HIGHLY TURBULENT TAYLOR-COUETTE FLOW

direct numerical simulations

Rodolfo Ostilla Mónico

Highly turbulent Taylor-Couette flow: Direct numerical simulations

Rodolfo Ostilla Mónico

Samenstelling promotiecommissie:

Prof. dr. Hans Hilgenkamp (voorzitter)	Universiteit Twente
Prof. dr. Detlef Lohse (promotor)	Universiteit Twente
Prof. dr. Roberto Verzicco (co-promotor)	Universiteit Twente/Tor Vergata
Dr. Chao Sun	Universiteit Twente
Prof. dr. Jaap van der Vegt	Universiteit Twente
Prof. dr. Gert-Jan van Heijst	Technical University of Eindhoven
Prof. dr. Javier Jiménez	Universidad Politécnica de Madrid

European Research Council

The work in this thesis was carried out at the Physics of Fluids group of the Faculty of Science and Technology of the University of Twente. It was funded by an European Research Council (ERC) advanced grant. Part of the results of this thesis have been achieved using PRACE Resources (projects 2012061135 and 2013091966).

Nederlandse titel:

Hoog turbulente Taylor-Couette stroom: Directe numerieke simulaties

Publisher:

Rodolfo Ostilla Mónico, Physics of Fluids, University of Twente,
P.O. Box 217, 7500 AE Enschede, The Netherlands
pof.tnw.utwente.nl

Cover: Constant azimuthal velocity contour in a Taylor-Couette simulation with $\eta = 0.909$, $Re_i = 10^5$ and resting outer cylinder.

© Rodolfo Ostilla Mónico, Enschede, The Netherlands 2014

No part of this work may be reproduced by print
photocopy or any other means without the permission
in writing from the publisher

ISBN: 978-90-365-3799-5 // DOI: 10.3990/1.9789036537995

HIGHLY TURBULENT TAYLOR-COUETTE FLOW: DIRECT NUMERICAL SIMULATIONS

PROEFSCHRIFT

ter verkrijging van
de graad van doctor aan de Universiteit Twente,
op gezag van de rector magnificus,
Prof. dr. H. Brinksma,
volgens besluit van het College voor Promoties
in het openbaar te verdedigen
op vrijdag 20 februari 2015 om 16.45 uur

door

Rodolfo Ostilla Mónico
geboren op 1 Augustus 1988
te Puerto de la Cruz, Spanje

Dit proefschrift is goedgekeurd door de promotoren:

Prof. dr. rer. nat. Detlef Lohse en Prof. dr. Roberto Verzicco

To my Family, with a capital letter

To maintain a cheerful attitude of mind in the midst of a gloomy and exceedingly responsible task, is no slight artistic feat. And yet, what could be more necessary than cheerfulness? Nothing ever succeeds which exuberant spirits have not helped to produce.

—FRIEDRICH WILHELM NIETZSCHE

Contents

1	Introduction	1
1.1	The problem of turbulence	1
1.2	The study of turbulence	4
1.3	The Taylor-Couette problem as a model for turbulence	6
1.4	Outline of the thesis	11
2	A pencil distributed code for strongly turbulent wall-bounded flows	13
2.1	Introduction	14
2.2	Rayleigh-Bénard convection and Taylor-Couette flow	15
2.3	Numerical scheme	17
2.4	Constraints on the timestep	21
2.5	Code parallelization	23
2.6	Computational performance	26
2.7	Summary and prospects	28
3	A multiple-resolution strategy for Direct Numerical Simulation of scalar turbulence	31
3.1	Introduction	32
3.2	Governing equations and numerical method	34
3.3	Pressure effects and heat transfer in RB	35
3.4	The multiple resolutions strategy	40
3.5	Results and computational performance	44
3.6	Summary and conclusions	48
4	Optimal Taylor-Couette: transition to turbulence	53
4.1	Introduction	54
4.2	Numerical method	60
4.3	Code validation	62
4.4	Global response	68
4.5	Characterization of the flow state	76
4.6	Local results	77
4.7	Summary and conclusions	91
5	Optimal Taylor-Couette: radius-ratio dependence	93
5.1	Introduction	94
5.2	Numerical method	95

5.3	Experimental set-up	100
5.4	Global response: Torque	102
5.5	Local results	112
5.6	Summary and conclusions	122
6	Boundary layer dynamics at the transition between the classical and the ultimate regimes	125
6.1	Introduction	126
6.2	Results	128
6.3	Summary and conclusions	140
7	Exploring the phase diagram of fully turbulent Taylor-Couette flow	143
7.1	Introduction	144
7.2	Explored parameter space	147
7.3	The inverse Rossby number dependence	151
7.4	The effect of radius ratio or the η -dependence	160
7.5	Dependence on number and size of rolls	165
7.6	Summary and conclusions	168
8	The effect of Taylor rolls on highly turbulent Taylor-Couette flow	175
8.1	Introduction	176
8.2	Numerical details	178
8.3	Results	180
8.4	Summary and conclusions	189
9	Effects of the computational domain size on the numerical simulations	193
9.1	Introduction	194
9.2	Results and analysis	196
9.3	Summary and conclusions	205
10	Turbulence decay towards the linearly-stable regime	209
10.1	Introduction	210
10.2	Numerical details	211
10.3	Results and analysis	213
10.4	Summary and conclusions	217
11	Conclusions and Outlook	221
	Summary	247
	Samenvatting	249

<i>CONTENTS</i>	iii
Publications	251
Acknowledgements	253
The life of Rodolfo	255

1

Introduction

1.1 The problem of turbulence

The mere mention of turbulence causes people to fasten their seat belts in aeroplanes and physicists to raise their eyebrows. Without turbulence, we would simply not exist. Stellar dust clouds would not collapse to form galaxies, solar systems and planets. Entire marine ecospheres would collapse. Our climate would drastically change. Juninho Pernambucano would not be able to perform his signature “knuckle-ball” free kicks. And worst of all, we would not be able to mix our coffee with milk by stirring a spoon, and we would be denied the pleasure of a “relaxing cup of café con leche in Plaza Mayor”.

Without a doubt, turbulence puzzles physicists. Horace Lamb famously said “I am an old man now, and when I die and go to heaven there are two matters on which I hope for enlightenment. One is quantum electrodynamics, and the other is the turbulent motion of fluids. And about the former I am rather optimistic”. It is widely regarded as the “most important unsolved problem of classical physics”. But what is exactly the problem of turbulence? What *is* turbulence? Etymologically, the more mundane meaning of turbulence, or something turbulent as something agitated and disturbed comes from the latin word “turba”, which refers to the disorderly motion of a crowd. To imagine how this relates to fluid flow, just think of a heavy metal music concert. Leonardo da Vinci was the first person to use the word in a scientific context in fluid mechanics. In this more precise manner, turbulence is harder to define. Wikipedia does it the following way: “Turbulence, or turbulent flow is a flow regime characterized by chaotic property changes. This includes low momentum diffusion, high momentum convection, and rapid variation of pressure and velocity in space and time” [1]. If one reads this definition from Wikipedia, turbulence is hard to picture. We appear to suffer from the “I-know-it-when-I-see-it” phenomena when attempting to define turbulence. We can easily recognize turbulence in every day life. Even more, to the formally trained person, the Wikipedia definition will seem

Figure 1.1: Left: Transition of laminar flow to turbulent flow in a rising plume in still air. Right: turbulent flow over a sphere. The Reynolds number of this flow is usually defined as $U_\infty D/\nu$, where ν is the kinematic viscosity of the fluid. The Reynolds number can be made larger by making the flow faster, the sphere larger or by reducing the kinematic viscosity of the fluid. Credits: Dr. Gary Settles / Wikimedia.

misleading and erroneous, and an attempt to more precisely define turbulence will lead to even more obscurity for the layman.

Therefore, let us begin by picturing (or alternatively, look at the left panel of figure 1.1) the rising smoke coming out from a candle, which first rises smoothly, and later in a less smooth and more *turbulent* way. By simple visual inspection of the photograph, it is quite clear to us that in the flow in the top part of the image is turbulent. Turbulence arises when the Reynolds number, that is, the ratio of inertia forces to viscous forces in a fluid becomes large enough. The right panel of figure 1.1 shows a way of understanding this Reynolds number. There are many kinds of flows in which turbulence plays a role. Internal flows, such as flows inside a pipe; external flows, such as atmospheric flow over the surface of the ocean, or “free” shear flows- an example of this being the rising plume of figure 1.1. The characteristics of turbulence in these flows, however different one flow is from the other, seem to be very similar. One of the first insights into turbulence was given by Richardson, immortalized as a poem: “Big whorls have little whorls, [t]hat feed on their velocity, [a]nd little whorls have lesser whorls, [a]nd so on to viscosity” [2]. Richardson essentially meant that large vortical structures in the flow, called eddies, which are of the size of the problem (e.g. the size of the sphere in figure 1.1), are unstable and break down into smaller and smaller vortices. Eventually, they become so small that viscosity, caused by the random motion of molecules, dampens them out. In various seminal works, Kolmogorov proposed a quantitative model for Richardson’s cascade at very large Reynolds numbers, the so-called K41 theory [3–5]. The K41 theory

postulates that the kinetic energy cascades from the largest, energy-containing scales to the viscous scales through a range of “inertial” scales. Small vortices in the inertial and viscous range of scales are universal, they “forget” where they come from, and do not feel the presence of the external flow. They do not have a preferred position, or a preferred orientation, i.e. they are homogeneous and isotropic, and only depend on two parameters, the amount of energy which is cascading through them and viscosity. Even more, inertial-range vortices are also too large to feel viscosity, and only feel the energy cascade. This insight led to several predictions for the velocity energy spectra, i.e. the kinetic energy at given length scale, which have been extensively validated by experiments [6]. Although K41 is not exactly correct, and corrections to it exist, it still is the most important model we have for understanding turbulence far away from walls, and the one taught in graduate-level turbulence courses.

As mentioned previously, turbulence is present in many different kinds of natural and man-made flows. The most commonly studied turbulent flows are free shear flows, such as wakes, jets and mixing layers. These flows are free from the influence of walls, and turbulence arises because of the velocity differences in the mean flow. Free shear flows are reasonably well understood from self-similarity arguments and due to the fact the K41 model works well away from walls [7]. However, turbulent flows near walls are more difficult to characterize, and wall-bounded flows seem to be harder to understand. One of the first key insights into the behaviour of turbulent boundary layers, i.e. the region of a turbulent flow very close to a wall, was provided by Prandtl [8], who, following a very similar line of argumentation as the K41 theory, proposed that there is a region in a turbulent boundary layer dominated by eddies which are too large to be damped by viscosity, and too small to feel that they are in the proximity of a wall which sets additional constraints. Prandtl postulated that the statistics of these eddies will only depend on the shear which is transported from the bulk to the wall, in a way which is analogous to the inertial eddies of Kolmogorov, which only depend on the energy cascade. As a result of this, turbulent boundary layers would have universal mean velocity profile, with a logarithmic-region, regardless of the flow in question. Prandtl’s model was extended in an aim to understand the cascades happening in the near-wall region. We highlight here the model based on a self-similar family of attached eddies by Townsend [9] and the eddy-forest model of Perry, Henbest and Chong [10]. Both models also predicted the shape of the velocity fluctuations near the walls, and the functional dependence of the velocity spectra. Many corrections to the original models exist, and more information can be found in reviews such as Refs. [11, 12], including discussions on the controversy between logarithmic behaviour versus power law behaviour of the mean velocity profiles.

1.2 The study of turbulence

Models for turbulence, commonly used in other areas such as geophysics (including climate science) [13] and engineering [14], are developed by studying turbulent phenomena from both experimental and computational perspectives. Experiments are usually very complicated, as to reach the very large Reynolds limit we need facilities which are extremely large and cost millions of Euros. Physicists also have to compete for the usage of some of these facilities. For example, high Reynolds number wind tunnels are also used by engineers attempting to tackle more mundane problems, such as the design of sports cars, an application which undoubtedly moves more money than fundamental research. Physicists can also study turbulence by looking at Earthly flows such as the atmosphere and oceans, and even to the Universe, inside stars or accretion disks, using them as extremely high Reynolds number experiments. However, it is hard to isolate single ingredients of these flows, as they often are a combination of many processes including radiation and magnetic fields, not just pure fluid dynamics. Furthermore, experimental measurements can be extremely challenging or even impossible, such as directly measuring the convective flows inside a star.

Turbulence can also be studied numerically through the use of direct numerical simulations (DNS). Simulations are usually constrained in the Reynolds number they can achieve. The reasons for this limitation will be revisited later. However, they have the immediate advantage of having access to the complete velocity and pressure field data at every time instant; something unthinkable in experiments. DNS means that the full Navier–Stokes equation for a fluid are advanced in time given the initial conditions and the boundary conditions. No turbulence modelling, such as those mentioned previously (and commonly found in engineering software packages such as FLUENT) is done. This is because our aim is to study the turbulence itself.

The Navier–Stokes equations for an incompressible fluid [15] are

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

plus the incompressibility condition

$$\nabla \cdot \mathbf{u} = 0, \quad (1.2)$$

where $\mathbf{u}$ is velocity, p is the dynamic pressure, ρ is density, ν is the kinematic viscosity of the fluid and t is time.

DNS attempts to resolve all scales in the flow, from the largest scales which are the size of the problem, to the small, viscosity-dominated scales. A snapshot of such a simulation is shown in figure 1.2. This leads to very high computational requirements for turbulent flows, as the scale separation between the largest and the smallest scale grows as $Re^{3/4}$ [7], where we now explicitly define the Reynolds number Re

Figure 1.2: Visualization of the turbulent structures in a two-dimensional cut through a simulation of turbulent Taylor-Couette flow, with a Reynolds number of 10^5 . Over a billion degrees of freedom were used in this simulation, adding up to a total of consumption of 1.2 million CPU hours for the whole simulation.

as $Re = UL/\nu$, with U a characteristic velocity of the system, and L a characteristic length of the system. The above relationship has as an immediate consequence that for a three-dimensional problem, the amount of grid points required to spatially discretize the system grows as $Re^{9/4}$. Furthermore, not only the dynamic range of length scales grows with increasing Reynolds number, but also the dynamic range of time scales (i.e. the scale difference between the fastest and the slowest time scales) grows. Once one adjusts also for the reduced time-step required to integrate the equations, a computational cost of $Re^{11/4}$ for simulating a characteristic time unit $\mathcal{O}(L/U)$ is obtained. This estimate is very optimistic, as it considers that the amount of computational work scales linearly with the amount of points and with the number of time-steps in the simulated system. This is usually not the case, and even larger scaling exponents relating the amount of work to the Reynolds numbers are seen in practice, such as Re^4 [7].

Even so, DNS remains a very versatile tool for the study of turbulence, and has led to many deep insights on its nature in the last years. The first numerical study of the full Navier-Stokes equations was the seminal work on homogeneous isotropic turbulence by Orszag & Patterson [16] in 1972. The computational requirements at that time allowed for a grids of $(32)^2$ and $(64)^3$, which took 30 seconds per time-step on a CDC6600 supercomputer. Computational power grows every day, and in the present, this computation could be performed in fractions of a second on any desktop,

or even smartphones. The current state-of-the art DNS for homogeneous isotropic turbulence (HIT), dating to 2003, was of the order of $(4096)^3$ [17]. Simulations using grids of $(12288)^3$ are being performed at the time of writing by the same research group [18].

DNS is, and will continue to be one of the most demanding computational problems. The 1972 DNS of HIT [16] reached a Taylor-microscale Reynolds number of $Re_\lambda = 35$, while the present effort (2015) by Ishihara *et al.* reaches $Re_\lambda \approx 2000$ [18]. This is less than a factor of one hundred increase in forty years. And even so, $Re_\lambda = 2000$ is not enough to achieve the desired scale separation. In the present day there seems to be a never-ending demand for faster computers with larger memory to perform larger and larger computations at higher and higher Reynolds numbers. Leading-edge codes (in 2013) have achieved scaling up to two millions of processors by joining two supercomputers together, in a simulation which simulates over 4.12 trillion degrees of freedom and uses 1.6 petabytes of memory [19].

1.3 The Taylor-Couette problem as a model for turbulence

In this thesis, we will focus on DNS of one canonical wall-bounded turbulent flow, namely Taylor-Couette (TC) flow to further understand its properties in relation to other canonical flows. TC flow, schematized in figure 1.3, consists of the flow in a fluid layer between two independently rotating cylinders. TC flow is among the most investigated systems in fluid mechanics, not only owing to its simplicity as an experimental model for shear flows, but also due to its applications in process technology. Examples of these applications are journal bearings, rotating machinery, viscometers, biological separation devices and (photo)chemical reactors [20–24].

The rotation difference between the cylinder shears the fluid and this drives the flow. The rotation rates of the cylinders can be expressed non-dimensionally as two Reynolds numbers, the inner cylinder $Re_i = r_i \omega_i d / \nu$, and the outer cylinder Reynolds numbers $Re_o = r_o \omega_o d / \nu$, where r_i and r_o are the radii of the inner and outer cylinder, respectively, ω_i and ω_o the inner and outer cylinder angular velocity, $d = r_o - r_i$ the gap width, and ν the kinematic viscosity of the fluid. The geometry of TC flow is characterized by two non-dimensional parameters, namely the radius ratio $\eta = r_i / r_o$ and the aspect ratio $\Gamma = L / d$, where L is the axial length of the cylinders. Thus, four parameters fully characterize the system. Instead of using the Reynolds numbers, one can also use two different parameters: the Taylor number $Ta = \frac{1}{4} \sigma (r_o - r_i)^2 (r_o + r_i)^2 \omega_i^2 \nu^{-2}$, which quantifies the shear-driving of the system, where $\sigma = [(r_o + r_i) / (2\sqrt{r_i r_o})]^4$ can be considered as a geometric quasi-Prandtl number [25], and the Rossby number $Ro = |\omega_o - \omega_i| r_i / (2\omega_o d)$, which quantifies the magnitude of the Coriolis force due to the rotation of the outer cylinder. A more detailed derivation of these alternative parameters is presented in §4.

Figure 1.3: Left: schematic of the Taylor-Couette system. The system consists of two coaxial cylinders, which have an inner cylinder radius of r_i and an outer cylinder radius of r_o . Both cylinders are of length L . The inner cylinder rotates with an angular velocity ω_i and the outer cylinder rotates with an angular velocity of ω_o . Right: Twente Turbulent Taylor-Couette (T³C) experimental facility with a dangerously-looking red button. This setup will be revisited in §5.

The investigation of the Taylor-Couette problem has a history of over three hundred years. Its study can be traced back to Newton, who touched on the subject of rotating circular motion in his *Principia Mathematica* back in 1687. The first experimental work dates back the end of the 1800s, when Maurice Couette [26] in France developed an apparatus which he could use to measure the viscosity of fluids. This was essentially a Taylor-Couette system in which only the outer cylinder rotated. His design was not completely original though, as he based himself on previous writings by Stokes [27] and Margules [28]. At around the same time in the United Kingdom, Mallock [29, 30] also built a device to measure torque between two coaxial cylinders, but he was able to rotate both cylinders. When only the inner cylinder was rotating, he entered the turbulent regime of TC flow. This was reflected on the torque-versus-velocity curve, which no longer followed a purely linear functional dependence, a result to be expected when only viscosity plays a role. Later work by Wendt [31] and Taylor [32, 33], greatly expanded the study of TC flow, the former measuring torques and velocities for several radius and rotation ratios in the turbulent case, and the latter being the first to mathematically describe the cells that form if the flow is linearly unstable. Taylor's studies were also pioneering in the sense that they convincingly demonstrated that the Navier-Stokes equations are the correct equations for describing fluid flow, and also that the correct boundary conditions to take at the boundaries (walls) were ones of no-slip, i.e. zero relative velocity between fluid and wall.

In the 1960s, TC flow was extensively discussed in Chandrasekhar's "Hydrody-

namic and Hydromagnetic Stability” [34], greatly detailing different variations to the basic TC problem. Experimental and numerical work on TC flow also continued over the years, cf. Refs.[35–45] at low and high Reynolds numbers and for different ratios of the rotation frequencies. This work has been complemented by numerical simulations, not only in the regime of pure inner cylinder rotation [46–50], but also for eigenvalue study [51] and counter-rotation at fixed rotation ratios [49]. For a broader historical overview on the field, which we have barely summarized here, we refer the reader to Donnelly’s review [52]. An outline of the parameter space explored at the time of beginning this thesis is shown in figure 1.4.

Following the tradition of Taylor, TC flow, and also Rayleigh-Bénard (RB) convection, the flow in a fluid layer heated from below and cooled from above, have been extensively used to explore new concepts in fluid mechanics, such as instabilities [34, 35, 57–61], non-linear dynamics and chaos [62–66], pattern formation [37, 67, 68], and turbulence [39, 69–73]. Studies of these systems have been performed both numerically and experimentally. The main reasons behind the popularity of these systems are, in addition to the fact that they are closed systems, as mentioned previously, their simplicity due to the high amount of symmetries present.

Interestingly enough, the existence of an exact mathematical analogy between TC flow and RB convection was shown by Eckhardt, Grossmann and Lohse [25]. The temperature difference between the plates in RB convection corresponds to the shear caused by the different rotation rates of the inner and outer cylinders in TC flow. One-to-one correspondence of the scaling laws relating driving and torque (or heat transfer for RB) are expected. This is of paramount interest, as Kraichnan [74] postulated that RB flow would reach an asymptotic or “ultimate” state if the driving is sufficiently large. In this ultimate regime, turbulence is fully developed throughout both bulk and boundary layers. In that regime, heat transport, which for less driving is limited by the laminar boundary layers, has no explicit, unit-wise dependence on viscosity. The ultimate-regime scaling laws relating heat transfer and flow driving are expected to be extendible to arbitrarily large Rayleigh numbers, like those present in both geo- and astrophysical flows.

