the isosurface and on the $y-z$ plane. (b) Symmetry-averaged ring-flow $\mathbf{u}_d(\mathbf{x}_s)$ on three cross-flow planes (left graphs) and three streamwise planes (right graphs). The in-plane components are represented by the arrows and the third component is quantified by the colors of the contours. The scale of the arrows in the right column is three times smaller than in the left column. The abscissa axes span one diameter in length.

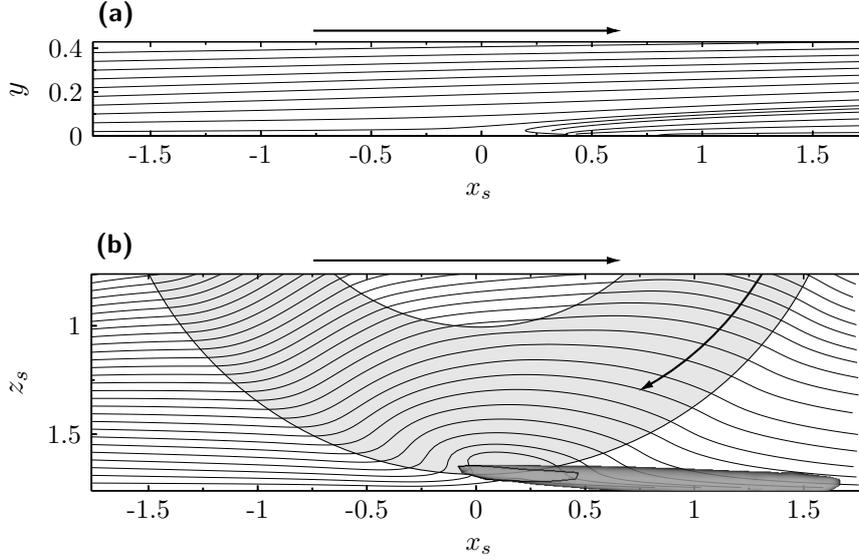


Figure 3.4: Rings-only flow (case R in Tab. 3.1). (a) Streamlines of $\mathbf{u}_d + \mathbf{U}$ in the $z_s \approx -D/2$ plane. (b) Streamlines of $\mathbf{u}_d + \mathbf{U}$ in the $y = 0.04$ plane. The dark gray three-dimensional structure is an isosurface where $|\mathbf{u}_d| = 0.14$ and $u_d < 0$.

u_d , forming the streamwise-elongated structures between the rings, depicted in Fig. 3.3a. The product $u_d v_d$ is thus negative where these structures exist, sharing the same sign of the turbulent Reynolds stresses $\overline{u_t v_t}$. Above the spinning portion of the rings, v_d acquires negative values because it balances the intense positive ejections at the sides of the rings and the weaker positive motion toward the portion of stationary wall above the center of the rings, as shown in Fig. 3.3b. These positive and negative regions of v_d cancel one another out over wall-parallel planes as $\langle v_d \rangle = 0$. This alternating suction/ejection pattern induces four coherent streamwise rolls, visualized by the $v - w$ vectors of Fig. 3.3b (left graphs). Two of these vortices are small and intense, appearing close to the blue region of negative u_d and to the red region of positive u_d , while the other two rolls are larger and weaker, existing over the central region of the rings and located at around $y = 0.5$. These larger rolls produce the weak upward motion directly over the centerline of the ring. All four rolls are persistent in the streamwise direction as they are observed in all the cross-stream sections of Fig. 3.3b (left graphs).

The upstream-rotating side of the ring is responsible for the creation of intense wall-normal gradients of u_d . Fig. 3.4a shows that, as the upstream-moving side meets the incoming mean flow, it gives rise to a local stagnation point in the vicinity of the wall where $U + u_d = 0$. This stagnation point moves away from the wall as the spanwise distance increases from the ring center and u_d grows from $-0.6W$ at the inner-ring edge to $-W$ at the outer-ring edge. Between these points and the wall a region of upstream motion thus occurs, while above the stagnation points the downstream flow is restored. The dark gray three-dimensional contour in Fig. 3.4b shows the region where u_d is most negative. This corresponds to the center of the upstream-flow region and occurs downstream of the ring.

a. Three-dimensional streamlines

Two-dimensional streamlines have been used in Fig. 3.4 for the sake of simplicity, integrating the flow field on plane sections. We report here a more complete three-dimensional visualisation of the

flow morphology at the disc upstream-rotating side. The 3D streamlines of the mean flow do not stay confined onto planes but bend in all three spatial directions. The four streamlines in Figure 3.5 represent the main types of behaviour found in the side region. They are integrated starting from four initial conditions located upstream of the ring, uniformly spaced between $y = 0.01$ and $y = 0$ and numbered in order of decreasing distance from the wall. Type I streamlines belong almost perfectly to a wall-normal, streamwise-aligned plane. Type II streamlines instead bend significantly in the spanwise direction. Type III and Type IV streamlines exhibit a more complex pattern with twists and reversals. Such topologies arise because the fluid trajectories become entrained into the near-wall layer where the swirling flow dominates. They undergo a partial rotation until the streamwise velocity component becomes large enough to deviate them across the fixed-wall region. At that point, their behaviour diverge. Type III finds itself sufficiently far from the edges and continues downstream, being only partially deflected by the swirling flow. Type IV instead re-enters the rotating region at a location where it is pushed backwards and sideways, lifted up and eventually attracted back to the same streamline bundle of Types I and II.

Fig. 3.5b depicts the streamwise symmetry-averaged velocity profiles at several spanwise positions in order to show the existence of a region of reverse flow at the ring upstream rotating edge. The red points are located at $y^{0+} = 11$ to visualise the wall-normal scale of the reverse flow region with respect to the viscous sublayer.

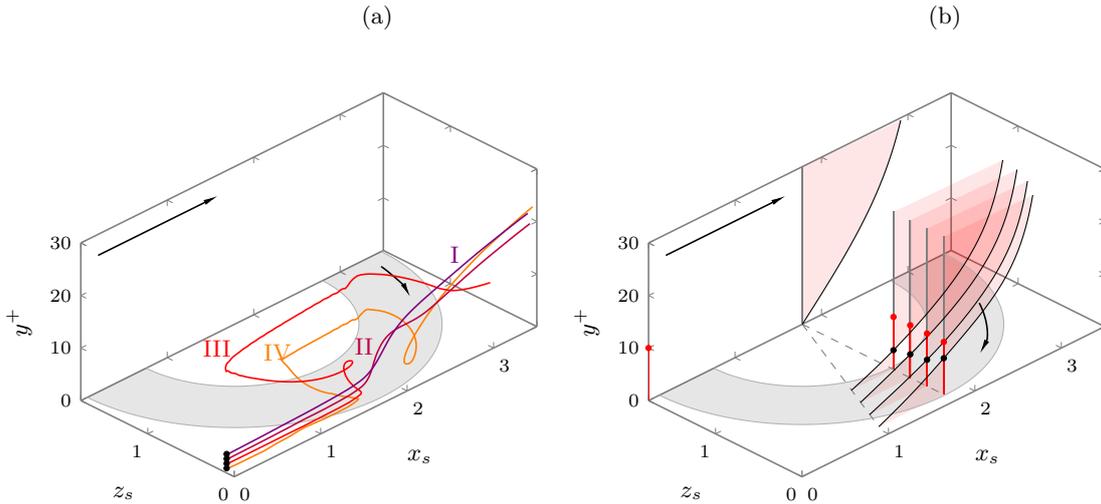


Figure 3.5: Three-dimensional organization of the flow at the upstream-rotating edge of a ring. (a) 3D streamlines integrated from four initial conditions uniformly spaced along the wall-normal direction. (b) Streamwise velocity profiles $U + u_d$ at the ring outer edge. Wall-normal scale is 9x exaggerated. The flow is the clockwise-rotating ensemble. Streamlines Type I (—), Type II (—), Type III (—) and Type IV (—).

As illustrated in Fig. 3.1, along the centreline of the rings the wall velocity distribution matches a triangular wave of spanwise velocity, with the exception of the stationary-wall central regions. This wave-pattern is analogous to that induced by the standing spanwise-velocity waves studied by Viotti et al. (2009) and Quadrio et al. (2009). For the rings, the equivalent wavelength is $2D^*$, which translates to about $1200\delta_\nu^*$. At comparable wavelength, Reynolds number, and forcing amplitude, the standing waves lead to their maximum drag reduction of about 45% (Viotti et al., 2009; Hurst et al., 2014). It is therefore useful to inspect the spanwise velocity

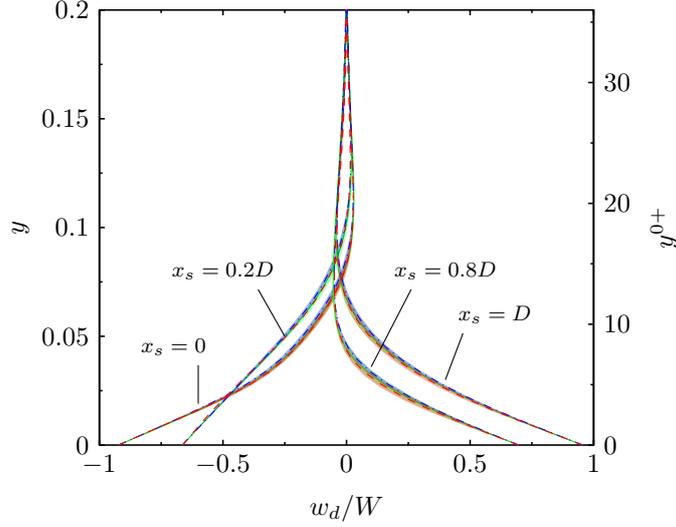


Figure 3.6: Profiles of the spanwise ring-flow velocity w_d at different x_s locations along the centerline of a ring. The shaded regions denote the uncertainty of the time-averaged field. Legend: — R case; - - - ROC case; - · - · RHS case.

profile w_d along the ring centreline at different streamwise locations. Figure 3.6 shows that these profiles qualitatively resemble the generalized Airy-function Stokes-layer profiles studied by Quadrio and Ricco (2011), and that, in the three control cases, they collapse on one another, irrespectively of the type of distributed forcing. This thus proves that the distributed control does not interfere with the beneficial drag-reducing action of the spinning rings because the viscous streamwise-modulated layer is unaffected when opposition control or hydrophobicity is imposed.

It is also relevant to quantify the thickness of this viscous layer. It has been shown by Quadrio and Ricco (2011) for the standing and traveling sinusoidal waves that the drag reduction scales linearly with the wall-normal penetration length, which quantifies the thickness of the boundary layer created by wall motion. In the present case, we define this length as $\delta = -\langle w_d(y=0) \rangle_{x,c} / w'_d(y=0)_{x,c}$, where $\langle \cdot \rangle_{x,c}$ denotes the spatial average along the centerline of the ring where $w_d(y=0) \neq 0$ and the prime symbolizes differentiation with respect to the wall-normal direction. We find that the penetration depth is $\delta^+ = 6.5$ in all the cases, which matches to the optimal value found by Quadrio and Ricco (2011) for the sinusoidal-wave forcing.

3.3.2 Turbulence statistics

a. Global statistics

Figures 3.7 and 3.8 present the wall-normal statistical profiles scaled in the reference wall units, i.e., by u_τ^* and δ_ν^* (superscript 0+), and scaled in native wall units, i.e., by u_τ^+ and δ_ν^+ of each drag-reduced flow (superscript +), respectively. Drag reduction is shown in Fig. 3.7a for the R and ROC cases ($\mathcal{R} = 20\%$ and $\mathcal{R} = 27.4\%$) as a lower streamwise velocity is computed in the viscous sublayer. The skin-friction coefficient also drops in the RHS case ($\mathcal{R} = 26.5\%$), but this is not immediately evident due to the wall-slip velocity. As the flow rate is constant, the R and ROC profiles display slightly larger velocity in the outer layer due to the near-wall reduced flow rate. The RHS profile shows lower velocities near the centreline, balancing the increased near-

wall velocity due to the hydrophobic slip. The inception of the logarithmic layer shifts upwards, denoting a thicker viscous sublayer. The mean profiles shift upward above the viscous sub-layer when scaled with their native wall-friction velocities, as shown in Fig. 3.8.

The turbulent fluctuations are significantly attenuated. Both the $[u_t u_t]_g^{0+}$ profiles in Fig. 3.7b and the $[uv]_g^{0+}$ profiles in Fig. 3.7c are reduced mostly up to $y^{0+} = 100$. ROC is the most effective in decreasing $[uu]_g^{0+}$ and $[uv]_g^{0+}$, although the latter increases in the viscous sub-layer because of the enhanced wall-normal velocity fluctuations. The wall-normal peaks of $[u_t u_t]_g^{0+}$ are shifted away from the wall by about $10\delta_v^*$ with respect to the reference case. The components of $[uv]_g^{0+}$, defined in (3.19), are shown in Fig. 3.7d. Since $[u_d v_d]_g^{0+}$ is negative over most of the channel height, the coherent streamwise and wall-normal ejections between spanwise-adjacent rings increase the drag. The turbulent Reynolds stresses $[u_t v_t]_g^{0+}$ are instead strongly suppressed, showing the beneficial drag-reducing influence of the rings. The cases with distributed forcing are observed to further weaken the total Reynolds stresses with respect to the rings-only case. The ROC $[u_d v_d]_g^{0+}$ profiles are weaker than in the rings-only case, whereas in the RHS case they are more intense than in the rings-only case outside the near-wall region. All the native-scaled profiles are lower than in the fixed-wall condition. This is due to the lower native Re_τ caused by drag reduction and to the wall-bounded flow being altered with respect to the reference flow.

b. Local Reynolds stresses

The wall-normal turbulent $[\overline{u_t v_t}]_s^{0+}$ profiles, shown in Fig. 3.9 at fixed x_s, z_s locations, are significantly non-uniform along z_s . The Reynolds stresses along the ring centreline are significantly diminished by the spanwise wave-like motion caused by the rings. The opposite behaviour characterizes the ring sides, where the strengthened Reynolds stresses reach peaks up to 150% of the reference case. These regions are narrower than those where the stresses are weakened, so the global effect is a reduction. The profiles (b) and (d) of Fig. 3.9 are representative of the central middle 60% of a ring. The two types of distributed control act differently and not uniformly, especially at location (a), i.e., over the stationary wall surface confined between four adjacent rings. Here ROC yields lower $[\overline{u_t v_t}]_s^{0+}$ peaks, while RHS is not very effective in reducing the fluctuations. In the region confined within the rings, the distributed controls offer little additional damping of the turbulent fluctuations compared to the R case as the turbulence intensity is already low.

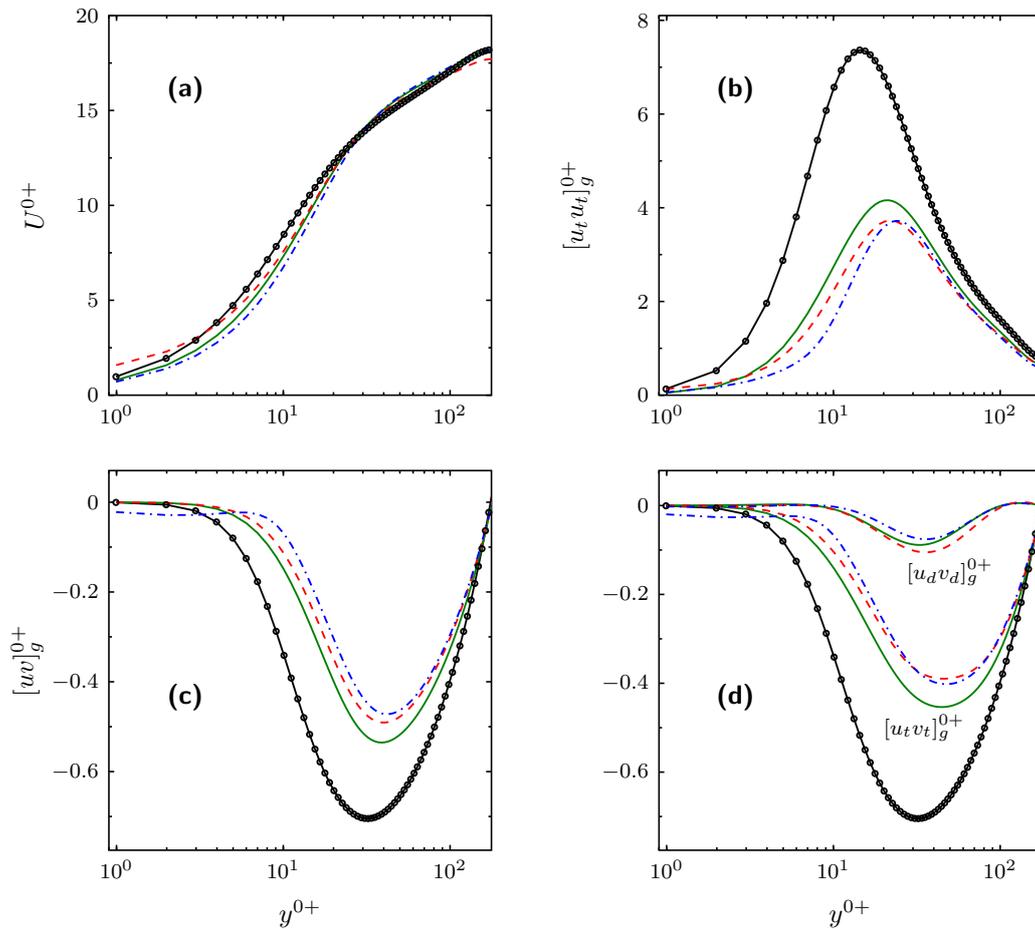


Figure 3.7: Wall-normal profiles of the turbulent statistics, scaled in wall units of the reference channel flow. (a) Mean velocity. (b) Streamwise turbulent intensity. (c) Total Reynolds shear stress. (d) Reynolds stress decomposition according to (3.19). Color key: \bullet — reference channel flow; — R case; - - - ROC case; - - - RHS case.

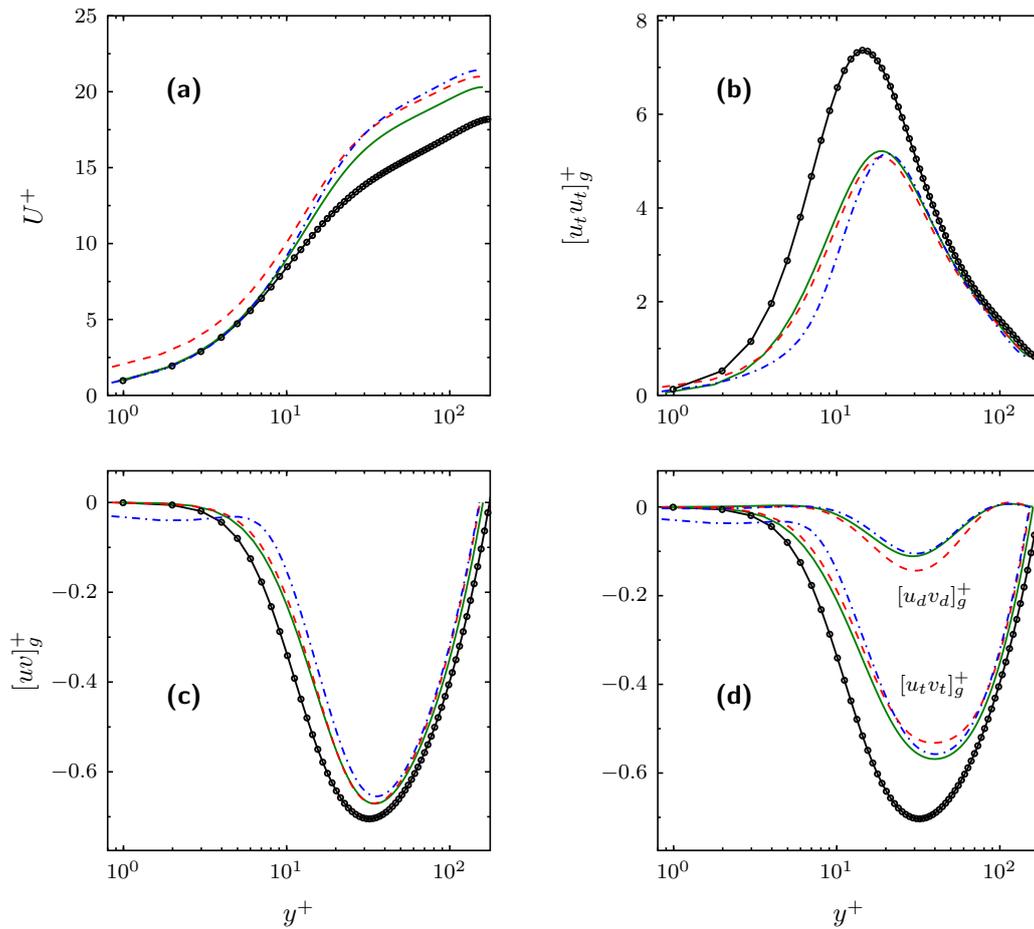


Figure 3.8: Wall-normal profiles of the turbulent statistics, scaled in native wall units. (a) Mean velocity. (b) Streamwise turbulent intensity. (c) Total Reynolds shear stress. (d) Reynolds stress decomposition according to (3.19). Color key: \ominus reference channel flow; — R case; - - ROC case; - · - RHS case.

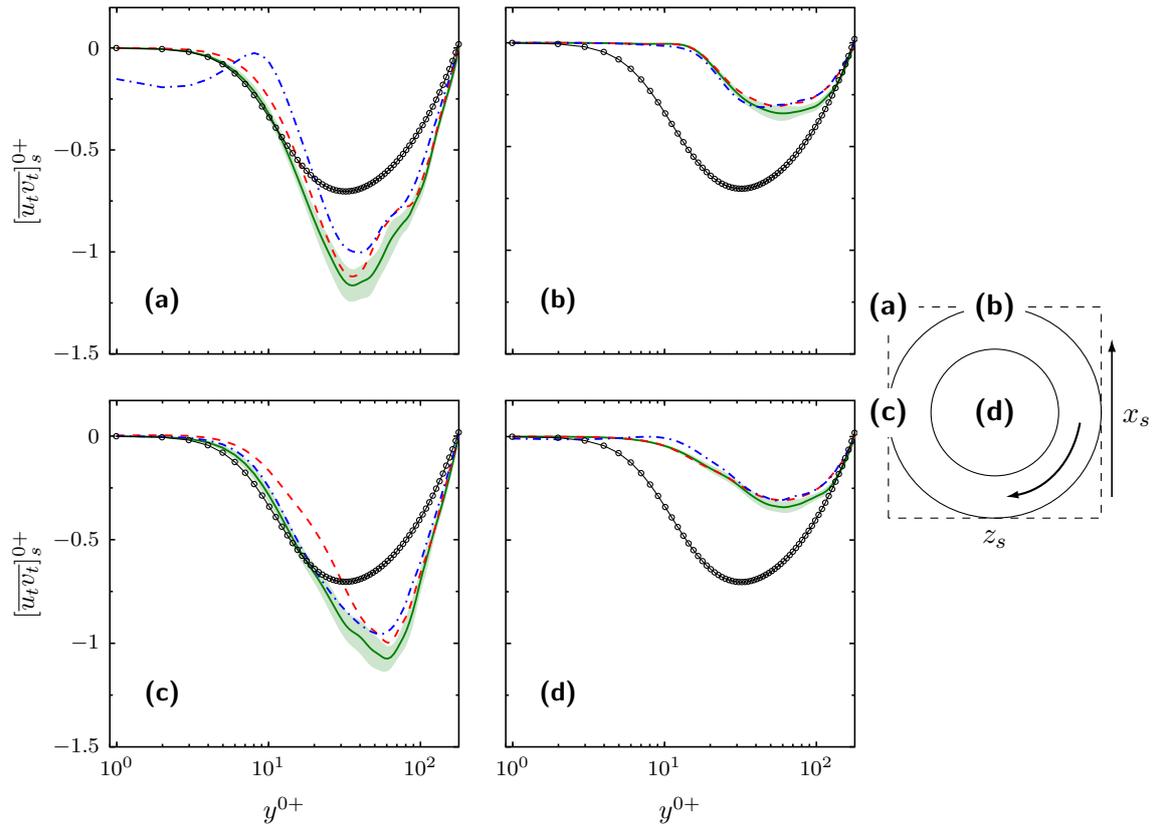


Figure 3.9: Turbulent Reynolds stress profiles at different positions in the symmetry-averaged frame. The variables are scaled with the wall units of the reference channel. The locations are letter-coded as in the sketch on the right. Color key: \ominus reference channel flow; — R case; -- -- ROC case; ⋯ RHS case. The shaded green regions denote the uncertainty of the time-averaged field.

c. Fine-scale vortical structures

Fig. 3.10 illustrates, via the λ_2 invariant (explained in Appendix 3.A), the fine-scale turbulent structures in the reference case and in the drag-reduced cases. The vortices are almost absent in the central region over the rings, in stark contrast with the more intense vortices on the streamwise-elongated structures between rings. These visualizations are consistent with the profiles of Fig. 3.9, which show the intense non-uniformity of the Reynolds stresses along the spanwise direction.

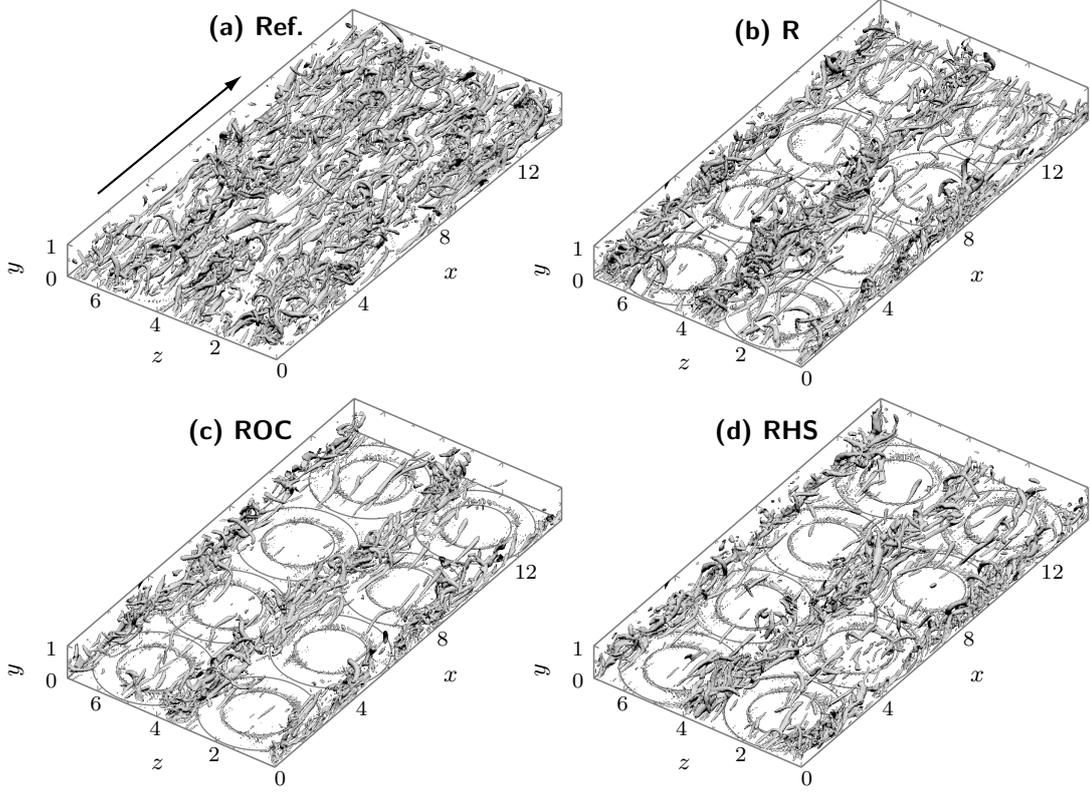


Figure 3.10: Instantaneous λ_2^+ -isocontours at $\lambda_2^{0+} = -0.007$. (a) reference flow, (b) rotating rings, (c) rotating rings and opposition control, (d) rotating rings and hydrophobic surface.

d. Turbulent kinetic energy

Due to the presence of \mathbf{u}_d , nonzero spatial fluxes of the turbulent kinetic energy $k = [u_t^i u_t^i]_s / 2$ are present in all three spatial directions. The general TKE evolution equation for the inhomogenous flow reads:

$$\frac{D_s}{D_{st}} k + \nabla \cdot \mathcal{T} = \mathcal{P} - \tilde{\varepsilon}, \quad (3.27)$$

where the terms are defined as follows, using the repeated-index notation:

$$\frac{D_s}{D_{st}} = (U^j + u_d^j) \frac{\partial}{\partial x_j} k, \quad \text{material derivative} \quad (3.28)$$

$$\mathcal{T}_j = \left(\frac{1}{2} [u_t^j u_t^i u_t^i]_s + [u_t^j p_t]_s + \frac{2}{Re} \frac{\partial k}{\partial x_j} \right), \quad \text{turbulent transport} \quad (3.29)$$

$$\mathcal{P} = [u_t^i u_t^j]_s \frac{\partial}{\partial x_j} (U^i + u_d^i), \quad \text{production} \quad (3.30)$$

$$\tilde{\varepsilon} = \frac{2}{Re_p} \left[\frac{\partial u_t^i}{\partial x_j} \frac{\partial u_t^i}{\partial x_j} \right]_s. \quad \text{pseudo-dissipation} \quad (3.31)$$

Fig. 3.11a shows that most of the TKE production is concentrated at the disc upstream-rotating edge. The same holds for the pseudo-dissipation (Fig. 3.11c). The collective transport term

show that in general, there is

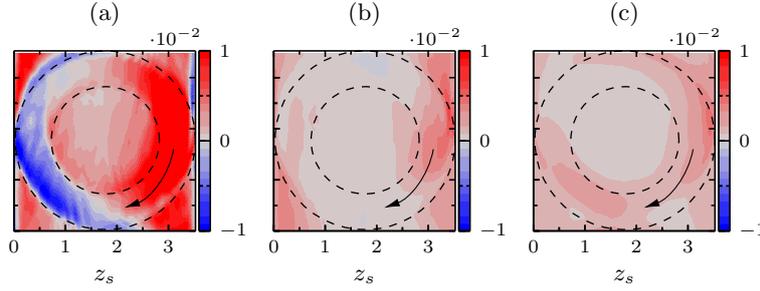


Figure 3.11: Turbulent kinetic energy budget at $y^{0+} = 11$. (a) mean production \mathcal{P} , (b) mean transport $\nabla \cdot \mathcal{T}$, (c) mean dissipation ε .

3.3.3 Turbulent drag reduction

In this section, the turbulent drag reduction generated by the spinning motion of the rings combined with the effects of the opposition control and the hydrophobic surfaces is examined, both globally and locally in space. The contribution of the different terms of the Navier-Stokes equations to the drag reduction is studied through an extended form of the Fukagata-Iwamoto-Kasagi identity (Fukagata et al., 2002).

3.3.4 Ideal prediction of combined-forcing drag reduction

We first predict the turbulent drag reductions \mathcal{R}_{id} generated by the combination of the rings and the distributed controls and compare these ideal predictions with the values computed via the direct numerical simulations. We use two assumptions:

0. For the first assumption, we consider the drag reduction generated by the distributed controls without the spinning rings. We assume that the reduction of the total drag is obtained by multiplying the fraction of actuated area by the drag reduction obtained if the whole surface area were altered by the distributed control, irrespectively of the geometrical arrangement of the controlled area.
1. For the second assumption, we first consider the wall-turbulent flow altered by the distributed actuation imposed over the area not occupied by the stationary rings. We assume that the action of the rings on this drag-reduced flow leads to the same amount of drag reduction \mathcal{R}_R as if the rings operated on the reference turbulent channel flow without control. This assumption is directly related to the spanwise ring-flow not being influenced by the distributed actuation, as demonstrated in Fig. 3.6.

It is useful to introduce the following definitions. We define $\mathcal{R}_{\text{dis-tot}}$ as the drag reduction via distributed control when the whole wall is actuated, $C_{f,1}$ as the skin-friction coefficient obtained via the distributed control activated over the area not occupied by the rings when the rings are not in motion, and $C_{f,2}$ as the skin-friction coefficient produced by the combination of the rings and the distributed control. The ratio between the combined skin-friction coefficient $C_{f,2}$ and the uncontrolled value C_f^0 is written as:

$$\frac{C_{f,2}}{C_f^0} = \frac{C_{f,2}}{C_{f,1}} \cdot \frac{C_{f,1}}{C_f^0}. \quad (3.32)$$

According to the first assumption:

$$\frac{C_{f,1}}{C_f^0} = 1 - \frac{\mathcal{R}_{\text{dis-tot}} S_{\text{dis}}}{S_{\text{tot}}}, \quad (3.33)$$

where S_{dis} is the area not occupied by the rings and $S_{\text{tot}} = L_x L_z$ is the total surface area. According to the second assumption:

$$\frac{C_{f,2}}{C_{f,1}} = 1 - \mathcal{R}_R. \quad (3.34)$$

Substitution of (3.33) and (3.34) into (3.32) leads to:

$$\mathcal{R}_{id} = \mathcal{R}_R + \frac{\mathcal{R}_{\text{dis-tot}} S_{\text{dis}}}{S_{\text{tot}}} - \mathcal{R}_R \frac{\mathcal{R}_{\text{dis-tot}} S_{\text{dis}}}{S_{\text{tot}}}. \quad (3.35)$$

The equation (3.35) shows that the predicted drag reduction is equal to the sum of the drag reductions produced by the two methods as if they operated separately minus the last term on the right representing the ideal loss. In our idealized model, this loss accounts for the drag reduction effect caused by the rings operating on a wall-bounded flow with a skin-friction coefficient that is lower than the reference one. We quantify the non-ideal interaction between the rings and the distributed control by defining the difference between the drag reduction computed via the direct numerical simulations and the ideal prediction given by (3.35):

$$\Delta \mathcal{R} = \mathcal{R} - \mathcal{R}_{id}. \quad (3.36)$$

The difference $\Delta \mathcal{R}$ arises mainly because the second assumption is not fully applicable, i.e., the effects of non-uniformity of the drag-reduced flow caused by the distributed forcing alone and by the ring-flow being altered by the distributed control between spanwise-adjacent rings.

Table 3.3: Drag reduction values obtained by the combination of the rings and the distributed controls. The full-disc drag reduction from Ricco and Hahn (2013) at the same diameter velocity are also added for comparison.

	Discs	Rings	OC	HS	ROC	RHS	OCR0	HSR0
$\mathcal{R}(\%)$	19.5	20.0	23.3	27.7	27.4	26.5	11.5	13.2
$\mathcal{R}_{id}(\%)$	-	-	-	-	29.3	31.1	11.7	13.8
$\Delta \mathcal{R}(\%)$	-	-	-	-	-1.9	-4.6	-0.2	-0.6

Table 3.3 reports the values of the numerically-computed \mathcal{R} , the predicted \mathcal{R}_{id} , and their difference $\Delta \mathcal{R}$. For the distributed control cases with stationary rings, OCR0 and HSR0, \mathcal{R}_{id} is calculated by adopting the first assumption, i.e., by multiplying the \mathcal{R} values given by the full-wall OC and HS controls by the fraction of actuated area $S_{\text{dis}}/S_{\text{tot}}$. These ideal values agree well with the numerically computed ones, i.e., within 95% confidence intervals, especially the opposition-control values. These agreements support the first assumption. In the combined cases, the drag reduction \mathcal{R} increases by a similar amount with respect to the rings-only case, i.e. 7.4% in the ROC case and 6.5% in the RHS case. However, neither fulfils the ideal expectation as $\Delta \mathcal{R} < 0$ in both cases. As these values are negative, the ideal estimate (3.35) represents an upper bound for the combined-forcing drag reduction. The $\Delta \mathcal{R}$ values in the ROC and RHS cases are comparable in magnitude and larger than the estimated uncertainties (refer to Table 3.2), but the agreement between the predicted and the numerical values is satisfactory. The discrepancy is larger in the hydrophobic-wall case than in the opposition-control case: the predicted drag reduction value is about 17% larger than the numerical value, whereas in the ROC case the prediction is about 6.5% higher than the computed value.

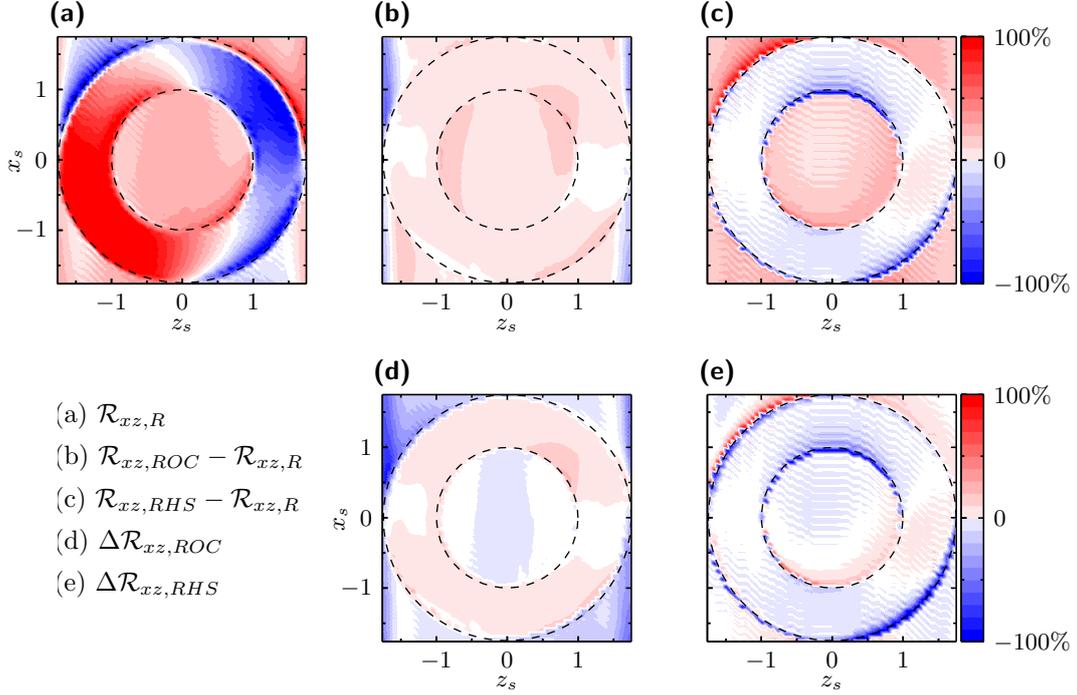


Figure 3.12: Spatial distribution of (a) the drag reduction \mathcal{R}_{xz} for the rings-only case, the difference between \mathcal{R}_{xz} in the rings-only case and the ROC and the RHS cases, (b) and (c) respectively, and the differences between \mathcal{R}_{xz} for the ROC and the RHS cases and the predicted $\mathcal{R}_{xz,id}$, (d) and (e) respectively.

3.3.5 Spatial distribution of drag reduction

The contour maps of Fig. 3.12 show \mathcal{R}_{xz} , defined in (3.24), for the rings-only case (Fig. 3.12a), the differences between \mathcal{R}_{xz} in the rings-only case and in the two combined-forcing cases (Fig. 3.12b,c), and the differences between the numerical \mathcal{R}_{xz} values in the combined-forcing cases and the predicted $\mathcal{R}_{xz,id}$, defined as

$$\mathcal{R}_{xz,id} = \mathcal{R}_{xz,R} + \mathcal{R}_{xz,OC/HS} - \mathcal{R}_{xz,R} \mathcal{R}_{xz,OC/HS}, \quad (3.37)$$

where $\mathcal{R}_{xz,OC/HS}$ is \mathcal{R}_{xz} generated by the opposition control or the hydrophobic control applied over the area not covered by the rings. These differences highlight the regions where the drag reduction is markedly altered with respect to the rings-only case. The spatial distribution of the time- and ensemble-averaged wall-shear stress reduction is strongly non-uniform and depends on the type of distributed forcing.

Fig. 3.12a shows that $\mathcal{R}_{xz,R}$ is determined by the rotation on the ring surface, thus negative on the upstream-rotating half (drag increase, blue color) and positive on the downstream-rotating half (drag reduction, red color). The red positive- \mathcal{R}_{xz} region on the rotating ring surface is larger than the blue negative- \mathcal{R}_{xz} region. The stationary inner-ring region benefits from levels of drag reduction as high as 40% and the wall-shear stress distribution is quite uniform. It is therefore not directly influenced by the ring sides spinning in opposite directions. The region outside of the ring is largely dominated by the combined effects of the radial swirling flow, the interaction between adjacent rings, and the adjustment of the flow from the rotating portion of the wall and the stationary surface. Drag reduction is observed downstream of the upstream-rotating part

of the ring, whereas drag increase occurs downstream of the downstream-rotating part of the ring. The drag is reduced upstream of the downstream-rotating side of the ring and increases upstream of the upstream-rotating side. On the lateral strips located directly underneath the ejections/low-speed regions, $\mathcal{R}_{xz,R}$ takes highly positive values. This proves that the streamwise-elongated structures have a direct but spatially limited beneficial effect on the wall-shear stress reduction, although their associated ring-flow Reynolds stresses are negative, as shown in Fig. 3.7d.

Fig. 3.12b shows that the combined forcing ROC yields about 10% more drag reduction in the inner-ring region than the rings-only case and has a similar beneficial effect on the large portions of the rotating rings, although it is not applied there. The bands between spanwise-adjacent rings experience a drag increase of up to 27% with respect to $\mathcal{R}_{xz,R}$, consistently with the weaker ring-flow structures in the ROC case, as displayed in Fig. 3.7d. Fig. 3.12c shows that the drag reduction \mathcal{R}_{xz} of the RHS case is less uniform than the ROC case. Therefore, their similar \mathcal{R} values, reported in table 3.3, do not reveal the markedly different physics. In the inner region, the RHS case delivers an improvement of about 15% drag reduction with respect to the rings-only case, which is much larger than in the ROC case. Outside of the ring, RHS has a uniform beneficial effect on the drag reduction. On the rotating region instead the average drag increases.

Several studies (e.g. Min and Kim (2004)) have established that in the reference channel flow the level of drag reduction depends on U_s^+ , the mean slip velocity scaled in native wall units. In the full-wall hydrophobic case, $U_s^+ = 3$. The global drag reduction achieved by the HSR0 case is very close to its \mathcal{R}_{id} , as shown in Table 3.4, and the mean slip velocity is very close to $U_s^+ = 3$ when scaled with the viscous length obtained from the averaged wall-shear stress over the hydrophobic surface. This supports the hypothesis that when the rings are not rotating, the ideal prediction is accurate. In the combined RHS case the slip velocity is smaller, i.e., $U_s^+ = 2.65$ in the central region, which is consistent with the intense reduction of wall-shear stress induced by the ring rotation.

The contours 3.12d and 3.12e depict the difference between the numerical \mathcal{R}_{xz} and the ideal $\mathcal{R}_{xz,id}$, defined in (3.37). In the ROC case, the \mathcal{R}_{id} overpredicts the numerical value over the regions not occupied by the rings, whereas the opposite occurs on the ring surface. In the RHS case, the difference is more non-uniform than the ROC case and the streamwise increase/reduction pattern along the streamwise direction is due to the local adjustment of the flow from the finite-slip over the hydrophobic surface to the no-slip condition on the ring surface and vice-versa.

3.3.6 Decomposition through the Fukagata-Iwamoto-Kasagi identity

In this section we examine the relation between the spatially-dependent skin-friction coefficient c_f defined in (3.25) and the terms of the Navier-Stokes equations to understand the origin of drag reduction and its spatial distribution. To this purpose, we apply the Fukagata-Iwamoto-Kasagi (FIK) theory (Fukagata et al., 2002) to our time- and ensemble-averaged streamwise momentum equation (3.21) for $\mathbf{U} + \mathbf{u}_d$. This leads to the following integral identity for the

spatially-dependent skin-friction budget:

$$\begin{aligned}
c_f(x_s, z_s) = & \underbrace{\frac{6}{Re_p U_b}}_{C_f^{\text{lam}}} - \underbrace{\frac{6U(0)}{Re_p U_b^2}}_{C_f^s} - \underbrace{\frac{6}{U_b^2} \int_0^1 (1-y) u_d v_d dy}_{c_f^d} - \underbrace{\frac{6}{U_b^2} \int_0^1 (1-y) [\overline{u_t v_t}]_s dy}_{c_f^t} \\
& + \underbrace{\frac{6}{Re_p U_b^2} \int_0^1 (u_d - u_{d,0}) dy}_{c_f^{ud}} - \underbrace{\frac{3}{U_b^2} \int_0^1 (1-y)^2 I_x'' dy}_{c_f^x} - \underbrace{\frac{3}{U_b^2} \int_0^1 (1-y)^2 I_z'' dy}_{c_f^z} \\
& - \underbrace{\frac{3}{U_b^2} \int_0^1 (1-y)^2 \mathcal{C}'' dy}_{c_f^{\mathcal{C}}} - \underbrace{\frac{3}{U_b^2} \int_0^1 (1-y)^2 \Pi_x'' dy}_{c_f^{\Pi}},
\end{aligned} \tag{3.38}$$

where c_f is defined in (3.25) and $u_{d,0} = u_d(x_s, 0, z_s)$. Following the notation adopted by Fukagata et al. (2002), the superscript '' indicates the difference between a quantity q and its bulk integral, as follows

$$q'' = q - \int_0^1 q \, dy. \tag{3.39}$$

The last four integrals in (3.38) are defined by grouping the terms in the streamwise momentum equation (3.21) that are null when averaged over x_s and z_s , as follows:

$$I_x = \frac{\partial}{\partial x_s} \left(u_d^2 + [\overline{u_t^2}]_s - \frac{1}{Re_p} \frac{\partial u_d}{\partial x_s} \right), \tag{3.40}$$

$$I_z = \frac{\partial}{\partial z_s} \left(u_d w_d + [\overline{u_t w_t}]_s - \frac{1}{Re_p} \frac{\partial u_d}{\partial z_s} \right), \tag{3.41}$$

$$\mathcal{C} = U \frac{\partial u_d}{\partial x_s} + v_d \frac{\partial U}{\partial y}, \tag{3.42}$$

$$\Pi_x = \frac{\partial p_d}{\partial x_s}. \tag{3.43}$$

The term C_f^{lam} represents the laminar part of the skin-friction coefficient. The term C_f^s in (3.38) is non-zero only in the RHS case: it is due to the finite average wall-slip velocity caused by the hydrophobicity of the surface (Aghdam and Ricco, 2016). Its magnitude is about 2% of C_f .

The terms $c_f^d(x_s, z_s)$ and $c_f^t(x_s, z_s)$ are the contributions of the two Reynolds-stress components $u_d v_d$ and $[\overline{u_t v_t}]_s$ to the skin-friction coefficient. The terms in the second and third lines of (3.38) are zero when (3.38) is averaged along x_s and z_s . The term c_f^{ud} originates from the spatial distribution of the ring-flow streamwise component u_d . The terms c_f^x and c_f^z are related to I_x in (3.40) and I_z in (3.41), respectively, and therefore to the x - x and x - z Reynolds and viscous stresses. The term $c_f^{\mathcal{C}}$ involves \mathcal{C} , defined in (3.42), and thus expresses the influence of the convective transport due to the interaction between the mean flow and the ejections at the sides of the rings. The term Π_x accounts for the non-uniform part of the pressure gradient.

By averaging (3.38) along x_s and z_s , the following expression for the skin-friction coefficient is found (Ricco and Hahn, 2013):

$$C_f = C_f^{\text{lam}} + C_f^d + C_f^t + C_f^s, \tag{3.44}$$

Table 3.4: Contributions of the Reynolds-stress components to the global skin-friction budget. The values of the partial skin-friction coefficients are multiplied by 10^3 and the values of the partial drag reduction values are percentage.

	C_f	C_f^{lam}	C_f^d	C_f^t	C_f^s	\mathcal{R}	\mathcal{R}_d	\mathcal{R}_t	\mathcal{R}_s
Reference flow	8.12	2.14	-	5.97	-	-	-	-	-
Rings	6.48	2.14	0.42	3.92	-	20.0	-5.2	25.2	-
ROC	5.88	2.14	0.36	3.38	-	27.4	-4.4	31.9	-
RHS	5.96	2.14	0.53	3.41	-0.12	26.5	-6.6	31.6	1.5

where $C_f^d = \langle c_f^d \rangle$ and $C_f^t = \langle c_f^t \rangle$. Table 3.4 reports the values of the terms in (3.44) for the drag-reduction cases, revealing that C_f^t exceeds C_f^d by one order of magnitude in all the cases. Table 3.4 also shows the partial drag reductions that add up to the total drag reduction, i.e., $\mathcal{R}_d + \mathcal{R}_t + \mathcal{R}_s = \mathcal{R}$, where $\mathcal{R}_d(\%) = -100 \cdot C_f^d / C_f^0$, $\mathcal{R}_t(\%) = 100 \cdot (C_f^0 - C_f^t) / C_f^0$, $\mathcal{R}_s(\%) = -100 \cdot C_f^s / C_f^0$, and $C_f^{t,0} = C_f^0 - C_f^{\text{lam}}$. The values of \mathcal{R}_d show the detrimental role of \mathbf{u}_d , which always increases the global drag. The analysis of the partial drag reduction budget also uncovers fundamental differences between the two combined control schemes. While both ROC and RHS can substantially reduce C_f^t to a very similar level, in the ROC case \mathcal{R}_d is more than 2% smaller than in the RHS case, while in the latter case C_f^d is larger than the rings-only case.

To improve the understanding of the flow, we average the terms in (3.38) along x_s , obtaining z_s -dependent quantities. This procedure leads to a limited loss of spatial information since the turbulence statistics are approximately uniform along the streamwise direction, as evident from Fig. 3.9. The budget of the skin-friction coefficient becomes

$$\langle c_f \rangle_x(z_s) = C_f^{\text{lam}} + C_f^s + \langle c_f^d \rangle_x + \langle c_f^t \rangle_x + \langle c_f^{ud} \rangle_x + \langle c_f^x \rangle_x + \langle c_f^z \rangle_x + \langle c_f^c \rangle_x + \langle c_f^{\text{II}} \rangle_x \quad (3.45)$$

The graphs in the first row of Fig. 3.13 show $\langle c_f^d \rangle_x(z_s)$ and $\langle c_f^t \rangle_x(z_s)$, thereby providing more insight into the values of Table 3.4.

The rings-only case is first examined (left column in Fig. 3.13). In the central region, the viscous boundary layer induced by the ring motion, shown in Fig. 3.6, reduces $\langle c_f^t \rangle_x$ (blue line) by 40%, i.e., at a level similar to that given by a sinusoidal wave at the same amplitude and wavelength, whereas at the ring sides, $\langle c_f^t \rangle_x$ has peaks at the ring sides exceeding the reference channel value $C_f^{t,0}$ (black dashed line) by 150%. This detrimental effect adds up to that of $\langle c_f^d \rangle_x$ (red line) that also peaks at the same locations. In all the configurations, the peak contributions to the global C_f originate from the sides of the rings where both $\langle c_f^d \rangle_x$ and $\langle c_f^t \rangle_x$ have a maximum. Since in the central region u_d is almost uniformly positive, the sign of $u_d v_d$ is dictated by the positive wall-normal component v_d , shown in the y - z plane of Fig. 3.3a. In the central region, the term $\langle c_f^d \rangle_x$ thus has a mild drag-reducing influence where ejections occur and a slightly more significant drag-increasing effect beneath the sides of the central rolls.