Transitions in the scaling laws to the ones expected in the ultimate state have recently been found in experiments for heat transport and wind in RB flow [75, 76] and angular velocity transport and wind in TC flow [41, 45, 77], confirming the existence of the ultimate regime, 50 years after Kraichnan’s prediction. RB and TC flow are expected to show the transition to the ultimate regime when the boundary layers are sheared strongly enough so that they undergo a shear instability and become turbulent [77–79]. The mechanical driving of TC flow is much more efficient in destabilizing the boundary layer, as it is directly sheared, than the one of RB convection. Thus, in the experiments of van Gils and coworkers [43, 77], only the ultimate regime of TC flow was seen, and the transition to the regime could not be captured. What is

Figure 1.4: State-of-the-art of the parameter space exploration of TC flow at the beginning of the thesis, including experimental data from Refs. [31, 33, 36–38, 40, 43, 44, 53, 54], and numerical data from [48–50, 55, 56]. The black square indicates the data points in the well-studied small-Reynolds-number regime of pattern formation and spatial temporal chaos.

happening at this transition, and where exactly does this transition take place across the whole parameter space is the first open question we focus on.

The second question we focus on this thesis is the lack of understanding of the phenomenon of optimal transport in TC flow. Recent experiments [43, 44, 80] revealed that for a given shear driving of the flow, a maximum in the torque-versus-rotation ratio curve was found for non-zero rotation rates. Further more, a deeper link between optimal transport, and the underlying local quantities was also found by van Gils *et al.* [77]. The action of a “neutral” surface, which separates regions of the flow which are unstable or stable to perturbations was found to play a crucial role in determining this optimum. However, the experimental capabilities of measuring flow quantities are much more limited than those of numerics. Simulations of TC flow can look in a more detailed way at the flow. Thus, in this thesis, we will use them to further look into the mechanisms behind optimal transport.

The next question we will focus on is on the universality of turbulence in TC flow. As mentioned in the introduction, turbulence is supposed to be universal. However, since the experiments by Lewis & Swinney [41], it is known that TC lacks the K41 inertial range, even at experiments conducted at very large Reynolds numbers $Re \sim \mathcal{O}(10^6)$. The lack of this inertial range was further confirmed by Huisman *et al.* [81]. The effect of the large scale circulation, i.e. the turbulent Taylor rolls, seems to play a crucial role in determining the bulk dynamics. However, recent experimental results point in the direction that the boundary layers in TC are very similar to those in channels and pipes [82]. Once again, we can use DNS, to look in more detail than experiments, and explore in depth the turbulent bulk and the turbulent boundary layer of TC flow. We can use this new data to make one-to-one comparisons to the available experimental and numerical data for canonical wall-bounded flows, and also to see how well the existing models for wall-bounded flows as those of Refs. [9, 10, 83] fare with TC flow.

Since the work of Rayleigh in 1917 [84], it is known that inviscid TC flow is linearly unstable as long as the angular momentum L decreases (in an absolute sense) with the radius [34]. This means that for certain control parameters, the whole flow is linearly stable even in the absence of viscosity. These parameters coincide with those which are of central relevance to astrophysics [85]. The final question we will focus on answering in this thesis takes place in this sub-critical regime. The question whether TC flow does indeed undergo a non-linear transition [86, 87] to turbulence at large shear Reynolds numbers has been probed experimentally by several authors with conflicting results [43, 44, 77, 88–93]. The discrepancies between the different experiments have been attributed to the end plates of the TC devices, i.e. the solid boundaries which axially confine TC flow. These cause secondary flows, known as Ekman circulation, which propagate into the flow and influence global stability. Avila [93], based on his numerical simulations with end plates imitating the Ji *et al.* [90]

and the Paoletti & Lathrop [92] experiments, concluded that both of these experiments were inadequate to study the non-linear transition due to the presence of end plates. DNS does not have these limitations, and can be run with periodic boundary conditions to completely prevent end-plate effects. Therefore, we will be able to the stability of sub-critical TC flow using DNS focusing on the “platonic” problem.

1.4 Outline of the thesis

This thesis attempts to answer the questions presented in the previous subsection by the means of numerical simulations. To perform these simulations, we first detail in §2 the development of a code to simulate high Reynolds number wall-bounded flows, which can be directly applied to TC flow. In §3, we expand this code with a multiple resolutions strategy for simulating scalars in turbulent flows, which proves advantageous even if the diffusivity of the scalar is the same as the momentum diffusivity.

In §4 we validate our computational code, and present in detail the rotating frame formulation of TC, and the analogy between RB and TC. Using these new tools, we attempt to further understand the phenomena of optimal transport. §5 presents an exploration of the dependence of optimal transport on the radius ratio η , combining DNS and experiments (provided by Sander Huisman and co-workers) to further understand the flow properties. §6 analyses in detail the boundary layer dynamics around the transition to the ultimate regime, and §7 explores the transition to the ultimate regime across all four dimensions of the TC flow parameter space. §8 studies in more detail the properties of the Taylor-rolls, attempting to characterize them, and drawing analogies to other kinds of wall-bounded flows. This is followed by a study on the effect of the computational domain on these statistics in §9, to check the applicability of previous results. In §10 a exploration of sub-critical TC flow using the rotating frame formulation is performed, using highly turbulent initial conditions to probe the stability. The thesis concludes with a summary of the results and an outlook for future work.

2

A pencil distributed code for strongly turbulent wall-bounded flows *

The unexamined life is not worth living for a human being.

—SOCRATES

We present a numerical scheme geared for high performance computation of wall-bounded turbulent flows. The number of all-to-all communications is decreased to only six instances by using a two-dimensional (pencil) domain decomposition and utilizing the favourable scaling of the CFL time-step constraint as compared to the diffusive time-step constraint. As the CFL condition at the wall is more restrictive at high driving, implicit time integration of the viscous terms in the wall-parallel directions is no longer required. This avoids the communication of non-local information to a process for the computation of implicit derivatives in these directions. We explain in detail the numerical scheme used for the integration of the equations, and the underlying parallelization. The code is shown to have very good strong and weak scaling to at least 64K cores.

*Based on: E. P. van der Poel, R. Ostilla-Mónico, J. Donners and R. Verzicco, *A pencil distributed code for strongly turbulent wall-bounded flows*, under review (arXiv:1501.01247). Parallelization and initial coding by van der Poel. Implementation of Poisson solver and new time-marching scheme by Ostilla-Mónico. Donners provided later help with optimizing the code. Writing and discussion done by everybody.

2.1 Introduction

Turbulence is known as the “last unsolved problem of classical physics”. Direct numerical simulations (DNS) provide a valuable tool for studying in detail the underlying, and currently not fully understood, physical mechanisms behind it. Turbulence is a dynamic and high dimensional process, in which energy is transferred (cascades) from large vortices into progressively smaller ones, until the scale of the energy is so small that these vortices are dissipated by viscosity. DNS requires resolving all of the flow scales, and to adequately simulate a system with very large size separation between the largest and the smallest scale, immense computational power is required.

The seminal works on homogeneous isotropic turbulence by Orszag & Patterson [16] and on pressure-driven flow between two parallel plates (also known as channel flow) by Kim, Moin and Moser [94], while difficult back then, could be performed easily on contemporary smartphones. Computational resources grow exponentially, and the scale of DNS has also grown, both in memory and floating point operations (FLOPS). In approximately 2005, the clock speed of processors stopped increasing, and the focus shifted to increasing the number of processors used in parallel. This presents new challenges for DNS, and efficient code parallelization is now essential to obtaining (scientific) results.

Efficient parallelization is deeply tied to the underlying numerical scheme. A wide variety of these schemes exist; for trivial geometries, i.e. domains periodic in all dimensions, spectral methods are the most commonly used [17]. However, for the recent DNS of wall-bounded flows, a larger variation of schemes is used. For example, in the present year, two channel flow DNSs at similar Reynolds numbers detailed DNSs were performed using both a finite-difference schemes (FDS) in the case of Ref. [95] or a more complex spectral methods in the case of Ref. [96]. FDS also present several advantages, they are very flexible, allowing for complex boundary conditions and/or structures interacting through the immersed boundary method with relative ease [97]. A commonly asserted disadvantage of low-order FDS is the higher truncation error relative to higher order schemes and spectral methods. However, this is only true in the asymptotic limit where the grid spacing becomes sufficiently fine, i.e. $\Delta x \rightarrow 0$, that is commonly not reached. Additionally, aliasing errors are much smaller for lower order schemes [98, 99]. Lower-order schemes have been shown to produce adequate first- and second-order statistics, but require higher resolution when compared to spectral methods for high order statistics [100–102]. Because lower-order schemes are computationally very cheap the grid resolution can in general be larger for the same computational cost compared to higher order schemes, although one has to consider the higher memory bandwidth over FLOPS ratio.

In this chapter, we will detail the parallelization of a second-order FDS based on the approach by Verzicco & Orlandi [103] to two wall bounded systems, Rayleigh-Bénard (RB) convection, the flow in a fluid layer between two parallel plates; one

heated from below and cooled from above and Taylor-Couette (TC) flow, the flow between two coaxial independently rotating cylinders; although our code can easily be extended to any flow that is wall-bounded in one dimension. This FDS scheme has already been used in pure Navier-Stokes simulations [103], Navier-Stokes simulations combined with immersed boundary methods [104], for Rayleigh-Bénard convection [105–111] and for Taylor-Couette flow [112, 113]. The numerical results have been validated against experimental data numerous times. We will exploit several advantages of the large Re regime and the boundary conditions to heavily reduce communication cost; opening the possibility to achieve much higher drivings.

The chapter is organized as follows: §2.2 describes TC and RB in more detail. In §2.3 the numerical scheme which is used to advance the equations in time is detailed. In §2.4 we show that in thermal convection, the Courant-Friedrichs-Lewy (CFL) [114] stability constraints on the time-step due to the viscous terms become less strict than those due to the non-linear terms at high Rayleigh (Reynolds) numbers. In §2.5 we detail a pencil decomposition to take advantage of the new time integration scheme and the choice of data arrangement in the pencil decomposition. Finally, in §2.6 we compare the computational cost of existing and the new approach and present an outlook of what further work can be done to combine this approach with other techniques.

2.2 Rayleigh-Bénard convection and Taylor-Couette flow

RB and TC are paradigmatic models for convective and shear flows, respectively. They are very popular systems because they are mathematically well defined, experimentally accessible and reproduce many of the interesting phenomena observed in applications. A volume rendering of the systems can be seen in figure 2.1. The Reynolds numbers in the common astro- and geo-physical applications are much higher than what can be reached currently in a laboratory. Therefore it is necessary to extrapolate available experimental results to the large driving present in stars and galaxies. This extrapolation becomes meaningless when transitions in scaling behaviour are present, and it is expected that once the Rayleigh number, i.e. the non-dimensional temperature difference, becomes large enough, the boundary layers transition to turbulence. This transition would most likely affect the scaling of interesting quantities. However, experiments disagree on exactly where this transition takes place [75, 76]. DNS can be used to understand the discrepancies among experiments. However, to reach the high Rayleigh numbers (Ra) of experiments new strategies are required. DNS must resolve all scales in the flow, and the scale separation between the smallest scale and the largest scale grows with Reynolds number. This means larger grids are needed, and the amount of computational work W scales approximately as $W \sim Re^4$ [7].

Figure 2.1: Left: RB flow for $Ra = 10^8$, $Pr = 1$ and $\Gamma = 2$ in Cartesian coordinates. The horizontal directions are periodic and the plates are subjected to a no-slip and isothermal boundary condition. Red/yellow indicates hot fluid, while (light) blue indicates cold fluid. The small heat carrying structures known as thermal plumes as well as a large scale circulation can be seen in the visualization, highlighting the scale separation in the flow. Right: TC flow with an inner cylinder Reynolds number $Re = 10^5$, a stationary outer cylinder, and a radius ratio $\eta = r_i/r_o = 0.714$. Green fluid has a high angular velocity while blue fluid has a low angular velocity. The smallness of the structures responsible for torque transport, and thus the need for fine meshes, can be appreciated clearly.

Simulations of RB commonly imitate the cylindrical geometry most used in experiments. Recently, a DNS with aspect ratio $\Gamma = D/L = 1/3$, where D is the diameter of the plates and L the height of the cell, reached $Ra = 10^{12}$ using 1.6 Billion points with a total cost of 2 Million CPU hours [115]. DNS in other geometries have been proposed, such as homogeneous RB, where the flow is fully periodic and a background temperature gradient is imposed. This geometry is easy to simulate [116], but presents exponentially growing solutions and does not have a boundary layer, thus not showing any transition [117]. Axially homogeneous RB, where the two plates of the cylinder are removed, the side-walls kept, and a background temperature gradient is imposed to drive the flow has also been simulated (e.g. in Ref. [118]). This system does not have boundary layers on the plates and does not show the sought after transition. Therefore, it seems necessary to keep both horizontal plates, having at least one wall-bounded direction. The simplest geometry we can simulate with these requirements is a parallelepiped box, periodic in both wall-parallel directions, which we will call “rectangular” RB for simplicity. Rectangular RB is receiving more attention recently [119–122], due to possibility to reach higher Ra as compared to more complex geometries. It is additionally the geometry that is closest to natural applications, where there are commonly no sidewalls.

For TC, we have one naturally periodic dimension, the azimuthal extent. The axial extent can be chosen to be either bounded by end-plates, like in experiments, or to be periodic. Axial end-plates have been shown to cause undesired transitions to turbulence if TC is in the linearly stable regime [93], or to not considerably affect the flow if TC is in the unstable regime [77]. Large Re DNS of TC focus on axially periodic TC, bounding the flow only in the radial direction [113, 123]. Therefore, the choice of having a single wall-bounded direction for DNS of both TC and RB seems justified.

2.3 Numerical scheme

The code solves the Navier-Stokes equations with an additional equation for temperature in three-dimensional coordinates, either Cartesian or cylindrical. For brevity, we will focus on the RB Cartesian problem, although all concepts can be directly translated to TC in the cylindrical coordinate system by substituting the vertical direction for the radial direction, and the two horizontal directions by the axial and azimuthal directions.

The non-dimensional Navier-Stokes equations with the Boussinesq approximation for RB read:

$$\nabla \cdot \mathbf{u} = 0, \tag{2.1}$$

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p + \sqrt{\frac{\text{Pr}}{\text{Ra}}} \nabla^2 \mathbf{u} + \theta \mathbf{e}_x, \quad (2.2)$$

$$\frac{\partial \theta}{\partial t} + \mathbf{u} \cdot \nabla \theta = \sqrt{\frac{1}{\text{PrRa}}} \nabla^2 \theta, \quad (2.3)$$

where $\mathbf{u}$ is the non-dimensional velocity, p the non-dimensional pressure, θ the non-dimensional temperature and t the non-dimensional time. For non-dimensionalization, the temperature scale is the temperature difference between both plates Δ , the length scale is the height of the cell L and the velocity scale is the free-fall velocity $U_f = \sqrt{g\beta\Delta L}$, where g is gravity and β is the isobaric expansion coefficient of the fluid. $\text{Pr} = \nu/\kappa$ is the fluid Prandtl number, where ν is the kinematic viscosity and κ is the thermal conductivity. The Rayleigh number is defined in this case as $\text{Ra} = g\beta\Delta L^3/\nu\kappa$. Finally, $\mathbf{e}_x$ is the unitary vector in the anti-parallel direction to gravity, which is also the plate-normal direction.

As mentioned previously, the two horizontal directions are periodic and the vertical direction is wall-bounded. The spatial discretization used is a conservative second-order centred finite difference with velocities on a staggered grid. The pressure is calculated at the center of the cell while the temperature field is located on the u_x grid. This is to avoid the interpolation error when calculating the term $\theta \mathbf{e}_x$ in equation (2.2), and thus to exactly conserve the transformation of potential energy into kinetic energy. The scheme is energy conserving in the limit $\Delta t \rightarrow 0$. A two-dimensional (for clarity) schematic of the discretization is shown in figure 2.2. For the case of thermal convection, an additional advantage of using FDS is present: the absence of pressure in the advection/diffusion equation for scalars causes shock-like behaviour to appear in the temperature (scalar) field [124] and low-order schemes fare better in this situation.

Time marching is performed with a fractional-step third-order Runge-Kutta (RK3) scheme, in combination with a Crank-Nicholson scheme [100] for the implicit terms. A second-order Adams-Bashforth (AB2) method is also implemented in the code. However, in all production runs the RK3 method takes precedence over AB2 even though the total RK3 time step includes three substeps as compared to one for AB2. The theoretical stability limit of AB2 and RK3 are CFL numbers lower than 1 and $\sqrt{3}$, respectively. In practice, the maximum CFL numbers of AB2 and RK3 are approximately 0.3 and 1.3, respectively. Because of three times higher amount of substeps in RK3, the computational cost is proportionally higher compared to AB2. Nevertheless, RK3 is more efficient as the progression in physical time per computational cost is better. In addition, even though the Crank-Nicholson integration with $\mathcal{O}([\Delta t]^2)$ error is the weakest link, the $\mathcal{O}([\Delta t]^3)$ error of RK3 decreases the total error significantly compared to the $\mathcal{O}([\Delta t]^2)$ error of AB2. In addition, RK3 is self-starting

Figure 2.2: Location of pressure, temperature and velocities of a 2D simulation cell. The third dimension (z) is omitted for clarity. As on an ordinary staggered scheme, The velocity vectors are placed on the borders of the cell and pressure is placed in the cell center. The temperature is placed on the same nodes as the vertical velocity, to ensure exact energy conservation.

at each time step without decreasing the accuracy and without needing additional information in the restart file. AB2 would require two continuation files per quantity.

The pressure gradient is introduced through the “delta” form of the pressure [125]: an intermediate, non-solenoidal velocity field $\mathbf{u}^*$ is calculated using the non-linear, the viscous and the buoyancy terms in the Navier-Stokes equation, as well as the pressure at the current time sub-step:

$$\frac{\mathbf{u}^* - \mathbf{u}^j}{\Delta t} = \left[\gamma_l H^j + \rho_l H^{j-1} - \alpha_l \mathcal{G} p^j + \alpha_l (\mathcal{A}_x^j + \mathcal{A}_y^j + \mathcal{A}_z^j) \frac{(\mathbf{u}^* + \mathbf{u}^j)}{2} \right], \quad (2.4)$$

where the superscript j denotes the sub-step, $\mathcal{A}_i$ is the discrete differential relationship for the viscous terms in the i th-direction, $\mathcal{G}$ the discrete gradient operator and H^j all explicit terms. The coefficients γ_l , ρ_l and α_l depend on the time marching method used. The pressure required to enforce the continuity equation at every cell is then calculated by solving a Poisson equation for the pressure correction ϕ :

$$\nabla^2 \phi = \frac{1}{\alpha_l \Delta t} (\nabla \cdot \mathbf{u}^*), \quad (2.5)$$

or in discrete form:

$$\mathcal{L} \phi = \frac{1}{\alpha_l \Delta t} (\mathcal{D} \mathbf{u}^*), \quad (2.6)$$

where $\mathcal{D}$ the discrete divergence operator, and $\mathcal{L}$ is the discrete Laplacian operator, $\mathcal{L} = \mathcal{D}\mathcal{G}$. The velocity and pressure fields are then updated using:

$$\mathbf{u}^{j+1} = \mathbf{u}^* - \alpha_l \Delta t (\mathcal{G}\phi), \quad (2.7)$$

and

$$p^{j+1} = p^j + \phi - \frac{\alpha_l \Delta t}{2\text{Re}} (\mathcal{L}\phi), \quad (2.8)$$

making $\mathbf{u}^{j+1}$ divergence free.

The original numerical scheme treats all viscous terms implicitly. This would result in the solution of a large sparse matrix, but this is avoided by an approximate factorization of the sparse matrix into three tridiagonal matrices; one for each direction [103]. The tridiagonal matrices are then solved using Thomas' algorithm, with a Sherman-Morrison perturbation if the dimension is periodic, in $\mathcal{O}(N)$ time. The calculation is thus simplified at the expense of introducing an error $\mathcal{O}(\Delta t^3)$. This method was originally developed and used for small Reynolds number problems, and without having in mind that data communication between different processes could be a bottleneck. The first parallelization scheme with MPI was a 1D-domain “slab” decomposition, visualized in the figure 2.3(a). The main bottlenecks were found in the all-to-all communications present in the pressure-correction step and the tridiagonal solver in the direction in which the domain is decomposed (cf. table 2.1 for more details). Slab decompositions are easy to implement, but are limited in two ways: First, the number of MPI processes cannot be larger than N , the amount of grid points in one dimension. A hybrid MPI-OpenMP decomposition can take this limit further, but scaling usually does not go further than $\mathcal{O}(10^4)$ cores. Second, the size of the “halo” (or ghost) cells becomes very significant with increasing number of cores. Halo cells are cells which overlap the neighbour's domain, and whose values are needed to compute derivatives. In the limit of one grid point per processor, halo cells are of the size of the domain in a second-order scheme, and even larger for higher order schemes.

2D-domain decompositions, also known as “pencil” decompositions, mitigate these problems. A schematic of this decomposition is shown in figure 2.3(b). To implement this decomposition, the 2DECOMP [126] library has been used, and extended upon to suit the specifics of our scheme. The limit on the amount of processes is now raised to N^2 , and the size of the halo cells on every core decreases with increasing amount of cores, so the amount of communication per core decreases. However, for a pencil decomposition solving all the tridiagonal matrices requires all-to-all communications for two directions, instead of one direction, as in the case of slab decomposition. As mentioned previously, solving the tridiagonal matrices involves large data communication, especially in the context of pencil decompositions. In this chapter we attempt to eliminate them as far as possible and at the same time implement an efficient pencil domain decomposition by arranging the data to gain

Figure 2.3: Left panel: Slab-type domain decomposition using four MPI processes. Right panel: Pencil-type domain decomposition using twelve MPI processes.

advantage of the inherent anisotropy with respect to the grid point distributions.

2.4 Constraints on the timestep

In the new scheme, solving the tridiagonal matrices in the horizontal directions is avoided by integrating not only the advection terms but also the viscous terms explicitly. A major concern is that this can cause the temporal stability issues that the implicit integration used to negate. In this section we argue that for high Ra , using the Courant-Frederich-Lewy (CFL) [114] time-step size constraint is sufficient. The CFL condition ensures the stability of the integration of the advection terms, and for high Ra it additionally ensures stability of the viscous terms. The grid point distribution in the wall-normal direction is different compared to the periodic directions. Namely, it is non-uniform in the wall-normal direction, with clustering of points near the boundaries in order to adequately resolve the boundary layers. As the periodic directions are homogeneous, no such clustering is required and the grid point distribution can be uniform. As a consequence, the minimum grid spacing in the wall-normal direction is much smaller than in the horizontal directions. Because of the scaling of both the time-step constraints, the viscous terms in the wall-normal direction do require implicit integration for all Ra .