Although most of the globally drag-increasing contributions (i.e., peaks of $\langle c_f^t \rangle_x$ and $\langle c_f^d \rangle_x$) originate from the side regions, in Fig. 3.12a the above-average values of $\mathcal{R}_{xz,R}$ occur between rings because the structures locally decrease the drag as u_d is negative between spanwise adjacent rings. The sum of all the terms of (3.45) (black circles in Fig. 3.13, second row) shows that the total $\langle c_f \rangle_x$ has a minimum on the right upstream-rotating side of the ring. This minimum is due to $\langle c_f^c \rangle_x$ (light blue line), which has drag-reduced negative peaks that are much more intense in absolute value than the local peaks of the drag-increasing $\langle c_f^t \rangle_x$ and $\langle c_f^d \rangle_x$ in the ring-side region. This explains the narrow region of large positive \mathcal{R}_{xz} at the right ring side. Once averaged along x_s , $\langle c_f^c \rangle_x$ consists only of the term $v_d \partial U / \partial y$, i.e., the wall-normal advection of U along v_d . Therefore $\langle c_f^c \rangle_x$ is maximum at the sides of the ring because of the lateral ejections

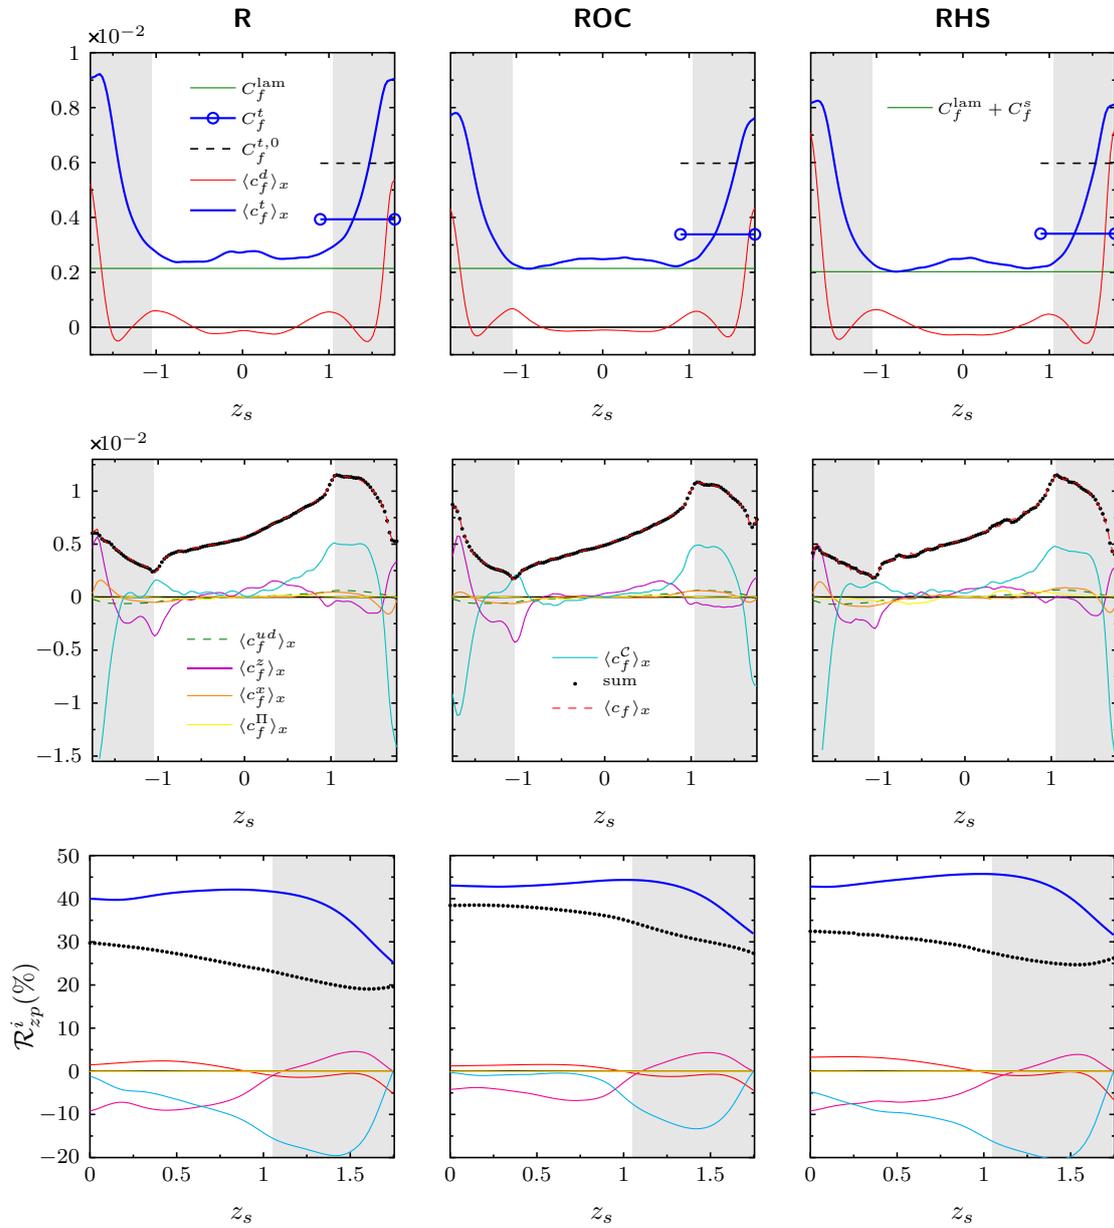


Figure 3.13: Spatially-dependent skin-friction budget defined in (3.45). The leftmost column is relative to rings-only control, the central column to ROC, and the rightmost column to RHS. The first row displays the first four terms of (3.45). The horizontal blue segments represent the global mean value C_f^t and the black dashed line denotes $C_f^{t,0}$. The second row displays the last three terms and the sum of all the terms of (3.45). The third row shows the progressively-averaged drag reduction, according to (3.46). The shaded areas represent the spanwise extension of the ring on a cross-flow section passing through its centre.

identified in §3.3.1. The term $\langle c_f^z \rangle_x$ (purple line) instead contributes detrimentally to the local skin friction by the intense peak at the sides. Most of the contribution to $\langle c_f^z \rangle_x$ is due to the Reynolds stresses $[\overline{u_i w_i}]_s$. The sum of the terms (black circles) coincides with $\langle c_f \rangle_x$ calculated from the wall-normal velocity gradient at the wall (dashed red line).

Further insight is gained by the bottom graphs of Fig. 3.13, which show

$$\mathcal{R}_{zp}(z_s) = \frac{100(\%) }{2z_s} \int_{-z_s}^{z_s} \frac{C_f^0 - \langle c_f \rangle_x(\widehat{z}_s)}{C_f^0} d\widehat{z}_s, \quad (3.46)$$

i.e., the progressive spanwise average of the local drag reduction from the ring centerline at $z_s = 0$ towards the edges. At $z_s = D/2$, the global drag reduction \mathcal{R} is retrieved by the integral (3.46). The partial spanwise-averaged drag reduction corresponding to the terms of equation (3.45), named $\mathcal{R}_{zp}^i(z_s)$ (where i denotes the corresponding superscripts of the terms), are found by substituting these terms into (3.46). The curves of $\mathcal{R}_{zp}^t(z_s)$ demonstrate that more than 20% of the global \mathcal{R}_R is lost in the side regions, while it remains essentially constant over the central region and of similar magnitude to the drag reduction produced by the standing waves (Viotti et al., 2009).

We now discuss how the wall-shear stress is modified in the ROC case by inspecting the graphs in the central column of Fig. 3.13. The action of OC results in a pronounced weakening of both the lateral peaks of $\langle c_f^d \rangle_x$ and $\langle c_f^t \rangle_x$, explaining the partial drag reductions \mathcal{R}_t and \mathcal{R}_d in Table 3.4. In the central region, $\langle c_f^t \rangle_x$ is slightly reduced with respect to the rings-only case. The progressively-averaged \mathcal{R}_{zp}^t curve of Fig. 3.13 (bottom central graph) confirms the additional 5% reduction of $\langle c_f^t \rangle_x$ with respect to the rings-only case in the central region. The weaker central rolls of \mathbf{u}_d , shown in Fig. 3.3b, are also attenuated as a consequence of the opposition control on v_d . The resulting reduction of $u_d v_d$ explains the slightly flattened $\langle c_f^d \rangle_x$ curve in the central region. In the ROC case, the negative peak of $\langle c_f^c \rangle_x$ is also diminished by more than 30% with respect to the rings-only case. This explains why the drag at the ring sides is higher than in the rings-only case despite $\langle c_f^d \rangle_x$ and $\langle c_f^t \rangle_x$ being much smaller in the same region. More insight into how $\langle c_f^d \rangle_x$ and $\langle c_f^c \rangle_x$ are affected by the OC is obtained by comparing the graphs 3.14-a (rings-only) and 3.14-b (ROC case), which show the flow field between two adjacent rings. At the ring sides the core of the negative- u_d region is located further away from the wall in the opposition-control case than in the rings-only case. This effect results in stronger gradients of $u_d + U$ at the wall with respect to the rings-only case. The wall-normal velocity v_d is also mitigated because of the direct opposing action of $v_d(y_d)$. This explains the lower peaks of $\langle c_f^c \rangle_x$ because the term $v_d \partial U / \partial y$ is less intense. The magnitude of $u_d v_d$ is also weakened and its minimum pushed away from the wall, resulting in the maximum of $\langle c_f^d \rangle_x$ being reduced.

We now examine the spatial skin-friction budgets of the RHS case from Fig. 3.13. On the side regions, $\langle c_f^d \rangle_x$ is increased with respect to the rings-only case. The peaks of $\langle c_f^t \rangle_x$ is reduced, but less than in the ROC case. In the central region the hydrophobic surface delivers a similar reduction of $\langle c_f^t \rangle_x$ as the ROC case, leading to an average 45% drag reduction. The progressively-averaged \mathcal{R}_{zp}^t curve of Fig. 3.13 (bottom right graph) also gives about an additional 5% reduction of $\langle c_f^t \rangle_x$ with respect to the rings-only case in the central region, similarly the opposition control case. The behaviour of the zero-global-mean, streamwise-averaged terms $\langle c_f^c \rangle_x$ and $\langle c_f^z \rangle_x$, shown in the second row of Fig. 3.13, remains largely unchanged.

These insights from the modified FIK analysis can help the future development of drag-reduction strategies, especially methods that require localized actuators. If, similar to our ring actuators, localized actuators interact detrimentally to induce additional Reynolds stresses that increase the wall-shear stress, wall-normal actuation confined between the actuators (wall transpiration in our case) is more efficient than streamwise wall-parallel forcing. We have also learned

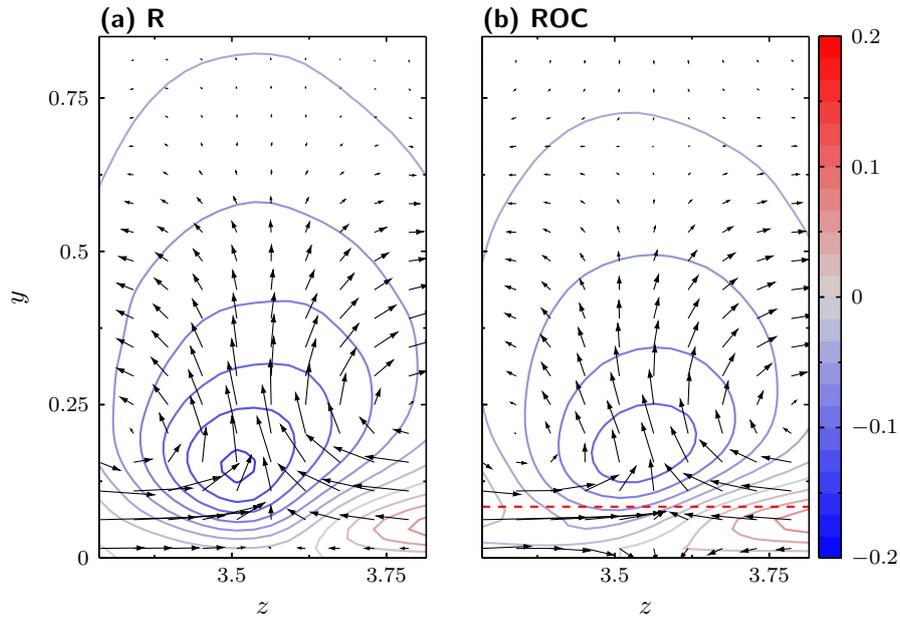


Figure 3.14: Action of the opposition control on \mathbf{u}_d , visualized on cross-stream sections at $x_s = D/2$. The contours depict the streamwise component u_d . (a) Rings-only forcing. (b) case ROC. The dashed red line marks the detection height $y_d^{0+} = 14.5$.

that the additional benefit of the distributed control in the central part of the ring is limited because the shearing drag-reducing action of the rings already operates effectively there.

3.3.7 Details of the skin-friction budget terms

Figure 3.15 depicts the complete local skin-friction budget on the wall surface. Each graph corresponds to one of the terms of Equation 3.38 for the symmetry-averaged flow unit. The three graphs in the top line are the terms that have a non-zero contribution to the global skin-friction. The remaining four terms do not contribute to the global skin-friction because of their symmetries that make their spatial average go to zero. As observed in Sec. 3.3.6, it can be noted that $c_f^t(x, z)$ and $c_f^d(x, z)$ are almost uniform along the streamwise direction but not along the spanwise direction. On the other hand, the term $c_f^z(x, z)$ is characterised by The other zero-mean skin-friction budget terms show a similar behaviour.

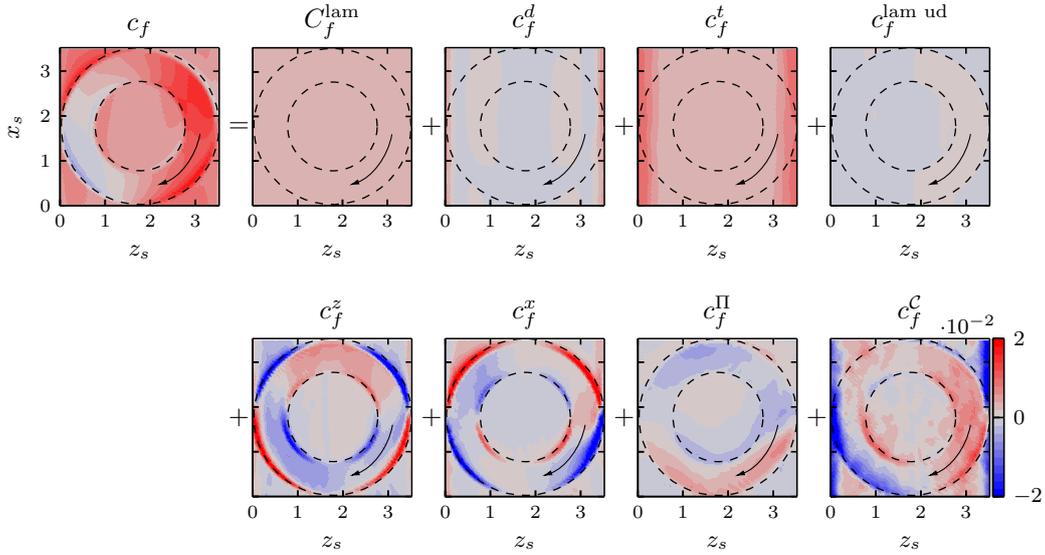


Figure 3.15: Full 2D maps of the skin-friction budget, forcing by rings only.

In order to avoid overcomplicated notation, some of the contributions of the individual terms of (3.21)-(3.22) to the skin-friction budget Eq. (3.38) were grouped into terms reflecting a common physical origin. The terms grouped under the name of c_f^C all belong to the convective transport of momentum \mathbf{u}_d by means of \mathbf{U} . The term c_f^z collects the spanwise fluxes of both the Reynolds and viscous stresses. Figure 3.16 represent the breakdown of these terms into their individual components.

The dominant component comes from the term $v_d \partial U / \partial y$, representing the transport of . The component $U \partial u_d / \partial x$, which Since no mean wall-normal velocity is present, the mean wall-normal convection contribution $V \partial u / \partial y$ is zero everywhere.

The contributions to the spanwise-inhomogenous component c_f^z show little variability across the different forcing configurations. The two major contributions come from the mean and the turbulent x, z -Reynolds stresses. In the lateral regions they have opposite sign, the positive (drag-increasing) turbulent stress contribution being larger in magnitude and leading to the maximum c_f^z .

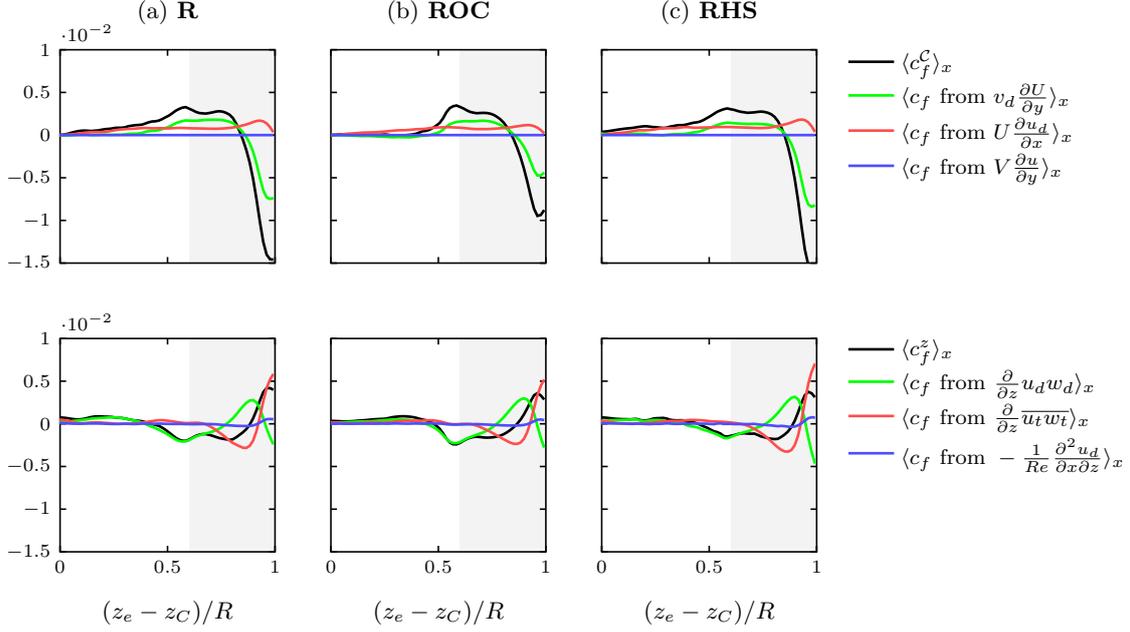


Figure 3.16: Detailed budget of the skin-friction contributions, averaged along x . Top row: term-by-term budget of the skin-friction contribution to c_f^c (Eq. (3.42)). Bottom row: term-by-term budget of the skin-friction contribution to c_f^z (Eq. (3.41)).

3.3.8 Detection height modification for case ROC

When the OC is applied on the entire wall surface with no wall motion (case OC), the drag reduction \mathcal{R}_{OC} is a function of y_d^+ only. If we consider a smaller value of Re_τ , raising y_d to maintain $y_d^+ = 14.5$ in the new wall units would preserve the same drag reduction value. Since the definition of \mathcal{R}_{id} assumes that y_d is rescaled, we expect that adjusting y_d to match Re_τ of the rings-only case would reduce $\Delta\mathcal{R}$. However, the failure of ROC to obey the wall-unit scaling is demonstrated in Fig. 3.17 where the results of four simulations with higher values of y_d are displayed. Starting from the left, the lowest height is the original configuration RHS with $y_d^{0+} = 14.5$ that has been discussed up to this point. The third value of y_d matches a detection height of $y_d \cdot 160 = 14.5$, preserving the wall-unit scaling for $Re_\tau = 160$ of the rings-only case. The two rightmost values (up to $y_d^{0+} = 26$) are also considered to test raising y_d based on the local friction at the ring centreline, which is lower than the global average. This attempt is motivated by the hypothesis that the canonical scaling is more likely to hold in the central region rather than at the heavily non-uniform ring sides. Any new y_d fails to reduce $\Delta\mathcal{R}$ or improve \mathcal{R} (Fig. 3.17, filled symbols), eventually resulting in drag increase for the largest detection height for which $\mathcal{R} < \mathcal{R}_R$ (dashed line). The same trend holds for the global contributions to the drag calculated at the ring centreline (Fig. 3.17, empty symbols).

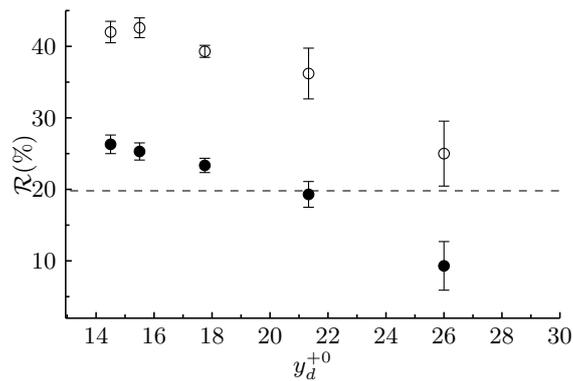


Figure 3.17: Effect of the variation of y_d on the drag reduction for the combined case ROC. Solid symbols (●) indicate the and empty symbols (○) refer to drag reduction calculated on the streamwise centreline of the disc.

3.4 Summary and conclusions

This study has examined two novel drag-reduction methods, i.e., the combination of rotating flush-mounted annular actuators and distributed control methods, i.e., opposition control and hydrophobic surfaces. The main conclusions of our work are outlined in the following.

(i) Rings-only flow

Despite the significant non-uniformity of the annular actuators, the resulting turbulence statistics are quasi-uniform along the streamwise direction. The ring-side streamwise-aligned structures, generated by the radial flow encountering stationary portion of the wall surfaces at the ring edges and interacting with the flow produced by neighbouring rings, carry a large fraction of the turbulent kinetic energy. While the local effect of these structures is to decrease the wall-shear stress in narrow streamwise-stretched bands between rings because of their negative near-wall streamwise velocity, they are responsible for additional negative ring-flow Reynolds stresses, which globally contribute detrimentally to the turbulent drag. The turbulence intensity at the ring side is larger than in the corresponding reference channel case, whereas the flow over the central part of the ring is almost relaminarized by the action of the spanwise triangular wave generated by the rotating rings. In this central region, the Stokes-type steady boundary layer produced by the rotation of the rings is responsible for the drag reduction and its thickness is equal to the optimal one proper of the streamwise-traveling waves of spanwise wall velocity studied by [Quadrio and Ricco \(2011\)](#). The damping of the turbulence in the central region occurs quite uniformly in the stationary ring core. The total influence of the rings on the turbulence intensity is an attenuation as the central low-turbulence region is much wider than the high-turbulence region of the narrow elongated structures. For the tip ring velocity and the diameter used in our study, the drag reduction produced by the rings is 20%. A spatially-dependent form of the Fukagata-Iwamoto-Kasagi identity has been useful to quantify the drag-altering effects of each term of the Navier-Stokes equation.

(ii) Combined flow: rings with distributed control

The combined control schemes deliver additional drag reduction with respect to the rings-only case: 7.4% in the case of the spinning rings with opposition control and 6.5% for the spinning rings with hydrophobic surface. We have used an idealized model, based on

reasonable assumptions, to predict the drag reduction produced by the combined actuations. The prediction is satisfactory and the ideal values are always slightly larger than the actual drag reduction, which renders the idealized value an upper bound. This simple model can be easily generalized to other combined drag-reduction techniques. The ring-Stokes boundary layer over the central ring region is unaffected by the distributed control schemes, which implies that the drag-reduction effect due to the rings is not altered in the combined cases.

Although the global drag-reduction performance of the two combined cases is similar, it is achieved via two different mechanisms. The modified form of the Fukagata-Iwamoto-Kasagi identity reveals that the attenuation given by the opposition control is more uniform than that of the hydrophobic surface. The rings-opposition control case exceeds the ideal performance in the side regions, weakening the Reynolds stresses carried by the streamwise-elongated structures at the ring sides. In the central region, the beneficial effect of the opposition control is almost negligible because the wall forcing is directly proportional to the near-wall wall-normal fluctuations, which are already strongly damped by the spinning rings. This view is endorsed by the intense attenuation produced by the opposition control in the central-ring region when the rings are not spinning. The rings-hydrophobic control case fares better in the central-ring region than at the ring sides, where the Reynolds stresses associated with the streamwise structures are strengthened by surface hydrophobicity while the turbulent Reynolds stresses are instead less affected by the surface hydrophobicity than in the opposition-control case. In the central-ring region, the hydrophobicity reduces the drag more than the rings-opposition control method and the flow globally benefits from the mean drag-reducing slip velocity. The distributed-control cases with fixed rings matches the ideal prediction very well, indicating that the additional actual losses are rooted in the interactions between the swirling flow and the distributed controls.

Experiments on wall turbulence over spinning discs and rings is certainly needed for validation of the numerical results. Given the highly three-dimensional character of these flows, the pressure-driven channel flow is certainly a preferred option over the free-stream boundary layer flow without a pressure gradient in order to measure the drag reduction by global pressure-drop measurements. The main reason for this choice is certainly the difficulty of measuring the highly spatially-dependent wall-shear stress through the wall-normal velocity gradient at the wall. Another very important topic of research is the effect of Reynolds number on the amount of drag reduction produced by the rotating rings and on the scaling of the optimum ring forcing parameters, i.e., the diameters defining the rings and the tip ring velocity, which deliver the maximum attenuation of wall-shear stress. It needs to be verified whether the rings, alone or combined with opposition control, can suppress effectively the large super-structures populating the outer part of high Reynolds number flows (Hutchins and Marusic, 2007). Another point of interest is the possible interaction of the ring-side structures, which scale in outer units (Wise and Ricco, 2014), with the super-structures at high Reynolds numbers.

We close our discussion with a short remark on the implementation of active control techniques in industrial scenarios. As amply discussed in Ricco and Hahn (2013), assuming that wall-unit scaling applies at high Reynolds numbers proper of commercial flight conditions, the optimal ring diameter would be of the order of 1mm and the typical period of rotation would be of the order of 0.1ms. These numbers render the practical realization of the rings method (and of any other active drag reduction technique) extremely difficult. Further complications obviously arise if feedback control is considered, even in the simplistic case of opposition control studied herein, because of the enormous database of the order of terabytes to be processed per second. Nevertheless, motivated by the enormous benefits from economical/environmental aspects, we

hope that research in turbulent drag reduction will continue to inspire generation of scientists and receive increased attention in order to achieve the ultimate goal of technological implementation in aircraft industry.

3.5 Further work

The flow decomposition Section 3.2.4 and the extended FIK identity (3.38) have proven critical tools to understand the fluid mechanics and the skin-friction of the disc-like actuators. In this section, we list a series of preliminary data and results about some research questions where the FIK analysis plays a central role and that may be considered as future research directions. We present three distinct lines of work:

- In Section 3.5.1 we argue that the same tools can be helpful in clarifying some issues that were left open in earlier work on rotating disc, namely the dependence on the parameters D and W and its origin in the spatially dependence FIK budget. Section 3.5.2 investigates qualitatively how differently the near-wall turbulent streaks are affected by homogeneous spanwise waves and by localised rings.
- Section 3.5.3 proposes a modification of the rotating-disc control based on sheltering the disc edges from the mean shear and reports promising results from a numerical simulation.
- Section 3.5.4 briefly examines a higher-Reynolds number flow case.

3.5.1 Dependence of \mathcal{R} on the control parameters

The original work on drag reduction by rotating disc by Ricco and Hahn (2013) left the open question of how the drag reduction \mathcal{R} depends on the control parameters W and D . In particular the authors note that, in contrast with a traditional spanwise wall forcing, the dependence of \mathcal{R} on the forcing amplitude W has a non-monotonic character, displaying a maximum at every value of D . The disc-flow decomposition and the spatially-dependent FIK identity help in answering this question.

Fig. 3.18a reveal that, while the disc flow generates drag increase that grows monotonically with W , the turbulent contribution \mathcal{R}_t is responsible for the convex character of \mathcal{R} and the existence of an optimal value of W .

Fig. 3.18b-c show how the spatially-dependent skin-friction budget changes with the control parameters. The turbulent contribution c_f^t increases with W at the sides while is decreases at the centre. An extreme increase at the sides is also seen in the disc-flow term c_f^d . We conjecture that the convex shape of $\mathcal{R}_t(W)$ is caused by the contrasting behaviour of c_f^t at the centre and at the sides, which would occur when the detrimental drag increase at the sides overtakes the increasingly marginal gains at the centre.

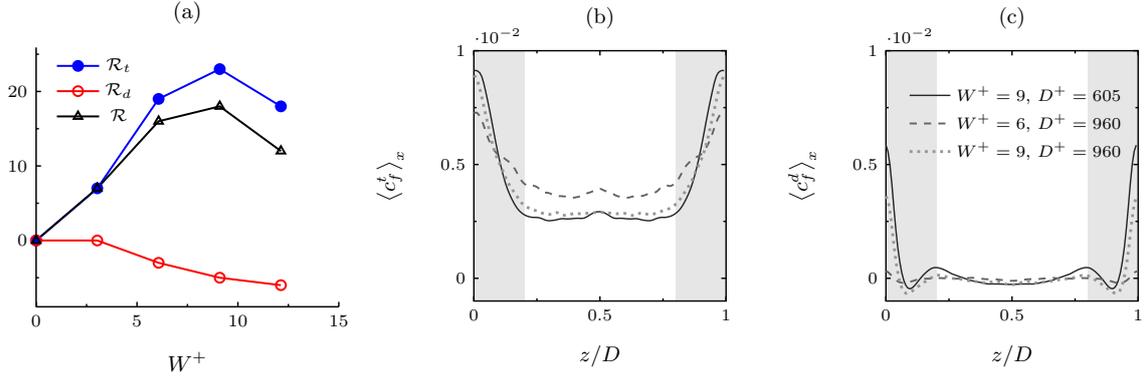


Figure 3.18: (a) Dependence of \mathcal{R} on W for discs with $D^+ = 605$, data from Wise et al. (2014). (b)-(c) Spatially dependent skin-friction components for three combinations of W and D for the ring forcing; the data are elaborated from the simulations performed by S. K. Aghdam (Khosh Aghdam, 2016).

3.5.2 Ring-like and wave-like boundary conditions

Since the rotating actuators create a partial triangular wave at their centreline, it is of great interest to compare the ring forcing with conventional spanwise wave forcing. To this purpose, a simulation of a channel flow was performed with standing triangular waves of spanwise velocity, having a wavelength $\lambda = 2D$ and amplitude W to match the corresponding parameters of the array of rings. The waves are enforced homogeneously along the spanwise direction. Another simulation has been carried out enforcing a traditional sinusoidal spanwise-velocity wave at the wall, also with the same parameters as the triangular wave. The numerical configuration of the simulations is the same for the wave and the ring simulations and can be found in Tab. 3.1.

Fig. 3.19 shows an instantaneous snapshot of the near-wall streaks over the periodic unit of the three flows. The streak organisation are known to be closely linked to the near-wall drag-reducing mechanism of the oscillating wall (Blesbois et al., 2013) and the analogous spanwise waves (Quadrio et al., 2009). The spanwise forcing gives rise to the pattern observed in Fig. 3.19b-c where the streaks are tilted to positive or negative angles depending on the forcing phase. A qualitatively similar pattern can be observed in the central region of Fig. 3.19a, although the streak organisation is heavily influenced by the spanwise inhomogeneity at the ring sides. Although the streak organisation in the central region is similar in all three cases, the triangular and sinusoidal waves appear to produce larger streak tilting angles than the ring forcing. The streak angle can be calculated as $\theta = \arctan(\lambda_z/\lambda_x)$ where λ_z, λ_x are respectively the spanwise and the streamwise length scales of the streamwise velocity (Blesbois et al., 2013). The length-scales can be computed for a given wall-normal distance by considering the maximum of the streamwise velocity spectrum along a given direction.

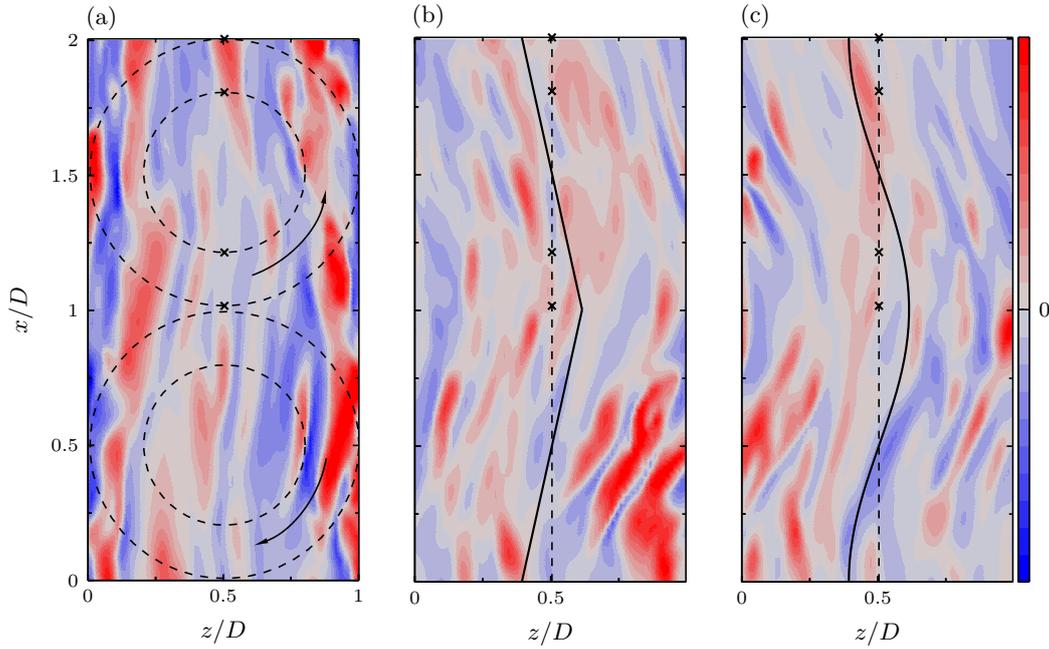


Figure 3.19: Near-wall streaks: instantaneous contours of u_t at $y^{0+} = 11$ for a) ring forcing, b) perfect triangular wave and c) sinusoidal wave. The shape of the spanwise velocity waves is sketched at the centre of each graph. The black crosses mark the location of the wall-normal profiles of Fig. 3.20.

In order to compare the mean flow of the ring forcing to the one produced by the perfect waves, the centreline spanwise velocity profiles are visualised in Fig. 3.20a.

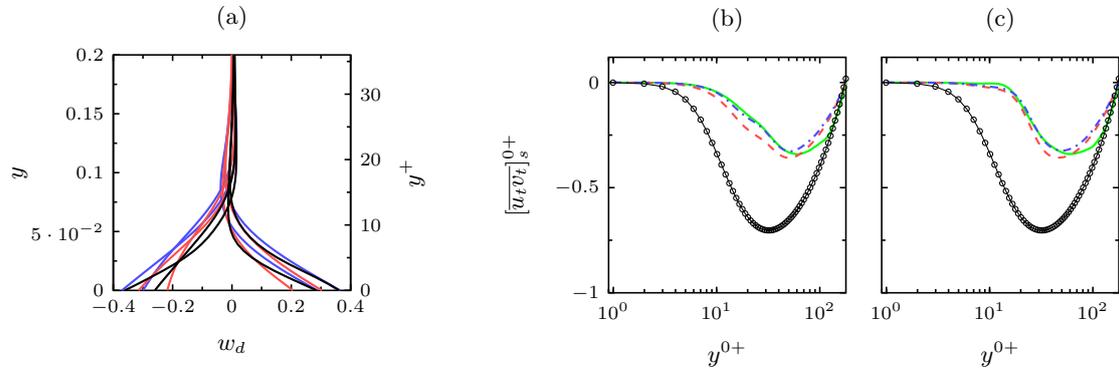


Figure 3.20: Flow statistics for rotating rings and perfect spanwise-velocity waves. The colour key is: R (—), triangular waves (—), spanwise waves (—). (a) wall-normal profiles of the mean spanwise velocity w_d . The profiles are taken at the positions marked by (x) in Fig. 3.19. Profiles at the outer edge (amplitude $\pm W$) and inner edge (amplitude $\pm 0.6W$) of the ring are represented and at the equivalent location in the perfect wave cases.

3.5.3 Disc edge sheltering

The modification to the rotating-disc control presented in this section is motivated by the finding that the upstream-rotating disc edges are responsible for a large portion (around 20%) of the skin-friction drag. suggests that preventing them to interact with the mean shear might produce a more orderly flow, possibly leading to higher drag-reduction. The proposed forcing configuration is illustrated in Fig. 3.21 and consists in fitting the disc array with platelets that suppress the direct interactions between the fast disc edge tips and the mean shear. The screens run continuously in the streamwise direction, they have a spanwise width of $B_{plat} = 0.1D$ and a thickness of $b_{plat}^+ = 1$. The platelet's lower surface is positioned at an height of $y^+ = 10$ above the channel wall. The platelets are modelled numerically via the Immersed Boundary Method of Parnaudeau et al. (2004) that is already integrated into Incompact3D's source code.

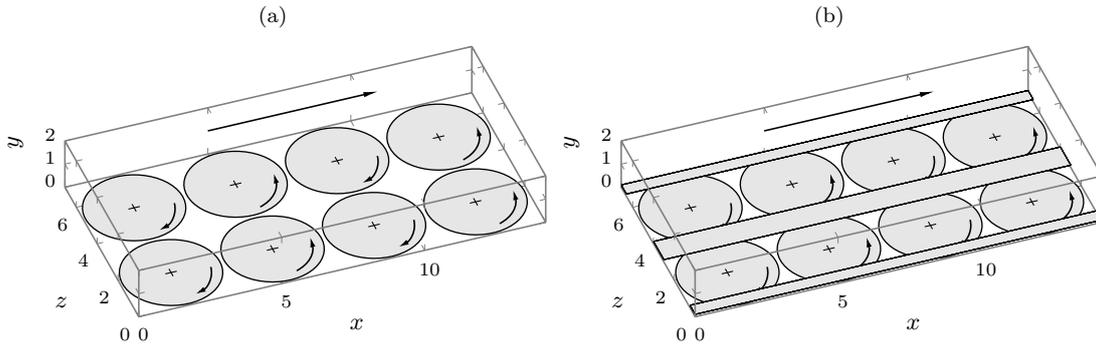


Figure 3.21: Arrangement of the disc-fitted channel, with and without edge sheltering: (a) discs, (b) discs and platelets. The sense of rotation of the discs is highlighted by the curved arrows, while the longer arrows indicates the mean-flow direction.

The fluid domain and the numerical parameters are identical to those used for the ring forcing of Sec. 3.2 and can be found in Tab. 3.6. Although the drag increase from the disc-flow contribution more than doubles with respect to the simple discs, a further 10% reduction of the turbulent contribution C_f^t is achieved, resulting in a more than 7% cumulative improvement in the performance of the scheme. More extensive analysis of the flow field and the skin-friction budget is still necessary to understand the drag reduction performance and speculate on the potential of this scheme.

Table 3.5: Contributions of the Reynolds-stress components to the global skin-friction budget. The values of the skin-friction coefficients are multiplied by 10^3 and the values of the partial drag reduction values are percentages.

Case	C_f	C_f^{lam}	C_f^d	C_f^t	$\mathcal{R}(\%)$	$\mathcal{R}_d(\%)$	$\mathcal{R}_t(\%)$
no-slip	8.12	2.14	0.00	5.97	-	-	-
discs	6.44	2.14	0.25	4.03	20.0	-3.1	23.1
discs+plats	6.10	2.14	0.64	3.10	27.3	-7.9	35.4

3.5.4 Higher-Reynolds flow

In this section we compare the effect of rotating discs in a channel flow at two values of the friction Reynolds numbers, $Re_\tau = 180$ and $Re_\tau = 400$. The work is motivated by the observation (Gatti and Quadrio, 2016), that the efficacy of the spanwise-velocity forcing methods declines with the

Reynolds number. A simulation of a no-slip and a controlled case is carried out at each value of Re_τ , for a total of four simulations. The parameters of the base channel flow and the forcing parameters are listed in Table 3.6. The values of the diameter D^+ and the disc-tip velocity W^+ are kept constant (in the wall units of the reference case at the same Re_p).

Table 3.6: Parameters of the higher-Reynolds channel-flow simulations with rotating discs.

Case	Re_p	Re_τ^0	$D^+ (D)$	$W^+ (W)$	L_x	L_z	$n_x \times n_y \times n_z$	Δy^+	N_D
D180	4200	180	600 (3.38)	9 (0.39)	4.5π	2.3π	$256 \times 129 \times 256$	0.9	2×4
D400	10500	400	600 (1.51)	9 (0.34)	2π	π	$256 \times 257 \times 256$	0.9	2×4

The results of the simulations are reported in Tab. 3.7. It must be noted that in a $Re_\tau^0 = 400$ flow, the turbulent skin-friction C_f^t contributes to an higher share of the total skin-friction than in a $Re_\tau^0 = 180$ flow. However the simulations indicate that at $Re_\tau^0 = 400$ the disc forcing is worse at reducing C_f^t than at $Re_\tau^0 = 180$, thus resulting in an around 4% smaller total drag reduction \mathcal{R} than the lower Reynolds case.

Table 3.7: Comparison of the drag-reduction performance of wall forcing through discs at $Re_\tau^0 = 180$ and $Re_\tau^0 = 400$. The values of the skin-friction coefficients are multiplied by 10^3 .

BC	Re_τ^0	C_f	C_f^{lam}	C_f^d	C_f^t	$\mathcal{R}(\%)$	$\mathcal{R}_d(\%)$	$\mathcal{R}_t(\%)$
no-slip	180	8.12	2.14	-	5.97	-	-	-
no-slip	400	6.35	0.86	-	5.48	-	-	-
D180	180	6.44	2.14	0.25	4.03	20.0	-3.1	23.1
D400	400	5.32	0.86	0.22	4.24	16.2	-3.5	19.5

The elements of the FIK skin-friction budget are visualised in Fig. 3.22, where each term is defined identically to (3.38) and they are normalised with the C_f of the no-slip channel at the corresponding value of Re_τ^0 . Compared to $Re_\tau^0 = 180$, the following differences can be noticed in $Re_\tau^0 = 400$: *i*) the streamwise Reynolds-stress term c_f^t displays an around 30% increase in the central region but remain largely unchanged at the sides, *ii*) the disc-flow term c_f^d is less intense and *iii*) the spanwise inhomogeneous term c_f^z is significantly stronger. The above points require a more extensive analysis of the simulation data to be explained, similarly to what undertaken in Sec. 3.3.1. It is also an open question whether the occurrence of *i*) is connected to the Reynolds-number effects reported in the literature to affect the spanwise-velocity drag-reduction methods.

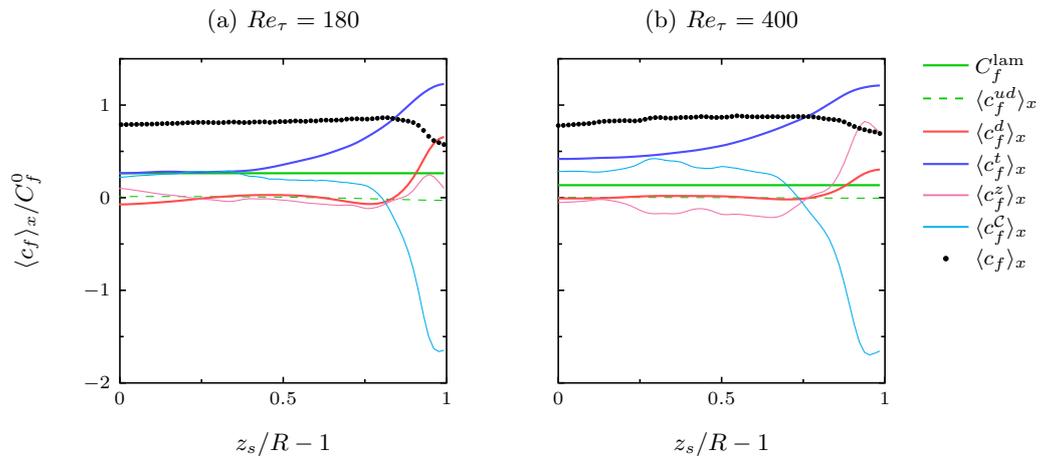


Figure 3.22: Skin-friction budgets of the ensemble-averaged flow for $Re_\tau^0 = 180$ and $Re_\tau^0 = 400$. The curves are normalised with C_f^0 i.e. the skin-friction coefficient of the no-slip case at the same Re_τ^0 .

Appendix to Chapter 3

3.A Identification of coherent vortical structures

The problem of identification of coherent vortices in a fluid flow is addressed in the seminal work of Jeong and Hussain, 1995 where the authors reviewed all the then-current methods and proposed a new one, called the “ λ_2 -method”, based on the spectral analysis of the velocity-gradient tensor. Recalling earlier vortex-identification criteria, the authors highlight three issues regarding the unambiguous identification of vortices:

The criterion that in a vortex the pressure tends to have a minimum on the axis of swirling motion, and that this inward pressure gradient is the centripetal constraint holding the vortex together, is not strictly true other than in steady inviscid planar flows.

Criteria based on streamline patterns, such as spirals or closed paths of revolution around a vortex core, do not satisfy the requirement to be Galilean invariant. The observation of streamlines in different reference frames does not yield the same curves. In addition, the trajectories may not close due to the topology of the vortex rapidly changing.

Isosurfaces of vorticity magnitude $|\boldsymbol{\omega}|$ are a better criterion compared to the previous two but still are not able to capture vortex cores correctly. For instance when shear, walls and vorticity sheets are present it is not able of telling true vortices from other non-vortical structures. The starting point for the λ_2 -method is the method of pressure minima and the observation that its failure is essentially due to two key facts: the pressure minima may be due to unsteady stretching, and the viscous effects may prevent a minimum to occur in the vortex’ core. A reliable criterion must overcome these two issues and must be Galilean-invariant. The local extrema of pressure are determined by the pressure’s Hessian $\partial_{ij}p$. Taking the gradient of the Navier-Stokes equation we have:

$$\partial_j a_i = -\partial_{ij}p + \frac{1}{\text{Re}} \partial_{jkk} u_i \quad (3.47)$$

where $\partial_j a_i$ is the acceleration gradient tensor which can be decomposed into a symmetric and an antisymmetric part, just like the velocity gradient. The symmetric part of (3.47) is:

$$D_t S_{ij} - \frac{1}{\text{Re}} \partial_{jkk} u_i + \Omega_{ik} \Omega_{kj} + S_{ik} S_{kj} = -\partial_{ij}p \quad (3.48)$$

where \mathbf{S} and $\boldsymbol{\Omega}$ are the symmetric and the skew-symmetric parts of the velocity-gradient tensor. The first two terms of (3.48) represent the contribution of the unsteady stretching and viscous effects to the pressure Hessian. Therefore by dropping them we obtain a kinematic characterization of the local extrema of pressure which excludes the two unwanted causes. In order to have a local pressure minimum on a plane, $\partial_{ij}p$ must have two negative eigenvalues:

Definition 1 (Vortex core) *Let the real numbers $\lambda_1, \lambda_2, \lambda_3$ be the three eigenvalues of the symmetric, real-valued tensor $\mathbf{S}^2 + \boldsymbol{\Omega}^2$, being $\lambda_1 \geq \lambda_2 \geq \lambda_3$. A connected region of the flow where $\lambda_2 < 0$ is defined as a vortex core.*

This condition is well-defined as it excludes the possibility of vortex cores on a wall, where $\mathbf{S}^2 + \boldsymbol{\Omega}^2$ vanishes.

Chapter 4

Passively moving discs for frictional drag reduction

4.1 Introduction

While the previous chapter the intent was to with the design and the analysis of better performing, robust and convenient methods, the remainder of this work will focus primarily on the energy efficiency aspect of wall-mounted rotating devices for drag reduction. To this purpose, the energy losses occurring when to actively drive the actuators need to be assessed accurately and realistically and, if possible, strategies to reduce them need to be devised. Furthermore, the ideal solution in terms of energy efficiency is to achieve drag reduction through a purely passive form of flow control, without any external power input. In this direction, the characteristics of passively moving wall mounted discs are studied. In the next Chapter a different path to higher energy efficiency is undertaken, namely the possibility of achieving higher power efficiency through a unsteady, closed-loop control of the disc velocity.

First, we carry out direct numerical simulations of a turbulent channel flow over flush-mounted rotating discs that are free to move under the action of the wall turbulence. We consider the full numerical modelling of the rotating actuators, taking into account the coupled dynamics of the turbulent flow and the discs, and the frictional losses occurring in a realistically designed disc housing.

Full modelling of the actuators in active control methods is critical to their success as it can produce better estimates of the power spent by the actuators, closer to those that would be observed in a real-world setting. This also addresses illustrates the wider task of modelling actuation devices together with the fluid mechanics (i.e. in a two-way coupling), which implies a more accurate prediction of the flow physics and therefore of the drag-reduction characteristics. In the case of the discs, the two-way coupling is able to capture the detailed physics of the stress and torque fluctuations to which a disc rotating under a turbulent flow is subjected. [Józsa et al. \(2019\)](#) simulated the coupled fluid/body dynamics of passive, shear-stress driven cylindrical actuators as a compliant-wall technique able to generate a finite in-plane wall velocity. They proposed an ideal model of the cylindrical actuators and considered the variation of three parameters governing their dynamics to produce drag reduction. [Mahfoze et al. \(2018\)](#) simulated suction and blowing control in a boundary layer and used an empirical model of the frictional losses in the electromagnetic speaker used to generate wall-normal blowing. To the authors' knowledge, no attempt exist yet to model the realistic losses and the coupled fluid/actuator dynamics for a spanwise in-plane wall forcing method.

The dynamics of freely-rotating full discs and half discs is studied. Freely-rotating discs are studied to determine the nature of their shear-driven motion, how it depends on the disc size, and whether it can modify the near-wall turbulence and influence the skin-friction. Half discs are realised by preventing the interaction of one half of a freely rotating disc (in the cross-flow direction) with the flow, thereby producing a non-zero average shear-stress torque on its surface. The half discs rotate with a non-zero velocity and introduce a slip velocity at the wall, potentially affecting the skin-friction. Two configurations of half discs have been the subject of an experimental boundary-layer study by Koch and Kozulovic (2013), who report a reduction skin-friction of up to 17% on the disc surface. The two types of freely-rotating discs can be classified as belonging to the class of compliant-surface methods.