In this semi-implicit method, both viscous terms and the non-linear terms are integrated explicitly. This requires two stability constraints on the time-step: one due to the non-linear terms, and one due to the horizontal viscous terms. For the non-linear terms, the definition of the CFL condition is given by:

$$\Delta t_{\mathbf{u}} \cdot \nabla \mathbf{u} \leq \mathcal{C}_1 \min_{\forall \mathbf{x} \in X} \frac{1}{\frac{|u_x|(\mathbf{x})}{\Delta x(\mathbf{x})} + \frac{|u_y|(\mathbf{x})}{\Delta y(\mathbf{x})} + \frac{|u_z|(\mathbf{x})}{\Delta z(\mathbf{x})}}, \quad (2.9)$$

where $\mathcal{C}_1$ is the integration scheme dependent Courant number, $\mathbf{x}$ is the position vector, X is the complete domain and $|\cdot|$ gives the absolute value. Here Δx gives the (non-dimensional) grid spacing in the x direction at position $\mathbf{x}$. The wall-normal direction is x and the wall parallel directions are y and z (cf. figure 2.2).

The additional constraint originates from the viscous terms, and is given by:

$$\Delta t_{\nu \nabla^2 \mathbf{u}} \leq \text{Re } \mathcal{C}_2 \min_{\forall \mathbf{x} \in X} (\Delta y(\mathbf{x}) + \Delta z(\mathbf{x}))^2, \quad (2.10)$$

where $\text{Re} = \sqrt{\text{Ra}/\text{Pr}}$ is the Reynolds number based on the freefall velocity, and $\mathcal{C}_2$ a number which depends on the integration scheme and the number of dimensions treated explicitly. This condition only needs to be satisfied in the horizontal directions, and not in the vertical direction, as the time integration of the vertical second derivatives is kept implicit. This is because $\Delta x(\mathbf{x})^2$ can be very small, and the resulting time-step would make the simulation infeasible.

We can now compare the two CFL constraints, and show that the non-linear constraint is more restrictive than the viscous constraint in the homogeneous directions. As $|u_z|$ and Δz are strictly positive, we have:

$$\frac{\text{Re } \mathcal{C}_2}{\frac{|u_x|(\mathbf{x})}{\Delta x(\mathbf{x})} + \frac{|u_y|(\mathbf{x})}{\Delta y(\mathbf{x})} + \frac{|u_z|(\mathbf{x})}{\Delta z(\mathbf{x})}} < \frac{\text{Re } \mathcal{C}_2}{\frac{|u_y|(\mathbf{x})}{\Delta y(\mathbf{x})} + \frac{|u_z|(\mathbf{x})}{\Delta z(\mathbf{x})}}. \quad (2.11)$$

Including the wall-normal grid spacing in the CFL condition gives a smaller time-step than only using horizontal spacing, and thus the expression on the right is an upper bound on the time-step. If we then use that the grid is uniformly spaced, and equally spaced in both horizontal directions, we can simplify the expression as $\Delta y(\mathbf{x}) = \Delta z(\mathbf{x}) = \Delta y$. We also know that the dimensionless velocity is $|u| \sim \mathcal{O}(1)$ by normalization. Using all of this, we obtain:

$$\Delta t_{\mathbf{u}} \cdot \nabla \mathbf{u} \sim \mathcal{O}(\Delta y), \quad (2.12)$$

for the non-linear CFL condition and

$$\Delta t_{\nu \nabla^2 \mathbf{u}} \sim \mathcal{O}(\text{Re } \Delta y^2), \quad (2.13)$$

from the CFL criterion for the viscous terms. If we assume $\text{Pr} \sim \mathcal{O}(1)$, we can get a bound on the viscous time-step as a function of Ra , $\Delta t_{\nu \nabla^2 \mathbf{u}} \sim \mathcal{O}(\text{Ra}^{\frac{1}{2}} \Delta y^2)$.

To compare both bounds, we need an estimation for Δy . For a resolved DNS, Δy should be similar to the smallest physical length scale in the system. Several length scales can be chosen in the thermal convection problem. The first choice stems from

homogeneous turbulence, where the most commonly used length scale that determines the numerical resolution is the Kolmogorov length scale, $\eta_K = \nu^{3/4} \varepsilon^{-1/4}$, where ε is the viscous dissipation rate. For RB, we can obtain an estimate for the Kolmogorov scale by using that the volume and time averaged viscous dissipation rate can be expressed directly as a function of Nu, Ra and Pr [69, 127]:

$$\langle \varepsilon \rangle_{V,t} = \frac{\nu^3}{L^4} \text{RaPr}^{-2} (\text{Nu} - 1). \quad (2.14)$$

For high Ra simulations, $\text{Nu} \gg 1$. Using again that $Pr \sim \mathcal{O}(1)$, we can obtain an estimate for η_K , and thus the grid spacing as $\Delta y = \eta_K/L \approx 1/(\text{RaNu})$. If we assume $\text{Nu} \sim \text{Ra}^\gamma$ we can now compare both CFL constraints on the time step, obtaining $\gamma \leq 1$ as a requirement for the non-linear CFL to be more restrictive on the time step than the viscous CFL constraint.

In RB convection, another restrictive length scale naturally arises, i.e. that of thermal plumes. These are conceptualized as detaching pieces of thermal boundary layers. The thickness of a thermal boundary layer can be approximated by $\lambda \approx 1/(2\text{Nu})$. Using this as an estimate for Δy in equations (2.9) and (2.10) give another bound: $\gamma \leq 1/2$. Trivial upper bounds in RB convection give a physical upper bound of $\text{Nu} \sim \text{Ra}^{1/2}$ [128], indicating that for the mild assumptions made, the criteria $\gamma \leq 1/2$ is always satisfied. This signifies that the scaling of $\Delta t_{\mathbf{u} \cdot \nabla \mathbf{u}}$ is more restrictive than $\Delta t_{\nu \nabla^2 \mathbf{u}}$, which results in that using only the non-linear CFL constraint in the time-marching algorithm, inherently satisfies the stability constraints imposed by the explicit integration of the horizontal components of the viscous terms. Including the vertical non-uniform grid in this derivation makes this statement even more valid, as the used CFL time step is based on this grid (cf. equation 2.9). Inherent to the big-O-notation is the absorption of the coefficients and offsets. This makes this derivation only valid for high Ra flows. For low Ra, the solver will be unstable as the viscous constraint is not satisfied in this regime.

In addition, we note that the previous analysis can be applied for the scalar (temperature) equation as long as $Pr \sim \mathcal{O}(1)$. If $Pr \gg 1$, which is the case in some applications, such as many oils or the Earth's mantle, the CFL constraint on the horizontal conductive terms becomes $\Delta t_{\kappa \nabla^2 T} \sim \mathcal{O}(\text{Pr}^{1/2} \text{Ra}^{1/2} (\Delta y)^2)$, which means a stricter restriction on the time-step than equation (2.13). This means that the Ra of the flow required to make $\Delta t_{\mathbf{u} \cdot \nabla \mathbf{u}} \leq \Delta t_{\kappa \nabla^2 T}$ will be higher.

2.5 Code parallelization

In the previous section, we reasoned that for large Ra the implicit integration of the viscous terms in the horizontal direction becomes unnecessary. The calculation becomes local in space as the two horizontal directions no longer require implicit

Figure 2.4: Domain decomposition of a $16 \times 12 \times 10$ grid using 12 distributed memory processes on a 4×3 process grid. Only data that is exclusive to one process is shown; i.e. the single gridpoint-sized halos are not shown in this figure. The pencils are a) x , b) y or c) z oriented.

solvers to calculate the intermediate velocity field $\mathbf{u}^*$. In this case it is worth decomposing the domain such that the pencils are aligned in the wall-normal (x) direction, i.e. that every processor possesses data from x_1 to x_N (cf. figure 2.4). Halo updates must still be performed during the computation of $\mathbf{u}^*$, but this memory distribution completely eliminates all the all-to-all communications associated to the viscous implicit solvers, as for every pair (y,z) , a single processor has the full x information, and is able to solve the implicit equation in x for the pair (y,z) without further communication.

All-to-all communications are unavoidable during the pressure correction step, as a Poisson equation must be solved. As the two wall-parallel directions are homogeneous and periodic, it is natural to solve the Poisson equation using a Fourier decomposition in two dimensions. Applying a Fourier transform to the variables ϕ and the right side in equation (2.5) reduces the pressure correction equation to:

$$\left(\frac{\partial^2}{\partial x^2} - \omega_{y,j}^2 - \omega_{z,k}^2 \right) \mathcal{F}(\phi) = \mathcal{F} \left[\frac{1}{\alpha_l \Delta t} (\mathcal{D}\mathbf{u}^*) \right] \quad (2.15)$$

where $\mathcal{F}(\cdot)$ denotes the 2D Fourier transform operator, and $\omega_{y,j}$ and $\omega_{z,k}$ denote the j -th and k -th modified wavenumbers in y and z direction respectively, defined as:

$$\omega_{y,j} = \begin{cases} \left(1 - \cos \left[\frac{2\pi(j-1)}{N_y} \right] \right) \Delta_y^{-2} & : \text{for } j \leq \frac{1}{2}N_y + 1 \\ \left(1 - \cos \left[\frac{2\pi(N_y - j + 1)}{N_y} \right] \right) \Delta_y^{-2} & : \text{otherwise} \end{cases} \quad (2.16)$$

and $\omega_{z,k}$ is defined in an analogous way. A modified wavenumber is used, instead of the real wavenumber, to prevent the Laplacian operator having higher accuracy in some dimensions. In the limit $\Delta y \rightarrow 0$, the modified wavenumbers converge to the real wavenumbers.

By using a second order approximation for ∂_x^2 , the left hand side of the equation is reduced to a tridiagonal matrix, and thus the Poisson equation is reduced to a 2D FFT followed by a tridiagonal (Thomas) solver. This allows for the exact solution of the Poisson equation in a single iteration with $\mathcal{O}(N_x N_y N_z \log[N_y] \log[N_z])$ time complexity. Due to the domain decomposition, several data transposes must be performed during the computation of the equation. The algorithm for solving the Poisson equation is as follows:

1. Calculate $(\mathcal{D}\mathbf{u}^*)/(\alpha_l \Delta t)$ from the x -decomposed velocities.
2. Transpose the result of (1) from a x -decomposition to a y -decomposition.
3. Perform a real-to-complex Fourier transform on (2) in the y direction.

4. Transpose (3) from a y -decomposition to z -decomposition.
5. Perform a complex-to-complex Fourier transform on (4) in the z direction.
6. Transpose (5) from a z -decomposition to a x -decomposition.
7. Solve the linear system of equations (2.15) with a tridiagonal solver in the x -direction.
8. Transpose the result of (7) from a x -decomposition to a z -decomposition.
9. Perform a complex-to-complex inverse Fourier transform on (8) in z direction.
10. Transpose (9) from a z -decomposition to a y -decomposition.
11. Perform a complex-to-real inverse Fourier transform on (10) in y direction.
12. Transpose (11) from a y -decomposition to a x -decomposition.

The last step outputs ϕ in real space, decomposed in x -oriented pencils, ready for applying in equations (2.7)-(2.8). Once the Poisson equation is solved, the corrected velocities and pressures are computed using equations (2.7)-(2.8), the temperature and other scalars are advected and the time sub-step is completed. The algorithm outlined above only transposes one 3D array, instead of three velocity fields, making it very efficient. Figure 2.4 shows a schematic of the data arrangement and the transposes needed to implement the algorithm. We wish to highlight that this algorithm also uses all possible combinations of data transposes. It can be seen from figure 2.4 that the x to z transposes and the z to x transposes need a more complex structure, as a process may need to transfer data to other processes which are not immediate neighbours. The non-overlapping of data before and after transposes is most striking for e.g. process 10 in figure 2.4 with no overlap at all between x and z oriented pencils. These transposes are absent in the 2DECOMP library on which we build. These transposes have been implemented using the more flexible all-to-all calls of the type ALLTOALLW, instead of the all-to-all MPI calls of the type ALLTOALLV used for the other four transposes. A complete list of the used libraries can be found in table 2.2 at the end of the chapter.

2.6 Computational performance

For optimal scaling conditions, each processor should have an equal amount of work, and the communication to computation ratio should be minimal. In our case, as we do not have iterative solvers, each grid point has the same amount of work, as long as the amount of points in every processor is the same there is perfect load balancing.

	Slab		Pencil	
	Halo	A2A	Halo	A2A
$\mathbf{u}^*$ computation	2	6	2	0
Pressure correction	8	2	10	6
Scalar equation	3	2	2	0
Total	13	10	14	6

Table 2.1: Number of communications necessary for the computation of all the terms per timestep of the different codes. Here, A2A is short for all-to-all communications. Halo updates involve updating all halo cells, which requires more, but smaller, communications in the case of the pencil code. The difference between the details of the halo and all-to-all for the slab and pencil codes have not been taken into account. It can be seen that the pencil code contains the majority of the communications in the pressure (Poisson) solver.

We also reduce the communication as far as possible. Not only the number of all-to-all communications are reduced, but also the halo communications as the halo is only one grid point wide. Table 2.1 presents the reduction in the number of communications when going from the slab decomposition with fully implicit viscous terms to the pencil decomposition with semi-implicit viscous terms. It is worth noting that the communications are not exactly the same: halo updates involve communications to four neighbours in the pencil decomposition, while only two neighbours are involved in the slab decomposition. However, the size of the halos is much smaller, so less data is transferred. Conversely, for the all-to-all communications, not all processes exchange data with each other in the pencil code, while all processes do so in the slab code. The total memory consumption of the code is approximately $M \approx 15 \times 8 \times N_x \times N_y \times N_z$. Here M is the total used memory in bytes of all processes without MPI overhead. The number of allocated 3D arrays is 14, with additional 1D and 2D arrays of which the memory consumption will not exceed that of one 3D array in the intended cases with large grids.

The left panel of figure 2.5 shows the strong scaling of the code for 2048^3 and 4096^3 grids on the Curie Thin Nodes system. Linear scaling can be seen up to 32K cores for the 4096^3 grid, with some scaling loss for 64K cores. The right panel of figure 2.5 shows the weak scaling of the code for 2^{23} points from 2 to 16K cores. The data in these plots is obtained by using only MPI parallelism, but hybrid OpenMP/MPI schemes are also available in the code. The choice between pure MPI or hybrid OpenMP/MPI depends heavily on the system on which the code is running. In addition, the presence of OpenMP will allow for a faster porting of the codes to a GPU architecture, in case it becomes viable for our application.

Finally, it must be noted that not only the scaling of the code is excellent, but also

the absolute performance. The computational cost per physical time step is very low. In the end, this is what counts.

Figure 2.5: (a) Strong scaling of the code for 2048^3 and 4096^3 degrees of freedom n_{dof} . Here n_c is the number of cores used. The dashed lines indicate linear behaviour. The walltime per timestep t_w (in seconds) is accounted for a full timestep, i.e. three sub-timesteps when using the RK3 integrator. (b) Weak scaling of the code for 8.3 Million (2^{23}) points per core. The dashed is a linear fit to the corresponding data points. Time is measured in seconds.

2.7 Summary and prospects

In this chapter, we have presented a parallelization scheme of a second-order centred finite difference method with minimal communication. Only six transposes are needed for every fractional time-step, and for large enough grids, the code's strong scaling is linear up to 32K cores, with slight performance loss from 64K cores. We emphasize that 64K cores is over half the total number of cores of the Curie Thin nodes system. The code will probably scale well for even larger grids (8096^3) on systems with larger amount of cores, as do similar codes based on the 2DECOMP library [126]. In addition, the absolute performance is excellent. The wall-time per physical time is very low and substantial progress can be made with few computational resources.

The performance of this code allows simulation of flows at high driving. For the application to Rayleigh-Bénard convection, we refer to figure 2.6(a), where the heat flux Nu as a function of the driving Ra can be seen. In this plot, Nu can be compared between the cylindrical, slab decomposed domain, the Cartesian, pencil decomposed domain and the theoretical prediction of the GL-theory [70]. The theoretical prediction is based on a fit to experimental data and thus implies an indirect comparison to experiments as well. The cylindrical domain in numerical simulations is used specifically to facilitate a comparison to experiments at the cost of increased time com-

Figure 2.6: a) The heat flux Nu as a function of the driving Ra for a cylindrical, slab decomposed domain and a Cartesian, pencil decomposed domain. In both cases $Pr = \mathcal{O}(1)$. The cylinder was simulated with an aspect-ratio of $\Gamma = 1/3$, while the lateral size in the periodic case was sufficiently large to approximate infinite aspect-ratio. The GL theory [70] is additionally shown with a solid line. The error bars are based on the temporal convergence of Nu . b) An instantaneous temperature field at height $x = 0.001$ for $Ra = 10^{11}$, where white and black indicate hot and cold fluid, respectively.

plexity of the pressure correction algorithm and the limitation of a one-dimensional domain decomposition (slab), as compared to the proposed Cartesian code. As a quantification of the difference in computational demands of these code: The highest $Ra = 10^{12}$ data point for the cylindrical and the Cartesian simulations have cost 5M and 1M CPU-hours, respectively. This difference, in favour of the Cartesian method, is amplified by the higher temporal convergence of the Cartesian simulation, judging from the smaller error bar size and the use of more degrees of freedom in the Cartesian geometry, as the system volume is slightly larger for identical Γ . The heat flux shows negligible differences, which shows that it is largely independent of the sidewall boundary conditions. Even though the impermeable no-slip wall in the cylindrical case differs largely from the lateral periodicity in the Cartesian case, the heat flux appears unaffected. This indicates that at least for the global quantities, there is no apparent reason to spend more computational resources on a cylindrical simulation, and one can safely resort to the proposed method while maintaining the possibility of comparing to experiments. Even without that possibility, the lateral periodicity is closer to natural applications of RB convection as it approximates infinite aspect-ratio, which by itself warrants the use of the Cartesian domain. The prospected analysis of RB convection is not limited to global quantities such as the heat flux. The highly parallel I/O and high resolution facilitates the study of local quantities. An unfiltered snapshot of the temperature field close to the lower boundary for $Ra = 10^{11}$ is shown in figure 2.6b, where the small scale temperature fluctuations that are required to be properly resolved, can be seen. These small scales can straightforwardly be

Purpose	Library
I/O	HDF5
FFT	FFTW (Guru interface)
Linear algebra	BLAS, LAPACK/MKL/ESSL/LibSci/ACML
Distributed memory parallelism	MPI + 2DECOMP
Shared memory parallelism	OpenMP

Table 2.2: The used libraries for the specified purposes are indicated in this table.

studied using using spectral analysis, or other techniques (cf. Ref. [115]).

The use of this code in Refs. [129, 130] has already allowed us to push the limits in Taylor-Couette simulations to $Re \sim \mathcal{O}(10^5)$, never simulated previously. Its use in future RB simulations is expected to allow us to achieve the large drivings required for entering the “ultimate” regime. The scheme, in combination to a multiple resolution strategy for the scalar field [124], has been used for simulating double diffusive convection [131], achieving the driving parameters relevant for oceanic convection. Due to the flexibility of finite difference schemes, we expect to be able to build further additions on to this code. The possibility of adding a Lagrangian phase, which can be either tracers, one-way or even two-way coupled particles is detailed in Ref. [132]. Other possibilities include adding rough walls using immersed boundary methods [133], or adding mixed Neumann-Dirchelet boundary conditions.

3

A multiple-resolution strategy for Direct Numerical Simulation of scalar turbulence *

When you have eliminated the impossible, whatever remains, however improbable, must be the truth.

—SIR CONAN BOYLE (SHERLOCK HOLMES)

A numerical procedure to simulate low diffusivity scalar turbulence is presented. The method consists of using a grid for the advected scalar with a higher spatial resolution than that of the momentum. The latter usually requires a less refined mesh and integrating both fields on a single grid tailored to the most demanding variable, produces an unnecessary computational overhead. A multiple resolution approach is used also in the time integration in order to maintain the stability of the scalars on the finer grid. The method is the more advantageous the less diffusive the scalar is with respect to momentum, i.e. particularly well suited for large Prandtl or Schmidt numbers. But even in the case of equal diffusivities the present procedure gives CPU time and memory occupation savings. The reason is that the absence of the pressure term in the scalar equation leads to much steeper gradients in the scalar field as compared to the velocity field.

*Based on: R. Ostilla-Mónico, Y. Yang, E. P. van der Poel, D. Lohse and R. Verzicco, *A multiple-resolution strategy for Direct Numerical Simulation of scalar turbulence*, in preparation. Initial proof of concept and coding by Ostilla-Mónico. Tricubic interpolation schemes and DDC implementation by Yang. Further debugging and RB implementation done by van der Poel. Writing and discussion done by everybody.

3.1 Introduction

Countless phenomena in Nature and technology involve one or more scalar fields that are advected and diffused by a turbulent flow. The dilution of pollution in the atmosphere [134], the transport of nutrients in oceans [135], the cooling or heating of devices [136] and the buoyancy-driven currents generated by natural- [69, 137] and double-diffusive [138, 139] convection are just few examples among many.

Numerical simulations have proven to be very helpful in unravelling the complex physics behind these phenomena [72] even if the calculations have shown to be more demanding than expected, taking up to millions of CPU hours in recent studies [108].

The common understanding of the problem is that in three-dimensional turbulent flows, there is a cascade from the largest towards the small spatial scales up to a lower limit that is determined by the diffusivity. As each field has its own diffusivity, these scales can have different magnitudes. In direct numerical simulation (DNS) the mesh size must be smaller than the smallest among them: This requirement quickly renders DNS infeasible. Denoting as η_K the smallest (Kolmogorov) scale of the momentum field, we can calculate the analogous quantity for a scalar field S as $\eta_B = \eta_K / Sc^{1/2}$, also called the Batchelor scale, with $Sc = \nu / \kappa_S$ the Schmidt number defined as the ratio of the kinematic viscosity ν and the scalar diffusivity κ_S , respectively. In some cases, like sugar in water, the Schmidt number exceeds 10^3 resulting in a Batchelor scale of $\eta_B \simeq \eta_K / 30$. With equal grid resolutions for the scalar and the momentum fields, this entails that the momentum field is overresolved by a factor of approximately 30 in each direction. The problem is exacerbated by the fact that a scalar is described by only a single quantity, while momentum is a vector field satisfying the incompressibility condition or other related constraints. This implies that the solution of the momentum alone generally takes an order of 90% of the total CPU time of a simulation and therefore resolving it on an unnecessary fine mesh is not desirable.