The numerical and statistical procedures that are shared by all the simulations are presented in Sec. 4.2. Sec. 4.3 describes the flow, the actuators and the detailed modelling of their dynamics. In Sec. 4.4 we study of the passive response of freely-rotating discs to the turbulent flow through numerical simulations. In Sec. 4.5, the simulation of an array of half-discs are discussed and compared to the existing experimental work. In Sec. 4.6, summarises the results and conclusions of the study.

4.2 Numerical and statistical procedures

4.2.1 Numerical solver

The open-source code Incompact3D (illustrated extensively in Chapter 1) has been used to simulate all the flows in the present work. The code has been and The freely rotating disc simulations of Section 4.4 have been run on the Cray XC30 ‘‘Archer’’ supercomputer of the National Supercomputing Service, using 1024 parallel computational cores in a 32×32 block decomposition. The half-disc simulations of Section 4.5 have been performed on the ‘‘ShARC’’ cluster at the University of Sheffield using 64 compute cores in an 8×8 block decomposition. All computations are initiated from the laminar Poiseuille channel flow between solid stationary walls, perturbed by random noise. The initial transient flow response to the modified boundary conditions is discarded to acquire meaningful statistics of the controlled flow.

The fluid interacts with the disc actuators via a Dirichlet boundary condition for the velocity on the walls. This means that the boundary conditions of the fluid velocity are updated at each time step in order to realise the wall-normal vorticity given by motion of the disc. The latter is itself determined by a rigid-body dynamical model that takes the instantaneous fluid shear-stresses on its surface as an input. The details of the disc dynamical model are given in Sec. 4.3.1. The time integration of the disc boundary conditions is performed using the same numerical scheme as the fluid solver and the same temporal resolution to preserve the same overall order of accuracy. Further details on the numerical procedures used for the disc boundary conditions are given in the next paragraph.

Numerics of the fluid-body coupling and torque calculation

The disc dynamical model is implemented as a rigid-body boundary condition on the wall. In the context of the fractional-step method, the problem is to update the Dirichlet boundary conditions of the first-step velocity $\mathbf{u}^* = \{u^*, v^*, w^*\}$ as follows:

$$\begin{aligned} u_{n+1}^*(y=0) &= \omega_{n+1} r \sin \theta + \Delta t (\nabla p_n)_x \\ v_{n+1}^*(y=0) &= 0 + \Delta t (\nabla p_n)_y \\ w_{n+1}^*(y=0) &= -\omega_{n+1} r \cos \theta + \Delta t (\nabla p_n)_z \end{aligned} \tag{4.1}$$

where ω_{n+1} is obtained by integrating the discrete-time version of (4.12) and the pressure correction is also applied. The 2nd-order Adams-Bashforth time advancement for (4.12) reads:

$$\omega_{n+1} = \omega_n + \Delta t \left(\frac{3}{2} RHS_n - \frac{1}{2} RHS_{n-1} \right) \quad (4.2)$$

$$\theta_{n+1} = \theta_n + \Delta t \left(\frac{3}{2} \omega_n - \frac{1}{2} \omega_{n-1} \right) \quad (4.3)$$

where:

$$RHS_n = (T_f^n + T_b^n + T_h^n)/I. \quad (4.4)$$

The chosen scheme preserves the same second-order accuracy of the main solver. The first integration step is performed using a simple Euler scheme.

The fluid torque T_f^n is obtained by applying the midpoint quadrature rule to Eq. (4.14) :

$$T_f^n = \frac{\Delta x \Delta z}{Re} \sum_{i,k} \left(x_i \frac{\partial}{\partial y} w_{i0k}^n - z_i \frac{\partial}{\partial y} u_{i0k}^n \right) \quad (4.5)$$

for i, k such that $\sqrt{(i\Delta x - x_C)^2 + (k\Delta z - z_C)^2} \leq D/2$

where (x_C, z_C) is the disc centre. The quadrature error for the midpoint rule is of the order $\mathcal{O}(\Delta x^2 \Delta z^2)$ for well-behaved functions.

The Power Spectral Densities of T_f and W are estimated via the SciPy implementation of the Welch algorithm (<https://github.com/scipy/scipy/blob/v1.2.1/scipy/signal/spectral.py#L294-L457>) applied to the time series. Hanning windowing is used, the segment length is set at 80000 points (i.e. 200 Poiseuille time units) and the overlapping at 40000 points.

4.2.2 Reference channel flow and non-dimensional units

We first describe the fixed-wall numerical case that is used as the reference flow to the other flow cases and to compute the drag-reduction performance. Unless explicitly stated otherwise, the same numerical set-up illustrated here is used for the freely-rotating cases described in Section 4.4.2 and the half-disc flow cases of Section 4.5.3.

An incompressible turbulent channel flow with a constant flow rate is studied via direct numerical simulation (DNS). The channel walls are separated by a distance $L_y^* = 2h^*$, where henceforth the * superscript indicates a dimensional quantity. The channel is assumed to be periodic in the streamwise (x^*) and spanwise (z^*) directions. The dimensions of the computational domain in the x^* and z^* directions are L_x^* and L_z^* , and the velocity components in the x^* , y^* and z^* directions are u^* , v^* , and w^* , respectively. The flow is driven by an uniform pressure gradient P_x^* along the x^* direction, which is adjusted at each time-step to maintain a constant flow rate per unit wall-area Q^* . The bulk velocity U_b^* is obtained by dividing Q^* , by the channel height as $U_b^* = Q^*/2h^*$. A schematic of the reference channel configuration is given in Fig. 4.1.

Non-dimensional quantities are obtained by choosing the channel half-height h^* as the unit of length, the fluid density ρ^* as the unitary density, and the unit of velocity as U_p^* , the centreline velocity of the laminar Poiseuille flow at the same flow rate, defined as $U_p^* = 3Q^*/4h^*$. The unit of time is thus h^*/U_p^* and the unitary stress is $\rho^* U_p^{*2}/h^*$. This normalisation choice defines the outer units of the flow and quantities expressed in outer units are written with no superscript. All simulations are conducted at a Reynolds number $Re_p = U_p^* h^*/\nu^* = 4200$, where ν^* is the kinematic viscosity of the fluid.

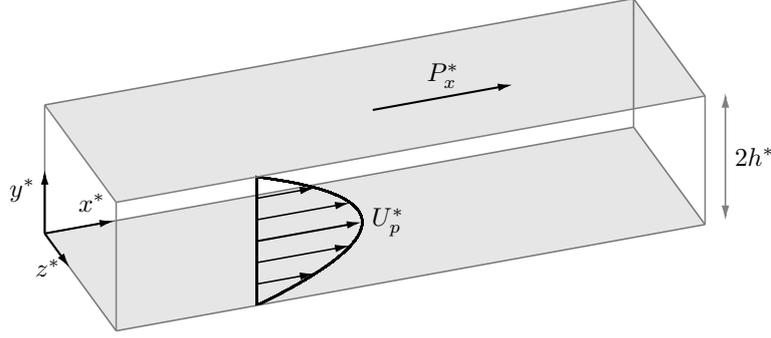


Figure 4.1: Schematic of the base plane channel used as the reference flow.

The reference stationary-wall friction Reynolds number for the simulations is $Re_\tau = u_\tau^* h^* / \nu^* = 180$, where $u_\tau^* = \sqrt{\tau_w^* / \rho^*}$ is the wall-friction velocity and τ_w^* is the mean wall shear-stress. Quantities scaled in viscous units, i.e. with respect to the scales that characterise the fluid motion near the wall, are denoted by a + superscript. The viscous length is $\delta_\nu = Re_\tau^- 1$ and the viscous time unit is $t_\nu = \delta_\nu / u_\tau$. Unless otherwise stated, scaling in viscous units is performed with respect to the reference stationary-wall case.

4.2.3 Averaging operators and performance quantities

The time-average of a variable f is indicated by \bar{f} . When estimating the mean of a disc variable, such as the disc-tip velocity W , we always average over each disc, i.e. $\bar{f} = \sum_i^{N_D} |\bar{f}_i| / N_D$ where \bar{f}_i is the signed time-averaged variable from the i -th disc, and N_D is the number of discs.

We give here the definition of the spatial period-averaging that is applied to the array of rotating half discs used in Section 4.5.3. For the flow field variables such as the velocity \mathbf{u} and pressure p , the periodic-averaging operator has to account for the spatially periodic symmetry induced by the array of rotating discs. Given the rectangular array of Fig. 4.19, the mean flow obeys the following linear transformation:

$$\mathbf{u}(x, y, z, t) = \mathbf{u}(x + D, y, z + D, t), \quad (4.6)$$

$$p(x, y, z, t) = p(x + D, y, z + D, t). \quad (4.7)$$

for any x, y, z .

The period-averaged flow field is therefore defined as:

$$\{[\mathbf{u}]_s, [p]_s\}(\mathbf{x}) = \frac{1}{N_D} \sum_{l=1}^{r_D} \sum_{m=1}^{c_D} \{\mathbf{u}, p\}(x + lD, y, z + mD, t). \quad (4.8)$$

where r_D and c_D are the number of rows and columns of an array of $N_D = r_D \times c_D$ discs. Using the definitions of the period-averaged flow field $\mathbf{u}_s = [\mathbf{u}]_s$ and $p_s = [p]_s$ we introduce the following flow decomposition:

$$\{\mathbf{u}, p\} = \{\mathbf{u}_s, p_s\} + \{\mathbf{u}_t, p_t\}. \quad (4.9)$$

where $\{\mathbf{u}_t, p_t\}$ represent the random fluctuations. The period-averaged component can be further decomposed into the spatially averaged flow $\{\mathbf{U}, P\}$ (the “velocity profile”) and the spatially inhomogeneous component $\{\mathbf{u}_d, p_d\}$ as:

$$\{\mathbf{u}_s, p_s\} = \{\mathbf{U}, P\} + \{\mathbf{u}_d, p_d\}. \quad (4.10)$$

The spatial average $\{\mathbf{U}, P\} = \{\langle \mathbf{u} \rangle, \langle p \rangle\}$ is obtained by integrating over x and z and dividing by the wall area $L_x L_z$, thus obtaining wall-normal profiles $\{\mathbf{U}(y), P(y)\}$.

The global skin-friction coefficient is defined by normalising the mean wall-shear stress τ_w with the bulk kinetic energy as $C_f = 2\tau_w/U_b^2$, where $\tau_w = dU/dy|_{y=0}/Re_p$. The percentage skin-friction drag reduction $\mathcal{R}(\%)$ is defined as:

$$\mathcal{R}(\%) = 100 \frac{C_{f,0} - C_f}{C_{f,0}}, \quad (4.11)$$

where the subscript “0” denotes the reference channel flow. It is also possible to define the local drag reduction $\mathcal{R}_{xz}(x, z)$ by using the spatially dependent, period-averaged wall-shear stress $\tau_w^s(x, z) = du_s/dy|_{y=0}/Re_p$.

4.3 Description and modelling of the actuators

One or more discs are flush mounted on the surfaces of the upper and the lower wall of the channel. The discs have radius $R^* = D^*/2$, are rigid and homogeneous, and are free to rotate about their central axis. The discs are characterised by the thickness b^* and material density ρ_s^* . The moment of inertia of a disc around the central axis is $I^* = \pi b^* D^{*4} \rho_s^*/32$. A thin gap of width $0.05D^*$ is modelled around the edge of each disc to simulate the clearance region between the disc and the stationary channel walls. Within the annular gap the velocity is modelled to decay linearly from its maximum value on the disc edge to zero on the no-slip wall. This treatment of the gap closely represents a real experimental set-up where such gaps would always be present, and it drastically reduces the numerical oscillations that would arise from a discontinuity in the velocity profile.

Each disc is fitted with a short shaft and housed in a cavity in the wall deep enough for the upper surface of the disc to be flush with the channel wall. A layer of fluid of thickness d_h^* exists between the lower surface of the disc and the bottom of the cavity. The shaft is supported by a bearing mounted into a purpose-built hole at the bottom of the disc housing. The disc housing has a similar design to those realised in the experimental study of [Koch and Kozulovic \(2013\)](#) and in a preliminary experimental apparatus built at the Deutsches Zentrum für Luft- und Raumfahrt (DLR) Göttingen (M. Rutte and U. G. Becker, private communication). A schematic of the disc housing is shown in Fig. 4.2.

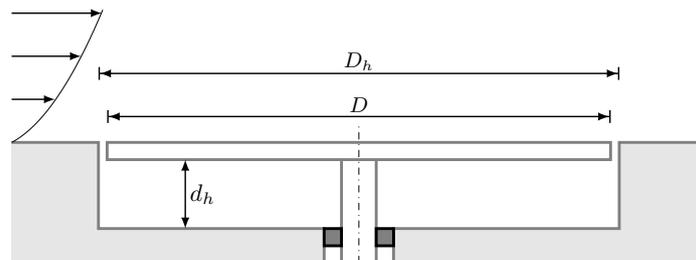


Figure 4.2: Schematic of the disc housing. The axis of rotation of the disc is indicated by a dash-dotted line.

In order to simplify the simulations, the fluid within the cavity will not be simulated, instead exact laminar models of the torque in a cylindrical cavity with a rotating lid will be used. The model is presented in Section 4.3.3, where it is also shown that the models are in realistic agreement with the flow regime that is expected to exist in the cavity in the range of disc velocities considered in this work.

4.3.1 Dynamics of the rotating discs

The motion of each disc at time t is described by its angular velocity $\Omega(t)$ or, alternatively, by its outer-edge or disc-tip velocity $W(t) = R\Omega(t)$. The disc-tip velocity evolves according to the following dynamical equation:

$$\frac{2I}{D} \frac{dW}{dt} = T_f + T_h + T_b \quad (4.12)$$

where T_f is the torque exerted by the wall-shear stresses of the channel flow on the flow-facing surface of the disc, T_h is the torque arising from the fluid contained in the housing beneath the disc and T_b is the frictional torque within the bearing. Using the non-dimensional units given in Section 4.2.1, torques are scaled with $\rho^* U_p^{*2} h^{*2}$ and the moment of inertia is scaled with $h^{*5} \rho^*$. The definitions and models for each torque component are detailed in the following sections.

4.3.2 Torque from fluid shear stresses

The fluid torque T_f on each disc results from the wall-shear stresses exerted by the turbulent flow over the surface \mathcal{S} of the disc and is expressed by the integral:

$$T_f(t) = \frac{1}{Re_p} \iint_{\mathcal{S}} \boldsymbol{\tau} \times (\mathbf{x} - \mathbf{x}_C) d\mathcal{S} \quad (4.13)$$

$$= \frac{1}{Re_p} \int_{-R}^R \int_{-\sqrt{R^2-(x-x_C)^2}}^{\sqrt{R^2-(x-x_C)^2}} [(z - z_C)\tau_x - (x - x_C)\tau_z] dz dx, \quad (4.14)$$

where R is the disc radius, $\mathbf{x}_C = (x_C, z_C)$ are the coordinates of disc centre and the components of the shear-stress $\boldsymbol{\tau}$ are defined as:

$$\tau_x = -z \frac{1}{Re_p} \frac{\partial u}{\partial y} \Big|_{y=0} \quad (4.15)$$

$$\tau_z = x \frac{1}{Re_p} \frac{\partial w}{\partial y} \Big|_{y=0}. \quad (4.16)$$

The integral in (4.14) is evaluated at each time-step and the boundary condition on the disc are updated accordingly, realising a fully coupled simulation of the body-fluid dynamics. More details on the numerical approximation to (4.14) are given in Appendix 4.2.1.

4.3.3 Resisting torques from disc housing

The motion of the fluid within the cavity results in a resisting torque T_h acting on the shaft, that can be calculated, knowing the flow field in the cavity, by applying (4.14) to the lower surface of the disc. For the aims of the present work, we want to use an approximate predictive model for T_h , thus avoiding the need to explicitly simulate the flow within the cavity.

To introduce the problem, we consider an idealised case where the disc angular velocity fluctuates following a monochromatic sinusoidal wave of zero mean and frequency f . The amplitude of the fluctuating angular velocity is denoted by Ω' . The dependence of the resisting torque T_h^* on the parameters of the flow within the cavity follows a non-dimensional law of the form:

$$T_h^* = \Omega'^{*3} D^{5*} \rho^* F_h \left(\frac{D^*}{d_h^*} d_h^* \sqrt{\frac{f^*}{\nu^*}}, \frac{\Omega'^* d_h^{*2}}{\nu^*} \right). \quad (4.17)$$

where F_h is an unknown function and d_h^* is the depth of the cavity. When the aspect ratio $D^*/d_h^* \gg 1$, the discs can be regarded as of infinite size and the effect of the cylindrical cavity

walls neglected. When the frequency parameter $d_h^* \sqrt{f^*/\nu^*}$ is much smaller than unity, the flow is in a quasi-steady regime (a similar definition for a different oscillating-flow configuration is found in [Barengi and Jones, 1989](#)). The last parameter is the cavity Reynolds number $Re_h = \Omega^* d_h^{*2}/\nu^*$ which determines the flow regime inside the cavity. If both the infinite size and quasi-steadiness can be safely assumed, the model (4.17) becomes much simpler, depending only on Re_h .

The infinite size is assumed and verified for most of the geometries discussed in this work. The extreme cases are $D^*/d_h^* = 2.8$ and $D^*/d_h^* = 67$. In two flow cases in Section 4.5 the assumption is purposefully violated for reasons that are made clear therein.

In the flow cases presented below, the disc velocity fluctuates randomly, i.e. it shows a continuous frequency spectrum. The quasi-steady analysis has to be adapted to this case by considering the power spectrum of the angular velocity and calculating the frequency parameter over the energy-containing frequency range. If a large fraction of the disc kinetic energy resides in frequencies f^* where $d_h^* \sqrt{f^*/\nu^*} \ll 1$, the quasi-steadiness condition can be considered satisfied. For freely-oscillating discs that have zero-mean angular velocity, our simulations show that this criterion is satisfied. In the half-discs cases, which rotate with a non-zero mean angular velocity, the quasi-steadiness assumption is always verified due to the standard deviation of the velocity fluctuations being much smaller than the mean component.

From the assumption of infinite disc moving quasi-steadily, it follows from (4.17) that the disc dynamics depends uniquely on the cavity Reynolds number. For randomly fluctuating full discs, Re_h is defined using the standard deviation of the angular velocity, as $Re_h = \Omega_{SD}^* d_h^{*2}/\nu^*$, whereas for the half-discs Re_h the mean angular velocity is used, i.e. $Re_h = \bar{\Omega}^* d_h^{*2}/\nu^*$. The steady flow between two coaxial infinite discs, one at rest and the other moving with angular velocity Ω , has been studied by [Stewartson, 1953](#) in the laminar regime of small to moderate Re_h . Within the laminar regime, for very small values of Re_h the equations of motion become linear and the azimuthal velocity u_θ is independent of the other two components. The solution takes the simple form:

$$u_\theta = \frac{r\Omega}{d_h} y. \quad (4.18)$$

where the polar coordinates r, θ, y with origin at the centre of the disc are used. The associated resisting torque, in non-dimensional form is:

$$T_h = -\frac{\pi D^4}{32 d_h Re_p} \Omega. \quad (4.19)$$

Even though the linear regime holds exactly for $Re_h < 1$, the linear solution describes reasonably well the cavity flow up to around $Re_h = 10$ ([Stewartson, 1953](#)). When $Re_h > 10$ the non-linear effects become increasingly important and the flow field is in a transitional regime where the velocity gradient on the disc at rest decreases and the one on the rotating disc grows. For around $Re_h > 40$ the swirling boundary-layer solution of [Von Kármán \(1921\)](#) for a disc beneath a semi-infinite fluid applies to the cavity flow. The non-dimensional torque of the Von-Kármán solution reads:

$$T_h = -\text{sgn}(\Omega) \frac{\pi G D^4}{32 Re_p^{1/2}} |\Omega|^{3/2}, \quad (4.20)$$

where $G = 0.6159$ is a numerically-determined constant.

In principle, the validity of the assumptions has to be verified by carrying out a full DNS simulation of the cavity flow for each case and measuring the resulting value of the frequency parameter and of Re_h . In this work we will assume quasi-steadiness in all the cases and restrict the validation to a self-consistency test. This means performing the simulations using one of the

approximate steady laminar models (4.19) or (4.20), and verifying that the resulting frequency parameter assumes sufficiently low values in the sense of the spectral analysis described above. In a similar way, the consistency of the laminar flow regime and the choice between the Von-Kármán or Couette model for T_h are verified by checking that Re_h , calculated with the disc velocity from the simulation, belongs to the appropriate range. In the freely rotating cases, the laminar regime and the Couette model are always found to be consistent with the assumptions, while for the half discs a comprehensive analysis of the results is found in Section 4.5.4.

a. Flow in a rotor-stator cavity

Several experimental and numerical studies exist on the resisting torque acting on finite discs that rotate in an enclosed cylindrical cavity with stationary walls. These works can be used to estimate the flow regime occurring in the cavity beneath the freely rotating (full and half) disc cases. The flow regime occurring in an enclosed cavity with a steadily rotating disc is classified, in the seminal work of Daily and Nece (1960), based on the two parameters $Re_R = \Omega^* R^{*2} / \nu^*$ and d_h^* / R^* , analogous to those used in (4.17). Some authors also stress that the presence of a central hub constitutes a third parameter that must be carefully considered when determining the flow regime (Lauder et al., 2010). However, data at very low Re_R are only available for hub-less cavities. Review of several studies (Daily and Nece, 1960; Lauder et al., 2010; Schouveiler et al., 2001) provides convincing evidence that in the freely rotating, and most of the half disc cases the cavity flow is correctly identified by the linear Stewartson solution; this also holds for low aspect ratios comparable the smallest freely rotating cases (Pao, 1972; Bertelá and Gori, 1982).

4.3.4 Frictional torque of a ball bearing

Each disc is supported by a ball bearing fixed onto its shaft, as shown in dark grey in Fig. 4.2. The resisting torque T_b of a rolling-element bearing is the combination of a number of phenomena such as lubrication, material deformation, thermal losses and others. Its complete modelling is complex and all the practical estimates are usually performed through the use of empirical formulae (Harris and Kotzalas, 2006).

The main component of T_b is the load-induced rolling friction, that arises from the pressure contact between the solid surface of the rotary elements and the metal grooves (the “races”), similar to train wheels on a rail track. In our application, the critical aspect is the capability of the bearing to support axial loads (i.e., the weight of the disc) without generating high friction torques. Radial loads such as the effect of the shear forces on the surface of the disc are orders of magnitude smaller.

Our estimates are based on the empirical formulae provided by the leading manufacturer of rolling-element bearings, the Swedish company SKF, which are freely available online (www.skf.com). This requires the selection of a real bearing model using dimensional quantities, obtained by considering the hypothetical water channel presented in Sec. 4.3.5. An example of a reasonable choice is represented by the SKF angular-contact thrust bearings of the series 7009, that produce a small amount of friction under high axial loads and are available in a size range that is compatible to the dimensionalised size of the disc. Using the bearing data-sheet parameters and an axial load corresponding to the weight of a disc of diameter $D = 5$ the SKF formulae return a friction torque that, translated back into non-dimensional units is $|T_b| = 0.0001$. In the same way, the values of $|T_b|$ for discs of different diameter are calculated using the same bearing and changing the axial load intensity. Using a value of $|T_b|$ that does not depend on W is justified because at the slow rotation rates found in the simulations, only rolling

friction occurs, that depends on the loading conditions but is nearly independent of the velocity. Accordingly, the bearing torque enters the RHS of (4.12) as $T_b = |T_b|\text{sgn}(W)$.

In a practical realisation of a water-immersed bearing, the presence of watertight seals might be unavoidable in order to protect the lubricant. Here however seals are neglected because, for a slowly rotating and lightly loaded bearing, lubrication plays a secondary role and can be avoided altogether with modern solutions, such as water-lubricated, ceramic-ball bearings. A velocity-independent model of the bearing frictional torque is also suggested in the experimental papers Koch and Kozulovic (2013) experimental work, where the bearing is not lubricated and thus subjected to pure rolling friction.

4.3.5 Flow parameters and dimensional reference

The flow statistics of the passively-rotating-disc channel flow, for example the skin-friction coefficient C_f and the disc-edge velocity W , can be written as a function of the parameters:

$$C_f = f(Re_p, D, d_h, T_b, \rho_s). \quad (4.21)$$

Tab. 4.1 shows the characteristics of the water channel used as a reference when dimensional quantities are needed. For the reference water channel in Tab. 4.1 the disc thickness translates to the realistic values of 0.5mm and the cavity depth to 10mm. The disc used in the bearing selection example above would result in $D^* = 0.5\text{m}$ and a weight of 0.27Kg. The material density ρ_s^* is assumed as that of aluminium, or 2.7 times the water density $\rho^* = 1000\text{kg/m}^3$, which translates into a non-dimensional disc density of $\rho_s = 2.7$. This allows us to calculate the dimensional disc weight and inertia needed for the bearing selection, and to determine the dimensional resisting torques.

The choice of $\rho_s = 2.7$ allows disc angular accelerations of a magnitude around 1000 times those possible in an air channel at the same Re_p , for which instead $\rho_s = 2200$. This is possible because the disc inertia is linearly dependent on ρ_s while the order of magnitude of the fluid shear-stresses on the disc surface is dictated by Re_p . This is relevant to this work because larger angular accelerations lead to higher oscillation amplitudes and therefore shorter simulation times with better statistical convergence, due to shorter integral time-scales of the disc variables. Moreover, this choice of dimensional units is also important for a practical experimental realisation of freely-rotating full discs, where one wants to produce measurable oscillations using a disc actuator that is constructively feasible with respect to the disc thickness and material choice.

Table 4.1: Physical characteristics of the water channel taken as the reference for dimensional quantities.

h^*	U_p^*	ν^*	ρ^*	Re_p	b^*	d_h^*
0.1 m	0.042 m/s	10^{-6} m ² /s	10^3 kg/m ³	4200	0.5 mm	10 mm

4.4 Freely rotating discs

This section presents the results of simulations involving freely rotating full discs, i.e. discs that are not driven by any external power source, but are set in motion by the action of the fluid shear stresses on their surface. The computational domain is depicted in Fig. 4.3 and it consists of two isolated coaxial discs located on the two opposite walls. This configuration allows a flexible modification of the disc diameter with little change to the numerical domain. The domain dimensions are $L_x = 4.53\pi$, $L_y = 2$ and $L_z = 2.26\pi$ and are solved on a discrete grid of size $256 \times 129 \times 256$.

For a single diameter ($D = 3.38$) the isolated-disc set-up was compared to a simulation of an array configuration, resulting in minor differences in the statistics. No drag reduction (local or global) is measured, neither in the isolated nor in the arrayed configuration. An isolated disc was also simulated within a wider channel, in order to make sure that no numerical issues related to the spanwise periodicity are present.

Only the diameter is varied, while the thickness $b = 0.005$ and the material density $\rho_s = 2.7$ remain constant. The resulting moments of inertia for each different diameter are listed in Table 4.2. The cavity depth is also kept unchanged at $d_h = 0.1$.

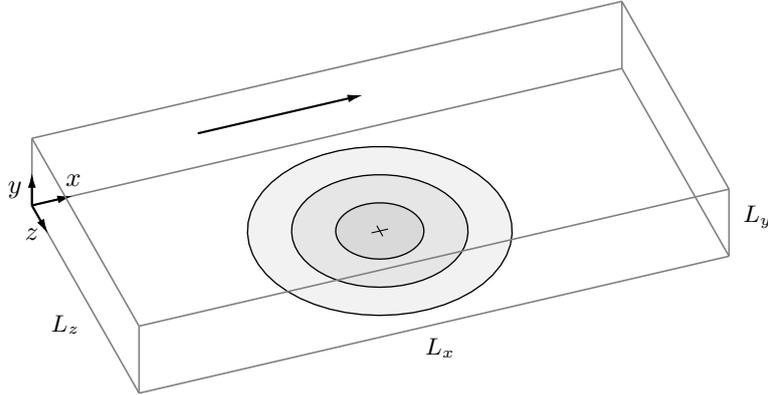


Figure 4.3: Schematic of the arrangement of the numerical channel used to simulate isolated, freely-rotating discs of increasing diameter. Only the disc on the bottom wall is shown. Three discs of different diameter are sketched.

We perform simulations of both the coupled and the uncoupled disc-fluid system, according to the framework introduced in Section 4.4.1. In Sec. 4.4.2 simulations of the coupled system are carried out for eleven different diameters, whereas in Sec. 4.4.3 the uncoupled dynamics of the discs is simulated independently of the fluid equations and compared to the coupled case. In Section 4.4.4, the system is analysed in the frequency-domain, especially the uncoupled case.

4.4.1 Description of the disc-fluid system

In this section we consider the dynamical system composed by the rotating disc and the fluid flow, and we develop some analytical tools that are applied in the following section to study the freely-rotating disc dynamics.

The dynamics of the disc obeys equation (4.12), while the fluid dynamics evolves according to the Navier-Stokes equations (NSE) and their boundary conditions (BCs). The two sets of governing equations are mathematically coupled because of the presence of two elements. First, the fluid torque T_f (4.14) couples the fluid to the disc, acting as an input term to the disc dynamics. Second, the disc boundary conditions to the NSE couple the disc to the fluid dynamics. On a disc centred at (x_C, z_C) and rotating with a disc-edge velocity W , the velocity boundary condition for the streamwise and the spanwise velocity components read:

$$u(x, 0, z, t) = W(t)(z - z_C) \quad (4.22)$$

$$w(x, 0, z, t) = -W(t)(x - x_C). \quad (4.23)$$

The system is visualised in Fig. 4.4 using a block diagram (Åström and Murray, 2008). A block can be viewed as a black-box function that takes one (or more) input and returns an output.

Each block in Fig. 4.4 corresponds to an element of the system, following the description outlined above. Since we are primarily interested in the motion of the disc, the diagram features the disc velocity as the main output. It can be noted that the system displays a closed-loop shape that is due to the two coupling elements, i.e. the fluid torque and the fluid BCs. A simple form of the response functions is difficult to calculate in the coupled case because they depend on their input and output through the dependence of the entire system on W .

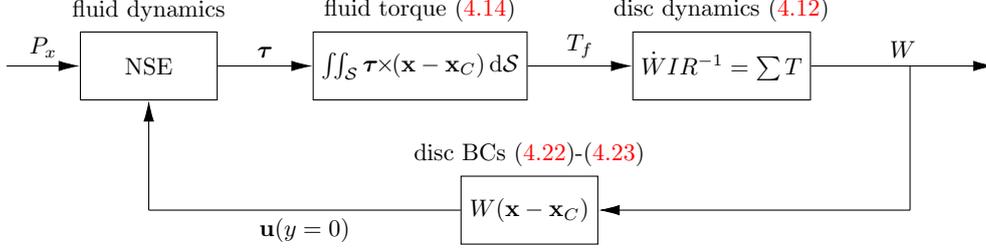


Figure 4.4: Block-diagram representation of the disc-fluid system as a feedback loop on the fluid.

Fig. 4.5 depicts the uncoupled system, where the discs are fixed ($W = 0$) and thus there is no feedback from the disc to the fluid boundary conditions. It follows that the system consists in a simple open-loop series, equivalent to considering the no-slip channel flow, calculating the shear-stress torque on a disc-shaped region of the wall, and independently solving the disc dynamics.

The uncoupled model has a number of advantages. Firstly, contrary to the coupled case, it does not require to perform a new DNS simulation for each disc of different diameter, but it is sufficient to simulate a no-slip channel flow once and extract the fluid torque on many disc-shaped wall patches simultaneously. This allows the study of a larger number of cases and benefits from potentially longer integration times to collect statistics. Secondly, when the system is studied in the frequency domain (Section 4.4.4), linearity makes it possible to calculate the frequency response of the disc to the fluid torque. Third, the uncoupled results can be compared to those of the coupled simulations, assessing the impact of the dynamical coupling on the dynamics.

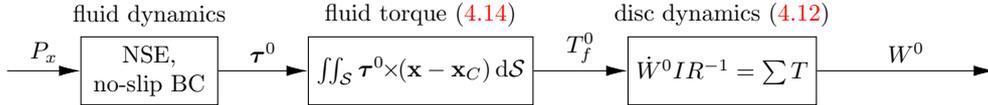


Figure 4.5: Block-diagram representation of the uncoupled system.

The uncoupled system of Fig. 4.5 can be characterised, in a more rigorous way, as the zeroth-order term of a regular expansion of diagram 4.14a in powers of W . For example, expanding the wall-shear stress τ with respect to the disc velocity around $W = 0$ gives:

$$\tau(\mathbf{x}, t; W) = \tau^0(\mathbf{x}, t) + \mathcal{O}(W), \quad (4.24)$$

meaning that for very small values of W , the fixed-wall stress τ^0 is a zeroth-order approximation to the fully coupled stress τ . The feedback loop vanishes because the zeroth-order term of the disc boundary conditions (4.22)-(4.23), which are linear in W , is zero. To emphasise the reduced-order view, the elements of the uncoupled system are indicated by “0” subscripts and superscripts.

4.4.2 Numerical simulation of coupled discs-fluid dynamics

The dynamical equation for the disc motion used in the simulations assumes the linear housing torque solution (4.19). Its non-dimensional form for the disc-tip velocity W reads:

$$\frac{dW}{dt} = \frac{D}{2I}T_f(t) - \frac{C_h}{I}W(t) - T_b, \quad (4.25)$$

where $C_h = \pi D^4/(32Re_p d_h)$ is the constant from the linear torque model (4.19). The bearing torque T_b is typically one order of magnitude smaller than T_h and therefore can be considered as a minor source of friction. Eleven diameters in the range $D^+=50$ to $D^+=1000$ are considered. Three time-histories of the disc-tip velocity, each 400 time units long, for three values of the diameter are shown in Fig. 4.6a. The simulations establish that for all the diameters considered, the discs oscillate randomly around $W = 0$. The variance of W is larger at small diameters, the oscillations of the disc becoming less intense as the diameter increases. The torque time-series on the same time interval are shown in Fig. 4.6b and they reveal an opposite trend, the variance of T_f growing rapidly with the diameter.

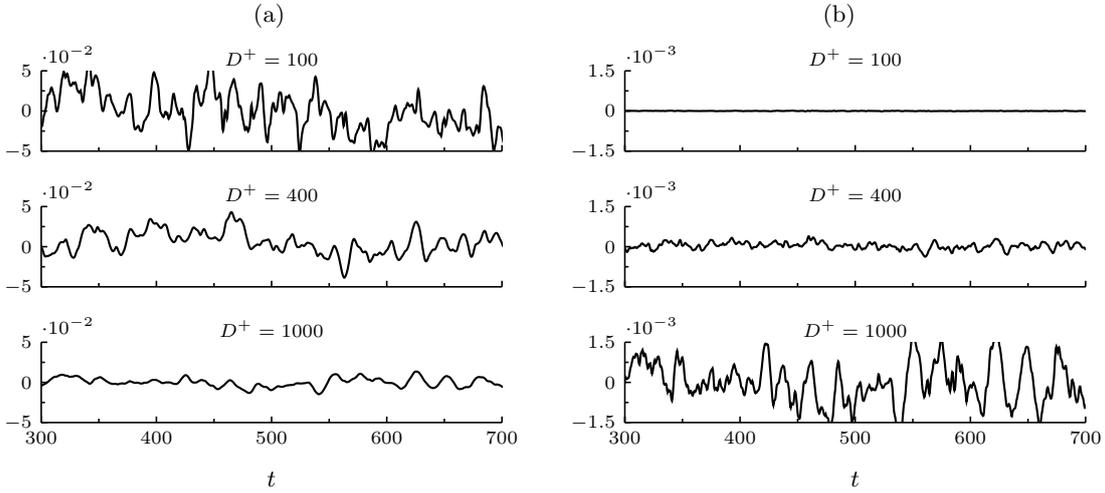


Figure 4.6: Extracts of the DNS time-series for three discs of increasing diameter. (a) disc-tip velocity W . (b) fluid torque T_f .

The r.m.s. estimates of the standard deviation of W and their chi-squared confidence intervals are reported in Table 4.2 for all the diameters. No reduction of the frictional drag is obtained. The work of Ricco and Quadrio (2008) and Quadrio and Ricco (2011) for transverse wall motions and Ricco and Hahn (2013) for rotating discs established that a minimal value of the velocity at the wall of $W = u_\tau$ is required for the mean wall gradient to be affected by the wall motion. The r.m.s. of the disc-tip velocity fluctuations W^+ remains always < 1 and does not affect the near-wall skin-friction significantly.

Table 4.2: r.m.s. of $W(t)$ from the simulations of the fully coupled freely rotating discs. The confidence intervals shown are at the 95% level. The two rightmost columns report the main results in the dimensional the water channel case of Tab. 4.1.

D	D^+	$I \times 10^3$	rms $W \times 10^2$	rms W^+	CI on rms $W \times 10^2$	D^* [mm]	W^* [m/s]
0.28	50	0.01	1.94	0.45	+1.57, -1.16	2.8	0.081
0.56	100	0.1	2.31	0.54	+1.87, -1.39	5.6	0.097
0.84	150	0.7	2.01	0.47	+1.63, -1.21	8.4	0.084
1.13	200	2.10	1.83	0.43	+1.48, -1.10	11.3	0.077
1.69	300	10.8	1.51	0.35	+1.22, -0.91	16.9	0.063
2.25	400	34.2	1.23	0.29	+1.00, -0.74	22.5	0.052
3.38	600	173	0.95	0.22	+0.77, -0.57	33.8	0.040
4.44	800	515	0.79	0.19	+0.64, -0.48	44.4	0.033
5.00	900	828	0.70	0.16	+0.56, -0.42	50.0	0.029
5.60	1000	1350	0.60	0.14	+0.49, -0.36	56.0	0.025
6.70	1200	2670	0.58	0.13	+0.47, -0.35	67.0	0.024

Figs. 4.7a-b depict the r.m.s. values of W and T_f for all the diameters considered in the simulations. The blue series represents the data from Table 4.2. The red series are the results from the uncoupled simulations and will be discussed in Sec. 4.4.3.

In the range $D^+ < 100$, the r.m.s. W follows a rapidly increasing regime, the maximum being located at around $D^+ = 100$. It can be noticed that the maximum response of the disc is for $D^+ = 100$, which is comparable with the characteristic spanwise spacing of the low-speed streaks. The r.m.s. of W scales approximately as $D^{-0.55}$ in the the $100 < D^+ < 900$ range.

The r.m.s. of $T_f(D)$ is a monotonically increasing function of D^+ . In the range $0 < D^+ < 100$, $T_f(D)$ shows a positive slope, that is steeper at lower D^+ . The r.m.s. T_f is proportional to $D^{2.2}$ in the range $100 < D^+ < 900$, as obtained by logarithmic regression.

For very small D , the torque is too small to exert the required power to move the discs, while for large discs, the inertia thwarts the action of the torque. The maximum r.m.s. W occurs at the disc size where the action of T_f is most effective. These qualitative considerations are explained mathematically by the frequency-domain analysis of the disc dynamics carried out in Section a..

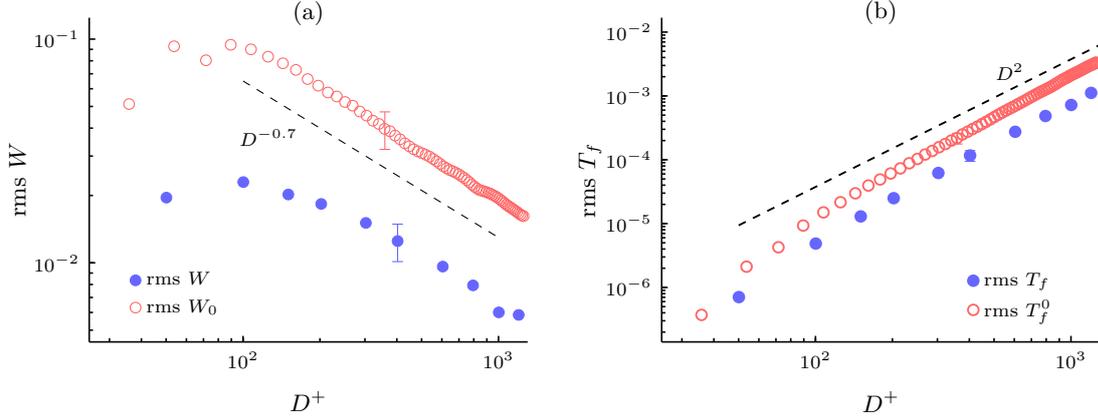


Figure 4.7: Root mean square values of (a) the disc-tip velocity and (b) the fluid torque as a function of the disc diameter. The time-averaging interval is $2500h^*/U_p^*$ or $19300t_\nu$ for each simulation. The confidence intervals are constant for any D when plotted in logarithmic coordinates, therefore only one is plotted for reference. The blue filled symbols are the coupled case data in Table 4.2, while the red circle symbols are the uncoupled case data i.e. W_0 and T_f^0 , which will be discussed in Sec. 4.4.3.

We also examine how the two components of the shear-stress contribute to the fluid torque on the disc surface. The streamwise component T_f^x corresponds to (4.14) with only the streamwise shear-stress term of the integrand:

$$T_f^x(t) = \frac{1}{Re_p} \int_{-R}^R \int_{-\sqrt{R^2-x^2}}^{\sqrt{R^2-x^2}} -(z - z_C)\tau_x dz dx. \quad (4.26)$$

where the point-wise torque contribution is $\tau_x = Re_p^{-1}\partial u/\partial y|_{y=0}$ and z_C is the spanwise position of the disc centre. The spanwise component T_f^z is defined analogously, with the corresponding point-wise contribution being $\tau_z = Re_p^{-1}\partial w/\partial y|_{y=0}$. It follows that $T_f = T_f^x + T_f^z$.

Fig. 4.8 visualises two instantaneous snapshots of the two point-wise torque contributions on the channel wall. As the moment arm grows linearly with the distance from the disc centre, the most intense contributions to the torque come from the edge regions.

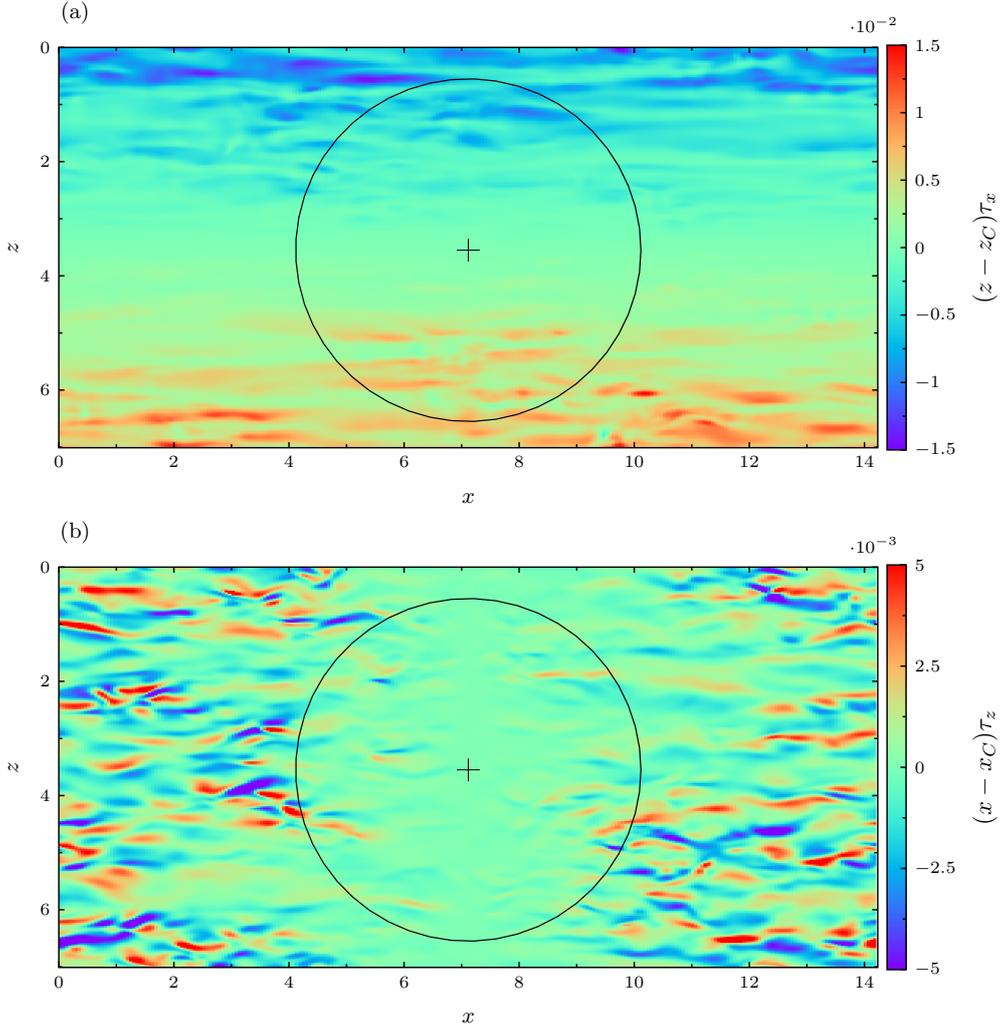


Figure 4.8: Instantaneous snapshot of the point-wise torque at the wall for a disc with $D = 6$: (a) contribution from the streamwise wall-shear stress, (b) contribution from the spanwise wall-shear stress.

Fig. 4.9a (bottom) reports a time-history snapshot of T_f^x and T_f^z for the coupled case, it can be seen there is a noticeable inverse correlation between T_f^x and T_f^z . From the definitions of the torque components, the torque variance is expressed as:

$$\text{Var } T_f = \overline{T_f^{x2}} + \overline{T_f^{z2}} + 2\overline{T_f^x T_f^z}, \quad (4.27)$$

where the last term is the covariance of the streamwise and the spanwise terms. The dependence of the right-hand side terms on the disc diameter are shown in Fig. 4.9b by the red unfilled symbols. All three components across both cases share a very similar power-law dependence on the disc diameter. The streamwise-stress components dominate the other two components, explaining more than 90% of the total variance. Additionally, the negative contribution from the correlation component is larger (in absolute value) than the contribution of the spanwise component.

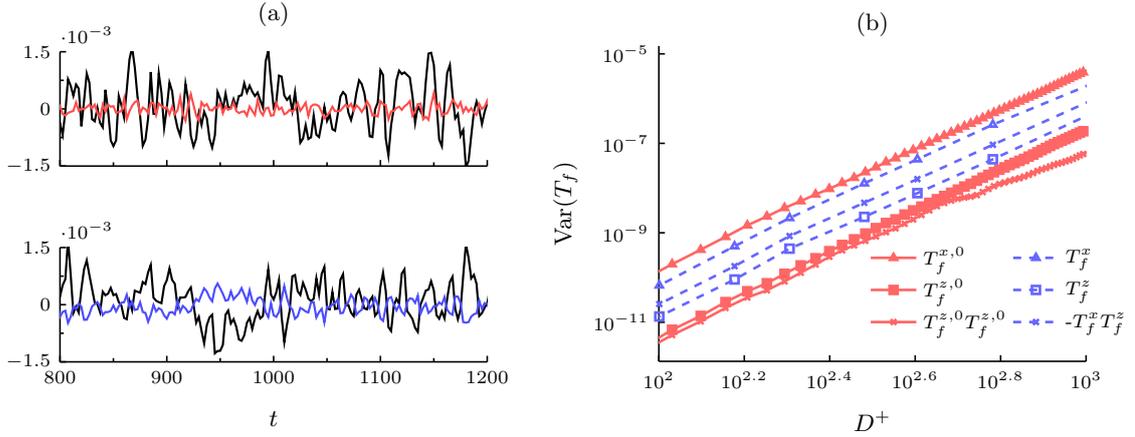


Figure 4.9: Spanwise- and streamwise-stress components of the fluid torque. (a) time-series of $T_f^x(t)$ and $T_f^z(t)$ for the uncoupled (top) and the coupled case (bottom). (b) mean-square values and covariance as a function of the disc diameter. The blue series correspond to the data from the coupled simulations, the red series are from the uncoupled simulations.

The standardised histograms for W and T_f are shown in Fig. 4.10 and compared to the standard normal pdf. The values of the disc-tip velocity W appear to be normally distributed for all the diameters. The tail behaviour of T_f instead indicates deviation from normality for the smaller diameters. The pdf of T_f for $D^+ = 50$ has heavy tails, with the estimated value of the standardised fourth moment being > 4 to a statistically significant level. A possible explanation of the tail behaviour is that, owing to the small size of the disc (smaller than the spanwise integral scale of the wall shear-stress of $\approx 100\delta_\nu$), the torque fluctuations depend on the localised extreme events that characterise the positive tail of the wall shear-stress distribution (Hu et al., 2006). For larger discs, the extreme shear-stress events are likely to play a more marginal role because they tend to be averaged out on a larger wall region.

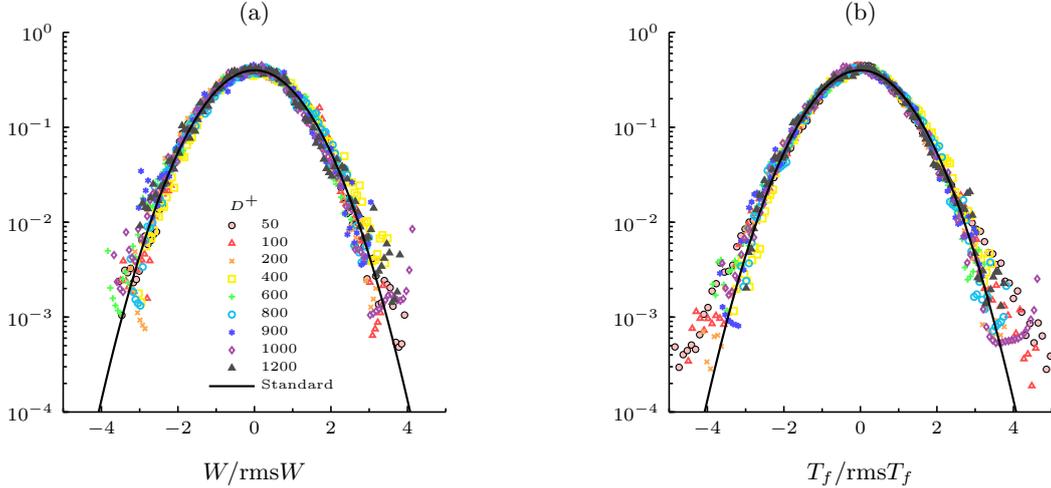


Figure 4.10: Coupled dynamics: Standardised histogram plots of (a) W and (b) T_f , compared to the standard normal distribution.

The Power Spectral Densities for the disc-tip velocity and the fluid torque are shown in Fig. 4.11. Normalising the PSDs with the total power does not result in a collapse of the curves for all diameters. At smaller diameters, the higher frequencies carry a larger contribution to the total power, progressively shifting to lower frequencies as D increases. For small D , the power contributions are spread over a broad high-frequency range, while for moderate values of D they are concentrated around $f = 0.04$, that is an oscillation period of almost $200t_\nu$. Such behaviour is observed in the PSDs of both W and T_f . The peak frequency and the shape of the spectra are determined by how the spatial and temporal scales of the wall shear-stress are filtered by the disc (Quadrio and Luchini, 2003). However, while T_f maintains a peak around $f = 0.04$, the spectral peak of W is rapidly lost, and energy is distributed more evenly at frequencies below $f < 0.03$.