The above scenario essentially derived from dimensional analysis, however, does not give the complete picture since it does not account for the structure of the equations. In fact, the naïve comparison between the Kolmogorov and Batchelor scales suggests that for $Sc \approx 1$, $\eta_K \simeq \eta_B$ although in practice the resolution requirements for the momentum and the scalar fields are not the same. Visual evidence of the latter statement can be obtained from the instantaneous snapshots of figure 3.1 showing horizontal cross-sections of temperature and vertical velocity in a thermally driven turbulent flow. In this flow, the fluid hotter than the average temperature (0.5 in nondimensional units) generates upward buoyancy and therefore positive vertical velocity (and vice versa). Although the two fields are very well correlated on the large scales, the sharp fronts of the scalar field do not have an analogous counterpart in the momentum distribution and this results in a different resolution requirement for scalar and momentum. We will see that the main reason for this difference is the presence of the pressure term in the momentum equation that makes the dynamics non-local and

tends to smooth the intense steep fronts. A problem related to the non-local equations for momentum is their high computational cost and the detrimental implications on the parallel performance. Therefore, the possibility of using different meshes for momentum and scalars opens the door to very large gains in performance, not only by reducing the amount of operations, but also the communication of data among processors and the total memory usage. This consideration motivates the present chapter that describes a strategy for efficiently simulating scalar driven turbulent flows with different spatial resolutions for momentum and scalar fields.

Figure 3.1: A horizontal plane halfway between the plates for a Rayleigh-Bénard simulation in a Cartesian geometry at $Ra = 10^{10}$ and Prandtl number $Pr = 1$. *a)* vertical velocity, red indicating rising fluid while blue indicates falling fluid, *b)* temperature, red indicating hot fluid and blue indicating cold fluid. Even though the Prandtl number is one, much sharper gradients can be seen in the right panel.

In this study, the multiple resolution strategy will be mainly applied to Rayleigh-Bénard (RB) convection, the flow of a fluid vertically confined by a top cold and a bottom hot plate. RB is a particularly suitable example for the present application since the flow is driven by the temperature (scalar) field whose diffusivity can be changed by the Prandtl number Pr . In addition, analytical exact relations are available for this problem that can be used to check the correctness of the numerical results. It is worth mentioning that in RB convection the forcing comes from the heated and cooled surfaces where viscous and thermal boundary layers develop. Since they become thinner as the forcing strengthens, the resolution requirements in these boundary layers become more stringent than in the bulk. An extensive analysis of the problem can be found in Ref. [140] where all the details, estimates and guidelines for numerical simulations are given. Here, it suffices to mention that a non-uniform mesh is required in the wall normal direction such to cluster the nodes within the boundary layers. Nevertheless, even if the grid spacing at the wall is much

finer than that in the bulk, the volume of fluid within these layers is at most a few per cent of the total and the nodes allocated there are only a small fraction of the whole mesh.

We will additionally show another numerical example. Namely, double diffusive convection (DDC), in which the flow is driven by two scalars with very different diffusivities and with opposite, stabilizing and destabilizing, effects on the flow. In this case the multiple resolution strategy is even more advantageous and it allows for the simulation of flow regimes that otherwise would be out of reach.

The organization of the chapter is the following. In the next section we describe the governing equations and the numerical method. §3.3 quantifies the differences of momentum and scalar gradients and presents some analytical exact relations for RB. The section closes with the results of reference simulations obtained on a standard single grid. In §3.4 the multiple resolution strategy and numerical details are explained. Finally, §3.5 discusses the results and the computational performance of the method for RB flow and double diffusive convection. Closing remarks are reported in §3.6.

3.2 Governing equations and numerical method

The incompressible Navier–Stokes equations with the Boussinesq approximation for thermal convection, in nondimensional form, read:

$$\nabla \cdot \mathbf{u} = 0, \quad (3.1)$$

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\partial_i p + \sqrt{\frac{Pr}{Ra}} \nabla^2 \mathbf{u} + \theta \mathbf{e}_z, \quad (3.2)$$

$$\frac{\partial \theta}{\partial t} + \mathbf{u} \cdot \nabla \theta = \sqrt{\frac{1}{RaPr}} \nabla^2 \theta, \quad (3.3)$$

where $\mathbf{u}$ is the velocity, p the pressure, θ the temperature (rescaled such that it is one at the hot plate and zero at the cold plate) and $\mathbf{e}_z$ the unitary vector in the anti-parallel direction to gravity, which is also the plate-normal direction. The Rayleigh number Ra is the non-dimensional temperature difference defined as $Ra = g\beta_T(T_h - T_c)L^3/\kappa_T\nu$, where ν is the kinematic viscosity, β_T the isobaric thermal expansion coefficient and κ_T the thermal diffusivity of the fluid, g is the gravity, L the distance and T_h and T_c the temperature at the hot and cold plates, respectively. The Prandtl number Pr , which is the temperature analogue of the Schmidt number, is defined as $Pr = \nu/\kappa_T$.

The integration of the above system is performed by a fractional timestepping [141] with the modifications proposed in Ref. [103]. In short, a provisional velocity

$\hat{u}_i$ is computed from the previous field u_i^n using the old pressure p^n

$$\frac{\hat{u}_i - u_i^n}{\Delta t} = -\partial_i p^n - N_i^{n+1/2} + D_i^{n+1/2}, \quad (3.4)$$

$N_i^{n+1/2}$ contains the non-linear terms and the temperature forcing while $D_i^{n+1/2}$ has the viscous terms: The former are computed explicitly in time, the latter implicitly. The flow incompressibility is then enforced by a pressure correction that takes the form of a Poisson equation $\nabla^2 \phi = \partial_i \hat{u}_i$ whose solution is the most computationally demanding step, especially on non-uniform grids. In addition, the Poisson equation is non-local and this has consequences on code parallelization, requiring the largest amount of communication. Once the scalar ϕ is obtained, the velocity $\hat{u}_i$ is projected onto the solenoidal field u_i^{n+1} and the new pressure p^{n+1} can be computed. The advancement of the temperature is performed directly through

$$\frac{\theta^{n+1} - \theta^n}{\Delta t} = -M^{n+1/2} + V^{n+1/2}, \quad (3.5)$$

where, again, $M^{n+1/2}$ contains the explicit nonlinear terms and $V^{n+1/2}$ the implicit diffusive terms.

All the variables are discretized by central second-order finite-differences on a staggered grid and the time advancement of the solution is obtained by a low-storage third-order Runge-Kutta scheme. Further details can be found in Ref. [103].

3.3 Pressure effects and heat transfer in RB

As mentioned in the introduction, the Prandtl number gives the ratio of momentum to thermal diffusivity, and even if Pr is of order unity, temperature and momentum do not have the same gradient magnitudes (cf. figure 3.1). We quantify this statement in figure 3.2(a) by showing instantaneous temperature θ and vertical velocity u_z profiles across a vertical line from a RB simulation: Much steeper gradients, almost ‘shock’ like, can be seen in the temperature field. These steep gradients are smoothed in the vertical velocity, owing to the pressure effects and this lowers the resolution requirements of momentum with respect to temperature. This observation is further corroborated by figure 3.2(b) showing the probability density functions of $\partial_z \theta$ and $\partial_z u_z$ computed in the bulk of the flow without the boundary layers. Extreme gradients can be seen to be more likely for θ thus evidencing a more intermittent behaviour. This behaviour has been extensively studied, and it is a well established fact that the intermittency corrections to the structure function exponents are much larger for scalars than for velocity [73, 142]. Indeed, these shock-like fronts become much sharper with increasing Reynolds number, and thus increasing small-scale intermittency [143–148].

Figure 3.2: *a)* Instantaneous θ and u_z profiles as a function of the vertical coordinate z/L for $Ra = 10^{10}$, $Pr = 1$ RB simulation. *b)* PDF of $\partial_z \theta$ and $\partial_z u_z$ for the same simulation in the bulk.

We have argued that the absence of the pressure term in the scalar equation is the reason for the localized steep gradients and this idea can be confirmed by the Burgers equation:

$$\partial_t u + u \partial_x u = \nu_B \partial_{xx} u, \quad (3.6)$$

which is often used to test numerical schemes since it is a simple one-dimensional partial differential equation that still retains the essential features of the more complex Navier–Stokes equations: The unsteady term, a quadratic non-linear term and a strictly dissipative viscous term [149].

A one-to-one comparison with equation (3.2), however, evidences the absence of a pressure term and this causes the solution $u(x, t)$ to develop ‘shock’-like discontinuities in a finite time and for finite values of the diffusivity ν_B .

As an example, in figure 3.3, a solution obtained for an initial condition $u(x, 0) = \sin x$ in the domain $-\pi \leq x < \pi$ and $\nu_B = 5 \times 10^{-6}$ is shown. The gradient at $x = 0$ becomes so steep that a fine nonuniform mesh clustered at the centre of the domain is necessary to properly capture the solution (figure 3.3(a)). For under-resolved meshes (figure 3.3(b)) the solution unphysically overshoots the initial extrema $u_{min} = -1$ and $u_{max} = 1$ and spurious oscillations are generated in the region of the steep gradient (figure 3.3(c)) resembling the behaviour of underresolved scalar fields. A very similar phenomenon is observed in a simulation with an underresolved temperature field as shown in figure 3.4.

In RB flow, one of the interesting quantities is the heat flux Q transferred from one plate to the other. In non-dimensional form this is the Nusselt number, $Nu = Q / \kappa_T (T_h - T_c) L^{-1}$. This is not only interesting from a physical point of view, but also as a monitoring variable since it has been observed [150] that when the separately

Figure 3.3: Spatial profiles at various times of the solution of the Burgers equation for $\nu_B = 5 \times 10^{-6}$: dashed $t = 0$, dotted $t = 0.2$, chain-dot $t = 0.4$ and solid $t = 0.8$. *a)* resolved simulation with $\Delta x = 10^{-4}$, *b)* underresolved simulation with $\Delta x = 4 \times 10^{-4}$, *c)* enlargement of *b)* in the region of the steep gradient.

calculated Nusselt number converge and are grid independent, at least all quantities up to second order statistics are properly resolved.

There are several ways to calculate Nu , either by measuring the convective heat transport in the system

$$Nu(z) = \sqrt{RaPr} \langle u_z \theta \rangle_{A,t} + 1, \quad (3.7)$$

or by using the exact relationships derived from global balances [127] of kinetic energy,

$$\varepsilon_v = \nu U_f^2 L^{-2} \langle [\partial_i u_j]^2 \rangle_{V,t} = \nu^3 L^{-4} (Nu - 1) RaPr^{-2}, \quad (3.8)$$

and thermal energy,

$$\varepsilon_\theta = \kappa_T (T_h - T_c)^2 L^{-2} \langle [\partial_i \theta]^2 \rangle_{V,t} = \kappa_T (T_h - T_c)^2 L^{-2} Nu. \quad (3.9)$$

Here the subscripts t , A and V denote, respectively, averages in time, horizontal homogeneous planes and the whole fluid volume, and U_f the free-fall velocity $U_f =$

Figure 3.4: *a)* Snapshot of an underresolved flow field for $Ra = 10^9$ and $Pr = 1$. The characteristic “wiggles” of the underresolved temperature field can be clearly appreciated (see arrows). *b)* Snapshot of a properly resolved velocity field.

$\sqrt{\beta g(T_h - T_c)L}$. Although equation (3.7) depends on z , once the equilibrium is reached its value becomes constant. This condition is requested to assess the statistical convergence of the results.

From here on, we denote Nu with a subscript, either $Nu_{u_z\theta}$, Nu_{ε_v} or Nu_{ε_θ} which specifies the particular equation, i.e. (3.7)–(3.9) respectively, with which Nu is calculated. In addition, we also denote the Nusselt number calculated by the temperature gradient at the wall as Nu_{θ_w} . All definitions are equivalent analytically, but they involve gradients or square gradients of the variables that, when calculated numerically, can deviate from each other if the simulations do not have enough spatial resolution to capture the smallest flow scales: Their comparison can thus be used as a test for the adequacy of the mesh [150].

Figures 3.5(a) and 3.5(b), show the ratios $Nu_{\varepsilon_\theta}/Nu_{u_z\theta}$ and $Nu_{\varepsilon_v}/Nu_{u_z\theta}$ for RB simulations performed on the same grid for momentum and temperature at $Ra = 10^9$ and $Pr = 1$ or $Pr = 10$. Resolutions between $96^2 \times 192$ and $384^2 \times 768$ were used, the larger number corresponding to the vertical (wall-bounded) direction. An aspect ratio of $\Gamma = 1$ was used in both directions, meaning that the computational domain is cubic. Points were clustered near the wall using a Chebychev-like distribution according to the prescriptions of Ref. [140]. The Kolmogorov scale is computed from the kinetic energy dissipation rate using $\eta_K/L = (\nu^3/\varepsilon_v)^{1/4}$ and the Batchelor scale $\eta_B/L = \eta_K/LPr^{-1/2}$. As mentioned above, the various expressions for Nu should be equivalent, their ratios however, approach the unity limit only when the normalized mesh sizes Δ/η_K and Δ/η_B decrease and they do not converge at the same rate. In particular it can be noted that $Nu_{\varepsilon_v}/Nu_{u_z\theta}$ tends to unity for larger grid spacings than $Nu_{\varepsilon_\theta}/Nu_{u_z\theta}$ even for $Pr = 1$ and this corroborates our hypothesis that a finer

resolution is needed for the scalar than for momentum. Using an identical mesh to spatially discretize both momentum and the scalars therefore produces an overhead in computational resources that is redundant.

At the highest resolution, all the definitions converge to the same value (within an uncertainty of 2–3%), therefore we will refer to it as Nu_{ref} without specifying the particular expression and figure 3.5(c) and 3.5(d) show the convergence of $Nu_{u_z\theta}$ to this asymptotic value.

Figure 3.5: (a) ratio between different ways of calculating Nu against grid size for a rectangular RB simulation with $Ra = 10^9$ and $Pr = 1$. Red squares are $Nu_{\epsilon_v}/Nu_{u_z\theta}$, blue circles are $Nu_{\epsilon_\theta}/Nu_{u_z\theta}$. (b) same as (a) for $Pr = 10$. Nu_{ϵ_θ} is plotted against max(Δ/η_B) while Nu_{ϵ_v} is plotted against max(Δ/η_k). For $Pr = 1$, $\eta_B \equiv \eta_K$. (c) convergence to an asymptotic value of $Nu_{u_z\theta}$ for $Ra = 10^9$ and $Pr = 1$. (d): same as (c) for $Pr = 10$.

3.4 The multiple resolutions strategy

3.4.1 Multiple resolution strategy in space

In this subsection we present a method to decouple the spatial discretization of the scalars and the momentum, which allows for large computational savings. This is achieved by refining every cell of a base mesh $\mathcal{M}^i$ times in each i -th direction. A simplified two-dimensional sketch of this procedure is shown in figure 3.6. On the left, the location of the scalars and velocities in the standard single mesh is shown for a staggered arrangement. There, the velocity components are at the centres of the cell faces, while the pressure and the scalar are discretized at the centre of the cell volume. The right panel shows a case with velocity and pressure on the base grid, and a doubly refined ($\mathcal{M}^x = \mathcal{M}^y = 2$) mesh for the scalar, which is temperature in the RB case.

The method works as follows. We first generate the refined mesh over which the scalar field is discretized and then a coarser mesh is obtained by taking only one out of $\mathcal{M}^i$ nodes in the i -th direction. Note that when the mesh is uniform in space this is equivalent to start from the coarse cells and split them into $\mathcal{M}^i$ identical parts. For a non-uniform mesh, however, this naive splitting would result in a staircase distribution of the metrics for the fine grid with constant coefficients within each coarse element and with jumps across its boundaries. These discontinuities would locally decrease the accuracy of the discretization to first order and introduce spurious oscillations in the resolved fields. The difference between the two methods is shown in figure 3.7(a) for a mesh obtained by a Chebychev-like distribution with 96 base nodes and a refinement factor of eight. In this chapter, most numerical examples are obtained using the same $\mathcal{M}^i$ in every direction therefore, unless otherwise specified, from here on we will use only $\mathcal{M}$ to indicate the isotropic refinement factor without specifying the direction. However, this is not required for the method to work and the same procedure can be applied to refinement levels different in each direction depending on the particular flow physics. An example will be shown in §3.5.2.

In order to advect the scalar, the velocity has to be interpolated from the base mesh onto the refined grid at the centres of the faces of the refined cells. This is achieved using a tri-cubic Hermite spline interpolation, taking a stencil of four points in every direction for a total of sixty-four points in three dimensions which, according to our numerical tests, is the minimum required. At the top and bottom boundaries, one-sided interpolation is used, which in principle is less accurate, but is performed on a much finer grid, and thus the amplitude of the error is much smaller than in the bulk. The accuracy of Hermitian interpolation has proven to be sufficient in turbulent flows, and it is comparable to that of B-splines [151]. Hermitian interpolation, however, is preferred in this method as B-splines are much more computationally expensive. Preliminary simulations have shown that a linear interpolation using a

Figure 3.6: Location of pressure, temperature and velocities of a 2D simulation cell. The third dimension (z) is omitted for clarity. As on an ordinary staggered scheme, the velocity vectors are placed on the borders of the cell and pressure is placed in the cell center. The temperature is placed on the same nodes as the vertical velocity, to ensure exact energy conservation.

two point stencil is not sufficient, because it results in a spatially interpolated velocity field which has equal fluxes inside every base cell, and different at the base cell boundaries (figure 3.7(b)). This lack of homogeneity results in spurious oscillations in the scalar field, in particular around local maxima and minima of velocity. These oscillations are visualized in figure 3.8.

Figure 3.7: (a) Normalized metric $\Delta x/N_x$, where N_x is the amount of grid points for a Chebychev-type grid clustering for a refined-mesh generated ($\mathcal{M} = 8$) from a base mesh by splitting the base cells (squares), and the base mesh generated from the refined mesh (diamonds). The first method causes a staircase-like metric, which leads to spurious oscillations. (b) Comparison on interpolated velocity v from a base mesh to the refined mesh ($\mathcal{M} = 8$) using linear interpolation (squares) and cubic Hermite splines (diamonds). Both interpolations coincide at the base mesh points. The underlying basis for the interpolation is also plotted.

The spatially interpolated velocity can then be used to advance equation (3.3),

Figure 3.8: (a) Pseudocolour plot of temperature at the mid-height when using a refinement of $\mathcal{M} = 3$ and linear interpolation for the velocities. Spurious oscillations every three points can be seen all over the domain. (b) Zoom-in of the region inside the black square of (a).

and compute the values of the scalars at the new time. If a scalar couples back to the momentum field, like in the case of RB flow, a spatial filter must be applied to calculate a “coarsened” scalar. In this case, an averaging using equal weights within each refined cell is used. This averaged scalar is then used in equation (3.2) to advance momentum and pressure. Notice that in this case the scalar field is interpolated from a fine mesh onto a coarser one, therefore the previous problem of having equal fluxes between neighbouring cells is not encountered.

We stress that the interpolation of a velocity field between two different grids is a very dangerous operation since its effect is equivalent to that of a low-pass filter, which usually leads to loss of energy and generation and destruction of information. This is catastrophic in the DNS of turbulent flows where the dynamics are based on the energy cascade through the scales. It has even more serious consequences in a RB flow where the balance between thermal (potential) and kinetic energy determines the heat transfer. Nevertheless, if the base mesh already captures the smallest momentum structures, the field is smooth at the scale of the grid cell and an interpolation kernel that is continuous enough does not alter the energy content of the field.

Indeed, the interpolated velocity on the refined grid is *not* solenoidal when the tricubic Hermitian interpolation is used[†]. There are several ways of addressing this problem. One could apply a “minimum energy” correction, by solving a Poisson equation for a velocity correction, and using the local divergence as source term.

[†]A linearly interpolated velocity does remain solenoidal on the fine grid. However, we have clearly seen that other spurious effects are generated, rendering it unusable.

However, this would involve solving a Poisson equation on the fine grid. While the multiple resolutions strategy would remain favourable, this option would negate one of the main advantages of the method. A second option is to remove the interpolation of the vertical velocity in the vertical direction. Instead, to obtain this velocity on the refined grid, the continuity equation (3.1) is integrated separately over each grid cell using the interpolated horizontal velocities for the horizontal derivatives and the vertical velocity on the base mesh as boundary condition. This would result in a solenoidal velocity field on the refined mesh with much less computational effort as the first option. The last option is to use the tricubic interpolation “as-is”. This non-solenoidal field has not resulted in apparent problems for the simulations in this chapter. The divergence is very small as long as we are in the DNS range. However, we anticipate that this could become a problem at higher drivings.

The previous discussion suggests that the proposed multiple resolution procedure can only work if the coarse mesh is fine enough to fully resolve the momentum field. In §3.5 numerical results will confirm this statement showing that when the base mesh adequately resolves the momentum field good results and CPU time saving can be obtained by refining only the grid for the scalar. On the other hand, if the coarse grid does not fully resolve the momentum field even very large values of $\mathcal{M}$ do not lead to correct results.

3.4.2 Multiple resolution strategy in time

The multiple resolution in space entails that the node spacing for the scalars (Δ_S) is smaller than that for the momentum (Δ_U) and this has immediate consequences on the stability of the time integration because of the explicit terms. Due to the Courant–Friedrichs–Lewy condition [114], $\Delta t \cdot \max[U/\Delta_S] \leq \mathcal{C}_{CFL}$, in fact, the time step must decrease by a factor equal to the refinement $\mathcal{M}$ because $\min[\Delta_S] = \min[\Delta_U]/\mathcal{M}$. As this small Δt is not needed by the base mesh for momentum and pressure, the usage of the current approach becomes disadvantageous very rapidly, especially in high Sc flows requiring high values of $\mathcal{M}$. However, a multiple resolution strategy can be applied also in time, by advancing the more expensive momentum and pressure with a larger time-step and the scalar with a smaller one, using a temporally interpolated velocity. In this way, the stability of the explicit terms in the scalar equation is retained without the penalization of an unnecessary small integration step for all the other equations.