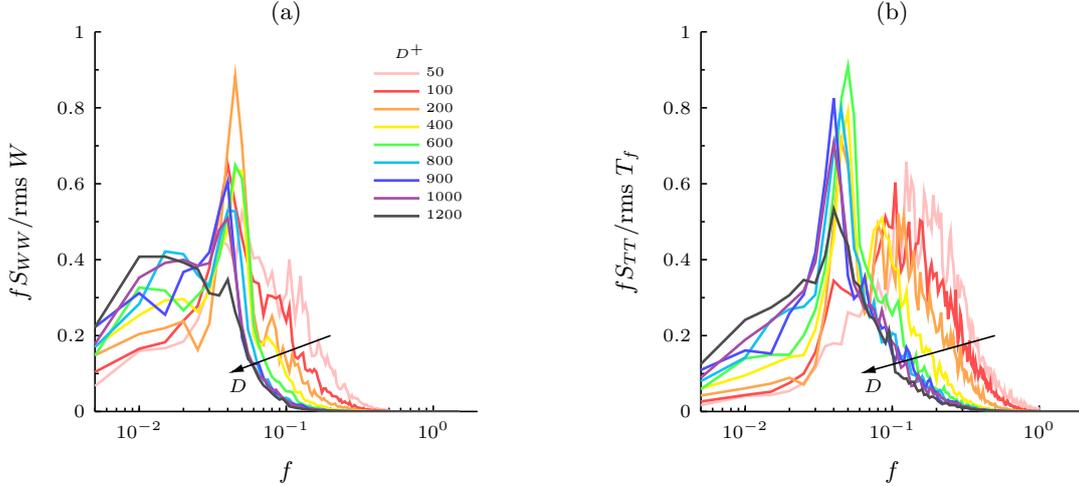


Figure 4.11: Coupled dynamics: PSDs normalised with the total power and shown in a pre-multiplied form. (a) disc-tip velocity. (b) fluid torque.

4.4.3 Uncoupled dynamics

In this section we study the dynamics of the discs subject to the fixed-wall fluid torque for a number of diameters. The disc dynamics is described by the following linear equation:

$$\frac{dW_0}{dt} = \frac{D}{2I} T_f^0(t) - \frac{C_h}{I} W_0(t) + T_b, \quad (4.28)$$

and is integrated independently of the equations of motion of the fluid.

The forcing torque T_f^0 is sampled from a single DNS of the fixed-wall channel at every time step, corresponding to a sampling frequency of $400U_p/h$. The time-series of T_f^0 is then used as an external forcing to numerically integrate (4.28) at the same temporal resolution. The integration is carried out through a second-order Adams-Bashforth scheme, identical to that employed to simulate the disc dynamics in the coupled simulations. Seventy diameters in the range $D^+ = 20$ to $D^+ = 1200$ have been considered, the discs are initially at rest, i.e. $W_0(0) = 0$ and the first 125 time units are discarded. Additionally, since (4.28) is linear, certain exact models can be derived for its average dynamics that provide a useful tool for understanding the dependence on D in both the uncoupled and the coupled cases. This will be further discussed in Sec. 4.4.4.

Fig. 4.7 (red symbols) shows that the r.m.s. of the uncoupled disc-tip velocity fluctuations is inversely proportional to the disc size for $D^+ > 100$, as it is in the fully coupled case. In that

range, the r.m.s. of W_0 scales as $D^{-0.7}$ as calculated by logarithmic regression. The fixed-wall fluid torque T_f^0 is also observed to follow a similar behaviour to the coupled case. The dependence of T_f^0 on D in the large- D range is also consistent with the $D^{2.2}$ scaling observed in the fully coupled case. Neither scaling holds in the range $0 < D^+ < 100$, where W_0 show an opposite trend and T_f^0 a steeper slope, as already found in the fully coupled dynamics.

The r.m.s. values of both T_f^0 and W_0 are between 2.5 and 3 times higher than the fully coupled case, regardless of D . This implies that the two-way interaction between the disc and the fluid produces an attenuating effect on the disc velocity and the torque. The r.m.s. intensity of this effect also scales approximately as $D^{2.2}$ in the large- D regime, as shown by the curves for the coupled and uncoupled cases remaining approximately parallel in log-log coordinates. At the present stage, the physical mechanism of the torque attenuation in the two-way coupling still remains an open question. Addressing this issue is complicated by the turbulence being locally modified by the disc motion.

The time-series of the fluid torque components $T_f^x(t)$ and $T_f^z(t)$ for the uncoupled case are shown in Fig. 4.9a (top). In comparison to the coupled case, the uncoupled simulations demonstrate a qualitatively different behaviour. In the uncoupled case there is no noticeable correlation between T_f^x and T_f^z , contrary to what found in the coupled simulations.

The contribution to $\text{Var}(T_f)$ from the streamwise, spanwise, and covariance components for the uncoupled case is shown in Fig. 4.9b (right) by the blue symbols. Similar to the coupled simulations, the streamwise-stress component dominates the other two components. The magnitude of the contributions from $\overline{T_f^{z,0}}$, and $\overline{T_f^{x,0}T_f^{z,0}}$ is much smaller than in the coupled simulations, and the contribution from the covariance term is > 0 .

Fig. 4.12 shows the standardised histograms of W_0 and T_f^0 for 9 diameters. The values of W_0 are normally distributed for all the diameters. As also observed in the fully coupled case, the tail behaviour of T_f^0 deviates from normality at small diameters. It follows that the coupling has a negligible effect on the shape of these distributions.

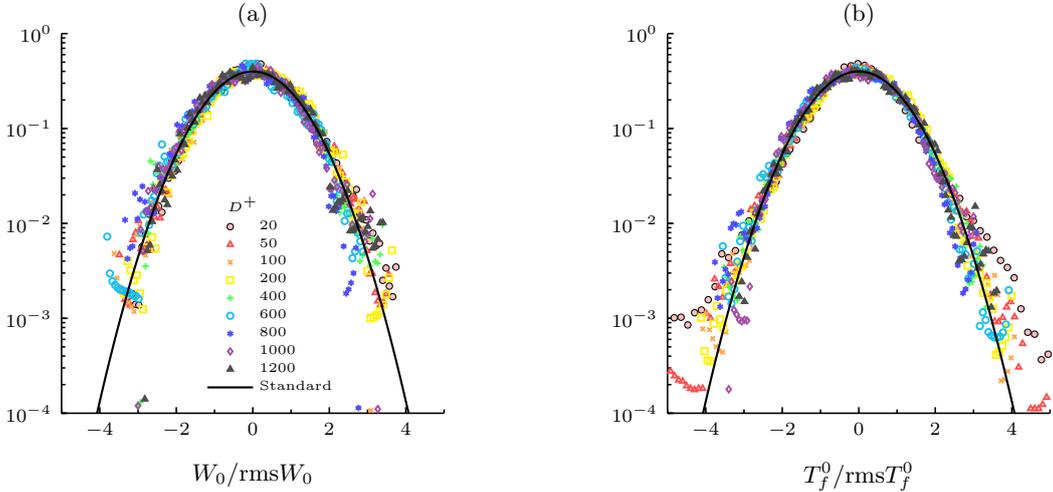


Figure 4.12: Uncoupled dynamics: standardised histogram plots of (a) W_0 and (b) T_f^0 , compared to a standard distribution.

The normalised PSDs of the disc-tip velocity and the fluid torque are shown in Fig. 4.13. Similar to that observed for T_f , the PSD of T_f^0 also presents a red shift for increasing D and a peak contribution at around $f=0.04$. The PSDs of W_0 also follow an analogous behaviour to the

coupled case, with the lower frequencies contributing more to the total power as D increases. In general the PSDs exhibit a flatter shape than in the coupled case, the energy being more evenly distributed with less pronounced peaks.

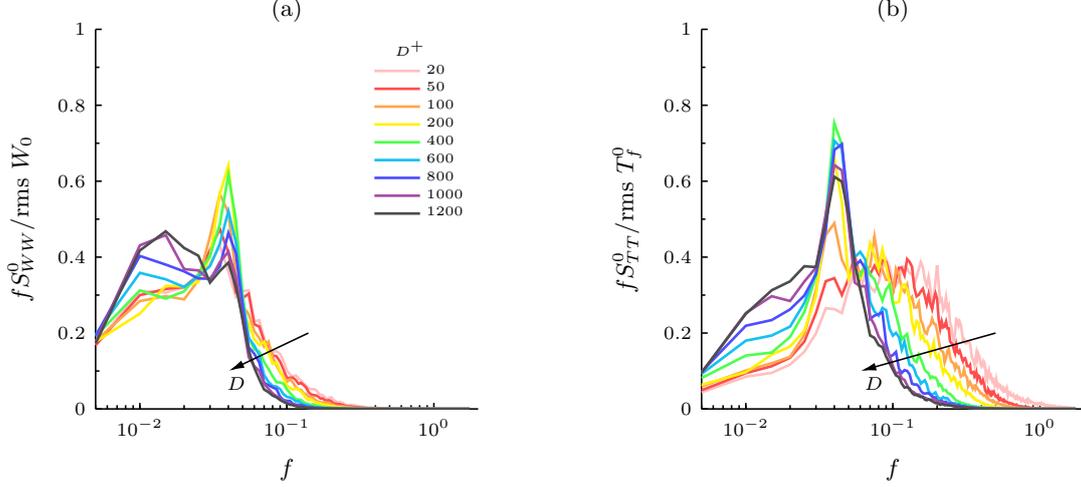


Figure 4.13: Uncoupled dynamics: PSDs normalised with the total power and shown in a pre-multiplied form. (a) disc-tip velocity PSDs, (b) fluid torque PSDs.

4.4.4 Frequency-domain analysis

The most natural approach to study the dynamics of freely-rotating discs under a statistically stationary turbulent flow is to investigate their behaviour in the frequency domain. Since the disc-fluid system is time-invariant and statistically stationary, the analysis is carried out through Fourier analysis.

The block diagrams in Fig. 4.14 provide a schematic view of the system in the frequency domain. The fully coupled system, represented by Fig. 4.14a is the frequency-domain equivalent of Fig. 4.4, the blocks representing each subsystem, namely the fluid governing equations (NSE), the torque induced by the wall-shear stresses on the disc, the dynamics of the rigid disc and the fluid boundary conditions. Each subsystem is characterised by a response function, i.e. a sort of spectral filter function that maps the Fourier transform of the input to that of the output.

The fluid governing equations with the boundary conditions produce the wavenumber-frequency FT of the wall-shear stresses $\hat{\tau}(\kappa_x, \kappa_z, f)$, where κ_x , and κ_z are the streamwise and spanwise wavenumbers, f is the frequency and the hat indicates the Fourier transform of a variable. In the time domain, \hat{G} corresponds to the fluid torque integral (4.14) and \hat{F} corresponds to the disc dynamical equation (4.25). The response function \hat{G} acts on $\hat{\tau}$ and returns the torque FT as $\hat{T}_f(f) = \hat{G}(\hat{\tau})$. The rigid-disc response function \hat{F} in turn characterises the rigid-disc velocity response $\hat{W}(f)$ to the torque \hat{T}_f , i.e. $\hat{W}(f) = \hat{F}(\hat{T}_f)$. The response functions are difficult to study in the coupled case because they depend on their input and output and the entire system is non-linear. In the uncoupled system of Fig. 4.14b, the disc dynamics is linear and thus its response function can be calculated analytically. Moreover, the analysis of the fluid torque in the frequency domain of is also more tractable. To emphasise the reduced-order view, the elements of the uncoupled system are indicated by “0” subscripts and superscripts.

In the next sections we apply the framework of Fig. 4.14 to understand the dependence of the variance of T_f and W on the diameter. Specifically, we construct models for \hat{G}_0 and \hat{F}_0 and

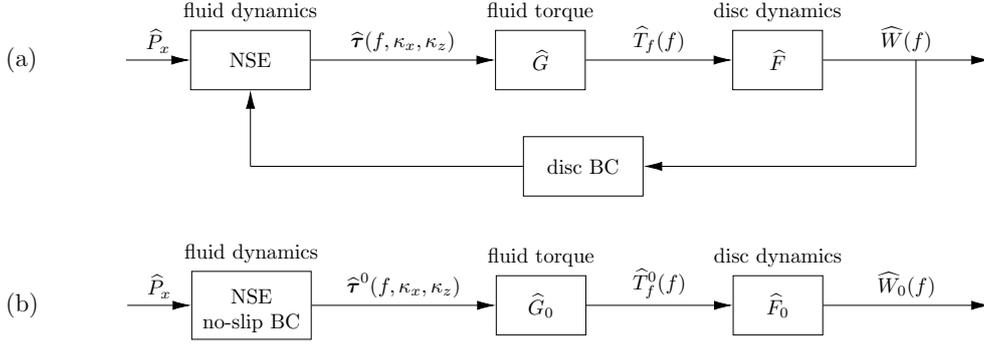


Figure 4.14: Block-diagram representation of the disc-fluid system in the wavenumber-frequency domain. (a) fully coupled case, (b) uncoupled case.

analyse how they depend on D , which is helpful in elucidating the dependence on D of the PSDs shown in Fig. 4.13 and the variance of T_f and W shown in Fig. 4.7 (i.e. is the integral of the PSDs). Sec. b. investigates the modelling of \hat{G}_0 and its application to the uncoupled dynamics. In Sec. a. we explain the r.m.s. response of the discs to the fluid torque by deriving a model of \hat{F}_0 . The results are discussed and compared with the coupled case.

a. Models for the disc dynamics

In this section we examine the frequency-domain response of the uncoupled disc to the turbulent fluid torque, and compare it to the fully coupled case. As (4.28) is linear and time-invariant, a steady-state frequency-domain solution can be found by applying the Fourier analysis. Taking the Fourier transform of both sides (excluding T_b that is negligible) and reordering the terms gives:

$$2\pi i f \hat{W}_0 = \frac{D}{2I} \hat{T}_f^0 - \frac{C_h}{I} \hat{W}_0 \quad (4.29)$$

$$\hat{W}_0 = \left(\frac{c_1}{D^3(2\pi i f + c_2)} \right) \hat{T}_f^0 \quad (4.30)$$

$$(4.31)$$

where i is the imaginary unit, f is the frequency, and the hat indicates the temporal Fourier transform of a variable. $c_1 = 16/\pi b \rho_s$ and $c_2 = 1/(Re_p d_h b \rho_s)$ are two constants. The expression in brackets in (4.31) is the linear transfer function for (4.28), that is defined as a function \hat{H}_0 such that $\hat{H}_0 = \hat{W}_0/\hat{T}_f^0$. Using the transfer function, we can write the exact linear form of the uncoupled response function as $\hat{F}_0(\hat{T}_f^0) = \hat{H}_0 \hat{T}_f^0$ and calculate the PSD response as:

$$S_{WW}(f) = |\hat{H}_0(f)|^2 S_{TT}(f), \quad (4.32)$$

where the transfer function magnitude, or gain, reads:

$$|\hat{H}_0| = \frac{c_1}{D^3 \sqrt{4\pi^2 f^2 + c_2^2}}. \quad (4.33)$$

Equation (4.33) characterises the frequency response of the uncoupled disc to the external input torque T_f^0 .

It is also useful for later discussion to derive a simplified model of the dynamics, where only the terms independent of W are kept on the right-hand side of (4.31):

$$\frac{dW_{0,s}}{dt} = \frac{D}{2I} T_f^0(t). \quad (4.34)$$

Applying Fourier analysis to (4.34) and calculating the modulus of the transfer function gives:

$$|\widehat{H}_{0,s}| = \frac{c_1}{D^3 2\pi f} \quad (4.35)$$

where $c_1 = 16/\pi b \rho_s$.

Fig. 4.15b depicts the reduced-order response magnitude $|\widehat{H}_0|$ scaled with D^3 . The fully coupled response functions \widehat{H} , also shown in Fig. 4.15b, are calculated from the coupled simulations according to the definition of linear transfer function i.e. $|\widehat{H}|^2 = S_{WW}/S_{TT}$. Since the system is not linear this is not a rigorous operation, however it provides a convenient tool to compare the coupled system to the reduced-order model and assess whether and when the second is applicable to the first. By inspecting (4.33) and its application to a number of diameters in Fig. 4.15b, a few facts can be outlined:

- i)* The reduced-order transfer function (4.33) reveals an exact D^{-3} dependence due to inertia, which is already present in the zeroth-order model (4.35). As \widehat{H}_0 exhibits an inverse dependence on ρ_s and to b , implying that a thinner disc or a lighter material would result in larger velocity amplitudes, all other parameters held constant.
- ii)* \widehat{H}_0 shows a f^{-1} slope at high frequencies and flattens out at low frequencies. These two features are also present in the fully coupled transfer functions \widehat{H} , although with quantitative differences. While the f^{-1} dependence is already present in the zeroth-order model (4.35), the flat shape at low frequencies is only explained after accounting for the disc housing torque, represented mathematically by the constant c_2 in the denominator of the augmented model (4.33).
- iii)* In general, the reduced-order model over-predicts the coupled response intensity i.e. $\widehat{H}_0(D) \geq \widehat{H}(D)$ for any D , with the difference increasing at lower frequencies and smaller D s. The difference is progressively reduced for increasing D , decreasing as $D^{-0.8}$. At large diameters, $\widehat{H}(D)$ is observed to converge to a $\propto D^3$ function that is very closely approximated by \widehat{H}_0 .
- iv)* As \widehat{H}_0 is proportional to D^{-3} and the r.m.s. of T_f^0 to $D^{-2.2}$, one may expect the r.m.s. of W_0 to scale as $D^{-0.8}$. However, as shown in Fig. 4.7, the r.m.s. of W_0 scales as $D^{-0.7}$. The discrepancy is explained by the spectral red shift of T_f^0 at large D (see Fig. 4.13). Since \widehat{H}_0 is a decreasing function of f , the red shift ultimately allows more power to contribute to the r.m.s. of W at higher D s. In a similar fashion, the gentler slope of the r.m.s. of the fully coupled W is associated with the lower intensity of \widehat{H} at smaller diameters.

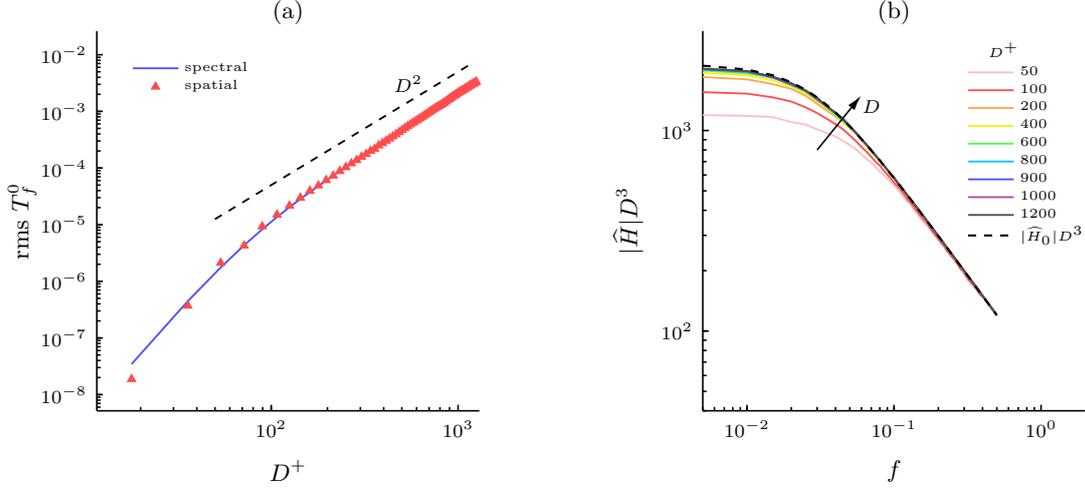


Figure 4.15: Uncoupled modelling. (a) Dependence of the r.m.s. T_f^0 on the disc diameter for the fixed-wall channel, as calculated from the spatial-domain shear stress fields and by the spectral-domain model, (b) Response functions from the uncoupled exact theory (4.33) (dashed lines) and calculated from the coupled DNS simulations (solid lines).

b. Models for the fluid torque

In this section we study how the statistical properties of the fluid torque on a fixed disc originate from the statistical properties of the wall-shear stress, and how they depend on D . In the framework of Fig. 4.14, this means studying the function \widehat{G}_0 . Examining the fixed-wall case is convenient because the wall shear-stress is statistically homogeneous and therefore analytically more tractable.

In order to keep the analysis simple and focused on the primary object of interest, that is the dependence of the torque (variance) on D , we introduce an additional assumption: exploiting the fact that the flow is statistically stationary, we neglect the dependence of T_f^0 on time, and study its variance as resulting from an ensemble of temporally uncorrelated realisations $T_f^0(t_i)$ where the $t_i^{N_t} = 1$ are sufficiently delayed to be temporally uncorrelated. This simplified approach is sufficient to study the dependence of $\text{Var}(T_f^0)$ (i.e. the integral of $S_{TT}^0(f)$) on D , but not of \widehat{G}_0 , for which it is necessary to consider full wavenumber-frequency spectra of the shear stresses. However, since only the spatial spectra are used, the more difficult computation of the wavenumber-frequency spectra can be avoided.

The analysis is carried out as follows: first, we derive the expression of a single realisation of $T_f^{0,x}$ – i.e. the torque induced by a single instantaneous snapshot of the wall-shear-stress field – and we study how it changes with the disc size; secondly, we calculate numerically $\text{Var}(T_f^{0,x})$, using the ensemble of uncorrelated snapshots mentioned above.

We recall that the instantaneous torque T_f^0 acting on the disc surface \mathcal{S} of radius $R = D/2$ at time t is given by (4.14), and can be decomposed into the two contributions from τ_x^0 and τ_z^0 , respectively called $T_f^{0,x}$ and $T_f^{0,z}$. For ease of explanation, in the following we will derive formulas for the streamwise-stress contribution $T_f^{0,x}$ only. The spanwise contribution is analogous and will be discussed afterwards.

The fluid torque $T_f^{0,x}$ acting on a disc of radius R at time instant t_i can be expressed as:

$$T_f^{0,x}(t_i; R) = \iint_{\mathcal{S}} z \mathcal{F}_{xz}^{-1} \{ \hat{\tau}_x^0(\kappa_x, \kappa_z) \} dx dz, \quad (4.36)$$

where $\hat{\tau}_x^0$ is the spatial Fourier transform of $\tau_x^0(x, z)$ and $\mathcal{F}_{xz}^{-1}\{\cdot\}$ its inverse. By explicitly writing the inverse Fourier transform and swapping the order of integration ($\hat{\tau}_x^0$ being spatially homogeneous) we obtain:

$$T_f^{0,x}(t_i; R) = \iint_{\mathbb{R}^2} \hat{\tau}_x^0(\kappa_x, \kappa_z) \left(\frac{1}{4\pi^2} \iint_{\mathcal{S}} z e^{i\kappa_x x} e^{i\kappa_z z} dx dz \right) d\kappa_x d\kappa_z \quad (4.37)$$

$$= \iint_{\mathbb{R}^2} \hat{\tau}_x^0 \Psi_R^x d\kappa_x d\kappa_z. \quad (4.38)$$

where the expression in brackets in (4.37) corresponds to the function $\Psi_R^x(\kappa_x, \kappa_z)$ that encodes the filtering effect of the disc size and the moment arm z on the shear-stress spectrum. Its definition can be simplified (refer to Appendix 4.A for the calculations) to the following self-similar form:

$$\Psi_R^x(\kappa_x, \kappa_z) = R^3 \Psi(R\kappa_x, R\kappa_z), \quad (4.39)$$

where Ψ is defined (using the generic variables κ_1, κ_2) as follows:

$$\Psi(\kappa_1, \kappa_2) = \frac{1}{2\pi^2 \kappa_1} \int_{-1}^1 z \sin(\kappa_1 \sqrt{1-z^2}) e^{i\kappa_2 z} dz. \quad (4.40)$$

No closed-form solution is found for the remaining integral, which is therefore evaluated numerically.

We note that since the integral in (4.38) is real-valued (because $T_f^{0,x}$ is real number), the integrand can be reduced to its real part $\mathbf{Re}\{\hat{\tau}_x^0 \Psi_R^x\} = |\hat{\tau}_x^0| |\Psi_R^x| \cos(\Phi_\tau + \Phi_\psi)$, where Φ_τ and Φ_ψ are the phase angles of $\hat{\tau}_x^0$ and Ψ_R^x respectively. In the following paragraphs, the magnitudes $|\hat{\tau}_x^0|$ and $|\Psi_R^x|$ are studied, because they are determined numerically to be the largest contribution to $\text{Var}(T_f^{0,x})$.

The phase spectra Φ_τ and Φ_ψ also contribute to $\text{Var}(T_f^{0,x})$, even though only marginally compared to the magnitudes. The fixed-channel simulation data establish that Φ_τ is a uniformly distributed random phase. Φ_ψ (shown in Fig. 4.16c) takes discrete values of either $-\pi/2$ or $\pi/2$, the sign being dictated by the same radial period observed for the oscillations of $|\Psi_R^x|$.

Fig. 4.16a shows the magnitude $|\Psi_R^x|$ of the spectral filter function, calculated by solving the integral in (4.40) numerically. The axes use the reciprocal of the wavelength (i.e. $1/\lambda_x = \kappa_x/2\pi$) normalised with the disc diameter in order to visualise the filtering behaviour of Ψ_R^x relative to the disc size. $|\Psi_R^x|$ is an even function with respect to both wavenumbers and thus only the first quadrant is shown. A series of regularly-spaced oscillations can be observed, centred at the origin and decreasing in intensity at higher values of the radius-wavelength ratio. The period of these oscillations, calculated along a radial direction centred at the origin, is $D/\lambda=1$ where $\lambda = \sqrt{\lambda_x^2 + \lambda_z^2}$. $|\Psi_R^x|$ privileges spanwise modes, gradually decreasing at low spanwise wavenumbers and vanishing completely for purely streamwise modes. The maximum of $|\Psi_R^x|$ is determined to be a purely spanwise mode such that $D/\lambda_z = 0.732$. The largest individual contribution to the torque is given by the spanwise mode of τ_x^0 whose wavelength is 1.37 times the diameter.

The map of Fig. 4.16a depict $|\Psi_R^x|$ when $D = 1$. When the disc size is increased to values > 1 , the coordinates scale linearly with D , therefore in absolute terms the maximum of $|\Psi_R^x|$

moves (linearly with D) closer to $\kappa_z = 0$, the period of the oscillations decrease linearly and areas decrease quadratically, generating an effect of “shrinking” centred at the origin. Similarly, when $D < 1$, $|\Psi_R^x|$ undergoes an “expansion” according to the same linear proportionality to D . The second effect of varying the disc size is that $|\Psi_R^x|$ is uniformly amplified by a factor of R^3 , according to (4.39). Larger discs therefore produce a very strong, localised filtering near the origin while smaller discs have a weaker, spread-out filtering across a broad wavenumber range.

Fig. 4.16b displays the square root of the PSD of the streamwise shear-stress. The PSD shows where the contributions to $|\widehat{\tau}_x^0|$ are located on average. The axes use the reciprocal wavelengths without normalisation (absolute units) as the wall shear-stress does not depend on D . As τ_x^0 is a real function, the amplitude of its PSD is even and thus only the first quadrant is shown. When $D = 1$, multiplying the map of Fig. 4.16a to the amplitude spectrum returns the mean amplitude of the integrand of (4.38). As $|\Psi_R^x|$ decreases and spreads out for increasingly large D , an attenuation of the entire shear-stress PSD is expected, leading to smaller torque values. Conversely, for sufficiently small D , Ψ_R^x has an amplifying effect, with the most amplified modes being those of small streamwise wavenumbers, which correspond to the energy-containing region of the shear-stress PSD. Therefore, we expect increasingly intense value of the torque as the disc size increases. This corresponds qualitatively to what observed in Fig. 4.7, the precise dependence on D being determined by the specific form of the shear-stress spectrum.

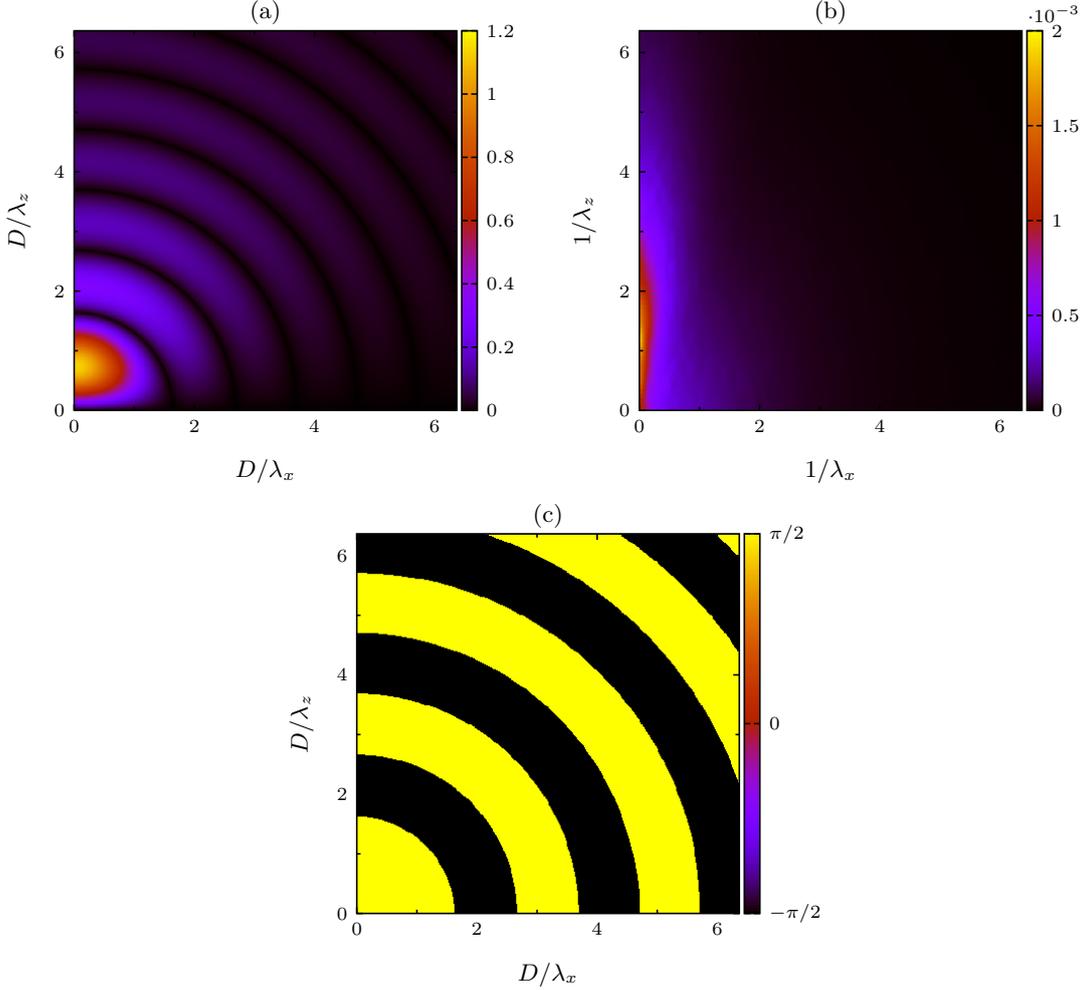


Figure 4.16: Filtering effect of the disc torque on the streamwise wall-shear stress. (a) Color map plot of $|\Psi_R^x|$ as a function of the streamwise and spanwise wavelengths, normalised with the disc diameter. This choice of coordinates emphasises how the relative scale of the shear-stress wavelengths to the disc size affects the filtering properties of Ψ_R^x . (b) Color map plot of the fixed-wall streamwise wall-stress PSD $|S_{\tau\tau}^{0,x}|^{1/2}$. (c) map of the phase Φ_Ψ of Ψ_R^x .

As illustrated in Sec. 4.4.2, the spanwise shear-stress torque $T_f^z(R)$ is much smaller than the streamwise contribution. A relation analogous to (4.38) can be derived for the spanwise-stress induced component $T_f^z(R)$ and reads:

$$T_f^{0,z}(R) = \iint_{\mathbb{R}^2} \hat{\tau}_z^0 \Psi_R^z d\kappa_x d\kappa_z. \quad (4.41)$$

The filter spanwise filter function Ψ_R^z is reported in Fig. 4.17a, together with the spanwise shear-stress mean amplitude in Fig. 4.17b. The difference between the spanwise filter function and the streamwise case is that $\Psi_R^z(\kappa_x, \kappa_z) = R^3 \Psi(R\kappa_z, R\kappa_x)$, i.e. the spanwise and the streamwise wavenumbers are exchanged. This can be shown by substituting z with x in (4.38) and calculating the spatial integral. As implied by its definition, Ψ_R^z privileges streamwise modes of $\hat{\tau}_z^0$ (as

opposed to spanwise for Ψ_R^z) and its maximum amplification happens for the purely streamwise mode such that $D/\lambda_x \approx 0.732$.

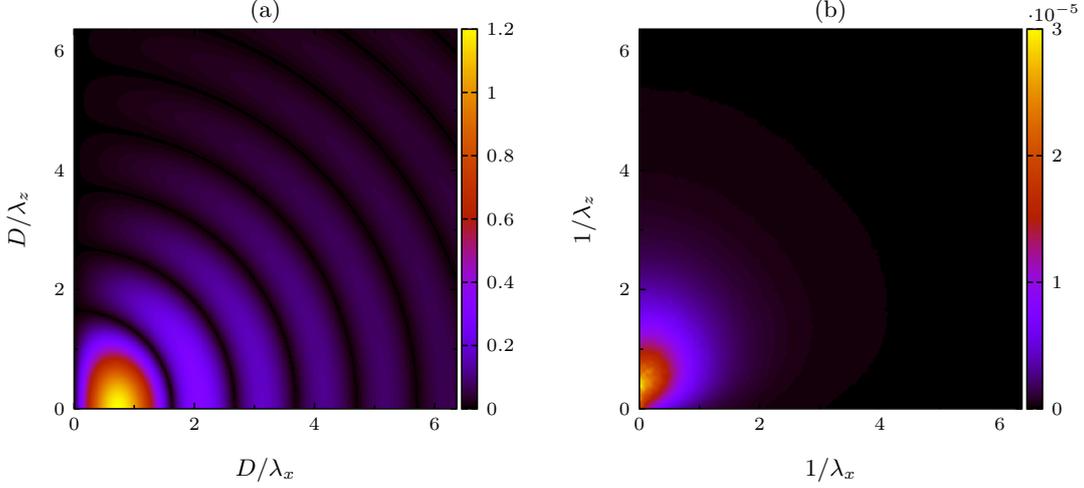


Figure 4.17: Filtering effect of the disc torque on the streamwise wall-shear stress. (a) Color map plot of $|\Psi_R^z|$ as a function of the streamwise and spanwise wavelengths, normalised with the disc diameter. (b) Color map plot of the fixed-wall spanwise wall-stress PSD $|S_{\tau\tau}^{0,z}|^{1/2}$.

Finally, the variance of T_f^0 is calculated by evaluating the variance of the right-hand side of (4.38) and (4.41) over the ensemble of uncorrelated shear-stress fields produced from the numerical simulation of the fixed-wall channel. After taking the square root, the resulting curve is shown in Fig. 4.15a alongside the r.m.s. $(T_f^0(D))$ curve measured directly from the spatial wall shear-stress fields as per (4.14), showing good quantitative agreement.

The slope change occurring at around $D^+=100$ originates from the scaling behaviour of the (4.39). A rough estimate of the slope turning point can be obtained by considering the diameter D_c at which the maximum of $|\Psi_R^x|$ and the maximum of $S_{\tau\tau}$ overlap. Knowing that the latter (Fig. 4.16b) is located around $1/\lambda_z = 1.2$, we have $D_c=0.732/1.2=0.61$, which in viscous units is $D_c^+ \approx 110$.

4.5 Freely rotating half discs

Passively rotating half discs have been studied experimentally by Koch and Kozulovic (2013) and Koch and Kozulovic (2014) as a method to reduce the skin-friction drag through a large-scale slip at the wall. The half discs consist of freely-rotating discs that have one half (in the spanwise direction) exposed to the mean shear, while the other half is covered by a solid surface and does not face directly the main body of the flow. The half disc therefore rotates with a non-zero-mean velocity, caused by the non-zero fluid torque ($\overline{T}_f \neq 0$) induced by the mean shear stress that acts on the exposed half of the disc. Since a steadily moving disc introduces a streamwise slip velocity at the wall, it is reasonable to expect a local flattening of the velocity profile and a reduction of the mean wall gradient.

In Section 4.5.1 we review and analyse the experimental results of Koch and Kozulovic (2013), in Section 4.5.2 we define the flow parameters of the half-disc channel flow. Section 4.5.4 describes the set-up and the results of the simulations, carrying out a comparison with the available experimental results.

4.5.1 Review of the experimental results

The original wind-tunnel experiments of Koch and Kozulovic, 2013 and Koch and Kozulovic, 2014 (KK13 and KK14 in the following) included half-discs of diameter $D^*=100$ mm in two different housing configurations. We focus on the more promising configuration, shown in Fig. 4.18a, that is the best performing of the two and is also more easily realised numerically. The experimental rig and the coordinate system used in the experiments is depicted in Fig. 4.18b. The centre of the disc is chosen as the coordinate origin.

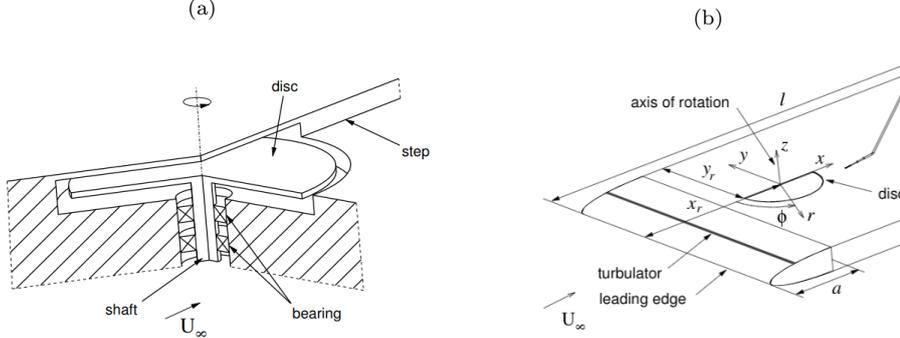


Figure 4.18: Experimental set-up of Koch and Kozulovic, 2013. (a) disc mount configuration. (b) schematic of the system.

KK13 and -14 estimate a drag-reduction rate of up to 17% on the exposed half of the disc. This number is not obtained from a direct measurement of forces or stresses but from an indirect estimate that proceeds as follows. The shear-stress distribution on the disc surface is not measured directly in the experiments of Koch and Kozulovic (2013) and Koch and Kozulovic (2014) but through an indirect estimate that proceeds as follows. First, the number of revolutions per second f^* of the disc is measured. Then, under the assumption that the boundary-layer thickness over the rotating half-disc is identical to the fixed-wall reference case, the wall-shear stress τ_w^* on the rotating half-disc surface is estimated according to the following empirical formulas:

$$\tau_w^*(x^*, z^*) = 0.0225 \rho^* [U_\infty^* - u_s^*(x^*, z^*)]^2 \left[\frac{\nu^*}{\delta_{99}^*(x^*, z^*) (U_\infty^* - u_s^*(x^*, z^*))} \right]^{1/4}, \quad (4.42)$$

$$\delta_{99}^*(x^*, z^*) = 0.37 (x^* - x_{t,v}^*) \left[\frac{\nu^*}{U_\infty^* (x^* - x_{t,v}^*)} \right]^{1/5} \quad (4.43)$$

where U_∞^* is the far-field velocity of the boundary layer and $u_s^*(x^*, z^*) = 2\pi f^* z^*$ is the streamwise slip velocity on the disc surface and δ_{99}^* is the 99% boundary-layer thickness. $x_{t,v}^*$ is the virtual origin of the turbulent boundary layer, which is extrapolated from the measurements of δ_{99}^* yielding $x_{t,v}^* = -118$ mm. Seven experiments (Table 4.3), each of increasing U_∞^* , are carried out in a closed-circuit wind tunnel. A kinematic viscosity $\nu^* = 1.5 \times 10^{-5} \text{ m}^2/\text{s}$ is assumed for air. The boundary-layer thickness is measured experimentally at three spanwise distances from the disc centre and shown to be in good agreement with the fixed-wall value, supporting the initial assumption. The authors thus conclude that the wall shear-stress value can be estimated via (4.42)-(4.43).

As the turbulent boundary-layer is a spatially developing flow, we choose the centre of the disc as the reference position for the purpose of calculating non-dimensional flow variables. For the boundary-layer flow, non-dimensional values are obtained by normalising lengths with δ_{99}^{0*} ,

velocities with U_∞^* and stresses with $\tau_w^*(0,0)$. Viscous units are defined in a similar way at the reference position.

The reference friction Reynolds number Re_τ , needed to match the experiment with the channel-flow simulations as required by the principle of near-wall similarity, is defined as $Re_\tau = u_{\tau,0}^* \delta_{99}^{0*} / \nu^*$, where $\delta_{99}^{0*} = \delta_{99}^*(0,0)$ and $u_{\tau,0}^* = \sqrt{\tau_w^*(0,0)} / \rho^*$.

For the sake of comparison with the numerical simulations and other passive drag-reduction methods, it is also convenient to model the freely-rotating half-disc as a slip-velocity surface, through the definition of the equivalent slip velocity $U_{s,D}$:

$$U_{s,D} \equiv U_D(y=0), \quad (4.44)$$

where U_D identifies the streamwise velocity, conditionally averaged on the exposed disc half (region D in Fig. 4.20) i.e. $U_D = \langle u_s \rangle_{\mathbf{D}}$. Knowing the mean disc edge velocity \overline{W} , the equivalent slip velocity is calculated exactly as $U_{s,D} = 4\overline{W}/3\pi$. The main performance metric used to discuss experiments and simulations is the half-disc drag reduction rate \mathcal{R}_D , defined as:

$$\mathcal{R}_D(\%) = 100 \frac{\tau_{w,0} - \tau_{w,D}}{\tau_{w,0}}, \quad (4.45)$$

where $\tau_{w,0}$ is the value obtained from formula (4.42) for $u_s = 0$ (fixed wall) and $\tau_{w,D} = \langle \tau_w(x,z) \rangle_{\mathbf{D}}$.

Table 4.3 contains the flow parameters estimated for the experimental turbulent boundary-layer according to the definitions given above.

Table 4.3: Flow parameters and results of the boundary-layer experiments from Koch and Kozulovic, 2013.

U_∞^* [m/s]	δ_{99}^{0*} [mm]	f^* [1/s]	Re_τ	D^+	\overline{W}^+	$U_{s,D}^+$	$\mathcal{R}_D(\%)$
20.32	3.98	9.94	276	6943	3.00	1.27	11.2
25.38	3.80	14.33	322	8481	3.54	1.50	12.6
30.41	3.67	19.01	366	9980	3.99	1.69	13.9
35.6	3.55	23.78	409	11500	4.33	1.84	14.7
40.7	3.46	28.78	449	12973	4.65	1.97	15.6
45.86	3.38	34.04	488	14444	4.94	2.09	16.3
50.96	3.31	39.04	525	15882	5.15	2.19	17.1

The experiments also find a persistent reduction of the momentum thickness of the boundary-layer downstream of the disc, according to measurements performed at three different spanwise locations. The reduction, of up to 5.3% in the best performing case, is present at least up to one disc radius downstream from the disc edge and therefore hints at the possible existence of a beneficial region where the skin-friction is persistently reduced by the disc rotation. The authors also observe that the streamwise extension of the beneficial region is proportional to the upstream extension of the disc, which is minimal at the disc edge and increases towards its centre.

4.5.2 Parameters of the channel flow with half discs

In this section we analyse the non-dimensional parameters of the half-disc channel flow, necessary to compare the numerical channel-flow simulations to the boundary-layer experiments. As explained in Section 4.3.5 any flow statistics can be expressed as a function of the reference (i.e. of the no-slip case) friction Reynolds number Re_τ , the diameter D , the disc inertial parameters

ρ_s, D, b and the parameters of the disc housing, namely d_h and T_b . The disc inertia can be neglected because mean disc-edge velocity \overline{W} (or alternatively, the angular velocity $\overline{\Omega}$) of a half disc is uniquely determined by the steady-state balance between the fluid torque and the frictional housing torques $\overline{T}_f(\overline{W}) = \overline{T}_h(\overline{W}) + \overline{T}_{min}(\overline{W})$. For half discs, it is convenient to reorganise the parameters introducing the equivalent slip velocity $U_{s,D}$ (defined in (4.44)) to obtain:

$$\mathcal{R} = f_{\mathcal{R}}(Re_{\tau}, D, U_{s,D}) \quad (4.46)$$

$$U_{s,D} = f_s(Re_{\tau}, d_h, D). \quad (4.47)$$

4.5.3 Numerical cases

The dependence of \mathcal{R} and $U_{s,D}$ on the parameters is investigated via a number of direct numerical simulations, each involving the variation of a parameter of the problem. All simulations (Tab. 4.4) are performed in a channel instead of a boundary-layer and at $Re_{\tau} = 180$, lower than the experiments of KK13. Although Reynolds similarity to the experiments is not enforced in the simulations, some considerations on the drag reduction can be made in the framework of a slip-velocity surface model. For slip-length surfaces, it is a widely accepted hypothesis that the drag-reduction performance scales in wall units (Min and Kim, 2004; Busse and Sandham, 2012) with very little dependence on the Reynolds number, at least for the range of Re_{τ} considered here. This fact endorses the idea of studying the half-disc drag reduction and its dependence on the equivalent slip velocity, scaled in wall units. The disc diameters are also smaller than the original experiment in terms of wall units. Larger values of D^+ are presently beyond our available computational resources: performing a simulation of a single half-disc at the experimental value of D^+ would require a grid size about 600 times that of HD2. In the channel-flow simulations, the flow in the housing cavity is not simulated explicitly, but rather modelled through one of the formulas (4.19), (4.20) that is pre-determined based on the assumed nature of the flow in the cavity (i.e. assuming a cavity Reynolds number Re_h). This means that the dependence of (4.47) on d_h is not natural but needs to be verified a posteriori by checking that the resulting Re_h is consistent with the assumption.

As explained in Section 4.3, the functional form of $\overline{T}_h(\overline{W})$ is determined by Re_h , therefore the choice of the cavity depth d_h must be considered carefully. Cases HD2 and HD3 assume $d_h = 0.1$, which is the same value that is found in the experiments when scaling with the boundary-layer thickness as d_h^*/δ_{99}^* , and make use of the linear T_h model (4.19). Cases HD4 and HD5 also use the linear T_h model but consider deeper cavities of $d_h = 1$ and 10, it is noted that the use of this model is not realistic for these cases, however it is attempted in order to obtain a lower friction torque and obtain a higher \overline{W} . Finally, HD6 instead uses the Von-Kármán solution (4.20) for T_h . Since \overline{W} and Re_h are mutually dependent variables, the consistency of these assumptions has to be verified a posteriori by measuring the actual Re_h produced in the simulations. In the experiments $d_h^* = 0.4\text{mm}$, giving a value of $Re_h = 3$ for the fastest rotating case and $Re_h = 1$ for the slowest, compatible with a Couette flow in the cavity.

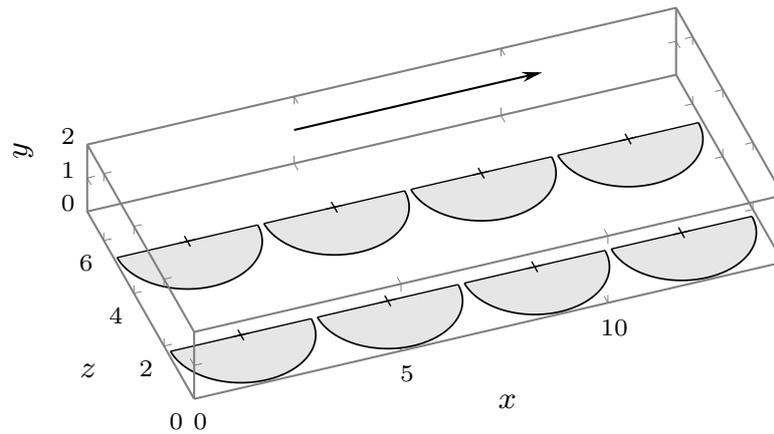


Figure 4.19: Schematic of the arrangement of the numerical channel of cases HD2,4,5,6.

Cases HD2 through HD6 (Fig. 4.19) simulate a set of simplified boundary conditions that does not explicitly simulate the sheltering structure that covers half of the disc. In these cases the half discs are modelled by simply eliminating the left-hand half of the disc and applying a fixed-wall condition. In cases HD2-6, the disc dynamical model is modified as follows: 1) T_f is only calculated on the existing disc half and 2) T_h is multiplied by a factor of 1.5 to account for the friction under the hypothetical plate that would cover the sheltered half of the disc. This operation implies that the flow under the plate is assumed to satisfy the same solution (i.e. Couette, Von Kármán, or turbulent) that is found in the housing cavity.

This simplified set-up is studied because it can be more easily adapted to study the dependence of the drag reduction on D (cases HD2 and HD3), the disc velocity (cases HD2 through HD6), and to decide which model of T_h is more realistic (cases HD5 and HD6).

Table 4.4: Numerical parameters of the passive half-disc flow simulations.

Case name	Re_p	L_x	L_z	grid size	D^+	N_D	T_h model
HD2,4,5	4200	4.53π	2.26π	$256 \times 129 \times 256$	605	8	Couette
HD3	4200	6.7π	3.33π	$256 \times 129 \times 256$	1210	2	Couette
HD6	4200	4.53π	2.26π	$256 \times 129 \times 256$	605	8	VK

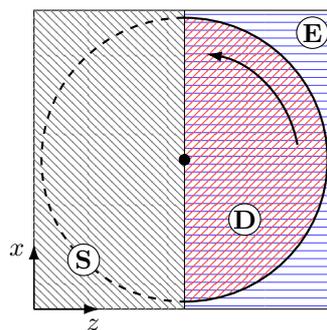


Figure 4.20: Nomenclature of the wall regions analysed in the text: D) exposed disc half, E) exposed wall region, S) sheltered wall region. Region D is a subset of Region E.

4.5.4 Simulation results

The main results of the simulations are shown in Tab. 4.5. In order to facilitate the comparison with the experiment, we compute the drag reduction \mathcal{R}_D on the exposed half of the disc defined in (4.45), as well as the global drag reduction \mathcal{R} .

The cases HD2 and HD3 return small drag reduction rates of around 2%, which becomes almost 5% when \mathcal{R}_D is considered. Cases HD5 and HD6 are the best performing configurations overall, they obtain the highest values of \overline{W}^+ and achieve drag reductions of 20% and 14%, respectively. The values of the mean slip velocity on the exposed half of the disc $U_{s,D}^+(0)$ are also shown.

The disc-tip velocity is identical for both HD2 and HD3, regardless of the differing disc diameters, that is the angular velocity decreases linearly with D . The equivalent slip velocity on the exposed disc half exhibits values similar to those needed to obtain drag reduction in the well-known slip-length modelling of hydrophobic surfaces (Min and Kim, 2004) and a direct correlation between the slip velocity and the drag reduction is also found.

The values of Re_h measured from the simulations remain well under 40 in HD2 and HD3, indicating consistency with the laminar Couette model of the cavity torque. However, the higher rotation velocity displayed by HD4 and HD5 means that the values of Re_h become of the order of 10^3 , which lie outside the Couette solution range and therefore are inconsistent with the model of T_h . Case HD6 instead achieves $Re_h = 45$ that is consistent with the Von Kármán torque model.

Table 4.5: Summary of the channel-flow simulation results.

Case name	Re_τ	D^+	\overline{W}^+	$U_{s,D}^+$	$\mathcal{R}(\%)$	$\mathcal{R}_D(\%)$
HD2	180	605	0.98	0.42	2.1	4.9
HD3	180	1210	0.98	0.42	2.1	4.9
HD4	180	605	1.42	0.60	2.8	8.6
HD5	180	605	4.85	2.06	5.6	19.9
HD6	180	605	2.49	1.05	5.1	13.9

a. Comparison and discussion

In this section we compare the values of \mathcal{R} and $U_{s,D}$ obtained in the experiment and the simulations, considering their dependence on the parameters specified in (4.46)-(4.47).

First, as noted in Section 4.5.4, the simulation of cases HD2 and HD3 shows that, all other parameters being equal, neither the disc velocity nor the drag reduction rate are influenced by the disc size. Although two data points do not provide conclusive evidence, we henceforth consider the dependence on D as negligible.

As the simulations are all carried out at a value of Re_τ different from the experiments, it is interesting to speculate on what would occur in the boundary-layer experiment if Re_τ were to match the lower value of the channel flow simulations. An estimate can be obtained as follows: first, we perform linear regression on the experimental data for the angular frequency as a function of the far-field velocity $f^*(U_\infty^*)$, which is shown in KK13 to be linear. Secondly, we extrapolate to smaller values of U_∞^* , and finally we use the extrapolated pairs U_∞^*, f^* in the empirical formulas (4.43)-(4.42). For example, by substituting $U_\infty^* = 1\text{m/s}$ and $f^* = 0.67\text{Hz}$ in the empirical formulas (4.42) and repeating the calculations illustrated in Sec. 4.5.1, we obtain $Re_\tau = 179$, $U_{s,D}^+ = 0.15$ and $\mathcal{R}_D = 1.4\%$. This estimation procedure can be carried out for an arbitrary number of velocities, four of which are shown in Fig. 4.21, identified by the (*)

symbols. At the same $Re_\tau = 179$ of the simulations, the extrapolated experimental disc rotates more slowly, producing lower values of $U_{s,D}^+$.

Fig. 4.21 shows the correlation between the drag reduction and the equivalent slip velocity in wall units. The simulations capture a wider range of slip velocities than the experiments. In both the experiment and the simulations, drag reduction is positively correlated to $U_{s,D}^+$. The experiments show lower values than the channel flow in the same range. The channel-flow curve seem to undergo a concavity change at low $U_{s,D}^+$, however the two data points are not sufficient to determine if a linear behaviour is also present in the channel-flow data. The latter are not captured by the empirical procedure used to calculate the drag reduction in the experiment, but can be very significant as shown in Fig. 4.22. At the same $Re_\tau = 179$ of the simulations, the extrapolated values have a smaller \mathcal{R}_D show a convex shape at low $U_{s,D}^+$.

The differences between simulations and experiments is due to either 1) Reynolds-number effects or 2) different physics in the boundary-layer related to the non-homogeneous flow phenomena.

The global drag reduction rate \mathcal{R} is depicted in Fig. 4.21a by the (Δ) symbols. There is very little increase in \mathcal{R} for the two largest disc tip velocities. The inspection of the spatial distribution of skin-friction on the wall is necessary to investigate the cause of this behaviour.

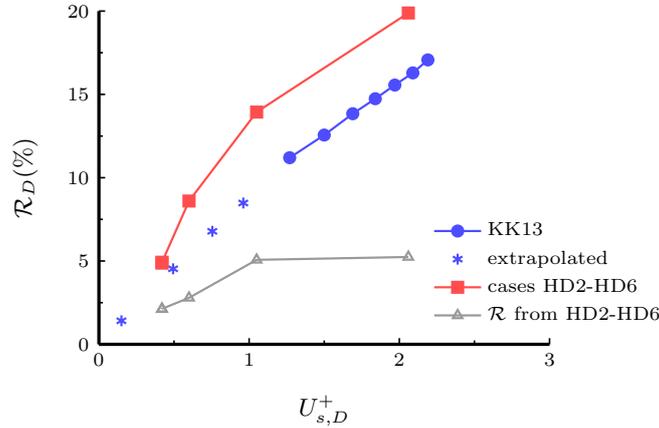


Figure 4.21: Dependence of the drag reduction rate on the exposed disc half (region D) on the slip velocity expressed in viscous units. Comparison between our simulations (HD2-HD6, Table 4.5) and the half-disc boundary-layer experiments (KK13, Table 4.3). The grey curve in is the global drag-reduction rate.

The maps of $\mathcal{R}_{xz}(x, z)$ on the wall are shown in Fig. 4.22b–c for the numerical cases HD5 and HD6 and reveal a starkly non-homogeneous spatial distribution, along both the x and z directions. In either case, the maximum drag reduction is encountered on the upstream rotating edge of the disc, while the maximum drag increase occurs in the region near the centreline, upstream of the disc centre. The latter appears to be an effect of the discontinuity that characterises the simplified half-disc boundary conditions. Although the simulations achieve a good level of drag reduction on region D, we do not find any region of persistent drag reduction downstream, as claimed by the experiments. On the contrary, an area of increased drag is observed immediately downstream of the disc, where the surface changes from a finite slip velocity to no-slip outside the disc.

Fig. 4.22a reports the spatial distribution of the drag reduction for the best-performing experiment of KK13. The drag-reduction distribution of the experiment is found in Fig. 4.22

and is calculated as:

$$\mathcal{R}_{xz}(x, z) = 100 \frac{\tau_w^{0*}(x, z) - \tau_w^*(x, z)}{\tau_w^*(x, z)} \quad (4.48)$$

where the shear-stress distributions are calculated from the empirical formula (4.42). No stream-wise non-homogeneous effect is captured by the experimental model, which simply reflects the proportionality to the slip velocity, increasing linearly along z . As no experimental measurements of the boundary-layer thickness are performed on the region of the disc near the wall or near the covering structure, no relevant experimental data exist for the spatial structure of the flow. It is also possible that the experimental finding on the constant boundary-layer thickness cannot be extended straightforwardly to the entire half-disc surface outside of the three spanwise position considered by the measurements.

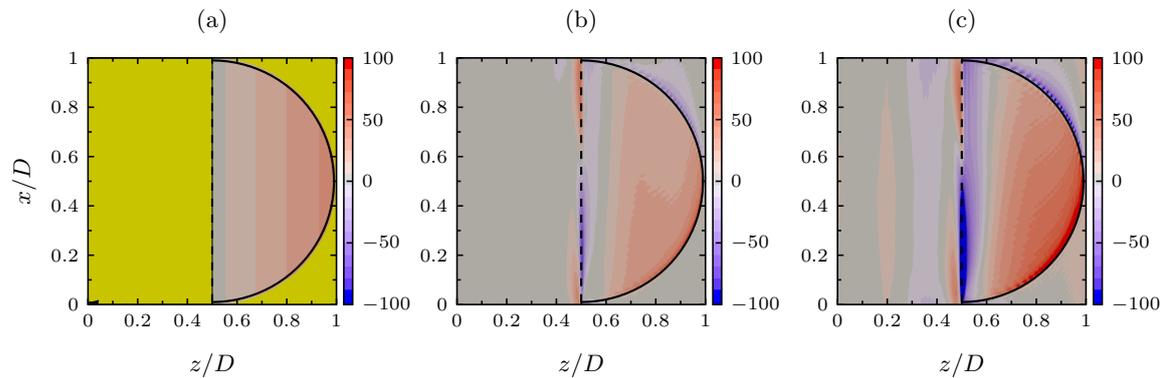


Figure 4.22: Spatial distribution of the drag reduction $\mathcal{R}_{xz}(x, z)$ on the wall, values in percent. A single, period-averaged flow unit is depicted. (a) experiment from KK13 (fastest rotating case): the yellow colour identifies the region where no data are available. (b) case HD6. (c) case HD5.

Fig. 4.23 provides a visualisation of the fine-scale vortical structures using the λ_2 method (Jeong and Hussain, 1995) for a single instant of time. Case HD5 generates a persistent reduction in the size and number of the vortex cores along the streamwise direction above the spinning region, in contrast with the no-slip part of the wall, which is also visible in HD5 although the effect is less pronounced.

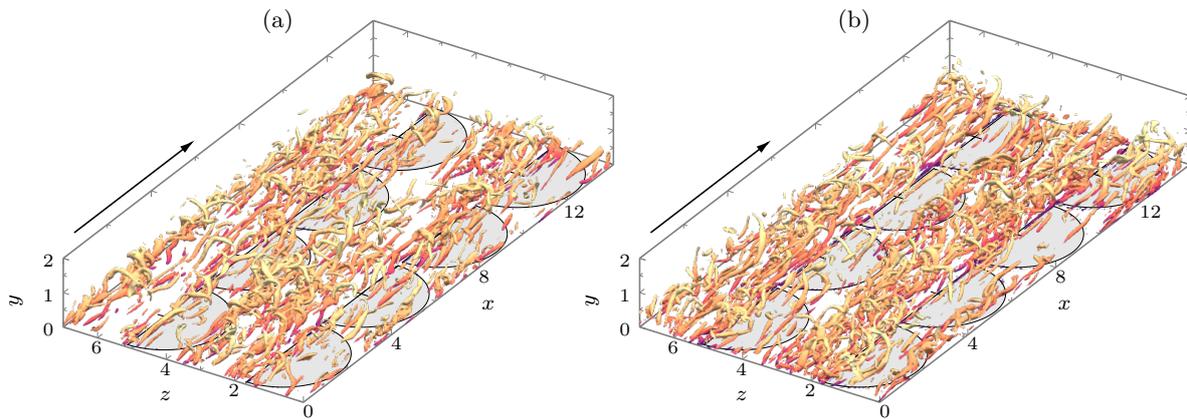


Figure 4.23: Instantaneous iso-surfaces at $\lambda_2^+ = -0.009$, coloured by the streamwise velocity component: (a) case HD6, (b) case HD5.

The wall-normal turbulent shear-stress profiles on region S (Fig. 4.24a-c) are virtually indistinguishable from the no-slip case, while those shown in Fig. 4.24b-d undergo significant weakening below $y^+ = 100$, with a 30% reduction at $y^+ = 30$. It can be noticed that, although the profile in Fig. 4.24b has reduced turbulent stress, drag increase is observed at the same position in Fig. 4.22. This suggests that the origin of the local drag increase is to be found in a laminar effect associated with the deceleration of near-wall high-momentum fluid on the no-slip region. The wall-normal profiles of the root-mean-square velocity also display a qualitatively analogous spatial dependence, being reduced on regions E and S and preserving the same intensity as the reference case on region S.

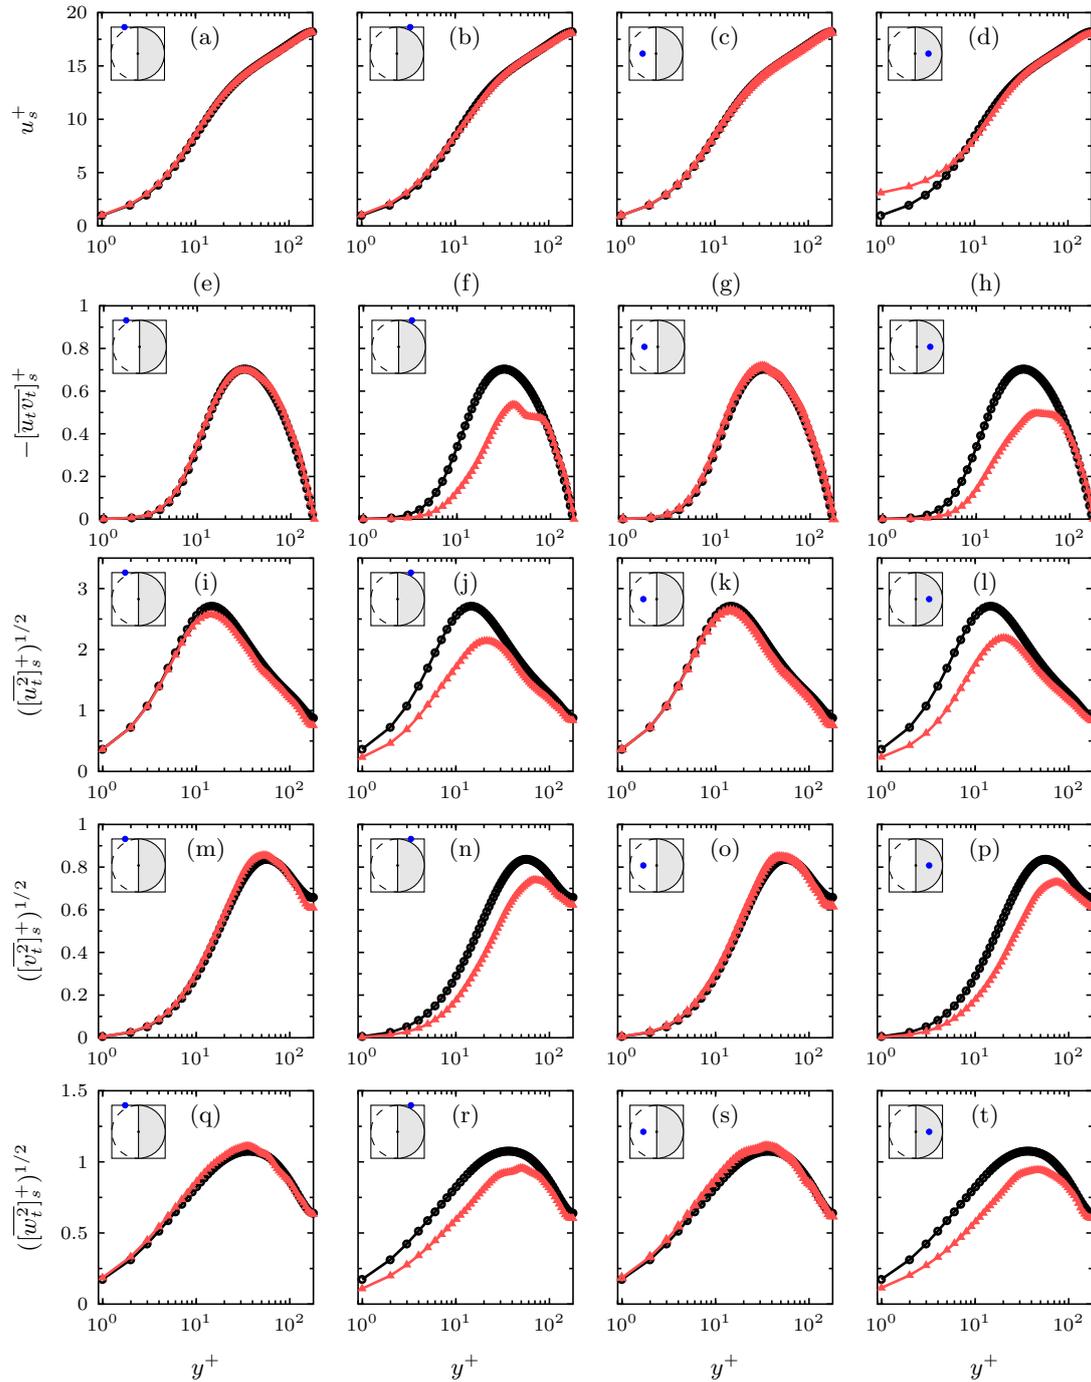


Figure 4.24: Wall-normal profiles of several flow statistics at four different positions on the wall: (a)-(d) the period-averaged streamwise velocity u_s^+ , (e)-(h) the turbulent Reynolds stress $[\overline{u_t v_t}]_s^+$, (i)-(l) the r.m.s. of the streamwise velocity. (m)-(p) the r.m.s. of the wall-normal velocity. (q)-(t) the r.m.s. of the spanwise velocity. The positions are identified by blue dots in the sketches. The statistics of Case HD5 (\blacktriangle) are compared to the reference channel-flow (\bullet). The variables are scaled with the wall units of the reference channel.

4.5.5 Experiments and simulations

In this section we analyse further the limitations of the half-disc numerical simulations in relation to experiments, focusing on the differences between TBL and channel flow, Reynolds numbers, and the uncertainties related to the experimental realisation of the actuators.

a. Boundary Layer or Channel flow

The validity of the results obtained in channel flow would rely on *i*) a perfect Re_τ similarity being preserved and *ii*) the same native-scaled disc velocity W^+ being maintained. The inhomogeneous distribution of the wall-shear stress in a the boundary layer flow implies that these conditions cannot be met exactly, irrespectively of the size or the position of the half-disc devices. In our analysis of the experimental data, we consider the value of the Reynolds number taken at the disc centreline as a reference. At the moderate values of Re_τ considered in our study this approach is sufficient to ensure close, given the difference between upstream and downstream Re_τ is of the order of a few units. At higher Reynolds numbers, as a consequence of the viscous length contraction, it is possible that more extended portions of the walls are subjected to sub-optimal slip velocities. A closer comparison to the experimental boundary-layer flow might be possible in the simulation programme advanced in Section 4.7 that tries to eliminate the free variable at the same Reynolds number of the channel flow. Furthermore, the interaction of a boundary-layer flow with the physical structure of the disc device might differ from that occurring in channel flow, thus introducing an additional source of uncertainty.

b. Effects of the Reynolds number on the disc rotation rate

In equations (4.46)-(4.47), the reference-case Reynolds number appears in the expression for both \mathcal{R} and U_s . In this and the following section, we give an overview of the effect of its influence on each one of the two variables and discuss how the combined effect affects the drag reducing capability at high Reynolds numbers.

Assuming the disc size and the flow regime in the housing cavity remain the same, it is expected that the disc angular velocity increases approximately as $\sim Re_p^{3/4}$ because, while T_f declines as $\sim Re_p^{-1/4}$ (being proportional to C_f , see (1.25)), the housing frictional torque (4.6) does so at the faster rate of $\sim Re_p^{-1}$. Since the wall-unit angular velocity u_τ/δ_ν also scales as $\sim Re_p^{3/4}$, the Reynolds scaling of W ensures that the native-scaled disc velocity – and thus U_s^+ – remains constant at any Re_p .

In this respect, an additional factor of uncertainty is related to the flow regime change within the cavity. Higher Re_p implies higher W , meaning that the cavity Reynolds number $Re_h = Wd_h Re_p$ soars at a faster rate than either one of its factors. While initially the laminar Couette-like state described by model (1.2) is preserved, when the cavity Reynolds number Re_h , undergoes a transition first to a Von-Karman-like laminar state and then to turbulence. This ultimately degrades the capability of the passive half-disc to produce an adequate value of U_s^+ and thus to deliver drag reduction.

c. Effects of the Reynolds number on \mathcal{R}

The dependence of the drag reduction \mathcal{R} on the Reynolds number is connected to the underlying mechanisms by which the actuators modify near-wall turbulence. Two known drag-reducing effects are understood to arise from the rotation of the half-discs: a streamwise, positive slip velocity at the wall and a spanwise periodic velocity pattern akin to spanwise velocity waves.

In either case the drag reduction performance can be parametrised, in the asymptotic high-Reynolds number limit, through the upward shift ΔB of the logarithmic velocity profile scaled in native wall units (Gatti and Quadrio, 2016; Rastegari and Akhavan, 2019). ΔB is taken as independent of the Reynolds number. In a constant-flow-rate turbulent channel, this approach gives the following relation for \mathcal{R} :

$$1 - \mathcal{R} = \left(1 + \left(\frac{1}{2\kappa} \log(1 - \mathcal{R}) + \Delta B \right) \sqrt{\frac{C_{f,0}(Re_p)}{2}} \right)^{-2} \quad (4.49)$$

where κ is the Von-Karman constant. This parametrisation allows the prediction of \mathcal{R} at arbitrarily high Reynolds number. Due to the presence of $C_{f,0}$ in relation (4.49), \mathcal{R} declines with Re_p for a given ΔB .

This theory has been applied in the literature to the prediction of high-Reynolds drag reduction for both spanwise velocity waves (Gatti and Quadrio, 2016) and slip-generating hydrophobic surfaces (Rastegari and Akhavan, 2019).

Among the half-disc cases discussed above, the best-performing one (case HD5) generates $\Delta B = 0.62$. Under the assumption that the ΔB measured at $Re_{\tau,0} = 180$ is already independent of the Reynolds number, Figure 4.25, following (4.49), shows that drag reduction at $Re_{\tau} = 10^6$ performance amounts to less than 40% of the low-Reynolds case. This however is a good performance given that such a decline takes places over four decades.

The assumption made in (4.49) that ΔB is independent of the Reynolds number only holds , as explained in the previous subsection. Since that is not likely to be the case at high enough Re_p , a further degradation of \mathcal{R} is foreseen.

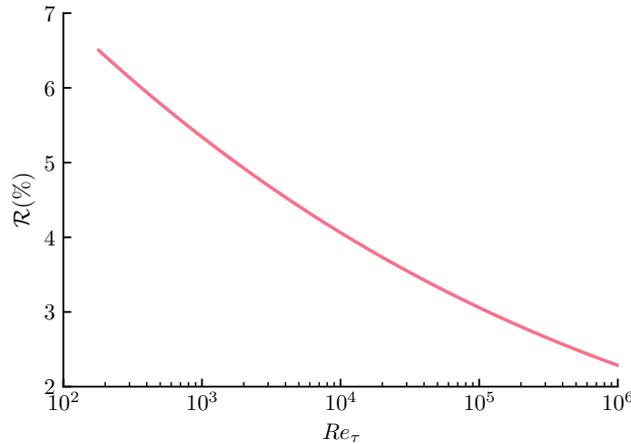


Figure 4.25: Predicted drag reduction for the half-disc configuration HD5 at high Reynolds number. The red curve is computed from (4.49) on the Reynolds number range 180 to 10^6 .

since many questions raised by the Reynolds number strongly encourages further research, and proposed in .

d. Device model uncertainty

A third aspect is related to the uncertainty inherent to the simplified device modelling necessary for the simulations and its comparison to the experimental set up, although the aim of our work

is not to replicate the experiments faithfully, but to provide a proof-of-concept of the mechanism and gain a broad understanding of the flow physics involved. The single biggest source of uncertainty in terms of actuator modelling are represented by the simplified housing with no disc shelter. More concerns are on the actuator's geometry and its constructive details, such as the flush, friction in bearing, the presence of nearby actuators, or additional sources of friction that may have been neglected. In general, since we have no access to the data of the exact components used in the original work, nor to the experimental uncertainties involved, a number of assumptions had to be made, except the housing geometry which was modelled as in the experiments. We are nonetheless confident of the quality of the models; a more nuanced comparison with the experiments is postponed after performing the simulations outlined in Section 4.7 that are intended to reproduce more closely the experimental conditions. More information on how the model compares with our experiment can be also found in Section 4.5.3.

4.6 Summary and conclusions

The key conclusions that can be drawn from the numerical study of passively rotating discs presented above can be summarised in the following points:

- (a) A dynamical model of the wall-mounted rigid disc has been developed based on a number of assumptions about its constructive and fluid-dynamical characteristics, grounded in existing experimental work on analogous actuators and an hypothetical experimental channel flow. The use of such model is necessary to simulate the real-time, two-way dynamic coupling between the turbulent flow and the disc and, consequently, the effect of the disc motion on the near-wall turbulence. The resisting torques arising from the disc bearing and from the motion of the fluid in the cavity underneath the disc are studied. In particular, the cavity torque is subject to case-specific evaluation as it is dependent on the flow regime that occurs in the cavity at different disc velocities. The disc model has been implemented numerically and integrated into the DNS solver. The disc-fluid coupling is realised by calculating, at each time-step, the torque T_f induced by the wall-shear stress on the disc surface, solving the disc dynamics and updating the boundary conditions of the fluid using the newly obtained angular velocity. The consistency of assumptions made about the fluid flow in the disc housing is verified *a posteriori* from the simulations.
- (b) The dynamics of rotating discs that oscillate passively under a turbulent channel flow has been studied via numerical simulations of discs of increasing diameter. The main objects of the analysis are the standard deviation of the fluid torque T_f and the standard deviation of the disc velocity W . The first is found to increase with the disc diameter as $D^{2.2}$ and the second is found to decrease as $D^{-0.5}$. The time-series and the spectra of both W and T_f have also been analysed, and found to undergo a shape change and a redistribution of the energy towards lower frequencies as the diameter increases. The amplitude of the disc velocity is small and no drag reduction can be measured. Motivated by the small intensity of the velocity amplitudes, the disc dynamics in the limit of asymptotically small oscillations (i.e. fixed with respect to the fluid) has also been studied and compared to the coupled case. This allowed the inexpensive simulation of a larger number of disc sizes using the fixed-wall channel flow data. We find that the disc velocity and the fluid torque depend on the disc diameter, following similar trends to those of the coupled case and around 30% smaller r.m.s. of both the torque and velocity oscillations. The qualitative behaviour of the velocity and torque spectra is also found to match the coupled case. Since the turbulent flow is uncoupled from the disc rotation, an exact spectral filter function can be derived that encodes how the fluid shear-stress fluctuations on the surface of the disc are converted

into the fluid torque fluctuations, thereby explaining the dependence of the r.m.s. T_f on D . As the uncoupled dynamics of the disc is linear, it is also possible to derive an exact response function for W that, given the fluid torque spectrum, explains *i*) the dependence of W on the diameter both in the uncoupled case and for most coupled cases and *ii*) the red shift observed in the disc velocity spectra.

- (c) We have reviewed and analysed the experimental results of the passively rotating half-discs of Koch and Kozulovic, 2013 in a turbulent boundary layer, conducted at seven increasing far-field velocities. The experiments claim drag reduction rates of up to 17% obtained through an estimation procedure based on the measurement of the BL thickness on the surface of the disc. The drag reduction rate grows with the disc velocity, which is proportional to the far-field velocity. The existence of a beneficial region on the fixed portion of the wall downstream of the disc is also reported.

In order to investigate the half-disc paradigm, five numerical simulations are performed in a channel flow fitted with half discs of two different sizes (by a factor of 2) and two different models of the housing torque (Von-Kármán and Couette solutions). The simulations are conducted at a lower value of the reference Re_τ^0 of the experiments. Rotational velocities similar (in wall units) to the experimental ones are obtained for the least intense values of the resisting torque produced in the disc housing. The drag-reduction rates on the exposed disc half \mathcal{R}_D and on the entire wall \mathcal{R} are evaluated. In the fastest rotating case, we find values of up to 20% and above 5%, respectively. Both drag reduction rates are observed to be positively correlated to the slip velocity that is introduced at the wall on the flow-facing half of the disc. The streamwise mean wall-shear stress has a non-trivial inhomogeneous distribution on the disc and a drag-increasing region downstream of the disc.

The simulation results make the case for passively rotating half-discs as a drag-reduction technique and offer a promising outlook on its potential improvement. Further investigation is necessary to clarify some crucial aspects, namely *i*) the added drag originating from the secondary flow phenomena on the disc superstructure; *ii*) the dependence of the disc velocity and the drag reduction on the friction Reynolds number, important for a closer comparison with the experiments.

4.7 Future research on half-discs

We propose a numerical study of a turbulent channel flow equipped with passively rotating rigid half-discs. The proposal is aimed at investigating the potential of the half-discs as a passive drag-reduction technique and expand the understanding of their flow physics, based on the results of Chapter 3 and the experiments of Koch and Kozulovic (2013) and Koch and Kozulovic (2014) (in short KK13 and KK14).

Tab. 4.6 lists the estimated numerical parameters for a series of DNSs that match the same values of the Re_τ^0 of the KK13 experiments as calculated in Tab. 4.3. The smallest domain size of $2\pi \times \pi$ is shown by Lozano-Durán and Jiménez (2014) to correctly reproduce the one-point statistics on numerical grids with a similar resolution. Nonetheless, a convergence study of the drag reduction with may be needed, as it is a known fact that using small domain sizes leads to overestimating drag reduction by in-plane wall motions (Gatti and Quadrio, 2016). The layout of the boundary conditions is identical to the one illustrated by Fig. 4.19 that uses a simplified model of the discs of KK13 and KK14. The modelling of the disc housing, the frictional torque T_h and its numerical implementation is explained in Sec. 4.3. Alternatively, we suggest the experimental correlation cited in KK14 to be studied and compared to the aforementioned models.

We propose the data analysis to follow what presented in Sec. 4.5, including *i*) the calculation of the equivalent slip-length U_D defined by (4.44) and the correlations $U_D(Re_\tau)$ and $\mathcal{R}(U_D^+)$, *ii*) the visualisation of the turbulent and disc-flow components and local statistics profiles, *iii*) the study of the spatially distributed skin-friction budget and its dependence on U_D and Re_τ .

Table 4.6: Numerical parameters of the planned channel-flow simulations with half discs. L_x is the streamwise domain length, L_z the spanwise width, D the disc diameter, N_D the number of discs and β is the wall-normal grid stretching parameter of Incompact3D. The values marked by * are estimates. Re_τ^0 is the nominal friction Reynolds number of the no-slip case at the same Re_p .

Case name	Re_p	* Re_τ^0	L_x	L_z	$n_x \times n_y \times n_z$	β	D (* D^+)	N_D
HD180	4200	180	4.5π	2.3π	$256 \times 129 \times 256$	0.25	3.38 (600)	4×2
HD280	6800	277	3π	1.5π	$256 \times 161 \times 256$	0.2	2.17 (600)	4×2
HD320	8100	320	2.5π	1.3π	$256 \times 161 \times 256$	0.18	1.88 (600)	4×2
HD365	9300	365	2.2π	1.1π	$256 \times 193 \times 256$	0.2	1.64 (600)	4×2
HD400	10500	405	2π	π	$256 \times 225 \times 256$	0.19	1.51 (600)	4×2
HD450	11900	450	2.2π	1.3π	$320 \times 257 \times 384$	0.2	1.33 (600)	5×3
HD525	14100	525	2.2π	1.1π	$320 \times 321 \times 384$	0.21	1.14 (600)	6×3

Appendix to Chapter 4

4.A Details of the reduced-order model from Sec. 4.4.4

Here we derive the expression of the spectral filter function Ψ_R^x . We start from the definition of Ψ_R^x and solve the double integral as follows:

$$\Psi_R^x(\kappa_x, \kappa_z) = \frac{1}{4\pi^2} \iint_{\mathcal{S}} z e^{i\kappa_x x} e^{i\kappa_z z} d\mathcal{S} \quad (4.50)$$

$$= \frac{1}{4\pi^2} \int_{-R}^R \int_{-\sqrt{R^2-z^2}}^{\sqrt{R^2-z^2}} z e^{i\kappa_x x} e^{i\kappa_z z} dx dz \quad (4.51)$$

$$= \frac{1}{4\pi^2} \int_{-R}^R z e^{i\kappa_z z} \left(\int_{-\sqrt{R^2-z^2}}^{\sqrt{R^2-z^2}} e^{i\kappa_x x} dx \right) dz \quad (4.52)$$

$$= \frac{1}{4\pi^2} \int_{-R}^R z e^{i\kappa_z z} \left(\frac{-i}{\kappa_x} \left[e^{i\kappa_x \sqrt{R^2-z^2}} - e^{-i\kappa_x \sqrt{R^2-z^2}} \right] \right) dz \quad (4.53)$$

$$= \frac{1}{4\pi^2} \int_{-R}^R z e^{i\kappa_z z} \left(\frac{-i}{\kappa_x} 2i \sin(\kappa_x \sqrt{R^2-z^2}) \right) dz \quad (4.54)$$

$$= \frac{1}{2\pi^2 \kappa_x} \int_{-R}^R z e^{i\kappa_z z} \sin(\kappa_x \sqrt{R^2-z^2}) dz. \quad (4.55)$$

We are unable to find an exact solution to the last integral, which is therefore evaluated numerically. The self-similar form (4.39) is obtained by considering the change of variable $z = z'R$.

Chapter 5

Feedback control of rotating discs

5.1 Introduction

In this chapter, we propose a simple on-off feedback-control scheme for rotating discs, based on the instantaneous angular velocity of the disc. The main purpose of introducing feedback control is to investigate whether it is possible to further improve the power efficiency of the base paradigm. Feedback control schemes have been applied to drag reduction in DNS studies [Bewley et al. \(2001\)](#); [Choi et al. \(1994\)](#); [Kim and Choi \(2017\)](#) in the context of the micro-actuators but not for in-plane wall motions. One disadvantage of these and other more advanced feedback-control schemes (e.g. optimal control) is that they require the measurement and processing of large amounts of flow data in real time.

The power efficiency of the feedback-control scheme is maximised through the computational optimisation of the control parameters. The optimisation is carried out using a powerful Bayesian optimisation framework that is able to select and propose new parameter combinations based on a surrogate model (i.e. a regression surface) of the objective function ([Schonlau, 1997](#); [Shahriari et al., 2016](#)). This technique is instrumental in reducing the number of computationally intensive direct numerical simulations needed for the convergence to the optimum. Variants of the Bayesian optimisation method have been long known in the engineering optimisation field, however only recently this technique has found application to enable DNS-based automatic optimisation studies. [Mahfoze et al. \(2018\)](#) simulated a suction/blowing device in a boundary layer, complete with an approximate model of the device's frictional losses, and used a variant of the Bayesian optimisation technique to study its drag-reduction performance.

This chapter layout is as follows: Section [5.2](#) gives an overview of the base flow case, the actuator design and the model used for their simulation. In Section [5.3](#) we introduce the feedback control scheme, develop the metrics used to evaluate its power efficiency and formalise the optimisation of the control parameters. The computational optimisation methodology is explained and motivated in Section [5.4](#). More technical information on Bayesian optimisation can be found in Appendix [5.A](#), along with the validation of the numerical code used. Section [5.5](#) describes the optimisation results and contains an extensive analysis of the simulation data, the main points of which are summarised in Section [5.6](#).

5.2 Modelling and simulation of the actuators

The flow case considered for feedback-controlled discs is analogous to the ring set-up of Chapter [3](#). Rotating discs are flush mounted on the surfaces of both walls of the channel in a orthogonal

array configuration, shown in Figure 5.1. Along the streamwise direction, the disc sense of rotation changes from one row to the next, as in Ricco and Hahn, 2013, but the angular velocity magnitude is dictated by the feedback-control mechanism introduced in Section 5.3. The disc geometry and the disc housing design are identical to those for the freely rotating discs given in 4.3.1.

The coordinate system and the definitions of the non-dimensional units are not changed from those defined in Chapter 3, Section 4.2 and 4.2. Non-dimensional quantities obtained by scaling in outer units, i.e. by normalising with the channel half-height h^* , the density of the fluid ρ^* , and the Poiseuille centreline velocity U_p^* at the same mass flow rate, are not indicated by any symbol. All simulations are conducted at a Reynolds number $Re_p = U_p^* h^* / \nu^* = 4200$, where ν^* is the kinematic viscosity of the fluid. The non-dimensional wall-shear stress is $\tau_w = (dU(0)/dy)|_{y=0} / Re_p$ and the bulk velocity is obtained by dividing the flow rate per unit wall area, Q , by the channel height as $U_b = Q/2h$. The stationary-wall friction Reynolds number of the reference, no-slip case is $Re_\tau^0 = u_\tau^* h^* / \nu^* = 180$, where $u_\tau^* = \sqrt{\tau_w^* / \rho^*}$ is the wall-friction velocity. Quantities scaled in viscous units are indicated with a + superscript.

The simulation parameters for the feedback-control cases are reported in Table 5.1, which is also the numerical configuration of the reference no-slip channel flow. In the following, quantities scaled in wall units are with respect to the reference stationary-wall case, unless explicitly indicated.

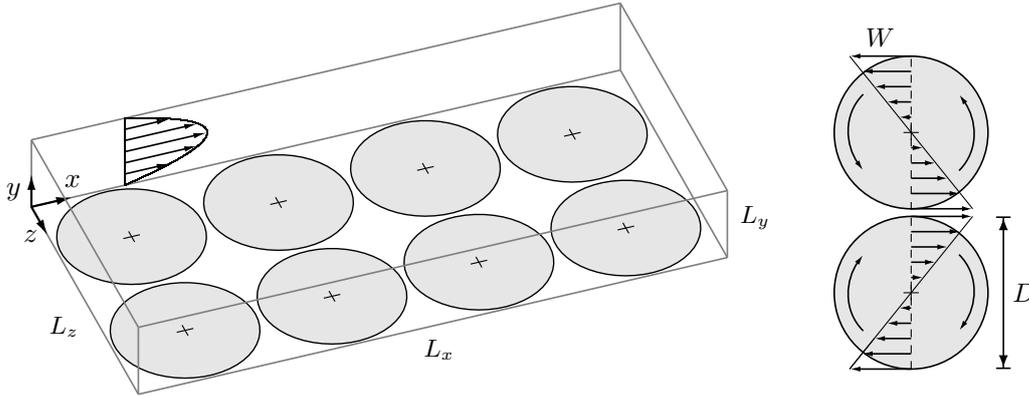


Figure 5.1: Schematic of the arrangement of the disc-fitted channel. Only the discs on the bottom wall are shown. The sense of rotation of the discs is highlighted by the curved arrows, while the longer arrow indicates the mean-flow direction. The sketch on the right illustrates the triangular-wave spanwise velocity pattern created along the centreline of the discs by two adjacent counter-rotating discs.

Table 5.1: Simulation parameters for the feedback-control flow cases. The number of grid points are along x , y , and z , respectively.

Case	Re_p	Re_τ^0	L_x	L_y	L_z	grid size	D^+	N_D
On-off control	4200	180	6.7π	2	3.33π	$378 \times 129 \times 320$	605	4×2

5.2.1 Averaging operators

a. Averaging operators

The time-average of a variable f is indicated by \bar{f} . Disc variables, such as the velocity W , are also always averaged across all the discs i.e. $\bar{f} = \sum_i^{N_D} |\bar{f}_i|/N_D$ where \bar{f}_i is the time-averaged variable from the i -th disc, N_D is the number of discs and the absolute value is used to correct for opposite senses of rotation.

It is also useful to define, for the flow field variables \mathbf{u} and p , the multi-disc or symmetry-averaging operator that accounts for the spatially periodic symmetry that exists across two adjoining counter-rotating rows of discs. The operator is defined identically to that detailed in Section 3.2.2.

Using the definition of the symmetry-averaged flow field $\mathbf{u}_s = [\mathbf{u}]_s$ the following flow decomposition can be introduced:

$$\mathbf{u} = \mathbf{u}_s + \mathbf{u}_t. \quad (5.1)$$

The symmetry-averaged component can be further broken down into the spatially-averaged flow \mathbf{U} (the ‘‘velocity profile’’) and the spatially inhomogeneous component \mathbf{u}_d as:

$$\mathbf{u}_s = \mathbf{U} + \mathbf{u}_d. \quad (5.2)$$

The spatial average of a flow field $\langle \mathbf{u} \rangle$ is obtained by integrating on x and z and dividing by the wall area $L_x L_z$, thus obtaining wall-normal profile $\mathbf{U}(y)$. The sampling uncertainty on time averaged or symmetry averaged variables is estimated via normal confidence intervals.

b. Performance metrics

The skin-friction coefficient is defined by normalising the mean wall-shear stress $\tau_w = Re_p$ with the bulk kinetic energy as $C_f = \langle \bar{\tau}_w \rangle / (U_b^2/2)$. The mean pumping power per unit wall-area that is needed to sustain a constant flow rate Q through a channel is $\mathcal{P}_p = 2\langle \bar{\tau}_w \rangle Q$. Using the definitions of the skin-friction coefficient and the bulk velocity, \mathcal{P}_p can be rewritten as $C_f U_b^3$. The percentage skin-friction drag reduction $\mathcal{R}(\%)$ is defined as:

$$\mathcal{R}(\%) = 100 \frac{C_{f,0} - C_f}{C_{f,0}}. \quad (5.3)$$

In a constant-flow-rate flow, the drag-reduction rate and the pumping-power reduction rate:

$$\Delta \mathcal{P}_p(\%) = 100 \frac{\mathcal{P}_{p,0} - \mathcal{P}_p}{\mathcal{P}_{p,0}} \quad (5.4)$$

are equivalent.

5.2.2 Model and dynamics of the rotating discs

The same disc actuators described in Chapter 4 are used with no modification, except the introduction of an electric motor as sketched in Fig. 5.2. The angular velocity of each disc at time t is $\Omega(t)$ and the disc-edge velocity $\Omega(t) = W(t)R$, where R is the disc radius. For each disc, the disc tip velocity evolves according to the dynamical equation:

$$\frac{2I}{D} \frac{dW}{dt} = T_f + T_h + T_{min} + T_M, \quad (5.5)$$

where the definitions and models adopted for the turbulent shear-stress torque T_f , the housing torque T_h and the minor frictional torques $T_{min} = T_b + T_{gap}$ are given in 4.3.1 and are used here with no modification. The motor torque T_M its the active forcing term and its modelling is introduced in the next section.

Some remarks need to be made on the validity, in the on-off controlled disc case, of the assumptions that are made in Section 4.3.3 to derive the models of T_h . In all the on-off flow cases considered below the value of cavity Reynolds number Re_h remains always < 1 even considering the worst case i.e. when the discs rotate at the upper velocity threshold. Therefore the flow is always in the laminar Couette regime and T_h is modelled by formula (4.19). The assumption of quasi-steadiness is not verified for the case where a disc is accelerated from or decelerated to the freely-rotating state by an external torque. In that case, it is necessary to quantify the rate of change of the angular velocity and show that it is such that it produces a small value of an acceleration parameter. The latter can be defined similarly to the frequency parameter, where the frequency f is replaced by the reciprocal of the transient duration and the velocity amplitude is replaced by the final value of the transient.

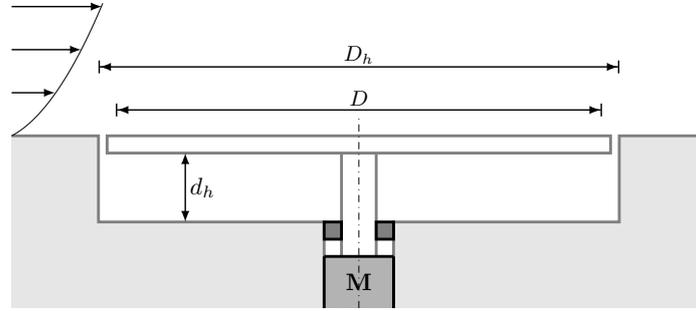


Figure 5.2: Schematic of the disc housing. The axis of rotation of the disc is indicated by a dash-dotted line.

5.2.3 Motor torque and active forcing

In general, the torque T_M exerted by an electric motor on the disc shaft depends on the angular velocity Ω and its precise functional form is a specific characteristic of the type and size of motor adopted. To model the motor torque T_M , we briefly review the procedure leading to the choice of the motor. We consider the use of a brush-less DC motor, which has small size, excellent controllability, low friction and low inertia. The motor inertia I_M , and its frictional torque T_{bM} , are neglected. The torque delivered by brushless motors is a linear function of the angular velocity $\Omega(t)$:

$$T_M(\Omega) = T_0 - \alpha\Omega. \quad (5.6)$$

A brushless motor can be easily controlled by proportionally increasing or decreasing the voltage at the motor terminals. The slope constant α and the feasible range of stall torques T_0 are determined by the motor type and power supply. Therefore, the motor model is chosen based on its capability to accommodate a range of expected loads, according to the motor-load matching explained in the following paragraph.

The motor constants are selected based on the characteristic curve $T_M(\Omega)$ matching the curve of the resisting loads at the steady-state velocity Ω_d . Since the resisting load of the spinning disc is the sum of the three contributions from the fluid, the housing, and the bearing, the matching equation reads:

$$T_M(\Omega_d) = T_f(\Omega_d) + T_h(\Omega_d) + T_{min}(\Omega_d). \quad (5.7)$$

The terms T_{min} and T_h are simple linear functions of Ω_d . T_f depends on the fully turbulent flow over the rotating disc and has, in principle, to be computed from a simulation with steadily rotating discs for each velocity Ω_d . Instead, we use a continuous approximation of $T_f(\Omega_d)$, obtained by quadratic interpolation of seven constant-velocity simulations performed at uniformly spaced values of Ω , thus avoiding a full-scale simulation each time a new Ω_d is needed. The approximation is justified because the power dissipated by the disc in the fluid can be reasonably well predicted (at low Ω_d) by the laminar solution for an infinitely extended disc (Ricco and Hahn, 2013). Using this method, we can determine the highest and lowest loads required by our range of Ω_d and choose a suitable motor model from manufacturer’s catalogue that is able to support the associated range of loads. This fixes α , while T_0 is the flexible control parameter that can be tuned to match precisely each load without exceeding the range specified by the manufacturer. Since we use the non-dimensional matching equation (5.7), the results are first dimensionalised using the parameters of the water channel described in Section 4.3.5, and then compared to the dimensional values given in real-world catalogues. In this way, we determine a realistic value of the slope constant to be $\alpha = -0.03$, which is used throughout this work, taken from the Faulhaber Series 1218 brushless DC servomotors (<https://www.faulhaber.com/en/products/series/1218b/>).

As a final note to this section, we note that this steady-state design methodology neglects the transient dynamics that brings the system from the initial velocity to Ω_d . During the transient, the resisting loads can exceed the steady-state curve due to inertial effects. Additionally, the motor itself can suffer performance losses due to transient electrodynamic phenomena. In a more refined model, the start-up dynamics and the transient loads would be considered during the motor selection process.

5.2.4 Dimensional reference

The dimensional water channel reference is identical to the one shown in Table 4.1.

An important effect of the choice of $\rho_s = 2.7$ for the on-off control is related to the time scales of the transients. In this case, a larger angular deceleration allows switch-off transient times one thousandth as long as those in an air channel. An unpowered disc decelerates from an initial angular velocity Ω_0 to rest in a time proportional to $\Omega_0 I / T_f$. In a water channel, ρ_s is around 1000 times smaller than it is in an air channel thanks to the much smaller fluid density. This implies that the decelerating transients have a duration about 1000 times shorter in the water-channel case. This allows much shorter simulation times and better statistical convergence of the flow variables.

In the on-off control, the simulations show that the disc transient time scale and the flow transient time scale are of a similar magnitude (Section 5.3.2). This is an advantage for the purposes of this study because both transients can be studied and compared with no need to perform longer and expensive simulations to resolve the slow disc transient.

5.3 On-off control

We propose an intermittent, or “on-off” forcing scheme as a basic form of feedback-based control of the discs. The motors are operated according to the following activation function:

$$T_M(t) = \begin{cases} T_0 - \alpha W(t) & t_i < t < t_f \\ 0 & t_f < t < t_i \end{cases} \quad \text{with } W(t_i) < W_m; W(t_f) > W_M, \quad (5.8)$$

that is enforced independently on each disc of the array. The sign of the motor torque T_M is determined by position of the disc in the array, complying with the wave-like pattern seen

in Fig. 5.1b. Upper and lower thresholds on the disc-tip velocity are represented by W_m and W_M , respectively. The real-world implementation of (5.8) is relatively effortless, relying on the easily accessible measurement of the instantaneous angular velocity. The on-off activation is studied in order to determine whether it is a more efficient alternative to the steady rotation of the discs. The best power-saving case from Ricco and Hahn (2013) is selected as the reference steady-rotation case and then a computational optimisation technique is applied to search for more efficient combinations of the on-off control parameters (W_M, W_m, T_0).

The broad aim of this approach is to test whether it is possible to enhance the power efficiency of the rotating-disc method of drag reduction by tuning the parameters of an heuristic control law. The underlying control principle is identified as the same as the steady disc operation, i.e. the generation of a sustained spanwise velocity wave.

It can be pointed out that the desired control objective, the skin-friction C_f , is never directly measured, instead the controller is tasked with regulating the disc velocity W . This operation is reasonable because the nature of the dependence of C_f on W is already known in the open-loop regime. Systems where corrections are applied to the control signal rather than the output of interest are typically classified as of the *feedforward* type. In feedforward systems, compliance to the control objectives relies on open-loop models of the response of the system to the control signal (Åström and Murray, 2008). This is distinct from *feedback* systems where the control objective is directly measured and regulated based on the departure from a determined objective. The system architecture is depicted in Figure 5.3a. The sub-system encircled by the dashed line can be regarded to as a self-contained feedback control loop on the disc velocity, with an external disturbance represented by the fluid torque.

The graphical representation of control law 5.8 given in Figure 5.3b highlights its hysteresis-like behaviour, where the control signal depends not only from the current state W but also on the past history – whether it belongs to an accelerating or decelerating transient. It can be also noted that the powered activation is only taking place in the upper branch of the hysteresis loop, whereas the lower branch (at $T_M = 0$) represents the natural dynamics of the disc dissipating its energy into the fluid.

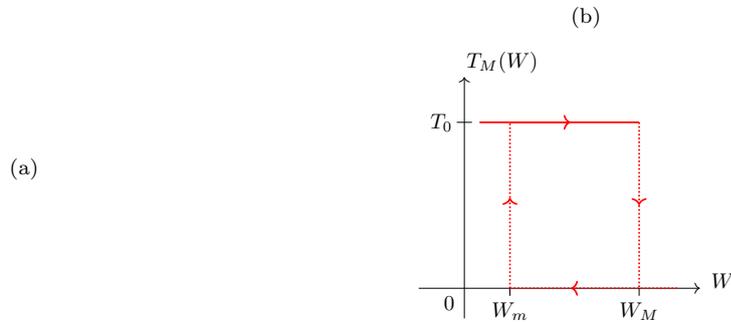


Figure 5.3: On-off feedforward control of the rotating discs. (a) Block diagram of the controlled disc-fluid system. (b) Response function of the on-off velocity controller ($T_M = T_0$ is assumed, neglecting the weak linear dependence of T_M on W).

The rest of the section is organised as follows: Sec. 5.3.1 describes the power budget of the on-off forcing and defines the efficiency metrics. Sec. 5.3.2 discusses the simulation of a controlled switch-on and switch-off transient. Sec. 5.3.3 formally defines the main control design problem.