The integration of the scalar equation is therefore performed in $\mathcal{L}$ sub-steps, and at each intermediate l time level the velocity is linearly interpolated through

$$\bar{q}^l = \frac{\mathcal{L} - l}{\mathcal{L}} \bar{q}^n + \frac{l}{\mathcal{L}} \bar{q}^{n+1}, \quad (3.10)$$

where $\bar{q}^n$ is the spatially interpolated velocity at time step n . A simple linear inter-

polation is used, which is enough to ensure correctness as is demonstrated in §3.5. This velocity is then used to advect the scalar(s) for every subimestep using equation (3.3) until the scalars have been advanced to the time t^{n+1} . Then, the velocity can be advanced a further time-step and the procedure repeated. If the maximum possible CFL is used for the velocity equations, which is usually the case, then $\mathcal{L} \geq \mathcal{M}$ must be satisfied to ensure stability.

3.5 Results and computational performance

In this section we will present the results of the multiple resolution method, and the associated saving in computational time and memory usage. A more extensive analysis will be performed in the first part for RB flow while, in the second part, additional results with anisotropic refinements will be shown for DDC flow.

3.5.1 Rayleigh–Bénard flow

A series of rectangular RB simulations with aspect ratio $\Gamma = 1$ were run to validate and to demonstrate the benefits of the described method. Meshes of $96^2 \times 192$, $128^2 \times 256$ and $192^2 \times 384$ (only for $Pr = 1$) were used for momentum, the Ra was kept constant at 10^9 while Pr was taken $Pr = 1$ or $Pr = 10$. In order to minimize the computational costs $\mathcal{L} = \mathcal{M}$ was always used except for a specific set of runs in which the effects of $\mathcal{L}$ were isolated.

Figure 3.9(a,b) shows the ratio between the different definitions of Nu and the maximum refinement level $\mathcal{M}$ while the bottom panels report the convergence of Nu_{u_θ} to the asymptotic reference value, calculated from a wiggle-free simulation such as the one seen in Figure 3.8(b). The raw numerical values can be found in table 3.2 at the end of the chapter.

As mentioned before, the multiple resolution strategy only works if the base mesh is sufficiently fine to fully resolve the momentum field. Looking at figure 3.5(a) we can see that at $Pr = 1$ the $96^2 \times 192$ mesh is not sufficiently fine, and the different Nu definitions do not show a monotonic convergence of the Nusselt ratios to the asymptote. It is interesting to note that for $\mathcal{M} = 4$ the Nu ratios get close to one, although the absolute values of Nu are wrong and do not indicate a convergence towards the reference grid independent value. On the other hand, the $192^2 \times 384$ grid yields a converged value of Nu_{ε_v} even though that resolution is not enough for the computation of Nu_{ε_θ} that involves squared temperature gradients. For this case a refinement factor $\mathcal{M} = \mathcal{L} = 2$ for the temperature gives an appropriate resolution as shown by the Nusselt numbers that converge to the reference value. Although a factor two in space and time might seem to produce only limited benefits, we should consider that it implies a grid for the momentum and pressure with 2^3 less elements than for the

Figure 3.9: (a) Ratio between different ways of calculating Nu against grid size for the two base meshes and increasing spatial refinement for $Pr = 1$. (b) same as (a) for $Pr = 10$. (c) convergence to an asymptotic value of Nu_{u_θ} for increasing refinement and $Pr = 1$. (d): same as (c) for $Pr = 10$. Circles are for base meshes of $96^2 \times 192$, while diamonds are for $192^2 \times 384$ base meshes. On the top panels, blue indicates the ratio $Nu_{\epsilon_\theta} / Nu_{u_\theta}$ and red the ratio $Nu_{\epsilon_v} / Nu_{u_\theta}$.

temperature. In addition, the momentum equations are solved only once every other time sub-step therefore, even if there is an overhead introduced by the interpolation of the fields, the CPU time savings are substantial. In our case the standard simulation on the single grid $384^2 \times 768$ required for the integration of dimensionless time unit a wall-time of 36.4 minutes on 96 processors, for a total of 58.2 CPU hours using 13 GB of RAM memory. Using $\mathcal{L} = \mathcal{M} = 2$, one dimensionless time unit required a wall-time of 36.7 minutes on 48 processors for a total of 29.4 CPU hours using 4.9 GB of RAM memory.

The method becomes even more advantageous as the Prandtl number increases. In fact, from equations (3.2)–(3.3) we see that the nondimensional diffusivities of momentum and temperature, read $Re = \sqrt{Ra/Pr}$ and $Pe = \sqrt{RaPr}$, respectively. Therefore for $Pr > 1$ the momentum field smoothens while the temperature field develops sharper gradients. This results in larger Kolmogorov η_K and smaller Batchelor η_B scales that need different meshes to be properly resolved. Figure 3.5(b) shows

that at $Pr = 10$ the increased momentum diffusivity makes even the relatively coarse grid $128^2 \times 256$ adequate for the description of the momentum field. On the other hand, the same mesh is clearly too coarse for the scalar field as the ratio $Nu_{u_z\theta}/Nu_{\varepsilon_\theta}$ strongly deviates from one. This grid, however, can be used to advance the momentum and to generate a refined mesh to advect the temperature. For this case the convergence of the Nusselt numbers to the reference value is obtained for $\mathcal{M} = \mathcal{L} = 3$ that yields a computational gain around a factor 7 and a reduction of RAM memory by a factor 3.5 when compared to the reference cases using a single mesh.

Before concluding this section we point out that for all simulations we have used a refinement factor for the time step $\mathcal{L} = \mathcal{M}$. Values of $\mathcal{L}$ smaller than $\mathcal{M}$ can be used provided the CFL number for momentum is reduced so that the scalar integration remains stable; this increases the number of time steps needed to advance the simulation over the same physical time resulting in an increased computational cost. On the other hand, further increasing $\mathcal{L}$ beyond $\mathcal{M}$ does not modify the results within statistical error and empirical evidence supporting this statement can be found in table 3.3 at the end of the chapter.

3.5.2 Double diffusive convection

Double-diffusive convection (DDC) is a system where two scalars with very different diffusivities are coupled to the flow field, one of which is stabilizing and the other destabilizing. A relevant example of this system is the ocean, where the scalars are temperature ($Pr \approx 7$) and salinity ($Sc \approx 700$). The former being warmer at the top surface and cooler at the bottom stabilizes the flow while the latter has the opposite effect because a higher salinity at the top boundary results in a denser, sinking fluid. A snapshot of the flow in a DDC system for a geometry similar to RB can be seen in figure 3.10. Very sharp gradients of salinity can be observed, while the temperature field is in a quasi-diffusive state.

In this case, the computational gains associated with a multiple resolution can be even larger than before, owing to the very large Sc number of salt in water. In the simulations of Ref. [131], the temperature is discretized without further refinement, for salinity a factor $\mathcal{M}^x = \mathcal{M}^y = 3$ is used for the horizontal directions while in the vertical direction it is $\mathcal{M}^z = 2$. The anisotropic refinement is necessary to cope with the unstable salinity field that is organized into thin vertical fingers with horizontal gradients steeper than in the vertical direction (figure 3.10(a)). Also the time step is refined by a factor $\mathcal{L} = 3$ in order to maintain the integration of the salinity equation stable without penalizing the evolution of momentum and temperature.

Similar to the RB case, the results in terms of heat transfer, salinity fluxes and turbulence statistics agree within the uncertainty of 2–3% with those obtained using a single refined grid for all the variables. Nevertheless, using the present numerical approach momentum and temperature equations are solved on a mesh with $3^2 \cdot 2$ less

Figure 3.10: Instantaneous snapshot of (a) salinity ($Sc_s = 700$), (b) temperature ($Pr = 7$) and (c) vertical velocity in a DDC simulation at drivings of $Ra_s = 5 \cdot 10^7$ and $Ra_T = 10^6$. This results in $Le = 100$ and $R_\rho = 2$.

nodes than that for salinity and 3 times less frequently in time; with respect to the single grid approach this leads to a reduction of CPU time by a factor of 5 for a grid of $144 \times 144 \times 144$ with refinement factors $\mathcal{M}^x = \mathcal{M}^y = 3$ and $\mathcal{M}^z = 2$, when compared to resolving the full flow field with a grid of $432 \times 432 \times 288$. A decrease of RAM memory usage by more than 50% is also achieved by using the strategy.

Before concluding this section it is worthwhile to discuss briefly why in DDC it is possible to simulate a flow with a Schmidt number of 700 without using a refinement factor $\sqrt{Sc} \simeq 27$. The DDC equations in nondimensional form read

$$\partial_t u_i + u_j \partial_j u_i = -\partial_i p + \sqrt{\frac{Pr_S}{Ra_S}} \partial_{jj} u_i + (R_\rho \theta - S) \delta_{iz}, \quad (3.11)$$

$$\partial_t \theta + u_j \partial_j \theta = Le \sqrt{\frac{1}{Pr_S Ra_S}} \partial_{jj} \theta, \quad (3.12)$$

$$\partial_t S + u_j \partial_j S = \sqrt{\frac{1}{Pr_S Ra_S}} \partial_{jj} S. \quad (3.13)$$

where the flow parameters are defined as $Ra_T = g\beta_T \Delta T L / (\nu \kappa_T)$, $Pr_T = \nu / \kappa_T$, $Ra_S = g\beta_S \Delta S L / (\nu \kappa_S)$ and $Pr_S = \nu / \kappa_S$, $Le = Pr_S / Pr_T$ and $R_\rho = Ra_T Le / Ra_S$.

It can be noted that the diffusivity of momentum $Re = \sqrt{Pr_S / Ra_S}$ increases with Pr_S . Therefore for large enough values the flow does not fully transition to turbulence. On the other hand, the equations for salinity and temperature are linear and they can not sustain the cascade without a turbulent velocity field. This is indeed the case for the flow parameters of the present numerical example (see also Ref. [131]) where the fully turbulent three dimensional cascade cannot be achieved and a factor $\mathcal{M} = \sqrt{Sc} \simeq 27$ is not required for salinity. Even so, the multiple resolutions strategy results in a substantial gain factor in computational time and RAM memory occupation when compared against the single grid strategy.

3.6 Summary and conclusions

In this chapter we have presented a numerical strategy for the direct numerical simulation of turbulent flows with active and/or passive scalar fields without over-resolving the momentum equation and its pressure correction. This is certainly the case of flows with scalar diffusivity smaller than the kinematic viscosity (Pr or $Sc > 1$). Substantial computational time and memory occupation savings are even obtained for equally diffusive fields with Schmidt numbers of order unity. The different requirements for spatial discretization of scalars with respect to momentum originate not only from the diffusivity but also from the pressure. Its non-local effect was found to smoothen the momentum gradients and thereby reduce the resolution requirements with respect to resolving the scalar field. This scenario modifies the picture obtained from dimensional analysis that compares only the Kolmogorov and the Batchelor scales.

To reduce computational costs, a multiple resolutions strategy was developed in which momentum is discretized on a base mesh while scalars are discretized on a refined mesh. To solve the scalar diffusion–advection equation, momentum is spatially interpolated onto the refined grid by tricubic Hermitian splines. The scalar is advanced in time, and if necessary, coarsened to couple it back to the momentum equations. Due to stability constraints on the non-linear terms, the scalar is advanced in time using a refined timestep. Velocity is linearly interpolated in time for all the

intermediate timesteps. The optimal amount of substeps $\mathcal{L}$ coincides with the grid refinement factor $\mathcal{M}$, when it is isotropic, or with $\max[\mathcal{M}^i]$ when it is anisotropic.

The method was applied to Rayleigh-Bénard convection, and decoupling the grid resolutions was found to result in computational speedups of around one-and-a-half to two times for Prandtl unity, and seven for $Pr = 10$. This strategy was also applied to high Sc flows, also resulting in computational advantages of approximately five. Due to the large costs, both in operations, memory usage and in communication associated to solving the Poisson equation, we expect the gains to increase for larger grids and larger drivings. This is because the Poisson solver is the most expensive part of advancing the Navier–Stokes equations in time, and this does not scale linearly with the amount of points, while the scalar diffusion-advection equations do. We expect gains of about three to four times for RB simulations at $Pr = 1$ and $Ra = 10^{12}$ with production grids of about 1 billion points and $\mathcal{M} = 2$, planned for the future. Also- the memory consumption is heavily reduced, by a factor 2.6x with a refinement of two, and this makes some simulations possible on supercomputers with a lower memory per core and decreases the dependence on high CPU–memory bandwidth.

It is crucial that the base mesh is fine enough to correctly resolve the momentum field. Adding more refinement to the scalar mesh when the velocity grid is insufficient does not give an improvement of the quality of the results, and it might even lead to the suppression of small scales that violate energy conservation. This method could in principle be additionally applied to flows with $Pr < 1$. Obviously, in this case the velocity field should be solved on a mesh finer than that of the scalar. Although explicit numerical tests have not been attempted, we expect that the computational overhead introduced by the interpolation and coarsening of the fields overcomes the advantages produced by solving the scalar equation on a coarser mesh.

$N_x \times N_y \times N_z$	Pr	$Nu_{u_z\theta}$	Nu_{θ_w}	Nu_{ε_v}	Nu_{ε_θ}
$96 \times 96 \times 192$	1	66.8	67.2	71.0	59.7
$128 \times 128 \times 256$	1	64.9	64.6	67.1	60.5
$192 \times 192 \times 384$	1	63.2	63.2	64.1	61.3
$256 \times 256 \times 512$	1	63.8	63.6	64.6	62.5
$384 \times 384 \times 768$	1	64.0	63.6	64.3	63.2
$96 \times 96 \times 192$	10	68.4	67.9	68.6	60.8
$128 \times 128 \times 256$	10	65.5	65.2	65.7	61.0
$192 \times 192 \times 384$	10	63.2	63.0	63.1	61.0
$256 \times 256 \times 512$	10	63.6	63.9	63.6	62.4
$384 \times 384 \times 768$	10	63.4	63.9	63.4	63.3

Table 3.1: Details of grid resolution used for standard single grid runs. Simulations were run until $Nu_{u_z\theta}$ achieved 1% temporal convergence. All the simulations are performed at $Ra = 10^9$ and $\Gamma = 1$. The first column shows resolution, the second Pr , while the other four show the results of the different definitions of Nu .

$N_x \times N_y \times N_z$	$\mathcal{M}$	Pr	$Nu_{u_z\theta}$	Nu_{θ_w}	Nu_{ε_v}	Nu_{ε_θ}
$96 \times 96 \times 192$	2	1	65.5	65.4	68.4	63.0
$96 \times 96 \times 192$	3	1	65.6	65.3	69.3	69.5
$96 \times 96 \times 192$	4	1	65.4	65.4	66.0	64.9
$128 \times 128 \times 256$	2	1	64.4	64.4	66.7	66.7
$128 \times 128 \times 256$	3	1	64.4	64.4	66.9	66.9
$192 \times 192 \times 384$	2	1	63.5	63.4	62.6	64.5
$96 \times 96 \times 192$	2	10	66.6	64.8	60.0	66.0
$96 \times 96 \times 192$	3	10	64.5	64.5	63.5	63.5
$96 \times 96 \times 192$	4	10	65.4	65.0	64.2	64.2
$128 \times 128 \times 256$	2	10	64.3	64.2	64.7	63.0
$128 \times 128 \times 256$	3	10	64.6	64.1	62.2	64.8

Table 3.2: Details of grid resolution used for multiple resolutions runs. Simulations were run until $Nu_{u_z\theta}$ achieved 1% temporal convergence. All the simulations are performed at $Ra = 10^9$ and $\Gamma = 1$. The first column shows resolution, the second shows the refinement of the scalar grid in all directions, the third Pr , while the other four show the results for the different definitions of Nu . For all simulations $\mathcal{L} = \mathcal{M}$.

$\mathcal{L}$	$\mathcal{C}_{CFL}$	$Nu_{u_z\theta}$	Nu_{θ_w}	Nu_{ε_v}	Nu_{ε_θ}
1	0.6	64.5	64.6	67.3	64.0
2	1.2	64.4	64.4	66.8	63.5
3	1.2	64.6	64.3	67.2	63.5

Table 3.3: Details of the testing for the temporal multiple resolutions. Simulations were run until $Nu_{u_z\theta}$ achieved 1% temporal convergence. All the simulations are performed at $Ra = 10^9$ and $\Gamma = 1$ on a grid $128 \times 128 \times 256$ with $\mathcal{M} = 2$. The first column shows the time refinement level $\mathcal{L}$, the second the maximum CFL computed on the momentum grid, while the last four show the results of the different definitions of Nu .

4

Optimal Taylor-Couette: transition to turbulence *

Scientists work from models acquired through education and through subsequent exposure to the literature often without quite knowing or needing to know what characteristics have given these models the status of community paradigms.

—THOMAS KUHN

We numerically simulate turbulent Taylor-Couette flow for independently rotating inner and outer cylinders, focusing on the analogy with turbulent Rayleigh-Bénard flow. Reynolds numbers of $Re_i = 8 \cdot 10^3$ and $Re_o = \pm 4 \cdot 10^3$ of the inner and outer cylinders, respectively, are reached, corresponding to Taylor numbers Ta up to 10^8 . Effective scaling laws for the torque and other system responses are found. Recent experiments with the Twente turbulent Taylor-Couette (T^3C) set-up and with a similar facility in Maryland at very high Reynolds numbers have revealed an optimum transport at a certain non-zero rotation rate ratio $a = -\omega_o/\omega_i$ of about $a_{opt} = 0.33$. For large enough Ta in the numerically accessible range we also find such an optimum transport at non-zero counter-rotation. The position of this maximum is found to shift with the driving, reaching a maximum of $a_{opt} = 0.15$ for $Ta = 2.5 \cdot 10^7$. An explanation for this shift is elucidated, consistent with the experimental result that a_{opt} becomes approximately independent of the driving strength for large enough Reynolds numbers. We furthermore numerically calculate the angular velocity profiles and visualize the different flow structures for the various regimes. By writing the equations in a frame co-rotating with the outer cylinder a link is found between the local angular velocity profiles and the global transport quantities.

*Based on: R. Ostilla, R. J. A. M. Stevens, S. Grossmann, R. Verzicco and D. Lohse, *Optimal Taylor-Couette flow: direct numerical simulations*, J. Fluid Mech., **719**, 14–46 (2013).

Figure 4.1: Geometry of the Rayleigh-Bénard (left) and the Taylor-Couette systems (right). The RB system consists of two cylindrical plates, a hot one at the bottom and a cold one at the top of diameter D separated by a distance L . The top plate is at a temperature T_0 and the bottom plate is at a temperature $T_0 + \Delta$, with Δ the temperature difference between the plates. The TC system consists of two coaxial cylinders of length L . The inner cylinder has the radius r_i and the angular velocity ω_i , while the outer cylinder has the radius r_o and the angular velocity ω_o .

4.1 Introduction

Taylor-Couette (TC) flow, i.e. the flow in the gap between two independently rotating coaxial cylinders, is among the most investigated problems in fluid mechanics, due to its conceptional simplicity and to applications in process technology. It was shown by Eckhardt, Grossmann & Lohse (2007) [25] (from now on referred to as EGL07) that TC flow has many similarities to Rayleigh-Bénard (RB) convection, i.e. the thermal flow in a fluid layer heated from below and cooled from above, which will be discussed in detail below. Figure 4.1 presents a schematic of both systems, side by side.

Both RB and TC flows have been popular playgrounds for the development of new concepts in fluid dynamics. Both systems have been used to study instabilities [34, 58–61], nonlinear dynamics and chaos [62–64, 66], pattern formation [37, 67, 68], and turbulence [39, 69–73]. The reasons that RB and TC are so popular include: (i) These systems are mathematically well-defined by the Navier-Stokes equations and the appropriate boundary conditions; (ii) these are closed systems and thus exact global balance relations between the driving and the dissipation can be derived; and (iii) they are experimentally and numerically accessible with high precision, thanks to the simple geometries and high symmetries.

The analogy between TC and RB may be better seen from the exact relations

(EGL07) between the transport quantities and the energy dissipation rates. For RB flow the conserved quantity that is transported is the thermal flux $J = \langle u_z \Theta \rangle_{A,t} - \kappa \partial_z \langle \Theta \rangle_{A,t}$ of the temperature field Θ , where κ is the thermal conductivity of the flow. The first term is then the convective contribution (u_z is the vertical fluid velocity component) and the second term is the diffusive contribution. Here $\langle \dots \rangle_{A,t}$ indicates the averaging over time and a horizontal plane. In the state with lowest thermal driving there is not yet convection. Therefore $J \equiv J_0 = \kappa \Delta L^{-1}$ and the corresponding dissipation rate is $\varepsilon_{u,0} = 0$ since $\mathbf{u} = 0$. In TC flow, the conserved transport quantity, which is transported from the inner to the outer cylinder (or vice versa) is the flux $J^\omega = r^3 (\langle u_r \omega \rangle_{A,t} - \nu \partial_r \langle \omega \rangle_{A,t})$ of the angular velocity field ω , where the first term is the convective contribution with u_r as the radial fluid velocity component and the second term is the diffusive contribution, cf. EGL07. In this case $\langle \dots \rangle_{A,t}$ indicates averaging over time and a cylindrical surface with constant radial distance r from the axis. In the state with lowest driving induced by the rotating cylinders and neglecting effects from the upper and lower plates (achieved in the simulations by periodic boundary conditions in the axial direction), the flow is laminar and purely azimuthal, $u_\theta(r) = Ar + B/r$, while $u_r = u_z = 0$. This flow provides an angular velocity current J_0^ω (called J_{lam}^ω in EGL07) and a non-zero dissipation rate, see equations (4.8) and (4.16) in table 1.