5.3.1 Power-budget analysis

The analysis of the power-budget components is essential to understand and potentially improve the power-saving performance of an active control scheme. We start by defining the terms of the mean power budget, defined on a single flow unit i.e. the system comprising a single disc and a fluid domain \mathcal{V} of dimension $D \times D \times h$, centred on the disc. The power needed to drive the controlled flow is supplied by two external sources. The first is the pumping power $P_p(t)$ spent to overcome the frictional drag and the other is the motor power $P_M(t)$ needed to move the disc. The instantaneous total input power is given by the sum $P(t) = P_p(t) + P_M(t)$. The mean power spent by the controlled flow can be calculated by averaging in time and across multiple discs, denoted as $\mathcal{P} = [P(t)]_s$. The mean power spent by the disc-fluid system is expressed by the following equations:

$$\mathcal{P}_p = C_f U_b^3 D^2, \quad (5.9)$$

$$\mathcal{P}_M = \overline{T_M \Omega} \quad (5.10)$$

The time-averaged (or averaged over multiple discs) net power \mathcal{P}_{net} (saved by using the control law (5.8) with respect to the steady activation) is calculated as:

$$\mathcal{P}_{net} = \frac{\mathcal{P} - \mathcal{P}_{ss}}{\mathcal{P}_0} = \underbrace{\frac{\mathcal{P}_p - \mathcal{P}_{ss}}{\mathcal{P}_0}}_{\equiv -\Delta\mathcal{R}} + \underbrace{\frac{\mathcal{P}_M}{\mathcal{P}_0}}_{\equiv \Delta\mathcal{P}_M}, \quad (5.11)$$

where \mathcal{P}_{ss} is the mean total power spent by the steadily-rotating discs case, and \mathcal{P}_0 is the mean power spent to drive the stationary-wall flow. \mathcal{P}_0 is equivalent to the mean pumping power of the stationary-wall flow i.e. $\mathcal{P}_0 = C_f^0 U_b^3 D^2$. The critical aspect of using an on-off control scheme is to determine whether it is a more efficient alternative to the steady operation of the disc array. If that is the case, \mathcal{P}_{net} must be a negative quantity.

A very useful tool for analysing how the actuator model affects the power-saving performance of the control scheme is to decompose the motor power into its components:

$$\mathcal{P}_M = \mathcal{P}_f + \mathcal{P}_h + \mathcal{P}_{min} + \dot{\mathcal{E}}_d \quad (5.12)$$

where \mathcal{P}_f is the mean power exchanged with the fluid through the disc surface, \mathcal{P}_h is the mean power dissipated into the housing flow, \mathcal{P}_{min} are the minor losses into the bearing and the disc gap, and $\dot{\mathcal{E}}_d$ is the mean change in the disc kinetic energy. The precise definitions and the derivation of the disc power-budget terms from (5.12) are given in paragraph b. below. Using definition (5.12), the net saved power can be decomposed into its budget contributions, becoming:

$$\mathcal{P}_{net} = \underbrace{\frac{\mathcal{P}_p - \mathcal{P}_{p,ss}}{\mathcal{P}_0}}_{\equiv -\Delta\mathcal{R}} + \underbrace{\frac{\mathcal{P}_f - \mathcal{P}_{f,ss}}{\mathcal{P}_0}}_{\Delta\mathcal{P}_f} + \underbrace{\frac{\mathcal{P}_h - \mathcal{P}_{h,ss}}{\mathcal{P}_0}}_{\Delta\mathcal{P}_h} + \underbrace{\frac{\mathcal{P}_{min} - \mathcal{P}_{min,ss}}{\mathcal{P}_0}}_{\Delta\mathcal{P}_{min}} + \dot{\mathcal{E}}_d. \quad (5.13)$$

$$\underbrace{\hspace{10em}}_{\Delta\mathcal{P}_M}$$

This identity will be used below to analyse the performance of the controlled flow.

a. Disc power budget

The motor power budget (5.10) is obtained by averaging in time the instantaneous power budget of the disc, that is derived by multiplying the disc equation (5.5) by Ω/R and reorganising:

$$P_M = P_f + P_h + P_{min} + \dot{E}_d, \quad (5.14)$$

where:

$$P_M(t) = \Omega(t)T_M(t) \quad (5.15)$$

$$P_h(t) = \Omega(t)T_h(t), \quad (5.16)$$

$$P_{min}(t) = \Omega(t)T_{min}(t), \quad (5.17)$$

$$\dot{E}_d(t) = I\Omega(t)\dot{\Omega}(t), \quad (5.18)$$

$$P_f(t) = T_f(t)\Omega(t), \quad (5.19)$$

where P_h is the power dissipated into the housing flow, P_{min} are the minor losses into the bearing and the disc gap, and \dot{E}_d is the change in the disc kinetic energy.

b. Fluid power budget

It is also interesting to model the power budget in the fluid, which can be derived by integrating the mean kinetic energy equation of the fluid on the flow unit volume \mathcal{V} and taking the time or multi-disc average. It reads:

$$\mathcal{P}_p + \mathcal{P}_f = -\mathcal{P}_\varepsilon \quad (5.20)$$

where the terms are defined as:

$$\mathcal{P}_p = C_f U_b^3 D^2 \quad (5.21)$$

$$\mathcal{P}_f = \frac{1}{Re_p} \iint_{\mathcal{S}} \left(\bar{w}(x, 0, z) \frac{\partial \bar{w}}{\partial y} \Big|_{y=0} - \bar{u}(x, 0, z) \frac{\partial \bar{u}}{\partial y} \Big|_{y=0} \right) dx dz, \quad (5.22)$$

$$\mathcal{P}_\varepsilon = \frac{2}{Re_p} \int_{\mathcal{V}} \overline{\mathbf{s}^2} dx, \quad (5.23)$$

where \mathcal{S} the disc surface and \mathbf{s} is the symmetric part of the velocity-gradient tensor.

It can be noted that the term \mathcal{P}_f is the coupling term, appearing in both the fluid budget (5.20) and in the motor budget (5.12). Equation (5.22) can be obtained from (5.19) by recalling that $w = \Omega(x - x_C)$, and $u = \Omega(z - z_C)$, where (x_C, z_C) is the disc centre and the definition of T_f from 4.14 is used.

The definitions are obtained by integrating the kinetic energy equation of the fluid on the flow unit volume \mathcal{V} . The evolution equation for the mean total kinetic energy of the fluid $E = [\overline{u_i u_i}]_s / 2$ reads (Pope, 2000):

$$\frac{DE}{Dt} + \frac{\partial}{\partial x_i} \mathcal{T}_i = -\varepsilon, \quad (5.24)$$

where the repeated-index notation for sums is used and for a lighter notation, the multi-disc symmetry average is not reported. The terms are defined as:

$$\frac{DE}{Dt} = \overline{u_i} \frac{\partial E}{\partial x_i} \quad (\text{material derivative}) \quad (5.25)$$

$$\mathcal{T}_i = \overline{u_i p} - \frac{2}{Re_p} \overline{u_j s_{ij}} \quad (\text{spatial fluxes}) \quad (5.26)$$

$$\varepsilon = \frac{2}{Re_p} \overline{s_{ij} s_{ij}} \quad (\text{dissipation}) \quad (5.27)$$

where $s_{ij} = (\partial_j u_i + \partial_i u_j)/2$ is the symmetric part of the velocity-gradient tensor. The term-by-term integration of (5.24) on the unit volume \mathcal{V} can be sketched as follows:

$$\int_{\mathcal{V}} \bar{u}_i \frac{\partial E}{\partial x_i} d\mathbf{x} = 0 \quad (5.28)$$

$$\int_{\mathcal{V}} \bar{u}_i \frac{\partial \bar{p}}{\partial x_i} d\mathbf{x} = P_x U_b D^2 = \mathcal{P}_p, \quad (5.29)$$

$$\frac{2}{Re_p} \int_{\mathcal{V}} \frac{\partial}{\partial x_i} \overline{u_j s_{ij}} d\mathbf{x} = \frac{1}{Re_p} \int_{\mathcal{S}} \left(\bar{u}_j \frac{\partial \bar{u}_i}{\partial x_j} + \bar{u}_j \frac{\partial \bar{u}_j}{\partial x_i} \right) n_i d\mathcal{S} = \mathcal{P}_f, \quad (5.30)$$

$$\int_{\mathcal{V}} \frac{2}{Re_p} \overline{s_{ij} s_{ij}} d\mathbf{x} = \mathcal{P}_\varepsilon, \quad (5.31)$$

where \mathcal{S} is the surface enclosing \mathcal{V} and n_i are the components of its normal vector. Exploiting the symmetries of the mean flow over a disc, the integrals (5.29)-(5.31) can be simplified to return the terms of (5.20).

5.3.2 Transient simulations

This section presents two preliminary simulations that address the dynamics of the switch-on and switch-off transients in a controlled fashion. The simulations are aimed at obtaining data on the disc velocity transient and the associated transient flow, focusing primarily the order of magnitude of their duration. The numerical set-up of the present simulations is identical to that illustrated in Sec. 5.2.

The first step is to assess the behaviour of a perfectly sharp transient i.e. such that the disc velocity follows the step-function activation $W(t) = W_{ss}$ when $t_i < t < t_f$ and $W = 0$ otherwise, where $W_{ss} = 0.26$ is the velocity of the best performing constant-speed case from Ricco and Hahn, 2013. The switch-on and switch-off times are respectively $t_i^+ = 100$ and $t_f^+ = 133$. The simulation begins at $t = 0$ from a fixed-wall turbulent flow snapshot. In this way, it is possible to observe the behaviour a transients flow with no influence that happen in a disc velocity transients. The second simulation encompasses the motorised disc transients, corresponding to applying the control law (5.8) using the same fixed values of t_i and t_f of the step-function transient. The motor stall torque, T_0 , is chosen such that $T_0 - \alpha \Omega_{ss} = T_{ss}$, where $\Omega_{ss} = W_{ss}/R$ and T_{ss} is the mean steady-state resisting torque. The value of T_{ss} is derived from the steady-state simulation by averaging the sum of the resisting torques $T_f + T_h + T_b$ over time and across all the discs. In this way, the motor accelerates the discs to the steady-state velocity W_{ss} following a natural transient.

Figure 5.4 depicts the evolution of some fluid and disc parameters for the step-function activation and the motorised activation of the discs. In the top graph of Fig. 5.4a the time evolution of the edge velocity $W(t)$ for each of the eight bottom-wall discs is compared to the step-function activation. Three distinct time-scales are observed in the simulations: one for the disc-tip velocity W , one for the skin-friction coefficient C_f , and a third one for the fluid kinetic-energy transients $\dot{E}_f \neq 0$ that is calculated as the time derivative of the spatial integral of the fluid kinetic energy $E_f = \int_{\mathcal{V}} \mathbf{u} \cdot \mathbf{u}/2 d\mathbf{x}$. In this case, the disc velocity reaches the steady state in about $700t_\nu$ for the start-up transient and $400t_\nu$ for the slow-down transients. The transients of the skin-friction and the fluid kinetic energy are however much longer, of the order of $1000t_\nu$. A lag is observed between the motorised activation of the disc and the modification of the skin-friction coefficient, which is not present in the sharp activation. Further numerical data (not shown) indicate that the two curves may present qualitative differences, however the order of magnitude of the time-scale of the fluid transient do not change significantly from the sharp

to the motorised activation. This is due to the fact that they always dominate because of the comparatively shorter (at least for this parameter combination) disc-velocity transient duration.

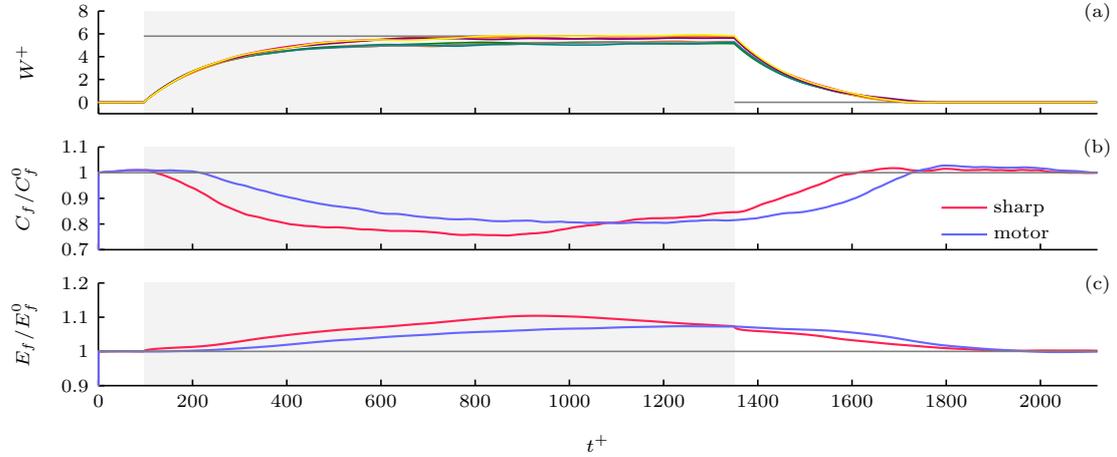


Figure 5.4: Flow statistics across an switch-on-switch-off transient activation of the discs. The shaded region identifies active control i.e. $T_M \neq 0$. Top: disc velocity, middle: skin-friction, bottom: total kinetic energy.

The penetration depth δ is a critical flow variable for the transverse-motion drag reduction paradigm, including the rotating discs or similar techniques (Chapter 2). It is therefore instructive to monitor its time evolution during the transients, implicitly assuming that the skin-friction modification can be attributed to the the (quasi-) steady realisation of a cross-flow motion and not to any transient effects. δ quantifies the thickness of the boundary layer created by the wall motion. The penetration depth is defined as $\delta = -\langle w(y=0)/w'(y=0) \rangle_{x,c}$, where $\langle \cdot \rangle_{x,c}$ denotes the spatial average along the streamwise symmetry plane of a disc and the prime symbolizes derivative with respect to y .

Figure 5.5 shows $\delta(t)$ across the perfectly sharp transients, revealing that in general it builds up and decays very rapidly, on a time scale close to that of W and much shorter than C_f and E_f . The decay of δ during the deceleration of the disc is extremely rapid, of the same order of the time resolution used in the plots.

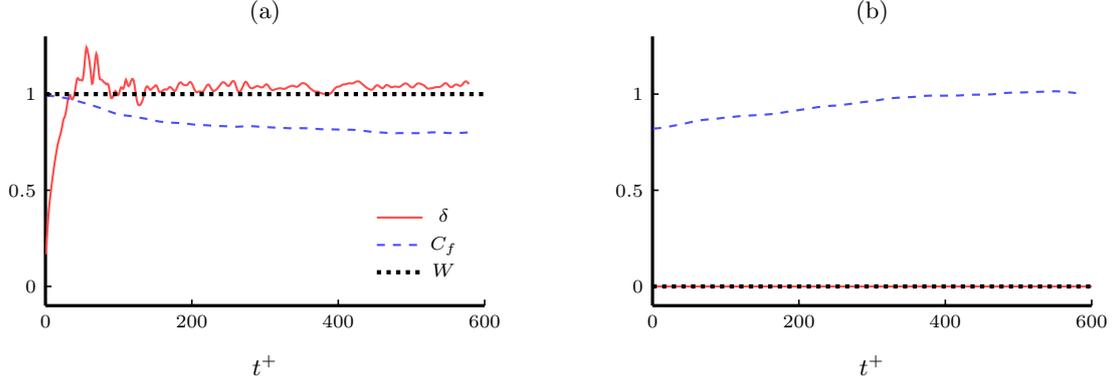


Figure 5.5: Time evolution of the perfectly sharp (a) switch-on and (b) switch-off transients, compared to the temporal evolution of velocity and drag-reduction. The instantaneous value of the Stokes layer thickness δ^+ , the skin-friction C_f and the disc-tip velocity W are normalised with their steady-state value in order to compare their transient time-scales effectively.

5.3.3 Optimisation of the total spent power

The net saved power \mathcal{P}_{net} for the on-off controlled channel is a function of the following parameters:

$$\mathcal{P}_{net} = f(W_M, W_m, T_0, \text{disc}, Re_p). \quad (5.32)$$

The disc parameters are kept constant by assuming the dimensionless thickness $b = 0.05$, cavity depth $d_h = 0.1$, material density $\rho_s = 2.7$ and fixing the disc diameter as $D = 5.07$. The value of $D=5.07$ is chosen as it is the optimal value for drag reduction found in [Ricco and Hahn \(2013\)](#). The Reynolds number is also kept unchanged. A given design of the control scheme is thus specified by the triplet $\mathcal{D}=(W_M, W_m, T_0)$. Since \mathcal{P}_{ss} and \mathcal{P}_0 are fixed, \mathcal{P}_{net} also depends on the same set of parameters of [\(5.32\)](#).

The best power-saving case from [Ricco and Hahn \(2013\)](#) provides the values of the steady-state velocity $W_{ss}=0.26$ and steady-state motor stall torque T_{ss} . For the sake of clarity, it is convenient to normalise the threshold velocities and the torque with W_{ss} and T_{ss} . For example, the point $\mathcal{D}_{ss}=(1, 1, 1)$ in the normalised parameter space corresponds to the initial design.

We set out to investigate the existence of a control design that is more efficient than the steady activation. Formally, the optimal design \mathcal{D}_{opt} of a control strategy is the solution to the problem:

$$\mathcal{D}_{opt} = \underset{W_m, W_M, T_M}{\text{arg min}} \mathcal{P}_{net} \quad (5.33)$$

The parameter space is subjected to the constraints:

$$T_0 > 1.05T_0(W_M), \quad (5.34)$$

$$W_M > W_m. \quad (5.35)$$

The problem [\(5.33\)](#)-[\(5.34\)](#)-[\(5.35\)](#) is referred to as ‘‘OPT1’’ in the following. The constraint [\(5.35\)](#) simply encodes the definition of upper and lower velocity thresholds. Since it is an exact function of the parameters W_M and W_m , any unfeasible parameter combination can be readily discarded without performing a simulation.

The constraint [\(5.34\)](#) excludes values of the motor torque smaller or equal to the torque $T_0(W_M)$ needed to steadily move the disc at W_M . This ensures that the disc can reach the velocity W_M in finite time, even considering instantaneous fluctuations of T_f as large as 5% of

the steady state torque $T_0(W_M)$. The function $T_0(W_M)$ is obtained by quadratic interpolation of 7 constant-velocity simulations performed at uniformly spaced values of W . The torque constraint (5.34) provides a one-sided bound on the parameter space, while both upper and lower bounds must be chosen for both W_M and W_m . These bounds were first set arbitrarily and then refined as preliminary results become available. The final search domain is reported in Tab. 5.2 and a three-dimensional visualization is also provided by Fig. 5.12.

Table 5.2: Absolute bounds on the search domain for the optimisation OPT1.

bound	W_M/W_{ss}	W_m/W_{ss}	T_0/T_{ss}
lower	0.1	0.05	$T_0(W_M)$
upper	1.50	0.95	$5 \cdot T_0(W_M)$

5.4 Bayesian Optimization

Each evaluation of the objective function (5.32) requires a costly DNS simulation and is subject to uncertainty. Furthermore, no information is available on the function's gradient or Hessian. Performing a systematic exploration of the three-dimensional parameter space would be inefficient in terms of the total computational cost and does not explicitly target the function's global minima. We instead adopt a Bayesian optimisation algorithm (Algorithm 3), that is well suited for optimising expensive, "black-box" functions (Shahriari et al., 2016; Jones et al., 1998) of a few variables. In this section, the abstract notation used in the BO literature is used. At the end of the section and in the result presentation we will clarify the correspondence of the theoretical notation to the elements of the physical problem (5.33).

A Bayesian optimisation algorithm makes a prediction about the location of a function's global minimum based on a probabilistic regression model of the available data (Schonlau, 1997; Jones et al., 1998; Shahriari et al., 2016). The data \mathcal{D} is the set of the known values of the true objective function $f(\mathbf{x})$. The set of points where the function is known is indicated by $X = \{\mathbf{x}_1, \mathbf{x}_2, \dots\}$ and the data is the set $\mathcal{D} = \{X, f(X)\}$.

a. Probabilistic regression

The probabilistic model of the function f is denoted by f^* , and the test input set, i.e. the points where the model is evaluated, is denoted by X^* (Williams and Rasmussen, 2006). f is modelled as one realisation of a Gaussian random process with mean $\mu(\mathbf{x})$ and covariance function $k(\mathbf{x}, \mathbf{x}')$. This type of regression is thus called Gaussian Process (GP) regression. The mean and the type of covariance function are parameters of the model and are discussed below. The model f^* calculated at the test input X^* given the data \mathcal{D} consists of a normal probability distribution:

$$p(f^*|X^*, \mathcal{D}) = \mathcal{N}(\bar{f}^*, \text{Var}(f^*)), \quad (5.36)$$

which represents the model prediction of the value of f at X^* and is called the "predictive distribution". The mean value \bar{f}^* and the variance $\text{Var}(f^*)$ are defined as:

$$\bar{f}^* = \mu + K(X^*, X)K(X, X)^{-1}(f - \mu), \quad (5.37)$$

$$\text{Var}(f^*) = K(X^*, X^*) - K(X^*, X)K(X, X)^{-1}K(X, X^*) \quad (5.38)$$

where μ is the expected value of the Gaussian process and $K(X, X')$ is the covariance matrix. The elements of the covariance matrix are defined as $K(X, X')_{ij} = k(\mathbf{x}_i, \mathbf{x}'_j)$ where $k(\mathbf{x}, \mathbf{x}')$ is the covariance function of the Gaussian process.

The expected value μ is typically a constant and can be chosen based on some information on f or set arbitrarily to zero. The type of covariance function must also be chosen. In this study, the Matérn 5/2 covariance function is used:

$$k_{\text{Matern}5/2}(\mathbf{r}) = \theta^2 \exp(-\sqrt{5}\mathbf{r})(1 + \sqrt{5}\mathbf{r} + 5\mathbf{r}^2/3) \quad (5.39)$$

where $\mathbf{r} = |\mathbf{x}'_i - \mathbf{x}_i|$. The choice of this kernel assumes f to be twice differentiable and is recommended in the literature for models of physical processes (Williams and Rasmussen, 2006). The parameters of the regression model (called "hyperparameters" in the literature) are the correlation length scales θ_i , one for each independent variable of the objective function. In practice, the model is calculated from the data by marginalising f^* over all possible values of the hyperparameters given the data. This operation incorporates in the model the uncertainty about the hyperparameters. The predictive distribution is calculated as:

$$p(f^*|X^*, \mathcal{D}) = \int p(f^*|X^*, \mathcal{D}, \boldsymbol{\theta})p(\boldsymbol{\theta}|\mathcal{D}) d\boldsymbol{\theta}, \quad (5.40)$$

where the distribution of the hyperparameters is defined according to Bayes' rule:

$$p(\boldsymbol{\theta}|\mathcal{D}) = \frac{p(f(X)|X, \boldsymbol{\theta})p(\boldsymbol{\theta})}{p(f(X)|X)}. \quad (5.41)$$

This is called the "posterior" distribution, as opposed to $p(\boldsymbol{\theta})$ which is the initial guess or "prior" distribution of the hyperparameters (typically an uniform distribution). The integral (5.40) is solved numerically, sampling from $p(f(X)|X, \boldsymbol{\theta})$ (called the likelihood) using a Monte-Carlo Markov Chain (MCMC) algorithm. At every iteration, or whenever new data become available, the model is updated by re-calculating (5.40) using the new data.

b. Minimum prediction and evaluation

Once the model is determined for the given \mathcal{D} , the next prediction of the minimum \mathbf{x}_{n+1}^* is made by maximising the Acquisition Function $a(\mathbf{x})$ according to:

$$\mathbf{x}_{n+1}^* = \arg \max_{\mathbf{x}^* \in X^*} a(\mathbf{x}^*) \quad (5.42)$$

The specific form of $a(\mathbf{x})$ has to be chosen, with many options being available in the literature for different problem requirements. The Expected Improvement (EI) acquisition function used in this study is a common choice in Bayesian optimisation literature and is the standard for engineering optimisation models (Shahriari et al., 2016). Given that f_{min} is the currently known minimum value of f , the notion of improvement is formalised by defining the following utility measure:

$$u(\mathbf{x}) = \begin{cases} f_{min} - f(\mathbf{x}) & \text{if } f(\mathbf{x}) < f_{min}, \\ 0 & \text{otherwise.} \end{cases} \quad (5.43)$$

The Expected Improvement is defined as the expected value of the above utility function i.e. $a_{EI}(\mathbf{x}) \equiv E[u(\mathbf{x})|\mathcal{D}]$. Maximising $a_{EI}(\mathbf{x})$ finds the point with the highest probability of finding a lower estimate of the global minimum. The EI acquisition allows the algorithm to perform both exploratory and exploitative evaluations of the objective function. An exploratory evaluation at \mathbf{x}_{min}^* occurs when \mathbf{x}_{min}^* is chosen, primarily because $f^*(\mathbf{x}_{min}^*)$ has high variance (i.e. high uncertainty) and therefore can be considered as an exploration of a region of the search space where f is little known. On the other hand, an exploitative evaluation occurs where the model

uncertainty is very low (i.e. f is known well) but the model mean is also low and therefore likely to be a better guess on the minimum. More exploratory evaluations are expected in the early stages of the optimisation, while they tend to decline as the algorithm nears convergence to the true minimum.

The function is then evaluated at \mathbf{x}_{n+1}^* and the new data $\mathbf{x}_{n+1}^*, f(\mathbf{x}_{n+1}^*)$ are used to construct an updated instance of the regression model for the next iteration. As the optimisation progresses, \mathbf{x}_{n+1}^* is expected to converge towards the true minimum \mathbf{x}_{min} . Exact convergence to the global minimum has been proven using the EI acquisition for any function that is exactly a Gaussian process (Bull, 2011). The pseudo-code in Algorithm 3 presents a convenient summary of the BO algorithm steps described above in more detail. A more technical discussion of GP regression and the acquisition function can be found in Appendix 5.A.

Algorithm 3 Outline of a Bayesian Optimisation algorithm, adapted from Shahriari et al. (2016).

```

initial data: evaluate  $f$  at  $m$  randomly selected  $\mathbf{x}$  and store in  $\mathcal{D}$ 
for  $n \leq N$  do
1: read data  $\mathcal{D}$ 
2: fit the probabilistic model parameters to  $\mathcal{D}$ 
3: calculate the acquisition function  $a_{EI}(\mathbf{x})$  on the model
4: determine  $\mathbf{x}_{n+1}^* = \arg \max a_{EI}(\mathbf{x})$ 
5: perform new evaluation at  $\mathbf{x}_{n+1}^*$ 
6: augment  $\mathcal{D}$  with the new point  $\{\mathbf{x}_{n+1}, f(\mathbf{x}_{n+1})\}$ 
end for

```

c. Practical considerations

Before running the optimisation algorithm, the objective function is evaluated at a number of locations $X_{ini}=\mathbf{x}_1, \dots, \mathbf{x}_m$ in order to obtain an initial set of data $\mathcal{D}=\{X_{ini}, f(X_{ini})\}$. These simulations are run before the beginning of the actual optimisation and are thereof independent. The initial locations are determined by a Latin Hypercube Sampling (LHS) algorithm that produces m space-filling random combinations of the design parameters (McKay et al., 2000). The initial data provide some initial information about the objective function in order to train the first instance of the regression model. LHS produces a set of points that is more evenly spaced than a purely random one, better representing the dependence of the unknown function on the parameters. This is a desirable property for estimating efficiently the function variability in high-dimensional spaces and for training surrogate models of the unknown function.

To clarify the abstract terminology used in this section, we remark that applied to (5.33) the objective function f corresponds to \mathcal{P}_{net} and the independent variable corresponds to the design parameters $\mathbf{x} = \{W_M, W_m, T_0\}$. The mean μ has been determined as the average \mathcal{P}_{net} obtained from the initial random simulations. A useful visualisation of the probabilistic model is given in Fig. 5.15a for a function of one variable, and Fig. 5.12 for a function of three variables. For more interesting visualisations and in-depth discussion refer to Chapter 2 of Williams and Rasmussen (2006).

The optimisation is carried out through an ad-hoc Python framework that is able to autonomously run, post-process and launch a sequence of DNS simulations at request of the Bayesian optimizer, the only hard limitation being the maximum allowed running time of 48 hours on the Archer system. The total integration time of each simulation is fixed at $750h/U_p$. The simulations are performed using 1152 parallel computational cores, resulting in a typical computation time of ≈ 6 hours. The open-source Python toolkit EmuKit (Paley et al., 2019)

provides the Bayesian optimizer and is easily coupled to the flow solver *Incompact3D* through the Python framework. We also collaborated to the development of the code, under the supervision of the EmuKit team. The code is freely available online at <https://github.com/amzn/emukit>.

5.4.1 Application to a closely related flow control problem

In this section, the Bayesian optimization framework is applied to the already available net-saved-power data of Ricco and Hahn (2013). The aim of this procedure is to validate the use of GP regression as a surrogate model for a problem as similar as possible to the net-saved-power optimization of the disc intermittent control. While not being an exact validation, it rather provides heuristic evidence in support of the applicability to the main research problem. The dots in Fig. 5.6a represent the combinations of the parameters (W, D) considered originally in the systematic study of Ricco and Hahn (2013). The colour maps instead visualise the mean and the variance of the GP regression performed over the data and the expected improvement acquisition function. The maximum expected improvement is very close to the current best estimate of the minimum (red dot). Fig. 5.6b shows the GP regression resulting from down-sampling the exact data from 49 to just 16. Both the structure and the location of the maximum EI do not depart significantly from the full dataset case, in spite of the available information about the objective function being just 30% of the full case.

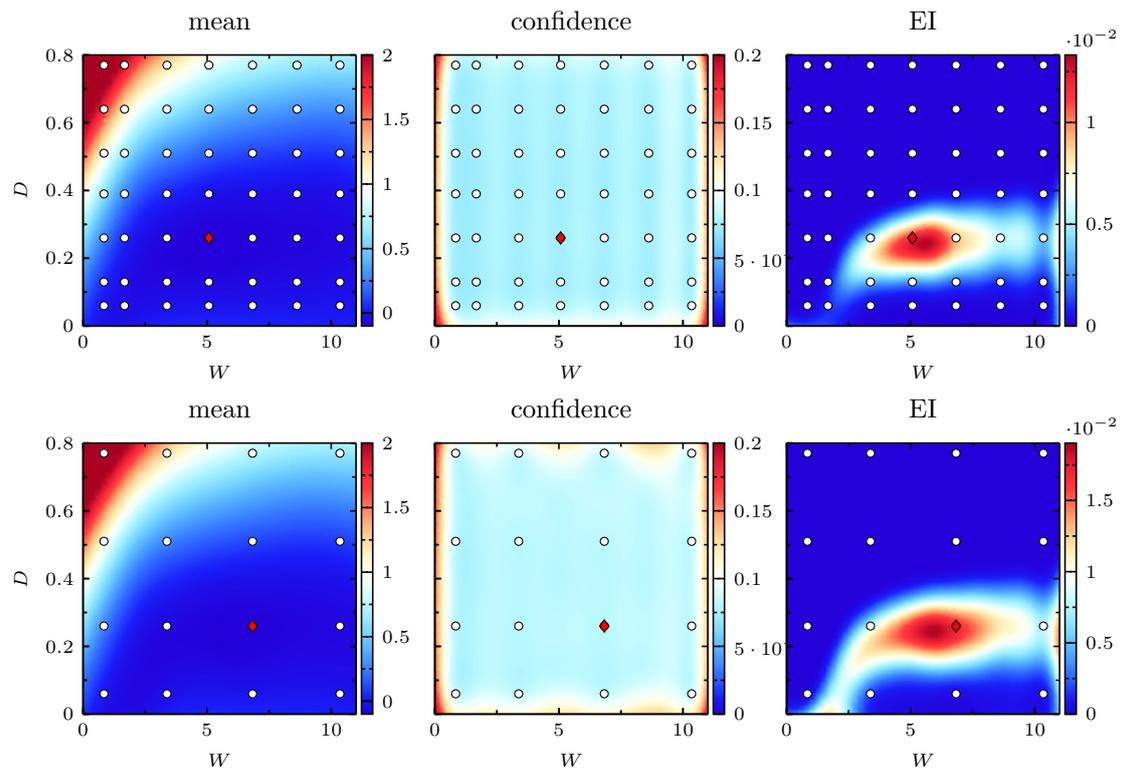


Figure 5.6: GP regression of $\mathcal{P}_{net}(W, D)$ using the systematic data from Ricco and Hahn (2013). The GP model uses the Matern52 kernel and full integration of the hyperparameter posterior.

5.5 Results and data analysis

This section contains the results of the on-off control design problem OPT1, defined in Sec. 5.3.3. Additionally, in Sec. 5.5.1, the on-off control is compared to an optimised version of the steady forcing and Sec. 5.5.2 examines the performance of the scheme relative to the motor power.

a. General results

Table 5.3 lists the main parameters of both the exploratory and optimisation simulations. The parameter-space exploration cases are letter-coded “E” and the optimisation cases “C”. The steady case is the initial design where the discs are constrained to rotate at the constant velocity W_{ss} in the alternating row pattern of Fig. 5.1 and provides the values of T_{ss} and \mathcal{P}_{ss} used in normalising the parameters. The fluid torque on the disc surface is measured at each time step, averaged and added to the exact housing torque (from (4.19)) to obtain T_{ss} . \mathcal{P}_{ss} follows by adding the steady-state motor power $\mathcal{P}_{M,ss}=2T_{ss}W_{ss}/D$ (and the minor losses) to the steady-state pumping power $\mathcal{P}_{p,ss}$ calculated from the skin friction. The initial $100h/U_p$ time units of each simulation are discarded in order get rid of the transient that initially steadily rotating discs to the on-off activation. The table also reports the results for the original simulation of Ricco and Hahn (2013) at the same diameter and velocity.

For each case, the Table 5.3 reports the parameter values, the total power difference \mathcal{P}_{net} and the partial power differences for the motor power ($\Delta\mathcal{P}_M$) and pumping power ($\Delta\mathcal{P}_p$).

No control design is found that can outperform the steady case in terms of the total saved power to a statistically significant level. Whenever a negative value of $\Delta\mathcal{P}_p$ is encountered, it is always balanced by a positive value of $\Delta\mathcal{P}_M$ and vice versa. Some of the designs (e.g. C05) display instead an enhanced power consumption in both terms. Conversely, none of the observed designs can achieve a simultaneous reduction of the motor and the pumping power. In all the parameter combinations that can achieve good performance we observe that $\Delta\mathcal{P}_M < 0$ and $\Delta\mathcal{P}_p > 0$ i.e. they save motor power but at the expense of a larger drag. This is expected as the on-off control concept is based on saving motor power without overly compromising the drag reduction. The two overall best control designs correspond to cases C12 and C13, which can be deemed to be very close to the total power consumption of the steady case.

In the final phase of the search, from design C09 onwards, the upper and lower velocity thresholds stabilise around $0.65W_{ss}$ and $0.60W_{ss}$ respectively. In this phase, the search reduces to determining the optimal value of T_0/T_{ss} , with values oscillating between around 1 to around 2. Interestingly, the two best designs entail very similar velocity thresholds (and therefore mean velocities) but a markedly different value of the motor torque parameter.

In the following paragraphs, the results are discussed with the help of more detailed statistics and visualisations and special attention is paid to designs C12 and C13.

b. Qualitative behaviour of the controlled discs

The three graphs of Fig. 5.7a show the disc-velocity time histories resulting from three different combinations of the control parameters. For each case, two $W(t)$ time series are plotted, one of a disc from the upper wall and one from the lower wall, in order to visualise the random nature of the inception and duration of the transients. The three designs E09, C13 and E08 are chosen because they are representative of the qualitative diversity of the outcomes. E09 has a relatively low value of the torque parameter and its upper and lower velocity thresholds differ by around 35% of W_{ss} . This produces long switch-on transients and comparatively short switch-off transients. Design C13 is instead characterised by the minimum possible (3% of W_{ss}) separation between upper and lower thresholds, and a relatively high torque ratio. This produces a series of

Table 5.3: Simulation results for the parameter combinations of the space exploration (E) and the optimisation OPT1 (C). The highlighted rows corresponds to the colour-coded cases discussed in the text, in Fig. 5.11 and below.

Case	W_M/W_{ss}	W_m/W_{ss}	T_0/T_{ss}	$\mathcal{P}_{net}(\%)$	$\Delta\mathcal{P}_p(\%)$	$\Delta\mathcal{P}_M(\%)$	\mathcal{G}	\bar{W}/W_{ss}
steady	1.00	1.00	1.00	± 0.87	0	0	1.22	1.00
RH13	1.00	1.00	1.00	N/A	N/A	N/A	2.19	1.00
E01	1.08	0.15	2.71	1.27 ± 1.40	10.29	-8.92	1.31	0.52
E02	1.38	0.96	2.51	7.25 ± 1.83	-1.96	9.11	0.85	1.16
E03	1.28	0.60	1.73	3.02 ± 1.32	1.73	1.10	1.03	0.93
E04	1.33	0.33	2.32	1.80 ± 1.23	4.95	-3.14	1.13	0.76
E05	1.03	0.42	1.54	0.45 ± 1.45	5.68	-5.10	1.27	0.72
E06	1.18	0.51	1.34	2.10 ± 1.40	2.95	-1.00	1.10	0.87
E07	1.23	0.24	2.90	1.18 ± 1.31	7.35	-6.15	1.24	0.64
E08	1.48	0.87	1.93	7.75 ± 1.29	-2.11	9.63	0.84	1.16
E09	1.13	0.78	1.15	3.56 ± 1.66	0.43	3.04	1.00	0.99
E10	1.43	0.69	2.12	4.93 ± 1.19	-0.22	4.95	0.93	1.04
E11	0.90	0.45	3.00	0.20 ± 1.50	7.49	-7.29	1.39	0.65
E12	1.42	0.81	3.38	6.12 ± 1.66	-0.81	6.93	0.88	1.09
E13	1.27	0.19	4.68	2.06 ± 1.30	8.70	-6.64	1.15	0.61
E14	1.10	0.50	4.03	1.44 ± 1.73	5.29	-3.85	1.17	0.77
C01	1.21	0.38	4.27	1.27 ± 1.42	5.70	-4.43	1.19	0.73
C02	1.21	0.34	4.20	0.76 ± 1.45	5.80	-5.05	1.26	0.70
C03	0.10	0.07	1.05	3.56 ± 1.54	18.82	-15.26	0.14	0.09
C04	0.25	0.22	3.82	2.52 ± 1.21	16.71	-14.20	1.72	0.25
C05	1.50	0.53	5.00	3.63 ± 1.34	1.27	2.36	0.99	0.94
C06	0.63	0.60	1.05	0.14 ± 1.41	9.62	-9.47	1.55	0.56
C07	0.73	0.31	2.35	0.75 ± 1.25	10.91	-10.15	1.50	0.51
C08	0.74	0.03	5.00	2.20 ± 1.32	15.32	-13.12	1.52	0.28
C09	0.63	0.60	1.05	0.08 ± 1.49	9.52	-9.45	1.56	0.57
C10	0.64	0.61	1.05	-0.19 ± 1.69	9.06	-9.25	1.58	0.58
C11	0.67	0.64	1.87	0.14 ± 1.27	7.62	-7.49	1.41	0.65
C12	0.68	0.65	1.05	-0.22 ± 1.71	8.30	-8.52	1.53	0.61
C13	0.63	0.60	2.17	-0.04 ± 1.56	8.26	-8.30	1.48	0.62
C14	0.63	0.60	2.10	-0.27 ± 2.74	8.01	-8.29	1.52	0.62
C15	0.65	0.62	1.88	0.04 ± 1.51	8.03	-7.99	1.45	0.63

high-frequency, low amplitude regular oscillations that closely resembles a steady control. The design E08 has similar velocity thresholds to E09 but a much larger torque parameter, resulting in faster switch-on transients whose duration is comparable to that of the switch-off phase.

There is also another qualitatively different class of behaviour that is encountered in designs such as C03 or C12. These are characterised by the minimum value of the torque parameter (5% above unity) and two very close velocity thresholds. The torque intensity is high enough to sustain $T_0(W_M)$ but not enough to give rise to strong accelerations, producing an extremely long switch-on transient. Therefore these designs result in a quasi-steady regime and do not reach the upper velocity threshold within the entire simulation time.

Fig. 5.7b represents the time-dependent power budgets for E08 according to (5.12), normalised with the reference channel pumping power \mathcal{P}_0 . The motor power time-history is taken from a single disc, while the pumping power is derived from the global average of C_f . It can be observed that $\mathcal{P}_M(t)$ is exactly zero during the switch-off phase and has a discontinuous switch-on due to the disc velocity being non-zero when the motor starts up. The transient term $\dot{\mathcal{E}}_d$ has a sharp peak at start-up due to high acceleration intensity, subsequently decaying to a nearly constant value. This type of behaviour is observed in every simulation with clearly defined cycles, with varying peak intensities and decay durations that depend on the specific characteristics of the transient. The time-dependent power budgets of every case of OPT1 can be seen in Fig. 5.8 and Fig. 5.9.

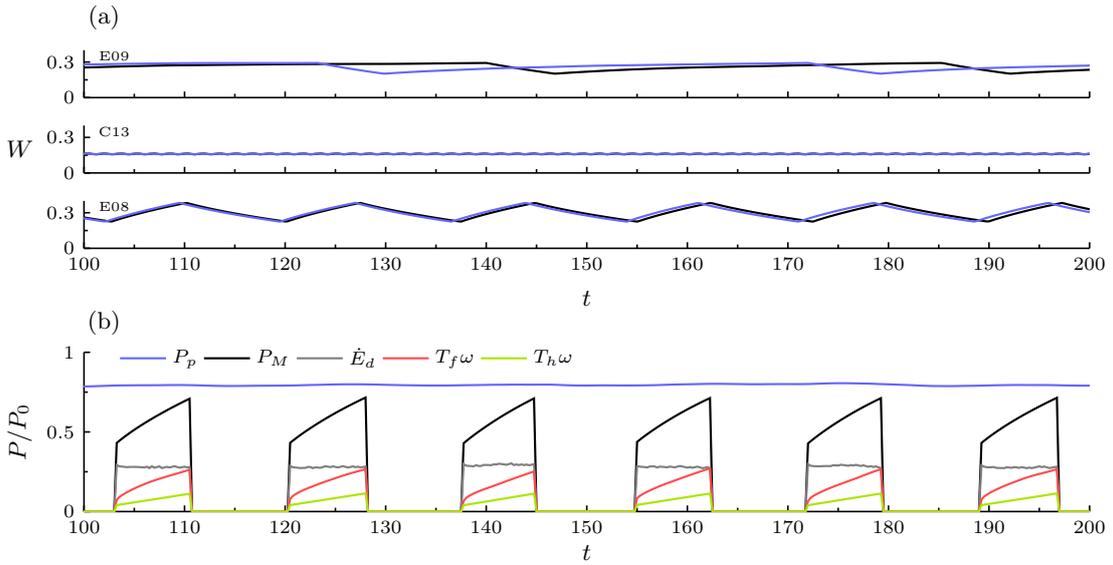


Figure 5.7: Disc-tip velocity and power budget time-series on an time interval of 100 viscous time units. (a) time series of the instantaneous disc-tip velocity $W(t)$ for three different control designs from Tab. 5.3. A disc from the lower and one from the upper wall are chosen to represent each case. (b) breakdown of the instantaneous motor power budget for case E08.

c. Power-budget data

The bar chart plots of Fig. 5.10 depict the time-averaged power-budget breakdown of the total power according to (5.12) for all the simulations of Tab. 5.3. The numerical values of the power-budget terms are reported in Tab. 5.4. The values of both \mathcal{P}_f and \mathcal{P}_p resulting from the

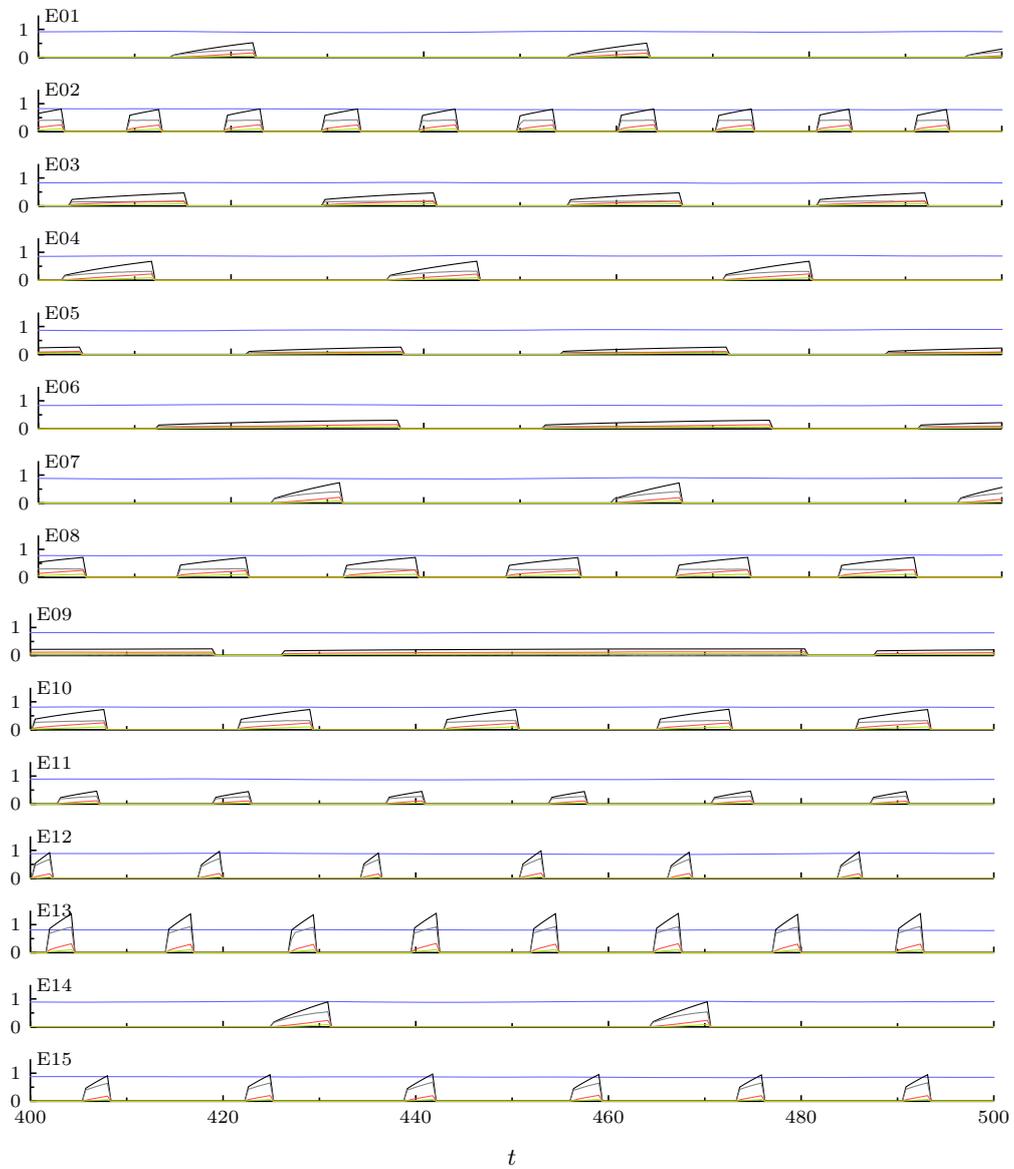


Figure 5.8: Time-series of the instantaneous power-budget terms for the exploratory simulations (E series in Tab. 5.3). A 100 time-units snapshot is represented.

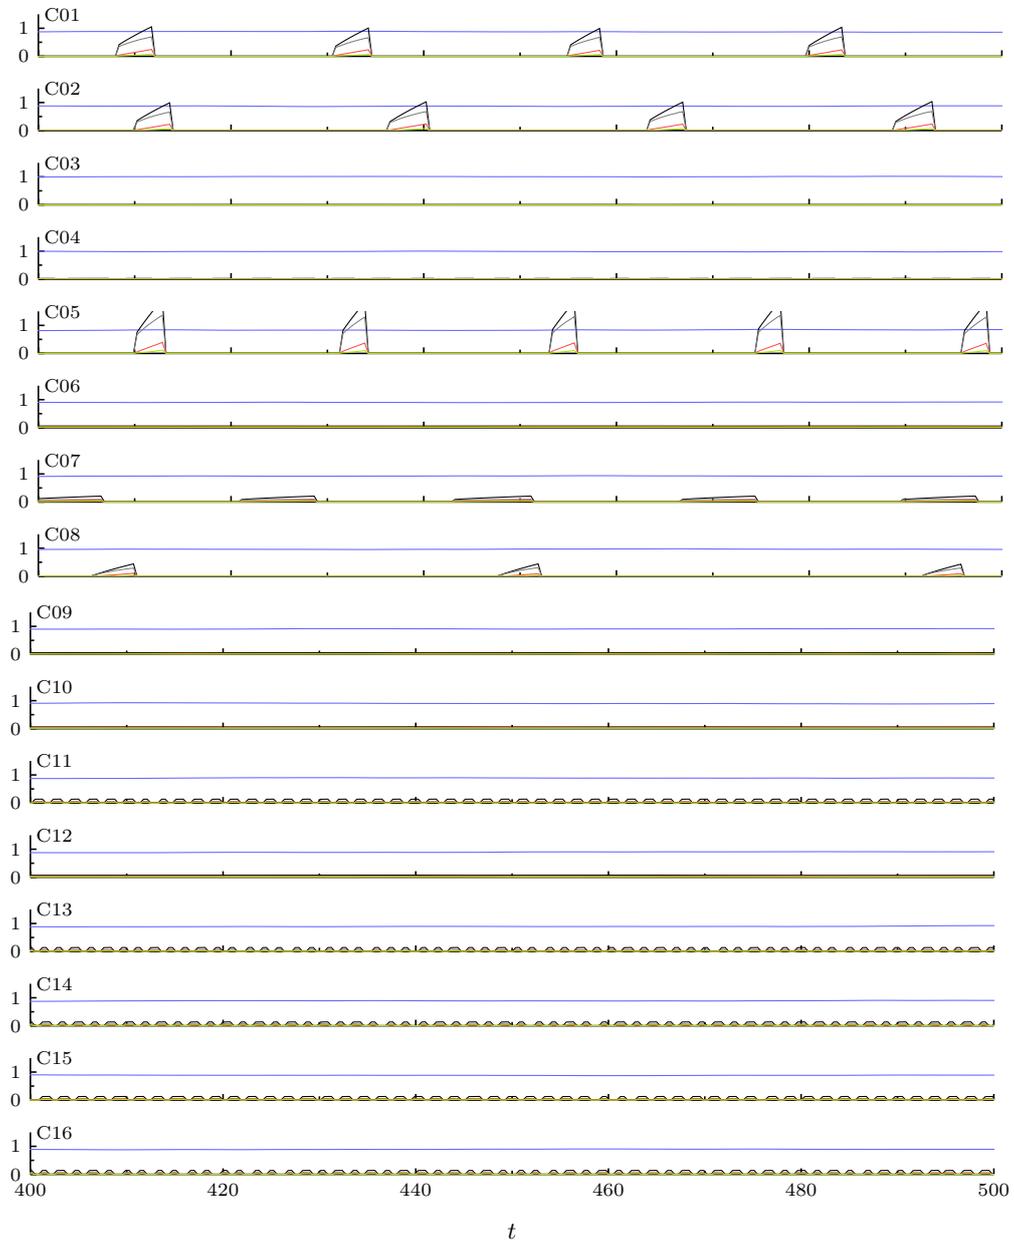


Figure 5.9: Time-series of the instantaneous power-budget terms for the OPT1 simulations (C series in Tab. 5.3). A 100 time-units snapshot is represented.

steady simulation are close to the values obtained in Ricco and Hahn (2013). The latter did not entail any modelling of the housing and bearing losses, together accounting for an additional 4.5% power spent w.r.t. the reference channel. It can be noted that, as the optimiser nears convergence (C09 to C15), the total power becomes almost constant and close to that of the steady control. Moreover, they all present a similar make-up of the different terms, with less than 90% of the power spent going into \mathcal{P}_p , < 3% into \mathcal{P}_f and 1% into \mathcal{P}_{min} . The only exception is the transient power $\dot{\mathcal{E}}_d$ that can take values from around 3% to 0% due to C09, C10 and C12 belonging to the quasi-steady regime described above. The quasi-steady cases show little to no transient power contributions from $\dot{\mathcal{E}}_d$, which is however balanced by higher values of \mathcal{P}_f yielding a total power budget close to the on-off activated cases. It must be noticed that only cases E02, E08 and E12 achieve a better drag reduction (smaller \mathcal{P}_p) than the steady case, however they use much more transient power which makes their total power budgets some of the largest of all cases. On the other hand, all the cases except for the same E02, E08 and E12 spend less motor power than the steady control.

The exploration-exploitation capability of the Bayesian optimiser is evident in Fig. 5.10b. In cases C01 through C08, significant variability of the parameters and highly inefficient designs are recorded, indicating an explorative (i.e. uncertainty-dominated) strategy. From case C09 onwards, the algorithm becomes purely exploitative (i.e. mean-dominated), with a more careful variation of the parameters needed to single out the best design in a narrower region of the search space.

The optimisation was stopped at C15 after observing that the 16th iteration of the optimiser did not entail any significant variation of the parameters with respect to cases C09 - C15 neither ‘decreasing’ values of the maximum expected improvement (Fig. 5.10c). The structure of the regression model of the data (Fig. 5.12) also provides evidence of convergence towards a region of the parameter space and is discussed below in more detail.

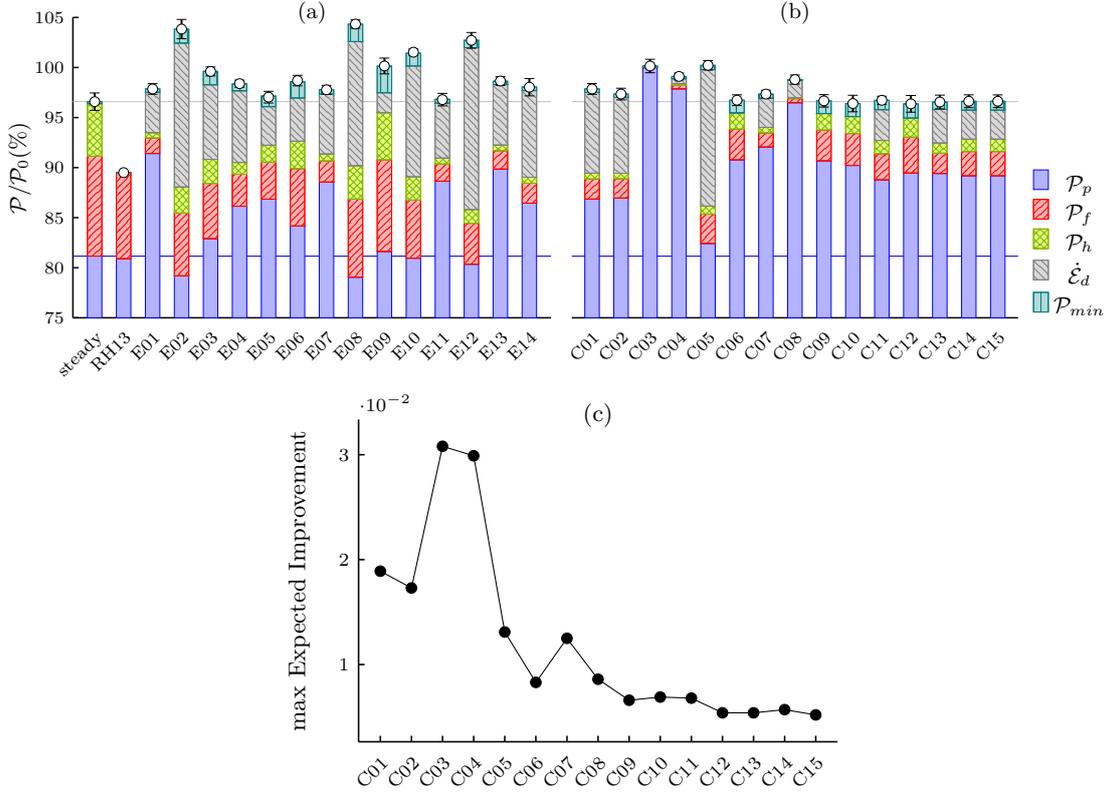


Figure 5.10: Optimisation of the mean power budget. (a) exploratory simulations. (b) OPT1 simulations. The power budget terms are indicated by different colors and hatching patterns and stacked on top of each other to return the total spent power. The power values are normalised with the total power \mathcal{P}_0 adsorbed by the reference uncontrolled flow. The error bars represent the 95% confidence intervals on the total power. In order to better visualise the minor budget terms, the y-axis starts at 75%. The numerical values are given in Table 5.4. The blue horizontal line marks the steady case pumping power and the grey horizontal line marks the total power budget of the steady case. (c) evolution of the maximum Expected Improvement over the course of the optimisation.

d. Correlation structure of the power-budget terms

Parallel coordinates are a useful tool to visualise correlation across many variables in high-dimensional data sets. The parallel coordinate plot of Fig. 5.11 has the three design parameters on the three leftmost axes, while the motor, pumping and total power are represented by the three central axes. The fourth axis report the values of the time-averaged disc velocity \overline{W} . A single control design is visualised by a discontinuous curve connecting each parameter value across the parallel axes, allowing a visually intuitive comparison of many different combinations. Although Fig. 5.11 shows all the designs from Tab. 5.3, we chose to highlight the best 10 cases in terms of \mathcal{P}_{net} (the last axis) plus case E13 for reasons that are discussed below. Inspecting the shape of the curves reveals that the following facts hold in general:

- i) the crossing pattern of the respective curves identifies an inverse correlation between the pumping power $\Delta\mathcal{P}_p$ and the motor power $\Delta\mathcal{P}_M$,

Table 5.4: Numerical values of the contributions to the mean power-budget of the cases in OPT1.

Case	$\frac{W_M}{W_{ss}}$	$\frac{W_m}{W_{ss}}$	$\frac{T_0}{T_{ss}}$	\mathcal{P}				
				\mathcal{P}_p	\mathcal{P}_M			
					\mathcal{P}_f	$\dot{\mathcal{E}}_d$	\mathcal{P}_h	\mathcal{P}_{min}
steady	1.00	1.00	1.00	81.16	9.98	-	5.25	0.20
RH13	1.00	1.00	1.00	81.20	8.60	-	-	-
E01	1.08	0.15	2.71	91.41	4.03	1.55	0.52	0.35
E02	1.38	0.96	2.51	79.19	14.38	6.27	2.60	1.38
E03	1.28	0.60	1.73	82.91	7.45	5.52	2.40	1.34
E04	1.33	0.33	2.32	86.14	7.16	3.20	1.18	0.69
E05	1.03	0.42	1.54	86.85	3.86	3.70	1.68	1.06
E06	1.18	0.51	1.34	84.17	4.32	5.74	2.73	1.61
E07	1.23	0.24	2.90	88.56	5.97	2.11	0.70	0.43
E08	1.48	0.87	1.93	79.05	12.40	7.81	3.33	1.74
E09	1.13	0.78	1.15	81.63	2.00	9.16	4.70	2.66
E10	1.43	0.69	2.12	80.95	11.07	5.81	2.33	1.27
E11	0.90	0.45	3.00	88.65	5.49	1.72	0.59	0.39
E12	1.42	0.81	3.38	80.35	16.17	4.09	1.38	0.74
E13	1.27	0.19	4.68	89.85	6.03	1.85	0.56	0.34
E14	1.10	0.50	4.03	86.45	8.65	2.00	0.60	0.36
C01	1.21	0.38	4.27	86.86	8.06	2.01	0.58	0.35
C02	1.21	0.34	4.20	86.96	7.58	1.92	0.55	0.34
C03	0.10	0.07	1.05	99.98	0.00	0.05	0.03	0.10
C04	0.25	0.22	3.82	97.87	0.50	0.35	0.16	0.23
C05	1.50	0.53	5.00	82.43	13.63	2.93	0.80	0.43
C06	0.63	0.60	1.05	90.78	-0.01	3.07	1.64	1.26
C07	0.73	0.31	2.35	92.06	2.90	1.41	0.55	0.42
C08	0.74	0.03	5.00	96.48	1.72	0.41	0.10	0.08
C09	0.63	0.60	1.05	90.68	-0.01	3.08	1.65	1.27
C10	0.64	0.61	1.05	90.22	-0.01	3.19	1.71	1.30
C11	0.67	0.64	1.87	88.78	3.08	2.61	1.32	0.94
C12	0.68	0.65	1.05	89.46	-0.01	3.59	1.92	1.41
C13	0.63	0.60	2.17	89.42	3.35	2.03	1.01	0.74
C14	0.63	0.60	2.10	89.17	3.22	2.10	1.05	0.77
C15	0.65	0.62	1.88	89.19	2.90	2.43	1.23	0.89

- ii) the mean disc-tip velocity \bar{W} is inversely correlated to $\Delta\mathcal{P}_M$ and directly correlated to $\Delta\mathcal{P}_p$,
- iii) consequently, a more complex correlation structure exists between \bar{W} and the total power \mathcal{P}_{net} .
- iv) a certain degree of direct correlation is observed between \mathcal{P}_f and \mathcal{P}_h .

Restricting the analysis to the best cases (dark grey and colour), it can be concluded that they share the following characteristics:

- i) increased pumping power (drag) w.r.t. the steady case with $7\% < \Delta\mathcal{P}_p < 10\%$,
- ii) decreased motor power w.r.t. the steady case with $-10\% < \Delta\mathcal{P}_M < -7\%$,
- iii) the values of the mean disc-tip velocity are closely clustered around $\bar{W} \approx 0.6W_{ss}$,

No straightforward correlation is observed between the torque parameter and the power outcomes. For example, it can be noted that the two cases C12 and C13 have analogous velocity thresholds while their torque parameters and transient duration are markedly different. Specifically, C12 has the minimum T_0/T_{ss} and the longest (quasi-steady) transients while C13 has a twice as large torque parameter and the shortest transients of any combination. Design E13 is noteworthy because its mean disc-tip velocity suggests it belongs to the best cases, while in fact its power efficiency is much worse (by $\approx 2\%$) due essentially to a higher value of $\Delta\mathcal{P}_M$. The most significant difference between E13 and the best cases is its much higher T_0 . In a similar way, design E07 is an outlier among the best cases because it is the only case with a relatively large T_0 which still returns a good \mathcal{P}_{net} . Following the above observations about E07 and E13 it can be conjectured that, even though \bar{W} is certainly the critical variable to the power efficiency, a minor role is also played by T_0 .

We also computed the median duration of the switch-on and switch-off transients and established that they are not correlated to any term of the power budget, with the exception a weak correlation to $\dot{\mathcal{E}}_d$.

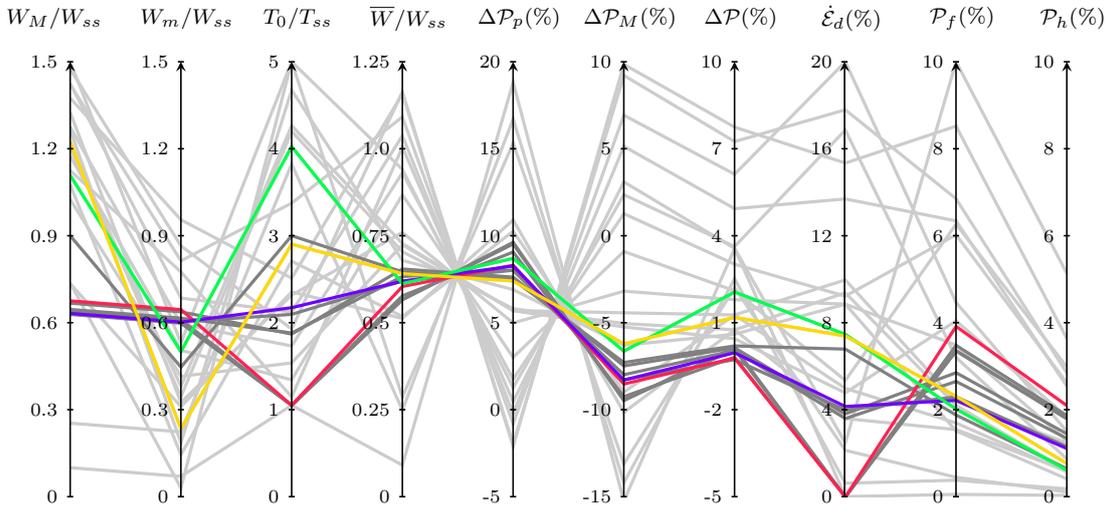


Figure 5.11: Parallel-coordinate visualisation of the results of OPT1. The numerical values are in Tab.5.3. The highlighted designs (both in colour and in dark grey) are the 10 most power-efficient cases, with the exception of E13 and E07. The other designs highlighted in colour correspond to cases C12 and C13 which are discussed in more detail in the text.

Fig. 5.12 presents the three main components of the probabilistic model for the objective function \mathcal{P}_{net} constructed by the Bayesian optimiser at the 15th iteration and therefore being the most complete available. The predictive distribution of \mathcal{P}_{net} , obtained by Gaussian Process regression of the DNS data of Tab. 5.3, is visualised in Fig. 5.12a-b through its mean and variance. The expected value of \mathcal{P}_{net}^* shows a markedly anisotropic structure, the gradient component along T_0/T_{ss} being much weaker than the gradient components along the two velocity thresholds. This fact implies that optimising the torque parameter is a harder problem than optimising the velocities, and explains why the algorithm converges more quickly to the optimal values of the velocity thresholds. It can be concluded that having a low value of T_0 is an advantage, however not as much as having the velocity threshold around 0.6. The variance of \mathcal{P}_{net}^* shows as expected small values in the region of space where the majority of the evaluations have taken place but is overall quite uniform. Fig. 5.12c shows the expected improvement, demonstrating that the minimum is most probably located in the region where the threshold velocities are ≈ 0.6 (red line) and the torque parameter < 2 . This region also has low variance, hence the EI essentially reflects the structure of the expected value as it would be expected from an exploitatively converged optimisation.

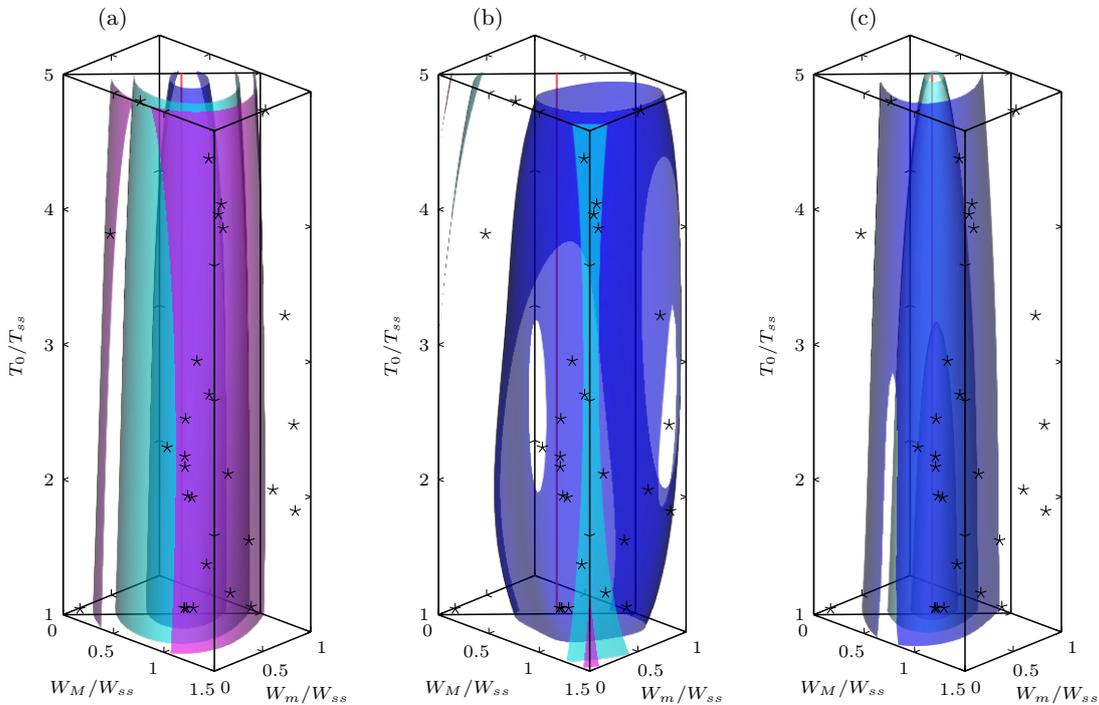


Figure 5.12: Isosurface visualisation of the probabilistic model for \mathcal{P}_{net} in the three-dimensional parameter space. The isosurfaces are taken at three equally spaced levels, increasing from blue to cyan to magenta. The diagonal plane represents the constraint $W_M > W_m$. From left to right the graphs show: (a) expected value, (b) $2\sigma(\mathbf{x})$ confidence interval and (c) expected improvement $EI(\mathbf{x})$. Each black star marker represent a simulation from Table 5.3. The red line indicates where the two velocity thresholds assume a value of 0.6, corresponding to the approximate vicinity of the optimal cases from OPT1.

e. Correlation of the power-budget terms to the mean disc velocity

Fig. 5.13a uses orthogonal axes to further inquire on how the mean velocity \bar{W} correlates with the power efficiency. It reveals that $\Delta\mathcal{P}_p$ is in an almost perfect linear inverse correlation to \bar{W} , with a slope coefficient of -21% estimated by least squares regression. Conversely, $\Delta\mathcal{P}_M$ shows a non-linear positive correlation structure. Their sum \mathcal{P}_{net} exhibits a concave shape, with the minimum \bar{W}_c located around $0.6W_{ss}$. At \bar{W}_c the extra pumping power requirement ($\Delta\mathcal{P}_p > 0$) and the motor power saving ($\Delta\mathcal{P}_M < 0$) balance out to give values close to zero. It can be noticed that the data are more densely clustered in the vicinity of the minimum, where all the late, exploitative evaluations are located as the optimiser nears convergence.

As previously highlighted, E07 and E13 are less efficient than the optimal cases even though they produce optimal values of \bar{W} . From Fig. 5.13a it can be concluded that the reason is that they spend more motor power. A closer look at the budget contributions to \mathcal{P}_M in Fig. 5.10 reveals that the reason E07 and E13 underperform with respect to the optimal cases can be tracked down to their much larger values of $\dot{\mathcal{E}}_d$. The same conclusion can also be drawn from Fig. 5.13b which also shows that individually the two components do not correlate well to the mean disc-tip velocity.

We also point out that the dependence structure of \mathcal{P}_p is completely explained by \bar{W} , whereas \mathcal{P}_M is well correlated to \bar{W} , but the presence of outliers (such as E13 and E07) suggests that some other variable must also play a role. In the next paragraphs we analyse more closely the dependence of \bar{W} , \mathcal{P}_f and $\dot{\mathcal{E}}_d$ on the control parameters.

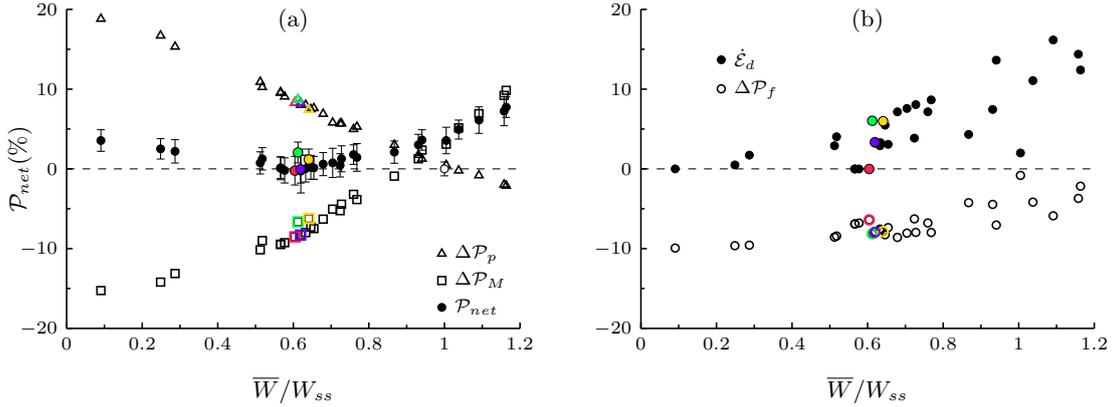


Figure 5.13: Correlation of two spent power metrics to the mean disc velocity. (a) correlation of \mathcal{P}_{net} and its two main components to the normalised mean disc velocity \bar{W}/W_{ss} . (b) correlation of \mathcal{P}_f and $\dot{\mathcal{E}}_d$ to the normalised mean disc velocity. The empty symbol identifies the steady control. The coloured symbols identify the designs C12, C13, E07, and E13. The confidence intervals on \bar{W}/W_{ss} are not shown because they are too small.

f. Modelling the power-budget dependence on the control parameters

In the first place we find that \bar{W} is remarkably well linearly correlated to W_M and W_m , and is almost completely uncorrelated to the torque parameter. Fitting a simple linear regression model to $\bar{W}(W_M, W_m, T)$ gives $\bar{W} = 0.47W_M + 0.58W_m$ with a determination coefficient $r^2 = 0.98$.

As \mathcal{P}_f and $\dot{\mathcal{E}}_d$ are non-trivial functions of the design parameters (Fig. 5.13b) and can't be explained through \bar{W} only, we used another strategy to investigate their dependence structure.

In this case, we solve a simplified model of the disc dynamics where the fluid torque T_f acting of the disc's flow-facing surface, instead of being simulated through DNS is substituted by an

exact model, either linear (i.e. Eq. (4.19)) or Von-Karman type Eq. (4.20). If a high degree of correlation is observed between the DNS and the approximate model results, we can think of studying the exact model in order to gain more insight about the underlying physics instead of solely relying on regression analysis of the DNS data. A certain degree of noise is nonetheless expected in the DNS data due to the randomness in the transient inceptions and ends. The models are used to predict a switch-on and a switch-off transient for the same control parameter combinations as from the data. The solution is obtained analytically in the linear case (Appendix 5.B) and numerically in the VK case. From the approximate solution, various mean quantities can be calculated such as \bar{W} and the mean $\bar{\mathcal{E}}_d$ across the transients. The analytical solution of the linear model and the derivation of an exact parametric formula for the mean disc-tip velocity can be found in Appendix 5.B.

Fig. 5.14 shows that the mean disc-tip velocity from both approximate models correlates remarkably well with the DNS results. Analysis of the exact formula for \bar{W} (5.71) reveals a weak dependence on T compared to the threshold velocities. The formula also demonstrates that while the simple linear regression of the data may provide an excellent fit, it does not reflect the actual physics encoded by formula (5.71). The other two quantities of interest, i.e. \mathcal{P}_f and $\bar{\mathcal{E}}_d$ also correlate well with the DNS data, encouraging the use of the present approach.

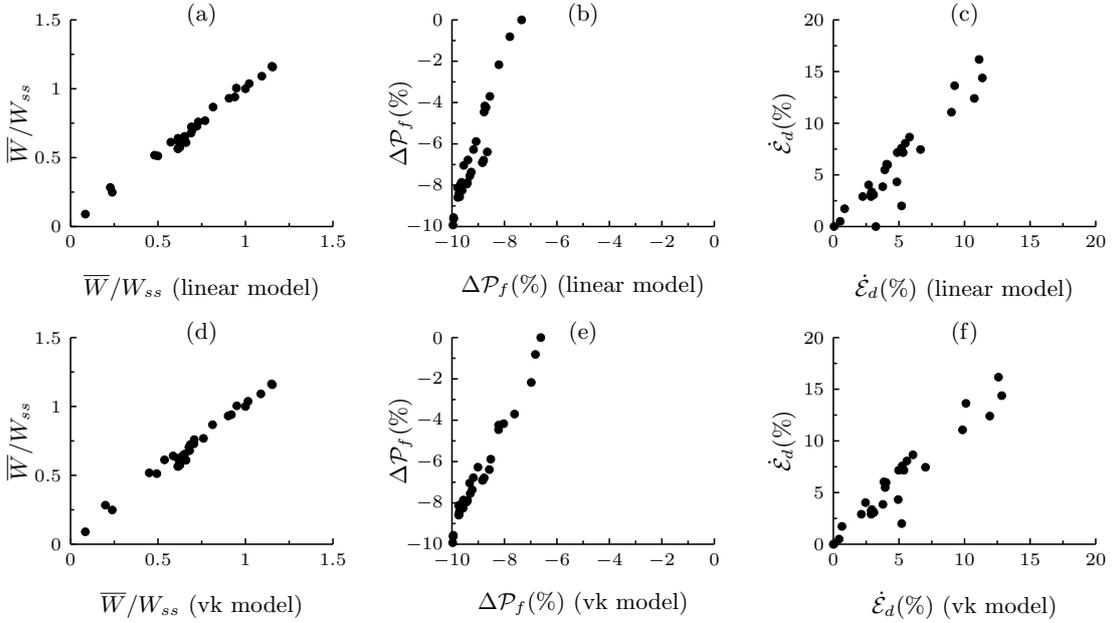


Figure 5.14: Comparison between the approximate models of the transients and the DNS simulations. Top row: linear model of T_f Bottom row: Von-Karman type model of T_f .

5.5.1 Steady forcing and optimisation

Since the power efficiency of the on-off controlled discs has been shown to be strongly correlated to the mean disc-tip velocity, it might be argued that an analogous trend exists for the constant-velocity activation of the discs. In the present section we investigate whether it is possible to improve the power efficiency of the steady forcing of the discs by reducing their rotational

velocity. This encompasses a simpler, one-dimensional problem in the velocity W :

$$\arg \min_W \mathcal{P}_{net}. \quad (5.44)$$

The search domain is restricted to the interval $0.4 < W/W_{ss} < 1.0$. The steady optimisation is realised using the EmuKit Bayesian optimiser through the same framework already described for the on-off control. The DNS numerical set-up does not change from what described for the on-off control, with the obvious exception that the disc velocity is kept constant. The results of the simulations are listed in Tab. 5.5. First, a preliminary LHS exploration of the parameter space (cases S01-S04) is performed, followed by the optimisation itself (S05-S08). A small number of evaluations is already sufficient to conclude that the RH13 steady forcing design utilises a suboptimal value of the disc-tip velocity. For $W = 0.721W_{ss}$ we obtain an almost 2% reduction of the total spent power with respect to the original case (here called “steady” for consistency with Tab. 5.3).

Table 5.5: Results of the DNSs for the parameter space exploration (S01-S04) and the optimisation (S05-) of the steady case.

Case ID	W/W_{ss}	$\mathcal{P}_{net}(\%)$	$\Delta\mathcal{P}_p(\%)$	$\Delta\mathcal{P}_M(\%)$	\mathcal{G}
steady	1.00	± 0.87	0	0	1.22
S01	0.925	-0.77 ± 1.64	1.71	-2.48	1.32
S02	0.625	-1.58 ± 1.65	7.95	-9.53	1.84
S03	0.475	-0.30 ± 1.21	11.73	-12.03	2.09
S04	0.775	-1.49 ± 1.66	4.87	-6.35	1.54
S05	0.400	0.64 ± 1.25	13.65	-13.01	2.14
S06	0.699	-1.23 ± 1.62	6.82	-8.05	1.64
S07	0.721	-1.84 ± 1.80	5.73	-7.58	1.67
S08	0.725	-1.81 ± 1.43	5.67	-7.48	1.66

The probabilistic model and the simulation data for $\Delta\mathcal{P}$ are visualised in Fig. 5.15a. It can be noticed that the objective function is not known exactly at the evaluation points because we fixed a minimum level of variance based on the median uncertainty obtained from the post-processing of the exploratory evaluations. This gives the model an enhanced capability of robustly managing outliers such as design S06. The best guess on the optimal value of W based on the available data corresponds to the maximum of the expected improvement function i.e. about 70% of the original W_{ss} .

Fig. 5.15b proves that the pumping power component of the power budget follows the same dependence on \bar{W} as the on-off control data. This supports the view that drag reduction depends uniquely on the mean rotational velocity of the discs irrespectively of its time history, implying that any difference in the power budgets of the constant-velocity and on-off forcing is to be found in the motor power budget terms. In Fig. 5.15c the fluid torque and housing torque contribution of the constant-velocity and on-off cases are compared, establishing that the on-off control always reduces both \mathcal{P}_f and \mathcal{P}_h with respect to a constant-velocity activation. Since for the constant-velocity forcing $\dot{\mathcal{E}}_d = 0$, the reason of the reduced efficiency of the on-off forcing can be tracked down to the transient power contribution.

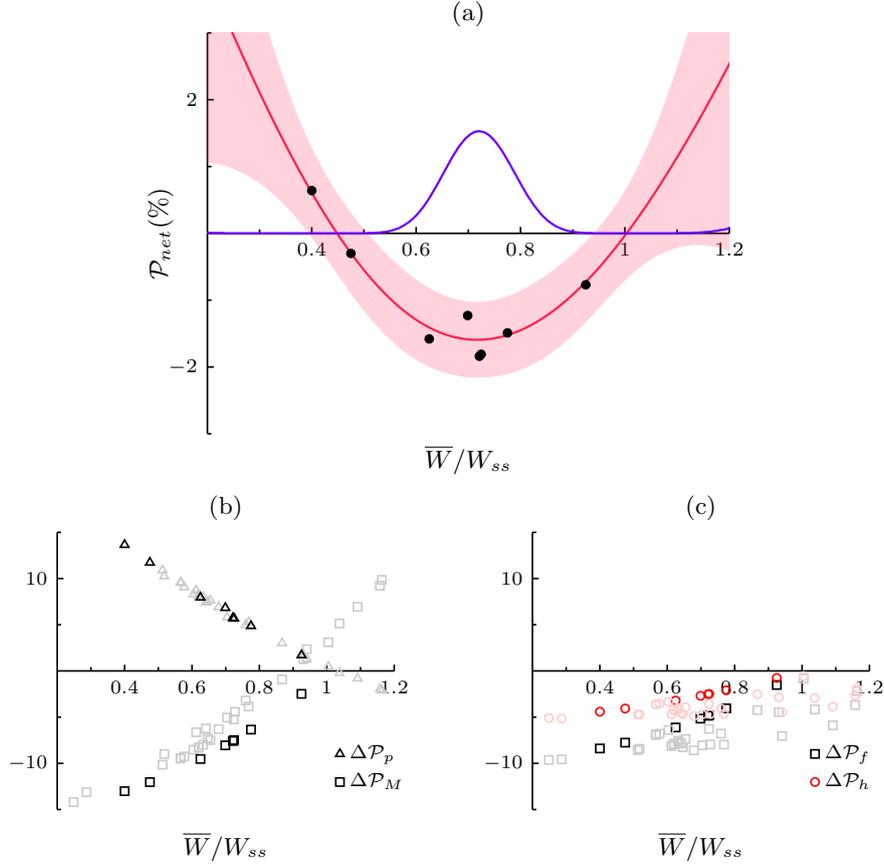


Figure 5.15: Results of the constant-velocity optimisation. (a) Visualisation of the probabilistic model of $\Delta\mathcal{P}$ for the steady optimisation. DNS evaluations (\bullet) from Tab. 5.5; expected value (---) with 2σ confidence intervals of the predictive distribution; Expected Improvement (---). (b) correlation of the pumping and motor power to the normalised mean disc velocity. The greyed out data are the on-off control. (c) main budget contributions to $\Delta\mathcal{P}_M$, correlated to the normalised mean disc velocity.

a. Comparison of the skin-friction budgets

Figure 5.16 presents a comparison of the effect of the on-off control on the skin-friction budget, compared to the steadily-rotating disc forcing. The same procedure used for the rings in Chapter 3 is applied. First, the flow is considered steady and time- and symmetry-averaged flow fields are calculated, averaging out the random on-off activation periods. This approach is justified provided that a sufficient number of forcing periods are resolved within the simulation time in order to achieve a reasonable level of confidence. Secondly, the FIK identity (3.38) is applied to the time- and symmetry-averaged flow fields, obtaining the skin-friction budget terms.

Figure 5.16 compares the streamwise-averaged FIK budget terms of the steadily rotating case S02 with that of the optimised on-off activated cases C12 and C13. For both cases the number of on-off periods resolved within the simulation time justifies the steady approach described above. The steady case S02 is chosen because its disc-tip velocity of $W = 0.625W_{ss}$ is very close to the mean velocity obtained in C12 and C13, respectively $W = 0.605W_{ss}$ and $W = 0.619W_{ss}$. No meaningful difference in the spatial structure of the skin-friction can be evinced from this data set.

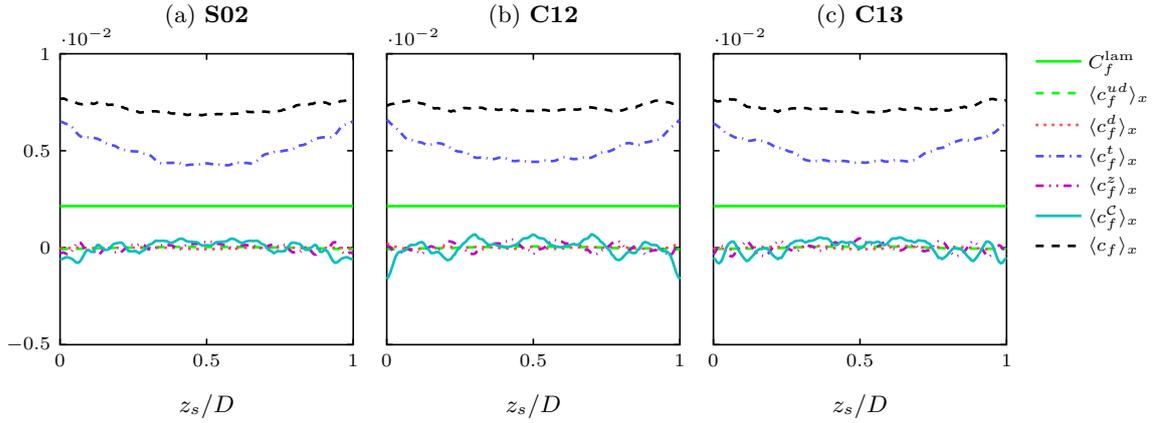


Figure 5.16: Budget of the skin-friction contributions for the on-off control. (a) steady activation. (b) on-off control, case C12. (c) on-off control, case C13.

5.5.2 Power gain and power gain optimisation

Beside the net power efficiency, another important performance indicator of an active flow control technique is the power gain, defined as $\mathcal{G} = (\mathcal{P}_0 - \mathcal{P}_p)/\mathcal{P}_M$ (Quadrio, 2011). \mathcal{G} provides a measure of how efficient the control scheme is relative to the external power needed to operate the actuators, which in this case is the disc motor power \mathcal{P}_M .

Some of the control designs evaluated during the total power efficiency optimisation display values of \mathcal{G} (Tab. 5.3) that improve up to 0.55 on the steady control. Moreover, the best cases in term of total power efficiency can be consistently placed among the best in terms of \mathcal{G} , surpassing the steady control of more than 0.3.

The optimised constant-velocity cases of Tab. 5.5 also show higher power gain than the original design. These observations are corroborated by the correlation of the power gain \mathcal{G} to \overline{W} shown in Fig. 5.17a. It can be pointed out that lower values of \overline{W} are positively correlated to a better power gain, and this is true for both the steady and the on-off controlled cases. This is explained by the fact that the steady forcing spends less motor power than an on-off control at the same mean disc velocity (Fig. 5.15b). For particularly low values of \overline{W} (i.e. C03) the power gain drops dramatically due to the drag reduction being almost zero. It can be therefore concluded that the improved power gain measured in some of the on-off control cases is an effect of the reduced disc velocity and not an intrinsic advantage of the control method.

Fig. 5.17c visually demonstrates that, for any desired value of \mathcal{R} , it is always convenient to prefer the more power-efficient steady forcing over the on-off activation. Conversely, given a certain $\Delta\mathcal{P}$, it is possible to find an on-off control that performs better than a steady one (this happens when the on-off case has a higher \overline{W}). However, the steady forcing still returns better combinations of \mathcal{R} and $\Delta\mathcal{P}$ that are not dominated by any on-off case with respect to neither performance metric.

The same conclusion remains true if we require to choose an optimal combination of \mathcal{R} and \mathcal{G} , as it is clearly seen in Fig. 5.17b where the constant-velocity cases dominate the others for any value of either variable.

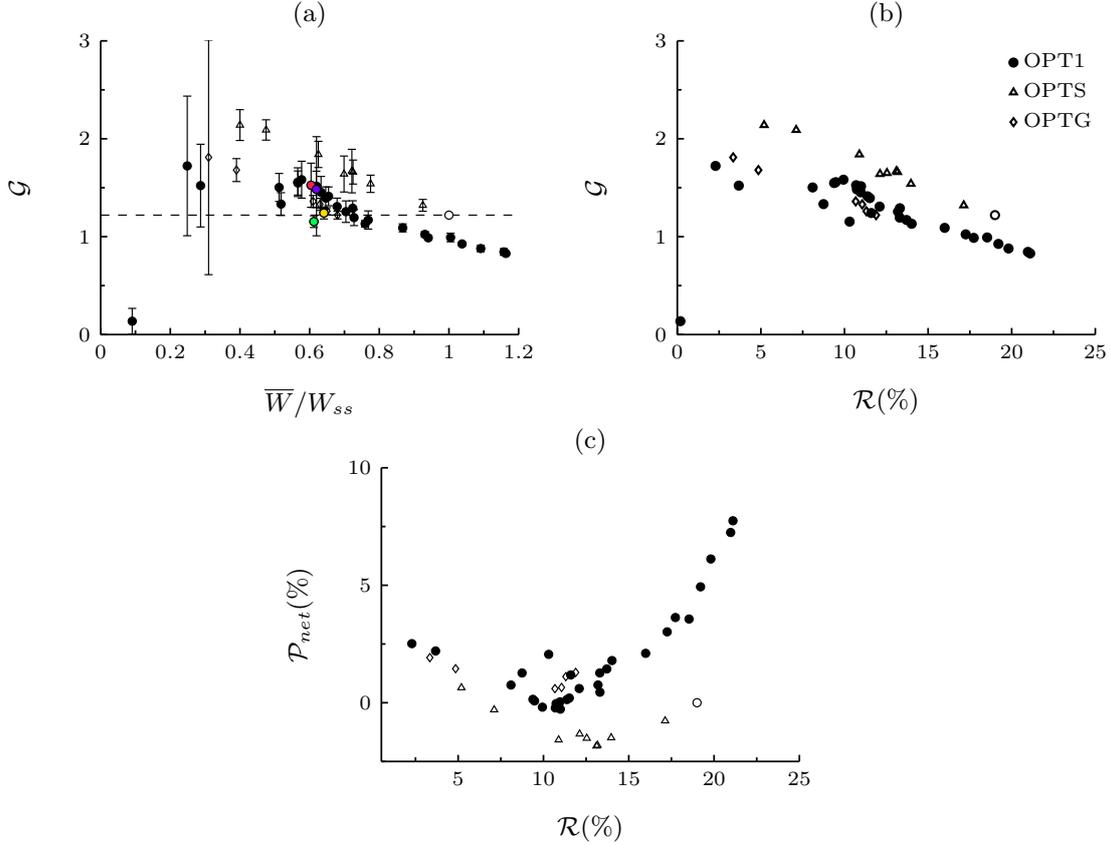


Figure 5.17: Power gain statistics of the on-off control. (a) Correlation of the power gain to the normalised mean disc velocity. The empty circle identifies the steady control. The coloured symbols are as in Fig. 5.13. (b) Correlation between the disc motor power and the drag reduction. A dashed line marks the steady forcing value of $\mathcal{G} = 1.22$.

In order to investigate whether it is possible to improve power gain of the on-off forcing, we carried out another optimisation setting \mathcal{G} as the objective to maximise. The power gain optimization problem (OPTG) is defined by:

$$\arg \max_{W_m, W_M, T_M} \mathcal{G} \quad (5.45)$$

and is subject to the constraints:

$$T_0 > T_0(W_M), \quad (5.46)$$

$$|1 - \mathcal{P}/\mathcal{P}_{ss}| < 0.01, \quad (5.47)$$

$$W_M/W_{ss} > 1, \quad (5.48)$$

$$W_m/W_{ss} < 1. \quad (5.49)$$

The constraint on the torque minimum value remains unchanged from OPT1. The second constraint guarantees that the total spent power does not depart more than 1% from the steady reference forcing. Additionally, we enforced an upper-bound constraint on the total spent power in order to prevent the optimiser to look for solutions which spend more than 1% more power

than the steady control. The constraint value is not known a priori from the control parameters and therefore needs to be evaluated from the simulations and incorporated into the acquisition function. The optimisation is carried out using the same BO framework as OPT1, with the exception that, in order to handle the unknown constraint, a modified acquisition function and the use of two probabilistic models (one for P and one for \mathcal{G}) is necessary. More details on BO with unknown constraints can be found in Appendix 5.A. In this case, no randomised exploration of the parameter space is performed because all the simulations from OPT1 are incorporated into the the first iteration of the GP regression model for OPTG.

The results of the power-gain optimisation (OPTG) are listed in Tab. 5.6. Although a few cases yield a good power gain (C02, C03), they have poor power efficiency because of their very low disc velocity that leads to a larger drag.

Table 5.6: Results of the simulations for the power-gain optimisation OPT-G.

Case ID	W_M/W_{ss}	W_m/W_{ss}	T_0/T_{ss}	$\mathcal{P}_{net}(\%)$	$\Delta\mathcal{P}_p(\%)$	$\Delta\mathcal{P}_M(\%)$	\mathcal{G}	\bar{w}/W_{ss}
G01	1.50	0.03	1.05	113.61	78.02	35.59	0.62	1.38
G02	0.31	0.28	4.98	98.51	96.66	1.85	1.81	0.31
G03	0.40	0.37	4.26	98.04	95.15	2.89	1.68	0.39
G04	1.11	0.31	3.77	97.70	88.69	9.01	1.26	0.65
G05	1.05	0.29	4.06	97.20	89.33	7.87	1.36	0.61
G06	1.08	0.30	3.85	97.24	88.94	8.29	1.33	0.63
G07	1.14	0.34	3.70	97.88	88.11	9.77	1.22	0.68

5.6 Summary and conclusions

This study has investigated a simple proportional feedback control for wall-mounted rotating discs in channel flow at low friction Reynolds number. The scheme consists in switching on and off the motor when a certain upper or lower threshold of the velocity is reached, realising acceleration and deceleration transients of the disc velocity. Models for the electric motor that powers the discs and the frictional losses in a realistic housing were implemented in the numerical model. The same models were also tested to replicate the steady open-loop control of Ricco and Hahn (2013), and compare its performance to the original study and to the closed-loop control. A Bayesian Optimisation framework has been implemented to perform an autonomous search of the optimal values of the control parameters of the control schemes with respect to the total power efficiency, with the goal to improve over the reference steady open-loop control of Ricco and Hahn (2013). For the closed-loop control, and optimisation of the power gain has also been run. The key conclusions can be summarised in the following points:

- i)* the Bayesian optimisation method has been successfully applied to the optimisation of the total power efficiency of the controlled flow, that consists in searching the three-dimensional parameter space of the two velocity thresholds and the torque magnitude. A fully parametric study would have required at least over a hundred evaluations instead of the around 30 simulations needed for the main optimisation to converge to a minimum with reasonable confidence. For the open-loop control, an optimal value of the disc-tip velocity could be located after performing 8 simulations, although a very narrow search domain was used. Only partial success has been achieved in case of the power-gain optimisation, likely because of the higher complexity of the optimisation in presence of constraints.
- ii)* no on-off control design is found that is more power-efficient than the reference steady forcing. The best guesses on the minimum correspond to values of the total spent power

that are indistinguishable from the steady forcing case, and the power budget of the more efficient cases is always characterised by higher (about +8%) skin-friction contributions and lower (about -8%) motor-power contributions compared to the steady forcing. In the more efficient parameter combinations, the control resembles closely a steady forcing, with high-frequency on-off cycles that results in small-amplitude fluctuations of the disc rotation velocity around a mean of around 60% of the reference steady forcing.

- iii) a more efficient constant-velocity activation of the discs than the reference was also found by the Bayesian optimiser. The optimal case corresponds to reducing the disc-tip velocity to around 70% of the reference steady forcing, a value closer to the optimal mean disc-tip velocity observed in the on-off forcing. The smaller efficiency of the on-off control is explained by the power needed to accelerate the discs that is then dissipated in the fluid during the deceleration transients.
- iv) the steady forcing and the on-off control are equivalent in terms of how well they reduce the skin friction (i.e. the pumping power). The drag reduction was shown to be perfectly correlated to the mean disc-tip velocity, regardless of whether the activation is the on-off or the steady type.
- v) the power gain \mathcal{G} i.e. the ratio of the drag reduction to the motor power needed to operate the discs the performance of the on-off control has also been assessed. The best cases display values of \mathcal{G} more than 30% larger than the steady activation while also maintaining a similar power efficiency. However, a better power gain can always be achieved by operating the discs at constant speed.
- vi) the disc power budget terms are shown to be reasonably well predicted by simple exact models, derived from simple dynamical models of the disc velocity transient. These models might be useful to unfold the physical meaning of the dependence of the disc power budgets on the control parameters and construct an approximate predictive model of the on-off control performance.
- vii) the motor and actuator models were constructed from a number of assumptions about the constructive and fluid dynamical characteristic of the disc and its housing and grounded in existing experimental work on analogous actuators. The use of models for the motor and actuator inherently introduces additional model uncertainty, however the precise modelling of active actuators is crucial in evaluating the power efficiency of such active forcing techniques because the variation of the non-ideal power losses in the actuator can be of a similar magnitude to the variation of the flow pumping power and can thus decide the viability of certain parameter combinations. For the reference steady forcing, a net power saving of around 3% can be achieved with respect to the no-slip channel flow with the realistic actuator model.

a. Sensitiveness of the power performance to the choice actuator model

It might be argued (Mahfoze et al., 2018) that the power efficiency of controlled flow is sensitive to the choice of the actuator model, and the absence thereof, or the choice of a different model, would lead to convergence to a different control-design optimum. In the closed-loop steady forcing case, we performed a simulation with the ideal actuator (open-loop, no frictional losses) at the new optimal rotation velocity $0.7W_{ss}$ of found by OPTS, finding a larger net power saving than the original reference case of Ricco and Hahn (2013). This is expected as both \mathcal{P}_p and

\mathcal{P}_f i.e. the only non-zero power-budget terms when using ideal actuators, are monotonic and perfectly correlated functions of the mean disc velocity.

For the feedback-controlled flow, an counterfactual in the ideal case is somewhat harder to obtain, because the disc-velocity time history depends on the parameters of the actuator model. For example, an optimisation could be performed using a prescribed velocity transient (for example a linear or a sharp transient with only W_m and W_M as control parameters) without the use of the disc dynamical model. The outcome of such optimisation could be compared to the fully modelled results.

Appendix to Chapter 5

5.A Bayesian Optimization

This appendix contains more technical details on some aspects of the Bayesian Optimisation algorithm. The first paragraph clarifies the Gaussian Process (GP) regression model and its properties. The second paragraph discusses in more detail the Expected Improvement and the third one discusses how Bayesian Optimisation is adapted to allow the use of constraints that are computationally expensive functions of the parameters.

a. Gaussian process regression

This section expands on the Gaussian Process regression modelling. We give here more precise definition of a Gaussian Process model and briefly illustrate the derivation of the predictive distribution following [Williams and Rasmussen, 2006](#). A Gaussian process $u(\mathbf{x})$ is a stochastic process for which the joint pdf of any finite sample is a normal distribution. Let us consider d variables, representing the values of the process sampled at d locations i.e. $\{u_i = u(x_i)\}_{i=1}^d$. Their joint normal distribution is defined as:

$$p(u_1, u_2, \dots, u_d) = \frac{1}{\sqrt{(2\pi)^d |K|}} \exp\left(-\frac{1}{2}(\mathbf{u} - \mathbf{m})K^{-1}(\mathbf{u} - \mathbf{m})\right) \equiv \mathcal{N}(\mathbf{m}, K), \quad (5.50)$$

where \mathbf{m} is the mean vector and K is the covariance matrix:

$$K = \begin{pmatrix} \sigma_{11}^2 & \text{cov}_{12} & \cdot & \cdot \\ \text{cov}_{12} & \sigma_{22}^2 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{pmatrix}. \quad (5.51)$$

The distribution [5.50](#) is therefore characterised by its mean vector and its covariance matrix. The entries of the covariance matrix can be thought of as the values of a function $\text{cov}(u_i, u_j)$ that given any couple of variables as input returns their covariance cov_{ij} . For a stochastic process $u(\mathbf{x})$, the covariance function is specified as a function of the underlying index variable \mathbf{x} such that:

$$\text{cov}(u(\mathbf{x}_i), u(\mathbf{x}_j)) = k(\mathbf{x}_i, \mathbf{x}_j). \quad (5.52)$$

Consequently, a Gaussian process is completely characterised by its mean vector $\mathbf{m}(\mathbf{x})$ and its covariance function $k(\mathbf{x}_i, \mathbf{x}_j)$.