The analogy between RB and TC (EGL07) is highlighted when the driving in TC is expressed in terms of the Taylor number Ta and the angular velocity ratio $a = -\omega_o/\omega_i$ of the cylinders, while the response is given by the dimensionless transport current density J^ω divided by the corresponding molecular current density of the angular velocity from the inner to the outer cylinder, called the “ ω -Nusselt number” Nu_ω . We recall that the Taylor number is defined as $Ta = \sigma(r_o - r_i)^2(r_o + r_i)^2(\omega_o - \omega_i)^2/(4\nu^2)$, or

$$Ta = (r_a^6 d^2 / r_o^2 r_i^2 \nu^2)(\omega_o - \omega_i)^2. \quad (4.1)$$

Here

$$\sigma = r_a^4 / r_g^4 \quad (4.2)$$

with $r_a = (r_o + r_i)/2$ the arithmetic and $r_g = \sqrt{r_o r_i}$ the geometric mean radii. $\omega_{o,i}$ are the angular velocities of the outer and inner cylinders, respectively; see also table 4.1 for definitions and relations.

TC flow has been extensively investigated experimentally [31, 33, 36–45] at low and high Ta for different ratios of the rotation frequencies $a = -\omega_o/\omega_i$, see the phase diagram in figure 4.3. However, up to now most numerical simulations of TC flow have been restricted to the case of pure inner-cylinder rotation [46–50], or eigenvalue study [51], or counter-rotation at fixed a [49]. Recent experiments [42–45] have shown that at fixed Ta an optimal transport is obtained at non-zero a : [43] obtained $a_{opt} = 0.33 \pm 0.05$, whereas Paoletti *et al.* [44] found $a_{opt} \approx 0.33$.

Table 4.1: Analogous relations between RB and TC flow, leading to the same effective scaling laws as derived by Eckhardt, Grossmann & Lohse [25]. In RB flow, the dimensionless control parameters are the Rayleigh number $Ra = \beta g \Delta L^3 / (\nu \kappa)$, the Prandtl number $Pr = \nu / \kappa$, and the aspect ratio $\Gamma = D/L$, where Δ is the temperature difference between the cold top and hot bottom, β the thermal expansion coefficient, g the gravitational acceleration, and κ the thermal diffusivity, see figure 4.1. The response of the system is the heat flux from the bottom to the top in terms of the molecular heat flux, known as the Nusselt number Nu . In analogy, for TC flow we define a Nusselt number Nu_ω as ratio of the total and the purely azimuthal and laminar angular velocity flow. $\tilde{\epsilon}_{u,0}$ is the dissipation in the purely diffusive state, equal to zero in RB flow, since the fluid velocity is zero and there is molecular heat transport only, while in TC flow $\tilde{\epsilon}_{u,0}$ is the purely azimuthal and laminar flow dissipation rate.

Rayleigh-Bénard	Taylor-Couette
Conserved: temperature flux	Conserved: angular velocity flux
$J = \langle u_z \Theta \rangle_{A,t} - \kappa \partial_z \langle \Theta \rangle_{A,t} \quad (4.3)$	$J^\omega = r^3 (\langle u_r \omega \rangle_{A,t} - \nu \partial_r \langle \omega \rangle_{A,t}) \quad (4.4)$
Dimensionless transport:	Dimensionless transport:
$Nu = \frac{J}{J_0} \quad (4.5)$	$Nu_\omega = \frac{J^\omega}{J_0^\omega} \left(= \frac{\tau}{2\pi L \rho J_0^\omega} \right) \quad (4.6)$
$J_0 = \kappa \Delta L^{-1} \quad (4.7)$	$J_0^\omega = \nu \frac{2r_i^2 r_o^2}{r_i + r_o} \frac{\omega_i - \omega_o}{d} \quad (4.8)$
Driven by:	Driven by:
$Ra = \frac{\beta g \Delta L^3}{\kappa \nu} \quad (4.9)$	$Ta = \frac{1}{4} \frac{\sigma (r_o - r_i)^2 (r_i + r_o)^2 (\omega_i - \omega_o)^2}{\nu^2} \quad (4.10)$
Exact relation:	Exact relation:
$\tilde{\epsilon}'_u = \tilde{\epsilon}_u - \tilde{\epsilon}_{u,0} \quad (4.11)$	$\tilde{\epsilon}'_u = \tilde{\epsilon}_u - \tilde{\epsilon}_{u,0} \quad (4.12)$
$= (Nu - 1) Ra Pr^{-2} \quad (4.13)$	$= (Nu_\omega - 1) Ta \sigma^{-2} \quad (4.14)$
$\tilde{\epsilon}_{u,0} = 0 \quad (4.15)$	$\tilde{\epsilon}_{u,0} = \frac{d^4}{\nu^3} \cdot \nu \frac{r_i^2 r_o^2}{r_a^2} \left(\frac{\omega_i - \omega_o}{d} \right)^2 \quad (4.16)$
Prandtl number:	Pseudo ‘Prandtl’ number:
$Pr = \nu / \kappa \quad (4.17)$	$\sigma = \left(1 + \frac{r_i}{r_o} \right)^4 / \left(4 \frac{r_i}{r_o} \right)^2 \quad (4.18)$
Scaling:	Scaling:
$Nu \propto Ra^\gamma \quad (4.19)$	$Nu_\omega \propto Ta^\gamma \quad (4.20)$

In this chapter we use direct numerical simulations (DNS) to study the influence of the rotation ratio a on the flow structures and the corresponding transported angular velocity flux for Ta up to $Ta = 10^8$. Our motivation is two-fold: as a first objective, we wish to further investigate the analogy between RB and TC flow by comparing the scaling laws of the global response across the different flow states. Our second objective is to study the optimal transport, which was recently observed in TC experiments [43, 44, 77], by using data obtained from DNS. In DNS we have access to the complete velocity field, which is not available in experiments, and this allows us to study this phenomenon in much more detail. At present, however, in DNS we are restricted to smaller Reynolds numbers compared to above-mentioned recent experiments.

Figure 4.2 shows the cases which are simulated, in the (Ta, a) , the $(Ta, 1/Ro)$ and the (Re_o, Re_i) parameter space. Note that a higher density of points has been used in places where the response (Nu_ω, Re_w) shows more variation (where Re_w is the wind Reynolds number defined below). All points have been simulated for fixed $\Gamma = 2\pi$ and $\eta = 5/7$ since these are very similar to the parameters of the Twente Turbulent Taylor-Couette (T³C) setup. There is a significant difference, however, as numerically we take periodic boundary conditions in the axial direction, while the T³C system is closed with solid boundaries at top and bottom which rotate with the outer cylinder.

In §4.2 we start with a description of the numerical method that has been used. In §4.3 we will discuss the validation and resolution tests that have been performed. In §4.4 the global response, in terms of Nu_ω and the wind Reynolds number Re_w , as functions of the angular velocity ratio a will be discussed. In order to understand the global system response we will analyse the coherent structures in §4.5 and the boundary layer profiles in §4.6, i.e. quantities that are difficult to analyse in experiments. This allows us to rationalize the position of the maximum in $Nu_\omega(a)$. We conclude with a brief discussion and outlook to future work in §4.7.

Brauckmann & Eckhardt [123] offer a complementary direct numerical simulation of turbulent Taylor-Couette flow: they employ a spectral code with periodic boundary conditions also in the axial direction and an aspect ratio $\Gamma = 2$ in the axial direction rather than $\Gamma = 2\pi$ as we do. Also they find a maximum in the angular velocity transport for moderate counter-rotation $a = -\omega_o/\omega_i \approx 0.4$, similar to the experiments by van Gils *et al.* [43, 77] and Paoletti & Lathrop [44] and in the present numerical simulations. So the result seems to be very robust, and at least does not strongly depend on Ta , Γ and other details of the flow. Brauckmann & Eckhardt [123] also offer an analysis of the probability density functions (p.d.f.s) of the local angular velocity fluxes in the different regimes for the inner and outer cylinder, similarly to what been done for the bulk in the experiments by Huisman *et al.* [45].

Figure 4.2: Control parameter phase space numerically explored in this chapter: (a) classical representation (Re_o, Re_i) , (b) (Ta, a) representation with $a = -\omega_o/\omega_i$, (c) $(Ta, 1/Ro)$ representation with Ro defined in equation (4.23). We fixed $\eta = 5/7$, $\Gamma = 2\pi$ and employed axial periodicity. The same color code, denoting the Taylor number, is maintained throughout the chapter. The shaded area outlines boundary conditions for which the angular momentum of the outer cylinder (L_o) is larger than the angular momentum of the inner cylinder (L_i). This causes the flow to have an overall transport of angular momentum towards the inner cylinder. In this region, the Rayleigh stability criterion applies, which states that if $dL/dr > 0$ the flow is linearly stable to axisymmetric perturbations [84].

Figure 4.3: Explored phase space (Re_o , Re_i) of TC flow with independently rotating inner and outer cylinders. Both experimental data [31, 33, 36–38, 40, 43, 44, 53] and numerical data [48–50, 55] are shown. Solid lines between markers represent a large density of experiments. The dashed lines are Esser and Grossmann’s (1996) estimate for the onset of turbulence with $\eta = 0.71$. The dark shaded area indicates the data points in the well-studied small-Reynolds-number regime of pattern formation and spatial temporal chaos (see e.g. Refs. [37, 58, 67]). The light gray area is the region shown in figure 4.2, covered by the present DNS.

4.2 Numerical method

Taylor-Couette flow was simulated by solving the Navier-Stokes equations in a rotating frame, which was chosen to rotate with $\Omega = \omega_o e_z$, where e_z is the unit vector in the direction parallel to the cylinders' axis. In this way the boundary conditions are simplified: at the inner cylinder the new boundary condition is $u_\theta(r = r_i) = r_i(\omega_i - \omega_o)$, while at the outer cylinder we have a stationary wall $u_\theta(r = r_o) = 0$. We can choose the characteristic velocity $U \equiv r_i|\omega_i - \omega_o|$ and the characteristic length scale d to non-dimensionalize the equations and boundary conditions. The characteristic velocity U can be written as

$$U = (v/d) \cdot [8\eta^2/(1+\eta)^3] \cdot Ta^{1/2}. \quad (4.21)$$

Up to a geometric factor, which is 0.810 for our choice of η , the characteristic velocity U is thus simply $Ta^{1/2}$, expressed in terms of the molecular velocity v/d . The non-dimensional variables will be labelled with a hat. In this notation, the non-dimensional inner-cylinder velocity boundary condition simplifies to: $\hat{u}_\theta(r = r_i) = (\omega_i - \omega_o)/|\omega_i - \omega_o|$. As $\omega_i - \omega_o$ is positive throughout the range covered in this work, in our coordinate system the flow geometry is simplified to a pure inner-cylinder rotation with the boundary condition $\hat{u}_\theta(r_i) = 1$. The effect of the outer cylinder is felt as a Coriolis force in this rotating frame.

The resulting Navier-Stokes equations in the rotating frame are now

$$\frac{\partial \hat{\mathbf{u}}}{\partial \hat{t}} + \hat{\mathbf{u}} \cdot \hat{\nabla} \hat{\mathbf{u}} = -\hat{\nabla} \hat{p} + \frac{f(\eta)}{Ta^{1/2}} \hat{\nabla}^2 \hat{\mathbf{u}} - Ro^{-1} e_z \times \hat{\mathbf{u}}, \quad (4.22)$$

where the Rossby number is defined as

$$Ro = \frac{|\omega_i - \omega_o|}{2\omega_o} \frac{r_i}{d} = -\frac{|1+a|}{a} \frac{\eta}{2(1-\eta)} \quad (4.23)$$

and $f(\eta)$ as

$$f(\eta) = \frac{(1+\eta)^3}{8\eta^2}. \quad (4.24)$$

Equation (4.22) is analogous to the Navier-Stokes equation for a rotating Rayleigh-Bénard system,

$$\frac{\partial \hat{\mathbf{u}}}{\partial \hat{t}} + \hat{\mathbf{u}} \cdot \hat{\nabla} \hat{\mathbf{u}} = -\hat{\nabla} \hat{p} + \left(\frac{Pr}{Ta} \right)^{1/2} \hat{\nabla}^2 \hat{\mathbf{u}} + \hat{\Theta} e_z - Ro^{-1} e_z \times \hat{\mathbf{u}}, \quad (4.25)$$

with the main difference that the Rossby number sign (carried by ω_o in equation (4.23)) is relevant in TC flow. As long as the transport of angular momentum takes place from the inner to the outer cylinder, i.e. $\omega_i > \omega_o$, Ro is always negative for counter-rotating cylinders and always positive for co-rotating cylinders. Therefore

the sign of Ro affects the flow physics, as it indicates the direction of rotation of the outer cylinder.

These equations were solved using a finite-difference solver in cylindrical coordinates. The domain was taken to be periodic in the axial direction. Coordinates were distributed uniformly in the axial and azimuthal direction. In the radial direction, hyperbolic-tangent-type clustering was used to cluster points near both walls. For spatial discretization, a second order scheme was used. Time integration was performed fractionally, using a third order implicit Runge-Kutta method. More details of the numerical method can be found in Ref. [103]. Large-scale parallelization is obtained with a combination of MPI and OpenMP directives.

In order to quantify the flow, it is useful to continue with the normalized radius $\tilde{r} = (r - r_i)/(r_o - r_i)$ and the normalized height $\tilde{z} = z/(r_o - r_i)$. As an aid to quantification, we define the time- and azimuthally-averaged velocity field as:

$$\hat{\mathbf{u}}(r, z) = \langle \mathbf{u}(\theta, r, z, t) \rangle_{\theta, t} . \quad (4.26)$$

This time- and θ -independent velocity is used to quantify the large-time-scale circulation through the wind Reynolds number:

$$Re_w = \frac{U_w d}{\nu} \text{ with } U_w = U \langle \hat{u}_r^2 + \hat{u}_z^2 \rangle_{r,z}^{1/2} . \quad (4.27)$$

As shown in equation (4.21), $U \propto Re_i - \eta Re_o$ scales as $Ta^{1/2}$; the non-dimensional transverse velocity fluctuations may or may not lead to corrections of this basic scaling.

The convective dissipation per unit mass can be calculated either from its definition as a volume average of the local dissipation rate for an incompressible fluid

$$\varepsilon_u = \frac{\nu}{2} \langle (\partial_i u_j + \partial_j u_i)^2 \rangle_{V,t} , \quad (4.28)$$

or from the global balance (EGL07)

$$\varepsilon_u - \varepsilon_{u,0} = \frac{\nu^3}{d^4} \sigma^{-2} Ta (Nu_\omega - 1) , \quad (4.29)$$

where $\varepsilon_{u,0}$ is the volume-averaged dissipation rate in the basic, azimuthally symmetric laminar flow, cf. equation (4.16).

In order to validate the code we have calculated ε_u from both equation (4.28) and equation (4.29) and checked for sufficient agreement. We define the quantity Δ_ε measuring the relative difference

$$\Delta_\varepsilon = \frac{\nu^3 d^{-4} \sigma^{-2} Ta (Nu_\omega - 1) + \varepsilon_{u,0} - \nu \langle (\nabla \mathbf{u})^2 \rangle_{V,t}}{\nu \langle (\nabla \mathbf{u})^2 \rangle_{V,t}} . \quad (4.30)$$

Re_i	Nu_ω (present study)	Nu_ω (Ref. [46])	Nu_ω (Ref. [50])
60	1.0005	1.0000	1.0000
68	1.0006	1.0000	1.0000
70	1.0235	1.0237	1.0238
75	1.0835	1.0833	1.0834
80	1.1375	1.1371	1.1372

Table 4.2: Nu_ω for low Re_i number and $Re_o = 0$, $\Gamma = 2$, $\eta = \frac{1}{2}$ and rotational symmetry of order 4.

Δ_ε is equal to 0 analytically, but will deviate when calculated numerically.

The strictest requirement for numerical convergence was that the radial dependence of $Nu_\omega(r)$ had to be less than 1%. This is a much harder condition to satisfy than torque equality at both inner and outer cylinders, which is satisfied if the Nu_ω at both cylinders is equal. Indeed, in many cases the torques were equal to within 0.01% but $Nu_\omega(r)$ was not constant within 1%, which meant either that a higher resolution had to be chosen or that the simulation had to be run for longer time. The time-average of the energy dissipation calculated locally (cf. equation 4.28) was also checked to converge within 1%; see §4.3.3 for more details.

4.3 Code validation

4.3.1 Validation against other codes at low Reynolds number

First of all, the code was validated against numerical results from Refs. [46] and [50]. This comparison was done through Nu_ω measurements at low Re_i , in the range between 60 and 80. Only a quarter of the TC system was used, assuming a rotational symmetry of order four. The aspect ratio Γ was taken as 2, the radius ratio η as 0.5. These geometrical parameters are different from the ones used in the rest of the chapter, but they are used here for validation. The resolution of the simulation ($N_\theta \times N_r \times N_z$) was taken as $32 \times 64 \times 64$, the same as in Ref. [50]. The results are shown in table 4.2. The values show a match up to three significant figures, or sometimes even higher.

For the two smallest Reynolds numbers, both references obtain the same result, while we obtain a slightly different result. This probably comes from the fact that they measure the torque directly at the inner cylinder, which we then convert to Nu_ω for comparison, while we measure Nu_ω by taking an average value of $J(r)$ and converting this to a value for the torque and thus Nu_ω . The difference between these approaches is probably the origin of the discrepancy. However, as it is very small (below 0.1%) it is not worrying.

Figure 4.4: (a) Absolute (Nu_ω) and (b) compensated ($Nu_\omega Ta^{-0.21}$) Nusselt numbers versus Ta for $\eta = 5/7$. Experiments (dots) and the present numerics agree in shape, but there is a slight shift between the data, which we attribute to the different boundary conditions in lateral direction.

4.3.2 Comparison with experiment

The code was also validated by comparing responses obtained at higher Taylor numbers versus data from Lewis & Swinney [41]. This was done through the Nusselt number for pure inner-cylinder rotation ($a = 0$) at fixed $\eta = 5/7$ and $\Gamma = 2\pi$. The overlap between the simulations and experimental data can clearly be seen in figure 4.4 for the higher Taylor number range which we have achieved with the numerics. The shift of about 5% might be attributed to the difference in both aspect ratio and boundary conditions at the top and bottom because of the vertical confinement in the experiment. As we also have an overlap at the lower Taylor number range with other numerical simulations as shown in §4.3.1, we feel sufficiently confident to proceed with our code.

4.3.3 Resolution tests

To achieve reliable numerical results, the grid's temporal and spatial resolution have to be adequate. The requirements for spatial resolution have been studied in Stevens *et al.* [150] for RB flow. There it was shown that the effect of under-resolved DNS is mainly visible in the convergence of the thermal dissipation rate $\varepsilon_\theta = \kappa \langle (\partial_i \Theta)^2 \rangle$, which in essence is the thermal Nusselt number Nu . The kinetic dissipation rate $\varepsilon_u = \nu \langle (\partial_i u_j)^2 \rangle$ turned out to be less sensitive to under-resolution. We note that even when the kinetic dissipation rate has converged within 1% the simulation can still be under-resolved. It is important to have grid lengths in each direction of the order of the local Kolmogorov or Batchelor lengths.

In TC the corresponding fields to Θ and u are the azimuthal velocity u_θ and the

perpendicular components u_r and u_z , respectively. But these are more closely related by the Navier-Stokes equations than the u , Θ fields in RB. Therefore we tested the grid spatial resolution at $a = 0$ by calculating Nu_ω beyond onset of Taylor vortices ($Nu_\omega > 1$) and Δ_ε from equation (4.30), which analytically is equal to 0, checking the (relative) difference between the transport (Nu_ω) and the dissipation rate (ε_u). Both were done for different grid resolutions with increasing Taylor number. For all these simulations we took $\Gamma = 2\pi$, $\eta = 5/7$ and $a = 0$. The results are shown in table 4.3.

Spatial convergence required more grid points than initially expected as satisfying the torque balance alone is a necessary but not a sufficient condition for grid resolution independence. Figure 4.5(a,b) shows graphs of the radial dependence of $Nu_\omega(\tilde{r})$ at $Ta = 2.44 \cdot 10^5$ ($Re_i = 400$) for an under-resolved case ($100 \times 50 \times 50$, $Nu_\omega = 2.70845$), a reasonably resolved case ($200 \times 100 \times 100$, $Nu_\omega = 2.76208$) and an extremely well resolved reference case ($300 \times 150 \times 150$, $Nu_\omega = 2.77855$). Nu_ω should not be a function of the radius as mentioned previously, but numerically it does show some dependence. For the under-resolved case we can see that the torque balance is satisfied very well (0.06%), even if other criteria are not satisfied, E.g. the peak-to-peak variation of Nu_ω is approximately 1% and the relative error in comparison to the reference case is 2.5%. The graph also shows that taking the value of Nu_ω at one of the cylinders gives a higher result for the transport current than taking the radial mean.

Figure 4.5(c,d) shows the same plots but for $Ta = 1.91 \cdot 10^6$ ($Re_i = 1120$) and the three cases: under-resolved ($192 \times 96 \times 96$, $Nu_\omega = 4.8354$), reasonably resolved ($256 \times 128 \times 128$, $Nu_\omega = 4.4600$) and extremely well-resolved reference ($384 \times 192 \times 192$, $Nu_\omega = 4.4776$). For this Taylor number, the under-resolved case shows a smaller deviation of Nu_ω from the mean value and the torque difference in comparison to the lower Taylor number case. However, the discrepancy in the mean value of Nu_ω between the under-resolved and the reference case is much larger (7.9%). For this Ta the value of Nu_ω at the cylinder walls is larger than the average value of Nu_ω , too.

If we look at the $Nu_\omega(r)$ profiles at given Taylor numbers Ta , they show similar radial dependences, whose magnitudes decrease with increasing resolution. However, the shape of this dependence is different for both Taylor numbers. The peaks of $Nu_\omega(r)$ are located close to the boundaries, indicating that they are probably produced by some boundary layer features and are not a systematic bias of our solver.

According to EGL07, dissipation should be equal, irrespective of the way in which it is calculated, either directly from its definition or indirectly via the Nusselt number balance, cf. (4.29). Stevens *et al.* [150] also mentioned the importance of the corresponding equality in RB flow, especially for low values of Pr , as a way to ensure that the flow field is sufficiently resolved and that the gradients are captured adequately. Under-resolving a flow in Taylor-Couette flow will result in a value of Δ_ε

Re_i	Ta	$N_\theta \times N_r \times N_z$	Nu_ω	Δ_ε	Case
160	$3.90 \cdot 10^4$	$128 \times 64 \times 64$	1.86927	0.0159	R
160	$3.90 \cdot 10^4$	$256 \times 128 \times 128$	1.85562	0.0074	S
260	$1.03 \cdot 10^5$	$160 \times 80 \times 80$	2.40536	0.0215	R
260	$1.03 \cdot 10^5$	$256 \times 128 \times 128$	2.40216	0.0322	S
400	$2.44 \cdot 10^5$	$100 \times 50 \times 50$	2.70845	0.0392	U
400	$2.44 \cdot 10^5$	$200 \times 100 \times 100$	2.76208	0.0102	R
400	$2.44 \cdot 10^5$	$300 \times 150 \times 150$	2.77855	0.0062	S
680	$7.04 \cdot 10^5$	$256 \times 128 \times 128$	3.49816	0.0147	R
680	$7.04 \cdot 10^5$	$256 \times 128 \times 128$	3.49816	0.0147	T
680	$7.04 \cdot 10^5$	$384 \times 192 \times 192$	3.51268	0.0056	S
1120	$1.91 \cdot 10^6$	$192 \times 96 \times 96$	4.83540	0.0949	U
1120	$1.91 \cdot 10^6$	$256 \times 128 \times 128$	4.46000	0.0174	R
1120	$1.91 \cdot 10^6$	$384 \times 192 \times 192$	4.47765	0.0065	S
1600	$3.90 \cdot 10^6$	$300 \times 144 \times 144$	5.42553	0.0216	R
1600	$3.90 \cdot 10^6$	$432 \times 216 \times 216$	5.37264	0.0063	S
2500	$9.52 \cdot 10^6$	$384 \times 192 \times 192$	6.42160	0.0168	R
2500	$9.52 \cdot 10^6$	$641 \times 321 \times 321$	6.34068	0.0078	S
3960	$2.39 \cdot 10^7$	$641 \times 321 \times 321$	7.46617	0.0161	R
5600	$4.77 \cdot 10^7$	$800 \times 400 \times 400$	8.76601	0.0166	R
5600	$4.77 \cdot 10^7$	$800 \times 400 \times 400$	8.78178	0.0167	T
8000	$9.75 \cdot 10^7$	$1024 \times 500 \times 512$	10.4360	0.0170	R

Table 4.3: Resolution tests for $\eta = 0.714$ and $\Gamma = 2\pi$. The columns display the inner Reynolds number, the Taylor number, the resolution employed, the calculated Nu_ω , the relative discrepancy Δ_ε between the two different ways of calculating the energy dissipation, and the ‘case’: (U)nder-resolved, (R)esolved, (S)patial error reference and (T)ime error reference, for which the maximum CFL has been halved. The resolution is sufficient for all cases, as the variations are small. Δ_ε turns out to be positive; thus the code gives for the dissipation rate a smaller value for the derivatives-squared-based definition than for the Nu_ω -based balance expression.

which is too large in magnitude. This can be seen in Table 4.3 for the under-resolved simulations at $Ta = 2.44 \cdot 10^5$ and $Ta = 1.91 \cdot 10^6$. But as was elaborated at the beginning of this subsection, we should consider the convergence of Δ_ϵ towards zero (becoming smaller than any chosen threshold) as a necessary but not as a necessarily sufficient way to ensure grid convergence. Besides Δ_ϵ being small enough, at least $Nu_\omega(r)$ must also be converged.

Sufficient temporal resolution is achieved by using an adaptive time step based on a CFL criterion. If the CFL number is too large, the code destabilizes and the velocities grow beyond all limits. The code used for the simulations employs a third order Runge-Kutta time marching algorithm, with a critical CFL number of $\sqrt{3}$. The implicit factorization of the viscous terms may lower this value, so a maximum CFL of 1.4 was taken for safety. This gives a typical time step of $\Delta t \sim 10^{-3}$ for the most refined grid, $1024 \times 500 \times 512$ at $Ta = 10^8$, and $\Delta t \sim 10^{-1}$ for the coarsest grid, $128 \times 64 \times 64$ at $Ta = 3.90 \cdot 10^4$.

In order to verify that a CFL number of 1.4 is indeed small enough to ensure that our results are independent of the time step used we performed two simulations with a maximum CFL of 0.7 (see table 4.3, case T), one in the time-independent regime and one in the time-dependent regime. In the time-independent regime the results obtained with this lower CFL number are identical to the results obtained when using a maximum CFL of 1.4. In the time-dependent regime the difference between the results obtained with maximum CFL of 0.7 and 1.4 is less than the statistically significant variations.

4.3.4 Dependence on initial conditions

For the lower Taylor numbers the flow was started from rest ($\mathbf{u} = 0$). The Taylor vortices start forming within a couple of revolutions. After enough time, a steady state with three pairs of Taylor vortices was reached. However, the simulations can also be started from non-resting conditions. Depending on these conditions a different number of vortex pairs can arise. This has a strong influence on both the global and the local response of the system. Once the vortices have formed, they are persistent in time during the simulation. Therefore, it is possible to bias a simulation through the initial conditions to have a higher or lower amount of vortex pairs, which results in a different response.

Although the importance of these coherent structures gets smaller and smaller with increasing Ta (§4.5), at lower Ta the number of vortex pairs must be fixed to determine the response. Brauckmann & Eckhardt [123] present an exhaustive study of the effects of the vortex wavelength and number of vortices on Nu_ω for pure inner-cylinder rotation. The study comes to two conclusions, namely that different vortex wavelengths can lead to variations of Nu_ω of up to 20% at low values of Ta and variations of up to 10% in Nu_ω at higher values of Ta and that for a fixed vortex

Figure 4.5: Radial dependence of $Nu_\omega(\tilde{r})$ for three different grid resolutions (see legends). (a,b) $Ta = 2.44 \cdot 10^5$ and (c,d) $Ta = 1.91 \cdot 10^6$. The figures on the left, (a,c), show the absolute values, an error bar indicating a 1% error for reference. The resolved cases lie within this error bar. The figures on the right, (b,d), show the curves normalized by their average value to compare the radial fluctuations of Nu_ω .

Figure 4.6: Dependence of $Nu_\omega - 1$ versus Ta for pure inner-cylinder rotation on the number of vortices. Although the values of Nu_ω are different, the scaling behavior with Ta is the same. The error bar indicates a 5% difference.

wavelength, the number of vortex pairs has no significant effect on Nu_ω .

The study of this parameter will not be repeated here. Instead, a base state, with three pairs of vortices, $N_{pairs} = 3$, has been chosen for all the simulations in this chapter. This choice is made to keep the aspect ratio of the vortices as close to 1 as possible ($\lambda_z = \Gamma/N_{pairs} \approx 2.10$). The effect of having 3 or 4 vortex pairs on the response for selected Ta and pure inner-cylinder rotation is shown in figure 4.6. A fixed offset can be seen between the two curves. It is important to note that this effect is different from the effect caused by neutral surface stabilization, which occurs when the vortices cannot penetrate the whole flow, and the number of vortices is changed as a result. That will be featured in more detail in §4.6.5.

4.4 Global response

In this section, the global response of the Taylor-Couette system is shown across the parameter space. First the onset of Taylor vortices is analysed. Then the scaling laws are revealed for pure inner-cylinder rotation. Finally, the effect of the outer-cylinder rotation on the scaling laws is investigated and an optimum value of Nu_ω as a function of a for given Ta is found, as has been reported for large Ta from experiments, cf. Ref. [77].

Figure 4.7: System responses $Nu_\omega - 1$ (left) and Re_w (right) as a function of Ta for pure inner-cylinder rotation near the onset of Taylor vortices. Onset in the present DNS occurs at $Ta \approx 1020$.

4.4.1 Transport Nu_ω and wind Re_w for pure inner cylinder rotation

The global response of the system is quantified through Nu_ω and Re_w . These two quantities measure two different flow responses. Nu_ω quantifies the transport of angular velocity and Re_w the “wind”, i.e. the additional velocity on top of the azimuthal flow. For the purely laminar-azimuthal flow $Nu_\omega = 1$ by definition, and $Re_w = 0$ as this laminar flow only has an azimuthal velocity component.

First of all we analyse how the onset of Taylor vortices is reflected in the global response quantities Nu_ω and Re_w . $Nu_\omega - 1$ is the additional transport of angular velocity on top of the laminar transport and the wind is the fluid motion on top of the purely laminar-azimuthal flow. Figure 4.7 shows the numerically calculated $Nu_\omega - 1$ and Re_w as functions of Ta close to onset of the Taylor vortex state. The critical Taylor number (Ta_c) for the onset of Taylor vortices is calculated to be around $Ta = 1020$ for our value of η . This DNS value can be compared with Ta_c as obtained from the analytical approximation of Esser & Grossmann [152], which is $Ta = 1038$ for the present η . The agreement is within 1.6%. Later on, we shall use these analytically calculated onset Taylor numbers.

After the Taylor vortices have appeared in the system, they are the dominating feature of the flow for several decades of Ta . Figure 4.8(a,b) shows the response of the system with increasing Taylor number in the case of resting outer cylinder and pure inner-cylinder rotation. We plot Nu_ω versus $Ta - Ta_c$ rather than versus Ta as it then shows a better scaling for the points at low Ta .

There seems to be a clear change in the scaling law of Nu_ω versus Ta , but not so in the scaling law for the wind Re_w (figure 4.8(c,d)). This change occurs between $Ta = 1.91 \cdot 10^6$ and $Ta = 3.90 \cdot 10^6$ and has been seen in other numerical

Figure 4.8: (a,b) The system response as a function of Ta with best fit lines for pure inner-cylinder rotation, $Nu_\omega - 1$ in (a) and Re_w in (b). (c,d) The data of (a,b) compensated by the scaling factor to test the quality of the scaling behavior. (e) The functional relation between the two responses $Nu_\omega - 1$ and Re_w . (f) The temporal dependence $Nu_\omega(t)$ for three different Ta . The time dependence can be seen to set in between $Ta = 1.91 \cdot 10^6$ and $Ta = 3.90 \cdot 10^6$, just where we see the change in the effective scaling in (a,c). The analytical approximation of Esser & Grossmann [152] is used, i.e. $Ta_c = 1038$ for the present η .

simulations too [47]. $Nu_\omega - 1$ scales as $(Ta - Ta_c)^{0.34}$ for $Ta < 2 \times 10^6$ and as $Nu_\omega - 1 \sim (Ta - Ta_c)^{0.21}$ for $Ta > 2 \times 10^6$. We attribute the change in scaling to the changes in the coherent flow structures that affect the angular velocity transport but not the global wind amplitude. As will be discussed later in detail, we expect coherent flow structures to lose importance for increasing Ta , see §4.5. We note already here that although the loss of influence of coherent structures in RB flow (for $Pr = 1$) and in TC flow sets in at similar values of Ra and Ta respectively, i.e. around 10^7 [153], there is a large difference in the shear Reynolds numbers Re_s of the boundary layers in these two systems. This will be discussed in the following section.

Re_w measures the amplitude (strength) of the Taylor vortices, which persist at long time scales. The non-dimensional characteristic speed of these vortices remains approximately constant with Ta , namely about 5-6% of the inner-cylinder velocity u_i , throughout the whole Taylor number range considered, and that is why we see a direct scaling law of $Re_w \sim U \sim (Ta - Ta_c)^{1/2}$, cf. (4.21).

The mutual functional dependence of the two responses Nu_ω and Re_w is presented in the bottom left panel of figure 4.8. As expected from figures 4.8(a,b), the relation between Re_w and Nu_ω also shows the change in the scaling. We interpret this as follows. Before the change, the Taylor vortices are mainly responsible for the additional transport. Beyond the change, some short-time-scale fluctuations appear, indicating other structures, which disrupt the flow and finally become its dominating features, while the Taylor vortices lose importance. In order to see these time scales, we show the temporal dependence of $Nu_\omega(t)$ in figure 4.8(f). The Nusselt number shows almost no time dependence for lower Taylor numbers. But it shows two different time scales at $Ta = 3.90 \cdot 10^6$. The short time scale gains much more importance for the highest Taylor numbers, causing fluctuations of about 10%.

4.4.2 The effect of outer cylinder rotation and optimal transport

In this section the effect of the outer-cylinder rotation on the global responses Nu_ω and Re_w will be studied. This effect is felt by the flow as a Coriolis force (cf. equation (4.22)), so plots in this section will be versus $Ro^{-1} \propto -a/|1+a|$ with $a = -\omega_o/\omega_i$:

$$Ro^{-1} = \frac{2d\omega_o}{U} = -2 \frac{1-\eta}{\eta} \frac{a}{|1+a|} = 2 \frac{1-\eta}{\eta} \frac{\omega_o}{|\omega_i - \omega_o|}. \quad (4.31)$$

The inverse Rossby number Ro^{-1} runs from $2\eta/(1+\eta)$ to -1 if ω_o runs from $\eta^2\omega_i$, the inviscid Rayleigh-line in the first quadrant of the (Re_o, Re_i) -plane, to $-\infty$. It is useful to note that for a given Ta constant Ro^{-1} means constant a and vice versa. Thus a seems the preferable parameter as it is more direct; Ro^{-1} will be used only when it provides a clear advantage in visualization or later in the chapter when we will trace back the occurrence of the maximum to the Navier-Stokes equation.

Figure 4.9: $Nu_\omega(Ta, 1/Ro)$. The black line shows the position of Ro_{opt}^{-1} , the value of Ro^{-1} for which Nu_ω is maximum at a given Taylor number

Figure 4.9 shows the complete results for Nu_ω as a function of Ro^{-1} and Ta . In order to better quantify the results from figure 4.9, cross sectional cuts are taken. By taking cross-sections of constant a , scaling laws can also be discovered for non-zero values of a , i.e. for co- and counter-rotation $\omega_o \neq 0$. This is shown in figure 4.10 for five different values of a , two under co-rotation $\omega_o > 0$, two under counter-rotation $\omega_o < 0$, and as reference case $\omega_o = 0$.

For counter-rotating cylinders and Taylor numbers below that at which the change in scaling happens, a universal scaling of approximately $Nu_\omega - 1 \sim (Ta - Ta_c)^{0.34}$ is seen. However, the change in scaling and its exponent happens at a lower Ta for $a = 0$, while the scaling prevails to larger Ta for the other a (0.2 and 0.4). The scaling is different for co-rotating cylinders, and is approximately $Nu_\omega - 1 \sim (Ta - Ta_c)^{0.27}$. $Ta_c = 1038$ has been subtracted as done previously so that the scaling is not lost for the first points.

The time independence of Nu_ω is broken for much smaller Ta if the outer cylinder is rotating. For both co- and counter-rotating cylinders time dependence can be noted to set in at Ta as low as 10^5 . Also, the scaling of Nu_ω with Ta is maintained throughout a much larger range of the Taylor number. Therefore, the breakdown of time independence can no longer be associated with to the change in scaling, as one could conclude when only considering pure inner-cylinder rotation, where the loss of

Figure 4.10: $Nu_\omega - 1$ versus $Ta - Ta_c$ for two values of co-rotating a (a) and pure inner-cylinder rotation and two values of counter-rotating a (b). Compensated plots are shown in (c,d). Numerical uncertainties are less than 1%. Ta_c depends on a and is determined respectively from the analytical approximation [152]; for $a = 0$ it is $Ta_c = 1038$.

time independence and change in scaling happened at about the same Ta .

Cross-sections of constant Ta are shown in figure 4.11. $Nu_\omega - 1 = 0$ indicates points for which the flow is purely laminar-azimuthal. For co-rotating cylinders, the maximum value of Ro^{-1} reaches the inviscid Rayleigh stability line $Ro_{Ra}^{-1} = 0.833$ for even the lowest values of Ta . On the other hand, the minimum Ro^{-1} which destabilizes the laminar state can be seen to decrease (become more negative) with increasing Ta .

In order to better visualize the results it is useful to define a normalized Nusselt number as $\bar{Nu}_\omega = (Nu_\omega(Ro^{-1}) - 1)/(Nu_\omega(Ro^{-1} = 0) - 1)$, which allows easier visualization of the dependence of Nu_ω on Ro^{-1} across the Ta range of interest. The numerator of $\bar{Nu}_\omega$ goes to zero, if Ro^{-1} becomes too large or too small, i. e. reaches the stability lines (where Nu_ω becomes 1, since the flow is laminar-azimuthal in the stable ranges), while the denominator is always larger than zero, as long as $Ta > Ta_c$.

If Ta is large enough the shape of the graph resembles two straight lines from the

Figure 4.11: $Nu_\omega - 1$ as a function of Ro^{-1} across the available Ta range. The advantage of plotting versus Ro^{-1} can be seen as some of the graphs show piecewise linear behaviour and the value of Ro^{-1} at the inviscid Rayleigh stability limit $a = -\eta^2$ (denoted as Ro_{Ra}^{-1}) is $Ro_{Ra}^{-1} = 0.833$, and appears for all Ta .

Figure 4.12: $\overline{Nu}_\omega$ versus Ro^{-1} . The shift of Ro_{opt}^{-1} with increasing Ta can be appreciated here.

Figure 4.13: Location of the optimal transport (hollow dots) in the (a) (Re_o, Re_i) phase space of figure 4.2(a), and (b) the $(Ta, 1/Ro)$ phase space of figure 4.2(c). In (a), the solid blue line represents the onset of instability according to Esser & Grossmann [152] and the thin dashed line is the line of equal distance to the left and right branch of the Esser-Grossmann instability line. The solid line is the bisector of the Rayleigh instability line ($a = -\eta^2$) and the line of pure outer-cylinder rotation ($a = \infty$).

maximum value of Nu_ω to Ro_{Ra}^{-1} and Ro_{min}^{-1} . These straight lines have already been seen when plotting Nu_ω versus a slightly different version of Ro^{-1} in Ref. [44] and are the reason we chose to plot Nu_ω versus Ro^{-1} in this section.

For smaller Ta an optimum transport for co-rotation can be seen, i. e., for positive values of Ro^{-1} . This holds for Ta less than the discussed change of the scaling behaviour of Nu_ω versus Ta . For the two lowest values of Ta the deviation of Ro^{-1} beyond 0 may still be within numerical uncertainties. However, Ro_{opt}^{-1} is definitely positive for $Ta = 7.5 \cdot 10^5$ and $Ta = 2 \cdot 10^6$ and seems to fit with the piecewise linear shape of the graphs. At around $Ta = 4 \cdot 10^6$, which is beyond the change in scaling, the maximum $\overline{Nu}_\omega$ begins to drift to negative Ro^{-1} , i. e., towards counter-rotation. This can be seen in figure 4.12. In figure 4.13 we plotted the positions of the optimum transport in the (Re_o, Re_i) phase space. Clearly, the curve does not have equal distance to the two instability branches of the Esser-Grossmann approximation, as was speculated in Ref. [77]. Another feature of the drift of Ro_{opt}^{-1} is the following: While the curve of $\overline{Nu}_\omega$ has a prominent peak at Ro_{opt}^{-1} for values of Ta of around $10^6 - 10^7$, this turns into a plateau for the highest value of Ta and Ro_{opt}^{-1} becomes hard to identify. For the highest value of Ta , the lower boundary Ro_{min}^{-1} already is beyond our parameter range of negative Ro^{-1} .

Experiments [43, 44] have found an optimum transport $a_{opt} \approx 0.33$, corresponding to $Ro_{opt}^{-1} \approx -0.20$ for Taylor numbers of the order of 10^{12} . Thus the position of the maximum shifts for higher Taylor numbers.

Figures 4.11 and 4.12 show some anomalous jumps in the graph around $Ro^{-1} \approx -0.2$ which corresponds to $a \approx 0.3$. These are caused by different vortical states as mentioned in §4.3.4. These may be present even if the simulations are started from the same initial conditions for different values of Ro^{-1} and Ta . If the number of vortices is higher, the vortices become stronger, and the value of Re_w , which measures their strength, also becomes higher. Since Nu_ω is monotonically related to Re_w , it also increases. We will further analyse this multi-vortex state in §4.6.5.

4.5 Characterization of the flow state

In this section we will analyse two characteristic Taylor number ranges in TC flow. The first, lower one, is the range in which the importance of coherent flow structures is lost, since these have become too small in size. In §5.4 we have observed a change in the scaling law for the angular velocity flux from $Nu_\omega \sim (Ta - Ta_c)^{0.34}$ to $Nu_\omega \sim (Ta - Ta_c)^{0.21}$. Although the Taylor number for this change coincides with the onset of time dependence for pure inner-cylinder rotation, when adding outer-cylinder rotation the onset of the time dependence is much earlier, and a transition in the scaling laws cannot even be seen.

Therefore, the onset of time dependence cannot be linked satisfactorily to the change in the Nu_ω -scaling. Another way of explaining this change is by estimating when the spatially coherent flow structures lose influence because their size becomes too small. We do this by defining an average, global coherence length in terms of the Kolmogorov length scale [153] resulting from the volume-averaged dissipation rate:

$$\ell_c = 10\eta_K = 10(v^3/\varepsilon_u)^{1/4} = 10d(\sigma^{-2}Ta(Nu_\omega - 1) + \hat{\varepsilon}_{u,0})^{-1/4}. \quad (4.32)$$

where equation (4.29) has been used for the second equality. We compare the global coherence length ℓ_c with the gap width d or, equivalently with the extent of the remnants of the Taylor vortices, whose length can also be estimated as d , since they tend to have a square aspect ratio.

Figure 4.14(a) shows the variation of ℓ_c/d with increasing Taylor number. The loss of importance of coherent structures happens in the range where ℓ_c/d is between 0.1 – 0.5, corresponding to $Ta \approx 10^6 - 10^7$. It is just within this Ta range where the change in the Nu_ω scaling law occurs. The graph is consistent with that change taking place at approximately the same Ta for different values of a , which is what is seen in figure 4.10. This transition is further analysed in §4.6.1.