In a GP regression problem, we are interested in the joint distribution of the training input vector $f(X)$ and test inputs vector $f^*(X^*)$, respectively defined as:

$$f(X) \equiv \{f(\mathbf{x}_1), f(\mathbf{x}_2), \dots\} \quad (5.53)$$

$$f^*(X^*) \equiv \{f^*(\mathbf{x}_1^*), f^*(\mathbf{x}_2^*), \dots\}. \quad (5.54)$$

According to the definition [5.50](#), their joint distribution reads:

$$p(f(X), f^*(X^*)) = \mathcal{N}\left(\mathbf{m}, \begin{pmatrix} K(X, X) & K(X, X^*) \\ K(X^*, X) & K(X^*, X^*) \end{pmatrix}\right), \quad (5.55)$$

where the correlation (sub)matrices are defined using the definition of correlation function. For example:

$$K(X, X^*) = \begin{pmatrix} k(\mathbf{x}_1, \mathbf{x}_1^*) & k(\mathbf{x}_1, \mathbf{x}_2^*) & \cdot & \cdot \\ k(\mathbf{x}_1, \mathbf{x}_2^*) & k(\mathbf{x}_2, \mathbf{x}_2^*) & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{pmatrix}. \quad (5.56)$$

The regression problem consists in calculating the probability distribution of f^* at the test inputs X^* , the so-called predictive distribution. Given the joint distribution, the predictive distribution is found by conditioning the joint distribution on the training inputs and the desired test input. For example, the predictive distribution at the single test input \mathbf{x}^* is $p(f^*|\mathbf{x}^*, f(X), X)$. Since the joint distribution is normal, it is possible to calculate the analytical solution shown in Eq. 5.38.

b. Acquisition function

In this section we give a more detailed background on the acquisition function, focusing on its useful properties in the context of optimisation under uncertainty. We recall the definition of the utility function:

$$u(\mathbf{x}) = \begin{cases} f_{min} - f(\mathbf{x}) & \text{if } f(\mathbf{x}) < f_{min} \\ 0 & \text{otherwise} \end{cases} \quad (5.57)$$

The Expected Improvement is defined as the expected value of the above utility function as:

$$a_{EI}(\mathbf{x}) \equiv E[u(\mathbf{x})|\mathcal{D}] = \int_{-\infty}^{f_{min}} (f_{min} - f)p(f^*|\mathbf{x}^*, \mathcal{D}) df \quad (5.58)$$

The point $\mathbf{x}_{n+1} = \arg \max a_{EI}(\mathbf{x})$ has the highest probability of finding a lower estimate of the global minimum and is therefore selected for the next evaluation. Substituting the closed-form expression of the predictive distribution (5.38) into (5.58), the EI acquisition takes the exact form:

$$a_{EI}(\mathbf{x}) = \underbrace{(f_{min} - \bar{f}^*(\mathbf{x}))\Phi(f_{min}, \sigma^*(\mathbf{x}^*))}_{a_1} + \underbrace{\sigma^*(\mathbf{x}^*)\mathcal{N}(f_{min}, \sigma^*(\mathbf{x}^*))}_{a_2}, \quad (5.59)$$

where Φ is the normal cumulative probability distribution. The term a_1 in (5.59) is the exploitative term, being proportional to the distance between the model mean and the minimum from the data. a_2 is the exploratory term, being proportional to the model variance. Optimising the exact acquisition function (5.59) is computationally inexpensive and can be done through efficient gradient-based algorithms such as L-BFGS.

c. Optimisation with unknown constraints

The situation addressed in this section arises when an optimisation problem is subjected to a constraint which is not known for any parameter combination but has to be evaluated similarly to the objective function. We do not know with certainty which regions of the parameter space contain infeasible designs (and can thus be avoided) before evaluating an expensive function. An unknown constraint can be defined in general by the inequality $\mathcal{C}(\mathbf{x}) < 0$, where $\mathcal{C}(\mathbf{x})$ is a function that requires an expensive evaluation at \mathbf{x} . Unknown constraints can be handled through a modified Expected Improvement acquisition following Gelbart et al., 2014. If the constraint and the objective are assumed to be independent, then we may define the Constrained Expected Improvement as:

$$a_{CEI}(\mathbf{x}) = a_{EI}(\mathbf{x})\Pr[\mathcal{C}(\mathbf{x}) < 0], \quad (5.60)$$

where $a_{EI}(\mathbf{x})$ is the Expected Improvement on the objective and $\Pr[\mathcal{C}(\mathbf{x}) < 0]$ is the probability of feasibility (PoF) i.e. the probability of satisfying the constraint. The PoF is computed by

building a second probabilistic model for \mathcal{C} (another GP regression) and evaluating the following integral:

$$\Pr[\mathcal{C}(\mathbf{x}) < 0 | \mathcal{D}_{\mathcal{C}}] = \int_{-\infty}^0 p(\mathcal{C}^* | \mathbf{x}^*, \mathcal{D}_{\mathcal{C}}) d\mathcal{C}. \quad (5.61)$$

where $\mathcal{D}_{\mathcal{C}}$ are the data on $\mathcal{C}(\mathbf{x})$.

d. Constrained optimization of a test function

This section illustrates an example where EmuKit's Bayesian optimization is applied to finding the global minimum of a multivariate test function with an unknown constraint, as done in [Gelbart et al. \(2014\)](#). The goal is to validate EmuKit and our optimisation framework using an exact function that is inexpensive to compute. The only difference with the research problem above is that the channel-flow DNS takes the place of the test function. The Branin-Hoo (BH) test function function is considered, that is defined as:

$$BH(x, y) = a(y - bx^2 + cx - r) + s(1 - t) \cos x + s \quad (5.62)$$

with $a = 1$, $b = 5.1/(4\pi^2)$, $c = 5/\pi$, $r = 6$, $s = 10$ and $t = 1/(8\pi)$. This function has three global minima, one of which is marked by hearts in [Fig. 5.18](#). Convergence to at least one of these minima will be considered a successful outcome of the optimization procedure. We also add a disc constraint function in the form:

$$\mathcal{C}(x, y) = \begin{cases} -1 & \text{if } (x - 2.5)^2 + (y - 7.5)^2 \leq 49 \\ 0 & \text{otherwise.} \end{cases} \quad (5.63)$$

The constraint function is unknown to the optimizer and has to be evaluated at each iteration together with the objective function, as explained in previous section. [Fig. 5.18](#) shows the results of a test optimization. The algorithm is run for 25 iterations, starting from an initial set of 8 space-filling (Latin design) evaluations.

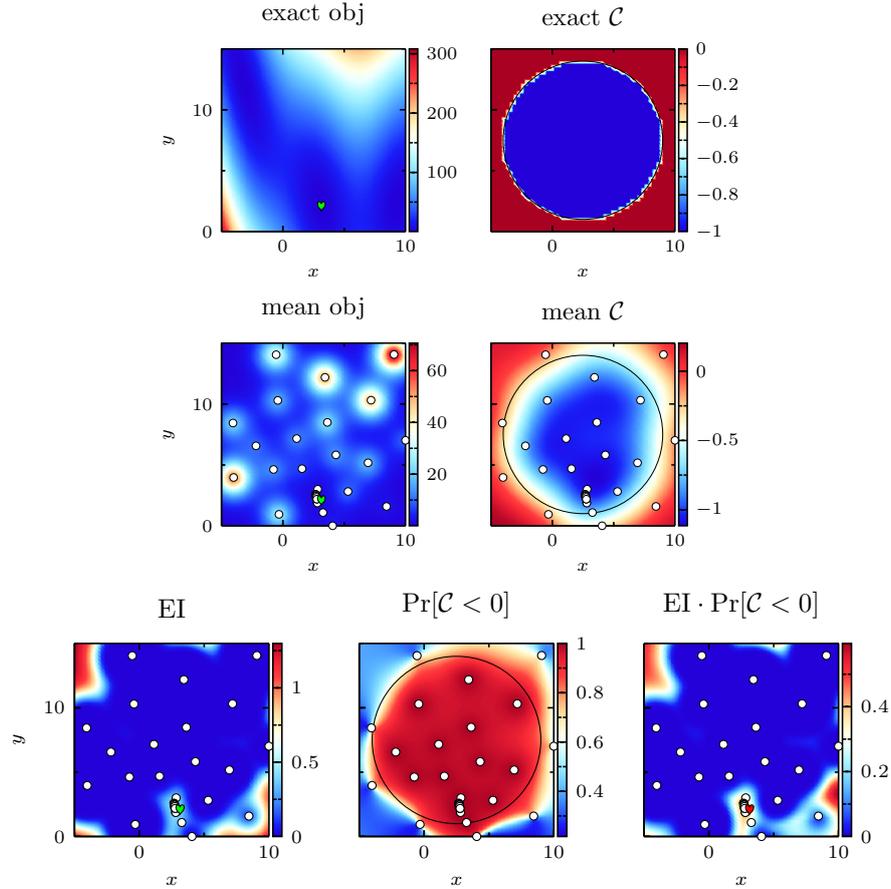


Figure 5.18: Constrained Optimization of the Branin-Hoo test function using EmuKit's Bayesian optimisation tools.

5.B Exact model of the transients

In order to derive an exact solution for the switch-on and the switch-off transients, we assume that the fluid stresses on the flow-facing surface of the disc follow the same linear Couette solution that holds inside the cavity.

a. Switch-on transient

The assumption leads to the linear disc dynamics during the switch-on transient being described by the following problem:

$$\begin{cases} \frac{dW_{on}}{dt} &= -\frac{2C_h}{I}W(t) + \frac{D}{2I}T_M \\ W_{on}(0) &= W_m \end{cases} \quad (5.64)$$

whose solution reads:

$$W_{on}(t) = \left(W_m - \frac{T_M D}{4C_h} \right) \exp\left(-\frac{2C_h}{I}t\right) + \frac{T_M D}{4C_h} \quad (5.65)$$

The switch-off time is found from the condition $W_{on}(t_{off}) = W_M$ which gives:

$$t_{off} = -\frac{I}{2C_h} \log \left(\frac{W_M - T_M D/4C_h}{W_m - T_M D/4C_h} \right). \quad (5.66)$$

b. Switch-off transient

Using the same set of assumptions, we can model the switch-off transient dynamics as:

$$\begin{cases} \frac{dW_{off}}{dt} &= -\frac{2C_h}{I} W(t) \\ W_{off}(0) &= W_M \end{cases} \quad (5.67)$$

The switch-on time is found from the switch-off condition $W_{off}(t_{on}) = W_m$ and it reads:

$$t_{on} = -\frac{I}{2C_h} \log \left(\frac{W_m}{W_M} \right) \quad (5.68)$$

c. Mean disc-tip velocity

From the solutions $W_{on}(t)$ and $W_{off}(t)$ the exact mean velocity \bar{W} can be calculated by averaging over the union of the two transients:

$$\bar{W} = \frac{(\bar{W}_{on} t_{off} + \bar{W}_{off} t_{on})}{t_{off} + t_{on}}, \quad (5.69)$$

where:

$$\bar{W}_{on} = \frac{1}{t_{off}} \int_0^{t_{off}} W_{on}(t) dt \quad \text{and} \quad \bar{W}_{off} = \frac{1}{t_{on}} \int_0^{t_{on}} W_{off}(t) dt \quad (5.70)$$

Proceeding from the above and substituting (5.65),(5.68),(5.66) it can be shown that:

$$\bar{W} = \gamma T_M \log \left(\frac{W_M - \gamma T_M}{W_m - \gamma T_M} \right) \log \left(\frac{W_m}{W_M} \frac{W_M - \gamma T_M}{W_m - \gamma T_M} \right)^{-1}, \quad (5.71)$$

where $\gamma = D/4C_h$. In a similar way, the mean value of any other quantity f over the exact period can be calculated as:

$$\bar{f} = \frac{(\bar{f}_{on} t_{off} + \bar{f}_{off} t_{on})}{t_{off} + t_{on}}. \quad (5.72)$$

Chapter 6

Conclusions and future directions

In this chapter we summarise the main outcomes of the dissertation and we recommend a few directions along which the work could be extended in the future, building on the findings reported in the previous Chapters.

6.1 General conclusions

This work represents a contribution to the study of rotating actuators as a practical device for skin-friction reduction in turbulent wall-bounded flow. On a fundamental level, we emphasised that rotating actuators are a practical and energy-efficient method to deliver two distinct drag-reduction paradigms: first, the generation of spanwise velocity waves by steadily rotating full discs or rings; second, the generation of streamwise slip through half discs. Our results endorse the idea that rotating rings or discs are an excellent practical way of generating a working standing wave of spanwise velocity, although with the side effect of generating detrimental secondary flow structures. The relative simplicity of the rotating disc and its scalability renders it superior in many aspects to other wave-generating techniques. The evidence collected in this work upholds the viability of rotating actuators for drag reduction either under active open-loop control, passive control or closed-loop feedforward control strategies. The suitability of open-loop, steadily rotating rings to be combined with other wall-based control methods was demonstrated in Chapter 3, resulting in methods that, albeit more complex, partially offset the shortcomings of each individual method and thus are not only better performing but also hold the potential to be more reliable in realistic applications.

The research that forms the second part of the dissertation, Chapters 4 and 5, was undertaken in the broad scope of addressing the energy efficiency of rotating devices. Realistic modelling of an entire reasonable rotating-disc apparatus, complete with a motor model and mechanical losses, was carefully performed for the first time and found to lead to encouraging results in terms of the energy efficiency of the steadily rotating case. Numerical experiments into passively and closed-loop controlled disc in Chapters 4 and 5 were performed to test more efficient activating strategies of rotating discs. Our most promising contribution to the passive skin-friction control is the half-disc array concept, that, inspired by already available experimental data, was proven to be an effective strategy that deserves deeper investigation. Regarding closed-loop control, a simple, pulsed regulation of the disc velocity was studied, with the characteristics of the on-off cycle being optimised to achieve the best energy efficiency. Although ultimately lacking optimal characteristics of efficiency compared to the steady activation, the work constitutes a first step for further developments in the feedback control of rotating discs and in the study of the transient

physics than may be leveraged for later developments.

6.2 Future outlook

In light of the findings reported in this dissertation, a number of directions can be envisioned for the future development of wall-mounted rotating discs, some of which are proposed in the following sections, tentatively listed from the more incremental to the more speculative.

a. Better understanding of the flow physics

A deeper understanding of the physics of steadily rotating discs under turbulent flow remains a critical objective. More insight would be beneficial into the mean-flow structures described in Chapter 3, such as rolls and ejections, and their dependence on the disc size, velocity and from the Reynolds number. Such developments would be important for obtaining more accurate drag-reduction performance predictions, especially at higher Reynolds number, and for exploring the application of the disc-based control to different flow scenarios, such as the separated flow described below.

Another topic worthy of further investigation is the physics of freely rotating discs. It can be speculated that using lighter discs and suitable numerical methods, stronger interactions with near-wall turbulence would occur. The usefulness of such passive devices to act as sensors – for e.g. shear-stresses – is also worth addressing.

b. Passive control

Even though small-scale, fast, sophisticated active actuators are technologically feasible (Keefe, 1997; Kasagi et al., 2009), a general trend in the field of skin-friction reduction research is, wherever possible, to gravitate towards passive forms of flow control, owing to their lower degree of operational complexity and cost effectiveness.

In this context, we believe that the research undertaken in Chapter 4 on passive half-discs deserves special attention for future development, especially due to its capability to generate both streamwise slip and a spanwise velocity pattern simultaneously. This may prove a precious upside, because the wall-wave forcing, which has limited flexibility due to the fixed actuator diameter, is complemented by the slip velocity generation that instead is able to adjust to off-nominal conditions. This can prove very helpful in adapting to changing flow regimes.

An outline of a series of numerical experiments at increasing velocities is given in Section 4.7, which would be important to paint a clearer picture of the efficacy of the method. Furthermore, an example of a possible alternative design is to use a staggered array configuration such as that depicted in Fig. 6.1B could be used in place of the regular array. The goal is to create an appropriate alternate spanwise motion near the wall and therefore combine the slip-length-like motion of the discs with the spanwise-velocity-wave paradigm through a purely passive technique.

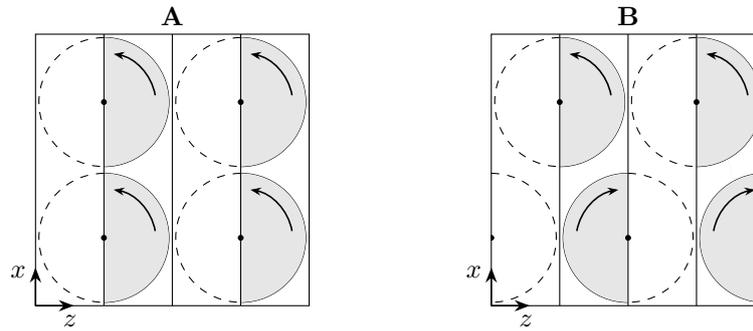


Figure 6.1: Disc mount arrangements type A (regular) and B (staggered).

c. Adaptive control

As it has been remarked in Chapter 3, perhaps the most promising direction is to imagine rotating discs as integrated into an adaptive flow-control framework, either in combination with other drag-reduction techniques – such as in Chapter 3 – or where the actuators can be leveraged to control other flow phenomena beyond pure skin-friction. It could also be interesting to elaborate on a closed-loop scheme where part of the devices are used for sensing and part is being used for actuation. The behaviour of passive half-discs could also be deemed as adaptive if one considers that it combines two effects, a spanwise motion and a streamwise slip, whose response to an increasing mean shear is different and could potentially result in an enhanced robustness of the scheme.

d. Separation control

Rotating discs may also prove a great tool to manipulate flow phenomena other than flat-plate skin-friction. An interesting perspective is to inquire whether the induced vortical motions can be applied to separation control, turning what is a drawback in wall-turbulence control into a convenient feature. The generation of streamwise-aligned, persistent vortices near the separation region of a stalling airfoil is a proven method to delay the onset of stall at high angles of attack (Lin et al., 1994). Such vortices promote the transport of streamwise momentum to the vicinity of the wall, flattening the mean velocity profile and thereby mitigating the effect of the adverse pressure gradient that occurs on the suction side of a stalling airfoil.

Both passive (Godard and Stanislas, 2006a) and active vortex generators have been investigated. Active VGs, typically in the form of skewed and pitched (Godard and Stanislas, 2006b) or swirling synthetic jets (Munday and Taira, 2018), have the additional advantages that they can be turned off with no parasitic drag and they can be controlled actively, adapting to a changing flow regime.

Given that rotating discs are wall-normal vorticity-generating devices, one may want to assess *i*) whether they are capable of generating appropriate streamwise-aligned vortices in the first place; *ii*) their attractiveness as an alternative to more conventional active devices such as jets.

Previous research has established – e.g. Section 3.3.1, Figure 3.3, and Wise and Ricco (2014) – that the interaction between the disc rotation and the mean flow results in the generation of streamwise vortex-like structures. Interestingly, the vortices' location and scale appear to be approximately matching those detected in experimental PIV measurements (Godard and Stanislas, 2006a) relative to the boundary-layer thickness, and they also display some degree of downstream persistence.

Based on these observations, Figure 6.2 sketches a minimal arrangement consisting of pairwise counter-rotating discs, leading to the establishment of counter-rotating vortex pairs. The dependence on the disc size and velocity of quantities crucial to successful separation control, such as vortex intensity, downstream persistence and wall-normal penetration, is currently an open issue.

Discs share the same upsides of other active VGs, while also offering the capability of being operated as drag-reduction control devices in the fully attached BL regime. The performance and the efficiency of disc VGs relative to conventional jets are, at present, a speculative subject.

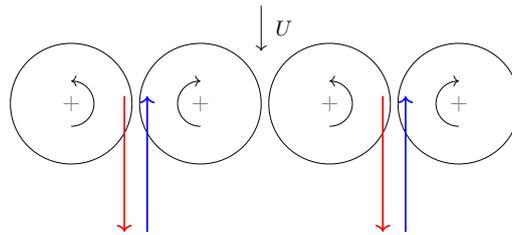


Figure 6.2: Hypothetical arrangement of wall-mounted rotating discs for counter-rotating vortex pair generation. The blue and red arrows represent the sign and the approximate location of the resulting vortex cores.

e. Closed-loop control

Hypothetically, more energy-efficient ways of delivering the same paradigms could be elaborated. While a rough attempt at this is represented by the on-off control from Chapter 5, more elaborate alternatives exist that are succinctly reviewed below.

The on-off control used in Chapter 5 is a basic form of active control that attempts to enhance the energy efficiency the simple open-loop forcing. Other possibilities exist in terms of control algorithms, mostly methods from linear control theory that are known to have been successfully applied to the reduction of skin-friction in turbulent flow (Brunton and Noack, 2015). Two possible general directions can be advanced for closed-loop control of rotating discs: first, maintaining a feedforward approach and implementing more advanced control algorithms; second, in a broader scope, pursuing a full feedback approach with respect to the skin-friction or some surrogate variable.

Proportional-Integral-Derivative (PID) control is a feedback algorithm where the control signal is a linear combination of the control objective, its finite-time integral and its derivative. PID controllers (usually free from the derivative term) are the most common control algorithms in industrial applications by an overwhelming margin (Åström and Murray, 2008). Proportional-integral control has also been successfully applied to wall-turbulence control via suction and blowing (Kim and Choi, 2017). PID control is an attractive alternative for feedforward skin-friction control via regulation of the disc velocity and represents an attractive evolution on the vanilla on-off control proposed in Chapter 5. To this purpose, a formulation of PID control for the disc velocity would entail a constraint on the mean velocity and parameter tuning would evolve according to the same energy efficiency optimisation criteria used in Chapter 5.

A Linear Quadratic Regulator (LQR) is the optimal controller for a linear system with respect to a cost functional that is defined as the integral of a quadratic function of the system's and the controller's state over a certain time horizon. In practical applications, a Kalman filter is used to approximate the full system state from a finite number of sensors. LQR has been applied to drag reduction via suction/blowing coupled with reduced-order modelling (Lee et al., 2001).

State-based control in general deals with systems driven by linear (or linearised) governing equations. Concerning a feedforward approach on the disc rotational dynamics, the disc governing equation, which is non-linear because the fluid torque depends on the disc velocity, would require linearisation. In this respect, the data discussed in Section f. are encouraging in that they show that the fully coupled non-linear transients are typically well correlated with the linear prediction. If full feedback control from the wall-shear stresses is desired, the flow equations need to be linearised and a suitable low-order representation identified to avoid intractable calculations. Many powerful tools can be drawn from the rich literature on reduced-order modelling of fluid flow, comprising analytical and data-driven techniques (Lee et al., 1997; Taira et al., 2017; Brunton and Kutz, 2019).

On a somewhat speculative note, an attractive property of state-based control is that it would open up the possibility of discovering new control principles for the rotating-disc array, analogously to what has been achieved for wall-normal suction/blowing (Kim, 2003). This is conditional on the existence of a control paradigm based on the wall-normal vorticity in the first place. This approach would see the application of feedback control to the entire actuator array, as opposed to Chapter 5's approach of using independent actuators. The idea of using part of the discs as sensors could also be integrated into this approach.

Bibliography

- Adrian, R. (2007). Hairpin vortex organization in wall turbulence. *Phys. Fluids*, 19(4):041301.
- Aghdam, S. K. and Ricco, P. (2016). Laminar and turbulent flows over hydrophobic surfaces with shear-dependent slip length. *Phys. Fluids*, 28(3):035109.
- Agostini, L. and Leschziner, M. A. (2014). On the influence of outer large-scale structures on near-wall turbulence in channel flow. *Phys. Fluids*, 26(7):075107.
- Aljallis, E., Sarshar, M. A., Datla, R., Sikka, V., Jones, A., and Choi, C.-H. (2013). Experimental study of skin friction drag reduction on superhydrophobic flat plates in high reynolds number boundary layer flow. *Phys. Fluids*, 25(2):025103.
- Åström, K. J. and Murray, R. M. (2008). *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton University Press.
- Auteri, F., Baron, A., Belan, M., Campanardi, G., and Quadrio, M. (2010). Experimental assessment of turbulent drag reduction by traveling waves in a turbulent pipe flow. *Phys. Fluids*, 22(115103/14).
- Barenghi, C. F. and Jones, C. A. (1989). Modulated Taylor–Couette flow. *J. Fluid Mech.*, 208:127–160.
- Batchelor, G. (1967). *An Introduction to Fluid Dynamics*. Cambridge University Press.
- Bertelá, M. and Gori, F. (1982). Laminar flow in a cylindrical container with a rotating cover. *J. Fluids Eng.*, 104(1):31–39.
- Bewley, T. (2009). A fundamental limit on the balance of power in a transpiration-controlled channel flow. *J. Fluid Mech.*
- Bewley, T. R., Moin, P., and Temam, R. (2001). DNS-based predictive control of turbulence: An optimal benchmark for feedback algorithms. *J. Fluid Mech.*, 447:179–225.
- Bidkar, R., Leblanc, L., Kulkarni, A., Bahadur, V., Ceccio, S., and Perlin, M. (2014). Skin-friction drag reduction in the turbulent regime using random-textured hydrophobic surfaces. *Phys. Fluids*, 26(8):085108.
- Bird, J., Santer, M., and Morrison, J. F. (2018). Experimental control of turbulent boundary layers with in-plane travelling waves. *Flow Turbul. Combust.*, 100(4):1015–1035.
- Blesbois, O., Chernyshenko, S., Toubert, E., and Leschziner, M. (2013). Pattern prediction by linear analysis of turbulent flow with drag reduction by wall oscillation. *J. Fluid Mech.*, 724:607–641.

- Brunton, S. L. and Kutz, J. N. (2019). *Data-Driven Science and Engineering*. Cambridge University Press.
- Brunton, S. L. and Noack, B. R. (2015). Closed-loop turbulence control: Progress and challenges. *Appl. Mech. Rev.*, 67(5).
- Bull, A. D. (2011). Convergence rates of efficient global optimization algorithms. *Journal of Machine Learning Research*, 12(Oct):2879–2904.
- Busse, A. and Sandham, N. (2012). Influence of an anisotropic slip-length boundary condition on turbulent channel flow. *Phys. Fluids*, 24(5):055111.
- Cain, A., Ferziger, J., and Reynolds, W. (1984). Discrete orthogonal function expansions for non-uniform grids using the fast Fourier transform. *J. Comp. Phys.*, 56(2):272–286.
- Cassie, A. B. D. and Baxter, S. (1944). Wettability of porous surfaces. *Transactions of the Faraday Society*, 40:546.
- Chan, T. F., Golub, G. H., and LeVeque, R. J. (1983). Algorithms for computing the sample variance: analysis and recommendations. *Am. Stat.*, 37(3):242–247.
- Choi, C. and Kim, C. (2006). Large slip of aqueous liquid flow over a nanoengineered superhydrophobic surface. *Phys. Rev. Lett.*, 96(066001).
- Choi, H., Moin, P., and Kim, J. (1994). Active turbulence control for drag reduction in wall-bounded flows. *J. Fluid Mech.*, 262:75–110.
- Choi, K.-S. (1989). Near-wall structure of a turbulent boundary layer with riblets. *J. Fluid Mech.*, 208:417–458.
- Choi, K.-S. (2002). Near-wall structure of turbulent boundary layer with spanwise-wall oscillation. *Phys. Fluids*, 14(7):2530.
- Choi, K.-S. and Graham, M. (1998). Drag reduction of turbulent pipe flows by circular-wall oscillation. *Phys. Fluids*, 10(1):7–9.
- Chorin, A. (1968). Numerical solution of the Navier-Stokes equations. *Math. comput.*, 22(104):745–762.
- Chorin, A. J. and Hald, O. H. (2013). *Stochastic Tools in Mathematics and Science*, volume 58 of *Texts in Applied Mathematics*. Springer New York, New York, NY.
- Chung, Y. M. and Talha, T. (2011). Effectiveness of active flow control for turbulent skin friction drag reduction. *Phys. Fluids*, 23(2):025102.
- Coles, D. (1956). The law of the wake in the turbulent boundary layer. *J. Fluid Mech.*, 1(02):191–226.
- Daily, J. W. and Nece, R. E. (1960). Chamber dimension effects on induced flow and frictional resistance of enclosed rotating disks. *J. Basic Eng.*, 82(1):217–230.
- Daniello, R. J., Waterhouse, N. E., and Rothstein, J. P. (2009). Drag reduction in turbulent flows over superhydrophobic surfaces. *Physics of Fluids*, 21(8):085103.
- de Giovanetti, M., Hwang, Y., and Choi, H. (2016). Skin-friction generation by attached eddies in turbulent channel flow. *J. Fluid Mech.*, 808:511–538.

- Dean, B. and Bhushan, B. (2010). Shark-skin surfaces for fluid-drag reduction in turbulent flow: a review. *Philos. Trans. R. Soc. A: Math. Phys. Eng. Sci.*, 368(1929):4775–4806.
- Dean, R. (1978). Reynolds number dependence of the skin friction and other bulk flow variables in two-dimensional rectangular duct flow. *J. Fluids Eng.*, 100:215–223.
- Dickey, D. A. and Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *J. Am. Stat. Assoc.*, 74(366a):427–431.
- Du, Y. and Karniadakis, G. (2000). Suppressing wall turbulence by means of a transverse traveling wave. *Science*, 288:1230–1234.
- Edinburgh Parallel Computing Centre (EPCC) (2016-2020). Training materials. http://www.archer.ac.uk/training/past_courses.php.
- Ferguson Jr., W. E. (1982). A simple derivation of Glassman’s general N fast Fourier transform. *Comput. Math. Appl.*, 8(6):401–411.
- Frigo, M. and Johnson, S. G. (2005). The design and implementation of FFTW3. *Proc. IEEE*, 93(2):216–231.
- Frisch, U. (1995). *Turbulence: The Legacy of AN Kolmogorov*. Cambridge University Press.
- Fukagata, K., Iwamoto, K., and Kasagi, N. (2002). Contribution of Reynolds stress distribution to the skin friction in wall-bounded flows. *Phys. Fluids*, 14(11):73–76.
- Fukagata, K., Kern, S., Chatelain, P., Koumoutsakos, P., and Kasagi, N. (2008). Evolutionary optimization of an anisotropic compliant surface for turbulent friction drag reduction. *J. Turbul.*, 9:N35.
- Fukagata, K., Sugiyama, K., and Kasagi, N. (2009). On the lower bound of net driving power in controlled duct flows. *Physica D*, 238(13):1082–1086.
- Garcia-Mayoral, R. and Jiménez, J. (2011). Hydrodynamic stability and breakdown of the viscous regime over riblets. *J. Fluid Mech.*, 678:317–347.
- Gatti, D. and Quadrio, M. (2016). Reynolds-number dependence of turbulent skin-friction drag reduction induced by spanwise forcing. *J. Fluid Mech.*, 802:553–582.
- Ge, Z., Holmgren, H., Kronbichler, M., Brandt, L., and Kreiss, G. (2018). Effective slip over partially filled microcavities and its possible failure. *Phys. Rev. Fluids*, 3(5):054201.
- Gelbart, M. A., Snoek, J., and Adams, R. P. (2014). Bayesian optimization with unknown constraints. *arXiv:1403.5607*.
- Ghebbali, S., Chernyshenko, S., and Leschziner, M. (2017). Can large-scale oblique undulations on a solid wall reduce the turbulent drag? *Phys. Fluids*, 29(10):105102.
- Godard, G. and Stanislas, M. (2006a). Control of a decelerating boundary layer. part 1: Optimization of passive vortex generators. *Aerosp. Sci. Technol.*, 10(3):181–191.
- Godard, G. and Stanislas, M. (2006b). Control of a decelerating boundary layer. part 3: Optimization of round jets vortex generators. *Aerosp. Sci. Technol.*, 10(6):455–464.
- Hamilton, J. M., Kim, J., and Waleffe, F. (1995). Regeneration mechanisms of near-wall turbulence structures. *J. Fluid Mech.*, 287:317–348.

- Hammond, E., Bewley, T., and Moin, P. (1998). Observed mechanisms for turbulence attenuation and enhancement in opposition-controlled wall-bounded flows. *Phys. Fluids*, 10(9):2421–2423.
- Harlow, F. H. (2004). Fluid dynamics in group T-3 Los Alamos National Laboratory. *J. Comput. Phys.*, 195(2):414–433.
- Harris, T. A. and Kotzalas, M. N. (2006). *Essential Concepts of Bearing Technology*. CRC Press.
- Hasegawa, Y., Quadrio, M., and Frohnapfel, B. (2014). Numerical simulation of turbulent duct flows with constant power input. *J. Fluid Mech.*, 750:191–209.
- Hu, Z. W., Morfey, C. L., and Sandham, N. D. (2006). Wall pressure and shear stress spectra from direct simulations of channel flow. *AIAA Journal*, 44(7):1541–1549.
- Hurst, E., Yang, Q., and Chung, Y. (2014). The effect of Reynolds number on turbulent drag reduction by streamwise travelling waves. *J. Fluid Mech.*, 759:28–55.
- Hutchins, N. and Marusic, I. (2007). Evidence of very long meandering features in the logarithmic region of turbulent boundary layers. *J. Fluid Mech.*, 579:1–28.
- Jacobson, S. A. and Reynolds, W. C. (1998). Active control of streamwise vortices and streaks in boundary layers. *J. Fluid Mech.*, 360:179–211.
- Jeong, J. and Hussain, F. (1995). On the identification of a vortex. *J. Fluid Mech.*, 285:69–94.
- Jones, D. R., Schonlau, M., and Welch, W. J. (1998). Efficient Global Optimization of Expensive Black-Box Functions. *J. Global Optim.*, 13(4):455–492.
- Józsa, T., Balaras, E., Kashtalyan, M., Borthwick, A., and Viola, I. (2019). Active and passive in-plane wall fluctuations in turbulent channel flows. *J. Fluid Mech.*, 866:689–720.
- Jung, W., Mangiavacchi, N., and Akhavan, R. (1992). Suppression of turbulence in wall-bounded flows by high-frequency spanwise oscillations. *Phys. Fluids A*, 4(8):1605–1607.
- Kasagi, N., Suzuki, Y., and Fukagata, K. (2009). Microelectromechanical systems-based feedback control of turbulence for skin friction reduction. *Annu. Rev. Fluid Mech.*, 41(1):231–251.
- Keefe, L. (1997). A normal vorticity actuator for near-wall modification of turbulent shear flows. In *35th Aerospace Sciences Meeting and Exhibit*. American Institute of Aeronautics and Astronautics.
- Keefe, L. (1998). Method and apparatus for reducing the drag of flows over surfaces. *United States Patent*, 5,803,409.
- Khosh Aghdam, S. (2016). *Turbulent drag reduction through wall-forcing methods*. PhD thesis, Department of Mechanical Engineering, University of Sheffield.
- Kim, E. and Choi, H. (2017). Linear proportional–integral control for skin-friction reduction in a turbulent channel flow. *J. Fluid Mech.*, 814:430–451.
- Kim, J. (2003). Control of turbulent boundary layers. *Phys. Fluids*, 15(5):1093–1105.
- Kim, J., Moin, P., and Moser, R. (1987). Turbulence statistics in fully developed channel flow at low Reynolds number. *J. Fluid Mech.*, 177:133–166.

- Koch, H. and Kozulovic, D. (2013). Drag reduction by boundary layer control with passively moving wall. In *ASME 2013 Fluids Engineering Division Summer Meeting*, pages V01BT15A004–V01BT15A004. American Society of Mechanical Engineers.
- Koch, H. and Kozulovic, D. (2014). Influence of geometry variations on the boundary layer control with a passively moving wall. In *AIAA SciTech Forum, 52nd Aero. Sc. Meeting*, page 0401.
- Kolmogorov, N. (1941). The local structure of turbulence in incompressible viscous fluid for very large Reynolds numbers. *The Royal Society*, 434(1890):9–13.
- Kravchenko, A. and Moin, P. (1997). On the effect of numerical errors in large eddy simulations of turbulent flows. *J. Comp. Phys.*, 131:310–322.
- Kristoffersen, R. and Andersson, H. I. (1993). Direct simulations of low-reynolds-number turbulent flow in a rotating channel. *J. Fluid Mech.*, 256:163–197.
- Kühnen, J., Song, B., Scarselli, D., Budanur, N. B., Riedl, M., Willis, A. P., Avila, M., and Hof, B. (2018). Destabilizing turbulence in pipe flow. *Nat. Phys.*, 14(4):386–390.
- Laizet, S. and Lamballais, E. (2009). High-order compact schemes for incompressible flows: A simple and efficient method with quasi-spectral accuracy. *J. Comp. Phys.*, 228:5989–6015.
- Laizet, S. and Li, N. (2010). 2DECOMP&FFT – a highly scalable 2D decomposition library and FFT interface. In *Cray User Group 2010 conference, Edinburgh*.
- Laizet, S. and Li, N. (2011). Incompact3d: A powerful tool to tackle turbulence problems with up to $O(10^5)$ computational cores. *Int. J. Numer. Meth. Fluids*, 67(11):1735–1757.
- Larsson, L. and Raven, H. C. (2010). *Ship resistance and Flow*. The Principles of Naval Architecture Series. Society of Naval Architects and Marine Engineers.
- Lauga, E. and Stone, H. (2003). Effective slip in pressure-driven Stokes flow. *J. Fluid Mech.*, 489:55–77.
- Lauder, B., Poncet, S., and Serre, E. (2010). Laminar, transitional, and turbulent flows in rotor-stator cavities. *Annu. Rev. Fluid Mech.*, 42(1):229–248.
- Lee, C., Kim, J., Babcock, D., and Goodman, R. (1997). Application of neural networks to turbulence control for drag reduction. *Phys. Fluids*, 9(6):1740–1747.
- Lee, K. H., Cortelezzi, L., Kim, J., and Speyer, J. (2001). Application of reduced-order controller to turbulent flows for drag reduction. *Phys. Fluids*, 13(5):1321–1330.
- Lee, M. and Moser, R. D. (2015). Direct numerical simulation of turbulent channel flow up to $Re_\tau \approx 5200$. *J. Fluid Mech.*, 774:395–415.
- Lele, S. (1992). Compact finite difference schemes with spectral-like resolution. *J. Comp. Phys.*, 103:16–42.
- Leschziner, M. A., Choi, H., and Choi, K.-S. (2011). Flow-control approaches to drag reduction in aerodynamics: progress and prospects. *Philos. Trans. Royal Soc. A*, 369(1940):1349–1351.
- Lin, J. C., Robinson, S. K., McGhee, R. J., and Valarezo, W. O. (1994). Separation control on high-lift airfoils via micro-vortex generators. *J Aircraft*, 31(6):1317–1323.

- Ling, H., Katz, J., Fu, M., and Hultmark, M. (2017). Effect of Reynolds number and saturation level on gas diffusion in and out of a superhydrophobic surface. *Phys. Rev. Fluids*, 2(12):124005.
- Lloyd, C. J., Peakall, J., Burns, A., Keevil, G., Dorrell, R., Wignall, P., and Fletcher, T. (2021). Hydrodynamic efficiency in sharks: the combined role of riblets and denticles. *Bioinspiration & Biomimetics*.
- Lozano-Durán, A. and Jiménez, J. (2014). Effect of the computational domain on direct simulations of turbulent channels up to $Re_\tau = 4200$. *Phys. Fluids*, 26(1):011702.
- Luchini, P. (2015). The relevance of longitudinal and transverse protrusion heights for drag reduction by a superhydrophobic surface. In *Proc. European Drag Reduction and Flow Control Meeting—EDRFMC 2015; March 23–26*, pages 81–82.
- Luchini, P. (2017). Universality of the Turbulent Velocity Profile. *Phys. Rev. Lett.*, 118(22):224501.
- Mahfoze, O., Laizet, S., and Wynn, A. (2018). Bayesian optimisation of intermittent wall blowing for drag reduction of a spatially evolving turbulent boundary layer. In *10th Int. Conference in Computational Fluid Dynamics, Barcelona, Spain*.
- McKay, M. D., Beckman, R. J., and Conover, W. J. (2000). A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 42(1):55–61.
- Meneveau, C. and Katz, J. (2000). Scale-invariance and turbulence models for large-eddy simulation. *Annu. Rev. Fluid Mech.*, 32(1):1–32.
- Min, T. and Kim, J. (2004). Effects of hydrophobic surface on skin-friction drag. *Phys. Fluids*, 16(7):L55–L58.
- Moin, P. and Mahesh, K. (1998). Direct numerical simulation: a tool in turbulence research. *Annu. Rev. Fluid Mech.*, 30:539–578.
- Moreau, E. (2007). Airflow control by non-thermal plasma actuators. *J. Phys. D: Appl. Phys.*, 40(3):605–636.
- Moser, R., Kim, J., and Mansour, N. (1999). Direct numerical simulation of turbulent channel flow up to $Re_\theta = 590$. *Phys. Fluids*, 11(4):943–945.
- Munday, P. M. and Taira, K. (2018). Effects of Wall-Normal and Angular Momentum Injections in Airfoil Separation Control. *AIAA Journal*, 56(5):1830–1842.
- Nakanishi, R., Mamori, H., and Fukagata, K. (2012). Relaminarization of turbulent channel flow using traveling wave-like wall deformation. *Int. J. Heat Fluid Flow*, 35:152–159.
- Navier, C. (1823). Mémoire sur les lois du mouvement des fluides. *Mémoires de l’Académie Royale des Sciences de l’Institut de France*, 6:389–440.
- Oak Ridge National Laboratory (ORNL) (1995). *High Performance Computer Architecture, Computational Science Education Project*.
- Oliver, T. A., Malaya, N., Ulerich, R., and Moser, R. D. (2014). Estimating uncertainties in statistics computed from direct numerical simulation. *Phys. Fluids*, 26(3):035101.

- Olivucci, P., Ricco, P., and Aghdam, S. K. (2019). Turbulent drag reduction by rotating rings and wall-distributed actuation. *Phys. Rev. Fluids*, 4(9):093904.
- Orlandi, P. (2012). *Fluid flow phenomena: a numerical toolkit*. Springer Science & Business Media.
- Orlandi, P. and Fatica, M. (1997). Direct simulations of turbulent flow in a pipe rotating about its axis. *J. Fluid Mech.*, 343:43–72.
- Orszag, S. A. and Patterson, G. S. (1972). Numerical simulation of three-dimensional homogeneous isotropic turbulence. *Phys. Rev. Lett.*, 28(2):76–79.
- Ou, J. and Rothstein, J. P. (2005). Direct velocity measurements of the flow past drag-reducing ultrahydrophobic surfaces. *Physics of Fluids*, 17(10):103606.
- Paleyev, A., Pullin, M., Mahsereci, M., Lawrence, N., and González, J. (2019). Emulation of physical processes with emukit. In *Second Workshop on Machine Learning and the Physical Sciences, NeurIPS*.
- Panton, R. L. (2007). Composite asymptotic expansions and scaling wall turbulence. *Philos. Trans. Royal Soc. A*, 365(1852):733–754.
- Pao, H.-P. (1972). Numerical solution of the Navier-Stokes equations for flows in the disk-cylinder system. *Phys. Fluids*, 15(1):4.
- Parnaudeau, P., Lamballais, E., Heitz, D., and Silvestrini, J. (2004). Combination of the immersed boundary method with compact schemes for dns of flows in complex geometry. In *Direct and Large-Eddy Simulation V*, pages 581–590. Springer.
- Pebay, P. P. (2008). Formulas for robust, one-pass parallel computation of covariances and arbitrary-order statistical moments. Technical report, Sandia National Laboratories.
- Pope, S. B. (2000). *Turbulent Flows*. Cambridge University Press.
- Quadrio, M. (2011). Drag reduction in turbulent boundary layers by in-plane wall motion. *Phil. Trans. Royal Soc. A*, 369(1940):1428–1442.
- Quadrio, M. and Luchini, P. (2003). Integral time-space scales in turbulent wall flows. *Phys. Fluids*, 15(8):2219–2227.
- Quadrio, M. and Ricco, P. (2004). Critical assessment of turbulent drag reduction through spanwise wall oscillations. *J. Fluid Mech.*, 521:251–271.
- Quadrio, M. and Ricco, P. (2011). The laminar generalized Stokes layer and turbulent drag reduction. *J. Fluid Mech.*, 667:135–157.
- Quadrio, M., Ricco, P., and Viotti, C. (2009). Streamwise-travelling waves of spanwise wall velocity for turbulent drag reduction. *J. Fluid Mech.*, 627:161–178.
- Quadrio, M. and Sibilla, S. (2000). Numerical simulation of turbulent flow in a pipe oscillating around its axis. *J. Fluid Mech.*, 424:217–241.
- Quarteroni, A., Sacco, R., and Saleri, F. (2010). *Numerical mathematics*. Springer Science & Business Media.

- Rastegari, A. and Akhavan, R. (2019). On drag reduction scaling and sustainability bounds of superhydrophobic surfaces in high reynolds number turbulent flows. *J. Fluid Mech.*, 864:327–347.
- Rathnasingham, R. and Breuer, K. S. (2003). Active control of turbulent boundary layers. *J. Fluid Mech.*, 495:209–233.
- Rebbeck, H. and Choi, K.-S. (2001). Opposition control of near-wall turbulence with a piston-type actuator. *Physics of Fluids*, 13(8):2142–2145.
- Rebbeck, H. and Choi, K.-S. (2006). A wind-tunnel experiment on real-time opposition control of turbulence. *Phys. Fluids*, 18(3):035103.
- Reholon, D. and Ghaemi, S. (2018). Plastron morphology and drag of a superhydrophobic surface in turbulent regime. *Phys. Rev. Fluids*, 3(10):104003.
- Reneaux, J. (2004). Overview on drag reduction technologies for civil transport aircraft. *European Congress on Computational Methods in Applied Sciences and Engineering - ECCOMAS 2004*.
- Reynolds, O. (1883). An experimental investigation of the circumstances which determine whether the motion of water shall be direct or sinuous, and of the law of resistance in parallel channels. *Phil. Trans. R. Soc.*, 35(224-226):84–99.
- Reynolds, W. and Hussain, A. (1972). The mechanics of an organized wave in turbulent shear flow. Part 3. Theoretical models and comparisons with experiments. *J. Fluid Mech.*, 54(2):263–288.
- Ricco, P. (2004). Modification of near-wall turbulence due to spanwise wall oscillations. *J. Turbul.*, 5 024.
- Ricco, P. and Hahn, S. (2013). Turbulent drag reduction through rotating discs. *J. Fluid Mech.*, 722:267–290.
- Ricco, P. and Quadrio, M. (2008). Wall-oscillation conditions for drag reduction in turbulent channel flow. *Int. J. Heat Fluid Flow*, 29:601–612.
- Rothstein, J. P. (2010). Slip on superhydrophobic surfaces. *Annual Review of Fluid Mechanics*, 42(1):89–109.
- Roy, C. J. (2005). Review of code and solution verification procedures for computational simulation. *J. Comput. Phys.*, 205(1):131–156.
- Ruelle, D. and Takens, F. (1971). On the nature of turbulence. *Comm. Math. Phys.*, 20(3):167–192.
- Russo, S. and Luchini, P. (2017). A fast algorithm for the estimation of statistical error in DNS (or experimental) time averages. *J. Comput. Phys.*, 347(Supplement C):328–340.
- Schmid, P. and Henningson, D. (2001). *Stability and Transition in Shear Flows*. Applied Mathematical Sciences, Springer, Vol. 142).
- Schonlau, M. (1997). *Computer experiments and global optimization*. PhD thesis, University of Waterloo.
- Schouveiler, L., Le Gal, P., and Chauve, M. P. (2001). Instabilities of the flow between a rotating and a stationary disk. *J. Fluid Mech.*, 443:329–350.

- Seo, J., García-Mayoral, R., and Mani, A. (2015). Pressure fluctuations and interfacial robustness in turbulent flows over superhydrophobic surfaces. *Journal of Fluid Mechanics*, 783:448–473.
- Seo, J., García-Mayoral, R., and Mani, A. (2018). Turbulent flows over superhydrophobic surfaces: flow-induced capillary waves, and robustness of air–water interfaces. *J. Fluid Mech.*, 835:45–85.
- Shahriari, B., Swersky, K., Wang, Z., Adams, R., and de Freitas, N. (2016). Taking the Human Out of the Loop: A Review of Bayesian Optimization. *Proc. IEEE*, 104(1):148–175.
- Skote, M. (2013). Comparison between spatial and temporal wall oscillations in turbulent boundary layer flows. *J. Fluid Mech.*, 730:273–294.
- Smits, A. J., McKeon, B. J., and Marusic, I. (2011). High–Reynolds Number Wall Turbulence. *Annu. Rev. Fluid Mech.*, 43(1):353–375.
- Stewartson, K. (1953). On the flow between two rotating coaxial disks. *Math. Proc. Cambridge Phil. Soci.*, 49(2):333–341.
- Taira, K., Brunton, S. L., Dawson, S. T. M., Rowley, C. W., Colonius, T., McKeon, B. J., Schmidt, O. T., Gordeyev, S., Theofilis, V., and Ukeiley, L. S. (2017). Modal analysis of fluid flows: An overview. *AIAA Journal*, 55(12):4013–4041.
- Thiébaux, H. J. and Zwiers, F. W. (1984). The Interpretation and Estimation of Effective Sample Size. *J. Climate Appl. Meteor.*, 23(5):800–811.
- Tomiyama, N. and Fukagata, K. (2013). Direct numerical simulation of drag reduction in a turbulent channel flow using spanwise traveling wave-like wall deformation. *Phys. Fluids*, 25(10):105115.
- Townsend, A. A. R. (1976). *The Structure of Turbulent Shear Flow*. Cambridge University Press, second edition.
- Tunncliffe Wilson, G. (1979). Some efficient computational procedures for high order arma models. *J. Stat. Comput. Sim.*, 8(3-4):301–309.
- Viotti, C., Quadrio, M., and Luchini, P. (2009). Streamwise oscillation of spanwise velocity at the wall of a channel for turbulent drag reduction. *Phys. Fluids*, 21(11):115109.
- Von Kármán, T. (1921). Über laminare und turbulente Reibung. *Z. Angew. Math. Mech.*, 1(4):233–252.
- von Storch, H. and Zwiers, F. W. (2001). *Statistical Analysis in Climate Research*. Cambridge University Press.
- Vreman, A. W. and Kuerten, J. G. M. (2014). Comparison of direct numerical simulation databases of turbulent channel flow at $re_\tau = 180$. *Phys. Fluids*, 26(1):015102.
- Wilcox, D. C. (2006). *Turbulence modeling for CFD*, volume 2. DCW industries La Canada, CA.
- Williams, C. and Rasmussen, C. (2006). *Gaussian Processes for Machine Learning*. MIT Press.
- Wise, D., Olivucci, P., and Ricco, P. (2018). Turbulent drag reduction through oscillating discs—corrigendum. *J. Fluid Mech.*, 856:1064–1066.

- Wise, D. and Ricco, P. (2014). Turbulent drag reduction through oscillating discs. *J. Fluid Mech.*, 746:536–564.
- Wise, D. J., Alvarenga, C., and Ricco, P. (2014). Spinning out of control: Wall turbulence over rotating discs. *Phys. Fluids*, 26(12):107–125.
- Wong, T.-S., Kang, S., Tang, S., Smythe, E., Hatton, B., Grinthal, A., and Aizenberg, J. (2011). Bioinspired self-repairing slippery surfaces with pressure-stable omniphobicity. *Nature*, 477(7365):443–447.
- Yamamoto, Y. and Tsuji, Y. (2018). Numerical evidence of logarithmic regions in channel flow at $Re_\tau = 8000$. *Phys. Rev. Fluids*, 3(1).
- Yao, J., Chen, X., and Hussain, F. (2019). Reynolds number effect on drag control via spanwise wall oscillation in turbulent channel flows. *Phys. Fluids*, 31(8):085108.
- Zhao, X. and He, G.-W. (2009). Space-time correlations of fluctuating velocities in turbulent shear flows. *Phys. Rev. E*, 79(4):046316.
- Zwart, S. P. (2020). The ecological impact of high-performance computing in astrophysics. *Nat. Astron.*, 4(9):819–822.