A second characteristic Taylor number is connected with the shear instability of the boundary layers (BL). Here the laminar Prandtl-type BLs become turbulent. Beyond that Ta value the flow is fully turbulent throughout and this state is known

Figure 4.14: (a) ℓ_c/d versus Ta for $\eta = 5/7$ and three values of a as shown in the legend for (b). The shaded region indicates the range of decreasing importance of the coherence structures. (b) The shear Reynolds number Re_s of the boundary layers versus Ta for the same three a values. In the shaded region we expect a shear instability of the boundary layers, in which the Prandtl laminar boundary layers become turbulent. The TC system then is in the so-called ‘ultimate’ regime, cf. Ref. [78].

as the ultimate state (cf. Ref. [78]). This happens if the BL shear Reynolds number becomes $Re_s > 250 - 420$, where Re_s is calculated from the shear velocity U_s as [77]:

$$Re_s = \frac{U_s \delta}{\nu} = a_{PB} \sqrt{Re_i - Re_w}. \quad (4.33)$$

The empirical constant a_{PB} is taken as 2.3 as in Ref. [77] for Prandtl-Blasius-type boundary layer in TC flow. This value is obtained by a fit to experimental data, detailed in Ref. [77].

For RB flow this transition is expected at $Ra \approx 10^{14}$ [72, 78, 154], while figure 4.14(b) shows that the transition in TC is expected for $10^8 \lesssim Ta \lesssim 10^9$. Recently, experiments have confirmed the ultimate scaling both for Nu_ω and Re_w . Huisman *et al.* [45] have shown that in TC flow $Nu_\omega \sim Ta^{0.38}$ and $Re_w \sim Ta^{0.50}$ when $Ta \gtrsim 10^9$. A confirmation of the analogy between RB and TC is obtained by the high- Ra experiments by He *et al.* [76] as they measured that $Nu \sim Ra^{0.38}$ and $Re_w \sim Ra^{0.50}$ for $Ra \gtrsim 5 \times 10^{14}$. These measured scaling exponents agree exactly with the predictions by Grossmann & Lohse [78]. In contrast to the experiments of Refs. [43, 77], in our present numerical simulations the ultimate state has not yet been achieved, as clearly seen from figure 4.14(b).

4.6 Local results

This section contains the analysis of local results. For convenience we skip in this section the “hat” on the dimensionless flow field variables, but still understand them

Figure 4.15: The shaded areas show where $\ell_c/d < 0.1$ for pure inner-cylinder rotation and Taylor numbers of (a) $Ta = 1.91 \cdot 10^6$ and (b) $Ta = 9.52 \cdot 10^6$. In these areas the breakup of the coherent structures is likely to occur. The arrows indicate $\mathbf{u}$ in magnitude and direction.

as being dimensionless. The angular velocity profiles are shown and the ratio of the BL thickness is calculated and compared with the theory of EGL07. The angular velocity profiles reflect the interplay of bulk and boundary layers and that of the mean flow and added perturbations. The importance of convective versus diffusive transport is quantified through the bulk slope of the angular velocity profile, and again we will find a maximum as function of a , which we will connect with the maximum in the angular velocity transport Nu_ω .

4.6.1 Local coherence length and vortex characterization

Figure 4.15 shows the local coherence length calculated from the local dissipation in analogy to equation (4.32). This figure adds details on where we expect the Taylor vortices to break down. At low Taylor number, the local coherence length is smaller than 0.1 only very near to the wall, where the highest local dissipation takes place. With increasing Taylor number, the highest local dissipation still is near the wall, but the dissipation rate is large enough in the whole domain to break up the dominance of the coherent structures, even if they do not fully disappear but become small enough.

From figures 4.14-4.15 we expect coherent structures to break up at Taylor numbers in the range of $Ta \approx 10^6 - 10^7$. This may at first sight contradict the earlier finding that the scaling of the wind remains constant across the whole Taylor number range studied (cf. §5.4), especially as the characteristic wind velocity is defined from a time-averaged field. One might expect that the perturbations destroy the large scale structures and as a consequence completely modify the wind. In order to analyse this transition in more detail, we investigate the instantaneous velocity fields before and after the breakdown of coherence. For this, vortices will be characterized employing the so-called λ_2 -criterion [155].

Figures 4.16(a) and 4.16(b) show full three-dimensional isosurfaces of λ_2 for two Taylor numbers, respectively $Ta = 7.04 \cdot 10^5$ before the transition, and $Ta = 2.39 \cdot 10^7$, after the transition. Figure 4.16(c,d) shows an azimuthal-cut contour plot of λ_2 for two Taylor numbers. The instantaneous “wind” is superimposed. It is important to note that for figure 4.16(c) time dependence has not yet set in, so the instantaneous and mean velocity fields are indistinguishable. In this panel we can see that the lowest values of λ_2 are located in the centre of the gap, coinciding with an area of positive u_r wind and almost no u_z wind. Structures of negative λ_2 occupy a significant portion of the space between the cylinders.

On figure 4.16(d), we can see a different picture. The structures of negative λ_2 are now much smaller, and no longer occupy a significant region of the domain, unlike in (c). These structures are also in a different place: clustered near the inner cylinder, from where they seem to originate. The instantaneous wind is superimposed on the contour plot. A similar structure as the one in (c) is seen, indicating that even though the coherent structures are no longer dominant, the underlying wind behaves in a similar manner. Indeed, once the velocity field is averaged in time, the large scale Taylor vortices are recovered. This result is consistent with the findings reported by Dong [48].

4.6.2 Angular velocity profiles

The angular velocity ω is the transported quantity in Taylor-Couette flow. Analysing the dependence of the $\omega(r)$ profiles on the driving parameters Ta and a seems useful to understand how transport takes place in the flow; $\omega(r)$ profiles are shown in figures 4.17 and 4.18. Beyond the breakdown of the laminar, purely azimuthal flow, three distinct regions in the gap appear. These are the inner and outer boundary layers (BL), in which the transport mechanism is dominantly diffusive, and a flatter bulk zone, in which the transport mechanism is dominantly convective, see figure 4.20(b) for a sketch, in which we approximate the profile of the mean azimuthal velocity $\langle \bar{u}_\theta \rangle_z$ by three straight lines, one for each boundary layer and one for the bulk. For the boundary layers, we calculate the slope of the lines by fitting (least mean square) a line through the first three computational grid points. For the bulk, we first force the

Figure 4.16: (a,b) Isosurfaces of λ_2 for pure inner-cylinder rotation and Taylor numbers of: (a) $Ta = 7.04 \cdot 10^5$ and (b) $Ta = 2.39 \cdot 10^7$. (c,d) Corresponding contour plots of λ_2 truncated for $\lambda_2 > 0$. The darkest shading represents the minimum value of λ_2 in each plot and white represents $\lambda_2 \geq 0$. The arrows indicate $\mathbf{u}$ in magnitude and direction. The topography of the negative- λ_2 areas changes from large coherent regions in the centre of the gap -indicating Taylor vortices- to small regions near the inner cylinder -indicating hairpin vortices. The underlying wind, however, remains unchanged.

Figure 4.17: The t, θ and also z -averaged angular velocity $\langle \bar{\omega} \rangle_z$ versus $\tilde{r}$ for increasing a and four Taylor numbers: (a) $1.91 \cdot 10^6$; (b) $3.90 \cdot 10^6$; (c) $9.52 \cdot 10^6$; (d) $2.39 \cdot 10^7$. The boundary layers of the ω profiles become thinner around slight counter-rotation than they are for higher values of a as well as for co-rotation as an indication of increased transport. For strong co-rotation $a = -0.4$ as well as high a , i. e., strong counter-rotation, at low Ta there is not yet a flat bulk zone since the flow is not yet turbulent enough.

line to pass through the inflection point of the profile (the nearest grid point). Then, its slope is taken from a least-mean-square fit using two grid points on either side of this inflection point. The respective boundary layer line will cross the bulk line at a point which then defines the thickness of that boundary layer.

In the next two subsections we will discuss the BL and bulk regimes separately.

4.6.3 Angular velocity profiles and resulting boundary layer thicknesses

With increasing Ta , in order to accommodate the increasing angular velocity transport, the boundary layers become thinner and the ω -slopes steeper. What is striking is the strong asymmetry between the inner BL and the outer BL, which is much

Figure 4.18: $\langle \bar{\omega} \rangle_z$ versus $\tilde{r}$ for increasing Ta and four values of a : (a) $a = -0.2$; (b) $a = 0$; (c) $a = 0.2$; (d) $a = 0.4$. The boundary layers for the ω transport become thinner with increasing Ta indicating increased angular velocity transport. The boundary layers get steeper with increasing Ta and the bulk region becomes more extended. For low values of Ta the bulk region is rather small. The $\langle \bar{\omega} \rangle_z$ profile in the center remains approximately unchanged with Ta .

thicker. Figure 4.19 shows the ratio between the outer and the inner BL thicknesses versus rotation ratio a for four Taylor numbers. This asymmetry is a consequence of the exact relation $\partial_r \langle \omega \rangle|_o = \eta^3 \partial_r \langle \omega \rangle|_i$, obtained from the r -independence of J^ω , (cf. equation 4.4); the slope at the inner cylinder is a factor of η^{-3} larger than at the outer one and thus the outer boundary layer is much more extended than the inner one. Since for the present Ta -range the shear Reynolds number Re_s is still below the threshold value range for the transition to turbulence in the boundary layers (cf. §4.5), we can compare the numerically obtained boundary thickness ratio with the one obtained by EGL07, which had been derived in the spirit of the Prandtl-Blasius (i.e. laminar-type) boundary layer theory, namely

$$\frac{\lambda_\omega^o}{\lambda_\omega^i} = \eta^{-3} \frac{|\omega_o - \omega_{bulk}|}{|\omega_i - \omega_{bulk}|}. \quad (4.34)$$

Here the value of ω_{bulk} is a characteristic bulk angular velocity chosen to be the value of ω at the inflection point of the ω -profile (see figure 4.20). It is calculated from the numerical simulations. The result for the BL thickness ratio $\lambda_\omega^o/\lambda_\omega^i$ is shown in figure 4.19. The agreement with the numerically obtained ratio is satisfactory for the counter-rotating a -cases, getting even better with increasing Ta . This is because the estimate is based on a flat profile in the bulk, and indeed the profile becomes flatter with increasing Ta . For co-rotation, formula (4.34) apparently fails. This had to be expected, because the approximation of the profile of $\langle \bar{u}_\theta \rangle_z$ by three straight lines, which was assumed in the derivation of equation (4.34), is then no longer appropriate.

4.6.4 Angular velocity profiles in the bulk

We now come back to the mean profiles in the bulk. As can be clearly seen from comparing figures 4.17 and 4.18, both the mean angular velocity and its slope are controlled by a (or Ro^{-1}) rather than by Ta . This behaviour can be understood from equation (4.22): The outer-cylinder rotation is reflected in that equation as a Coriolis force. This force is present in the whole domain, while Ta controls the strength of the viscous term, which is dominant in the boundary layer. Therefore the profile is controlled by the Coriolis force, i.e. Ro^{-1} or a , and not by Ta .

To further quantify this, the gradient of $\langle \bar{\omega} \rangle_z$ is calculated. This is done by fitting a straight line to $\langle \bar{\omega} \rangle_z(\tilde{r})$ at the point of the profile's inflection, numerically using the grid points around it. An example of how this is done can be seen in figure 4.20(a). The results for the profile slopes in the bulk as functions of Ta and a or Ro^{-1} are shown in figure 4.21.

The graphs collapse on each other for co-rotation ($a < 0$), which is what we expected from figure 4.18 and our previous analysis. An almost linear relationship between Ro^{-1} and the bulk slope of $\omega(r)$ is found. If pure inner-cylinder rotation ($a = 0$) is approached, the graphs for the various Ta start to diverge and reach a

Figure 4.19: Comparison between the present DNS and the theory of EGL07 for $\lambda_\omega^o / \lambda_\omega^i$ versus a for four values of Ta : (a) $1.91 \cdot 10^6$; (b) $3.90 \cdot 10^6$; (c) $9.52 \cdot 10^6$; (d) $2.39 \cdot 10^7$. The agreement between theory and simulation is better for counter-rotation and with increasing Ta , but does not match for co-rotation. A dashed line indicates the optimal rotation ratio a_{opt} .

Figure 4.20: Example of the two fitting procedures for $Ta = 1.91 \cdot 10^6$ and $a = -0.2$: (a) shows a linear fit to the bulk part of the angular velocity $\langle \bar{\omega} \rangle_z$, and (b) shows in addition the fit to $\langle \bar{u}_\theta \rangle_z$ for its boundary parts, which will be used in §4.6.3. Both bulk fits are done at the inflection point, but for different variables ($\bar{\omega}$ or $\bar{u}_\theta$), which gives a slight difference.

Figure 4.21: The average slope of the angular velocity $d\langle \bar{\omega} \rangle_z / d\tilde{r}$ versus a (left) and Ro^{-1} (right) in the bulk. For co-rotation and only slight counter-rotation, there is a nearly linear relationship linking Ro^{-1} with $d\langle \bar{\omega} \rangle_z / d\tilde{r}$. A black straight line has been added in this region to artificially extend this relationship up to $d\langle \bar{\omega} \rangle_z / d\tilde{r} = 0$. There is also a plateau of $d\langle \bar{\omega} \rangle_z / d\tilde{r}$ at counter-rotation. Here the r -slopes are smaller, i. e., the profiles flatter, for increasing Ta , reflecting an increased convective transport.

plateau. The absolute value of the slope decreases with increasing Ta . This is due to the increasing importance of convection at higher Ta . Note however that also for this counter-rotating case, for large enough Ta the center slopes again lose their Ta dependence, i.e. are again mainly controlled by Ro^{-1} and thus the Coriolis term.

We now come back to the co-rotating regime $a < 0$ and want to connect the numerically found approximately linear relationship between the slope of $\omega(r)$ in the bulk and Ro^{-1} with the dynamical equation (4.22), which for the θ -component of the velocity can be rewritten as

$$\partial_t u_\theta + u_r \partial_r u_\theta + \frac{u_\theta}{r} \partial_\theta u_\theta + \frac{u_r u_\theta}{r} + u_z \partial_z u_\theta = -Ro^{-1} u_r - \partial_\theta p. \quad (4.35)$$

The linear relationship can be obtained if we assume that the Coriolis force term $Ro^{-1} u_r$ and convective term $u_r \partial_r u_\theta + u_r u_\theta / r$ balance each other, i.e. we assume that the axial, azimuthal, and temporal dependences are small in equation (4.35), which then boils down to $u_r (\partial_r u_\theta + u_\theta / r) \sim -u_r Ro^{-1}$. Next, we use the fact that the radial velocity component u_r (the wind) in its non-dimensional form is constant in the present Ta -range (cf. §5.4, seen also in experiment of Huisman *et al.* [45]). Therefore, and as u_θ hardly depends on Ro^{-1} , an increased Coriolis force can only be balanced by a larger slope $\partial_r u_\theta$. The only alternative is that the wind u_r vanishes altogether, $u_r = 0$, and the flow state returns to the purely azimuthal, laminar case.

To further substantiate this argument, we now decompose the flow field into a (t, θ, z) -averaged mean azimuthal flow component U_θ , depending on the radial position r only, plus fluctuations u'_θ , as well as a decomposition of the pressure into a mean pressure P plus the pressure fluctuations p' . Inserting these Reynolds-type decompositions into the radial and azimuthal momentum equations, in which besides U_θ only its r -derivative survives and ignoring viscosity for now, we arrive at the following equations:

$$\begin{aligned} \partial_t u'_r + u'_r \partial_r u'_r + \frac{u'_\theta}{r} \partial_\theta u'_r - \frac{U_\theta^2}{r} - 2 \frac{U_\theta u'_\theta}{r} - \frac{u_\theta^2}{r} + u'_z \partial_z u'_r \\ = -\partial_r (P + p') + Ro^{-1} (U_\theta + u'_\theta), \end{aligned} \quad (4.36)$$

$$\begin{aligned} \partial_t u'_\theta + u'_r \partial_r U_\theta + u'_r \partial_r u'_\theta + \frac{u'_\theta}{r} \partial_\theta u'_\theta + \frac{U_\theta}{r} \partial_\theta u'_\theta + \frac{u'_r U_\theta}{r} + \frac{u'_r u'_\theta}{r} + u'_z \partial_z u'_\theta \\ = -Ro^{-1} u'_r - \partial_\theta p'. \end{aligned} \quad (4.37)$$

It is important to note that U_r and U_z are both equal to zero, so $u_r = u'_r$ and $u_z = u'_z$. As long as Ro^{-1} is larger than Ro_{opt}^{-1} , we assume that the mean flow contributions alone already balance in equations (4.36) and (4.37),

$$-\frac{U_\theta^2}{r} \sim -\partial_r P + Ro^{-1} U_\theta, \quad (4.38)$$

and

$$u'_r \partial_r U_\theta = u'_r \left(r \frac{d\langle \bar{\omega} \rangle_z}{dr} + \langle \bar{\omega} \rangle_z \right) \sim -u'_r \left(Ro^{-1} + \frac{U_\theta}{r} \right). \quad (4.39)$$

As in the bulk $r \approx r_a$ is almost constant, the linear relationship $\partial_r \omega \propto \text{const} - Ro^{-1}$ between Ro^{-1} and bulk slope $d\langle \bar{\omega} \rangle_z / d\tilde{r}$ is obtained.

Figure 4.21(b), displaying this linear relationship, can be used to obtain a quantitative estimate for optimal transport for large Ta . We can see two distinct features in the slope of this $d\langle \bar{\omega} \rangle_z / d\tilde{r}$ versus Ro^{-1} curve. There is a plateau, where the value of the slope is linked to Ta (and therefore to the viscous term in the equation of motion), and there is a line to the right of the plateau where the value of $d\langle \bar{\omega} \rangle_z / d\tilde{r}$ is independent of Ta and thus linked only to the Coriolis force. From the previous discussion and from the experimental evidence of Ref. [77] we know that optimal transport is linked to the flattest ω -profile. We can interpret the shift of Ro_{opt}^{-1} with Ta as that value of Ro^{-1} where the plateau meets the co-rotation linear relationship, i.e. the flattest possible ω -profile that does not break the large scale balance discussed before. If Ro^{-1} becomes more negative, the profile would have to become flatter to keep on satisfying the large scale balance. As this does not happen, the transport decreases for more negative Ro^{-1} .

With increasing Ta , the value of $d\langle \bar{\omega} \rangle_z / d\tilde{r}$ at the plateau increases, and the curves cross at a smaller value of Ro^{-1} , which corresponds to a shifted maximum. Eventually, the plateau value of $d\langle \bar{\omega} \rangle_z / d\tilde{r}$ will tend to zero as seen in the experiments of Ref. [77], and the co-rotation line will cross the plateau at the x-axis. We can extend the straight line to get an estimate for when this happens and obtain $Ro_{opt}^{-1}(Ta \rightarrow \infty) = -0.20$, corresponding to $a_{opt} \approx 0.34$, an estimate consistent with the experimental values $a_{opt} \approx 0.33 \pm 0.05$ of Refs. [44, 77].

If Ro^{-1} is too negative, $Ro^{-1} < Ro_{opt}^{-1}$, the large scale balance of equation (4.39) can no longer be satisfied. This can be seen as the Coriolis force now has values which would require a smaller (or even a negative) value of $d\bar{\omega}/d\tilde{r}$ for the balance to hold. Since this is not possible to accommodate, bursts will originate from the outer cylinder towards the inner cylinder, because the flow tries to accommodate a large Coriolis force. These bursts increase in importance until they end up stabilizing parts of the flow, or even the whole flow which will drastically reduces the transport. Therefore, a maximum transport is reached just when the Coriolis force balances the large scale convective term. If it is further increased, stabilized regions start appearing. This is linked to the appearance of a neutral surface, which is analysed in the next section.

The large scale balance cannot be satisfied either if Ro^{-1} is too positive. The Coriolis force then requires a value of $d\bar{\omega}/d\tilde{r} > 1$ to be balanced through the convective acceleration forces due to the average flow, and this cannot be accommodated. Unlike the previous case, the flow cannot be separated into stable and unstable re-

gions. Instead, this can be linked to the complete disappearance of the radial and axial components of the flow (the so-called wind). This causes Nu_ω to drop to the purely azimuthal value, as seen in figure 4.11.

4.6.5 Neutral surface

In this subsection, we leave the rotating reference frame, and return to the inertial reference frame to analyse the neutral surface which is defined as that surface where, in the inertial reference frame, ω is zero. This surface only exists for non-negative values of a and coincides with the outer cylinder in the case of pure inner-cylinder rotation. In an inertial reference frame, it marks the division between the Rayleigh (inviscid) stable and unstable regions. This means that this surface separates two regions, an unstable inner region and a stable outer region. In the stable region, perturbations to the azimuthal flow (both large scale wind as in Taylor vortices and small-scale perturbations such as plumes) cannot grow. Therefore we expect this surface to play a significant role in the behaviour of the flow. It was already shown to be important in controlling optimal transport by van Gils *et al.* [77].

In general, the position r_N of the neutral surface depends on the height z . Figures 4.22 and 4.24 show contour plots of the angular velocity with $r_N(z)$ indicated. Indeed, $r_N(z)$ shows a strong axial dependence, showing heights with positive or negative u_r , at which the neutral surface is pushed more outside or inside, respectively. This strong axial dependence of r_N is a measure for the vortex strength and becomes weaker, when the vortices lose importance at very high Ta . By comparing figures 4.22 and 4.24 the effect of the Taylor number on the position of the neutral surface can be seen. Its distortion happens at larger a for increasing Ta , as expected.

If a is large enough, the vortices are no longer able to penetrate the whole gap. There the neutral surface separates the Rayleigh stable and unstable regions. The vortices are mainly located in the unstable range, but partially enter the Rayleigh stable region. Being restricted to part of the gap, they also shrink in the horizontal direction. Because the vortices try to remain as square-like as possible, their height (wavelength) also shrinks, allowing new vortices to appear in the available given height. This is visible in figure 4.23(c). These vortices are also associated with a stronger wind. If the value of a is not very large, they thus will again fill the distance between the two cylinders but with a distorted aspect ratio. A zoom of this effect can be seen in figure 4.23(b). This causes both the rise in Nu_ω for positive a seen in figure 4.12 around $a \approx 0.3$ at low Ta and the crossovers seen in figure 4.25(a) below.

Next, in addition to the temporal and azimuthal average, we also average $r_N(z)$ in the axial direction and call this average $\bar{r}_N$. Figure 4.25 shows how $\bar{r}_N$ varies with a and Ta . The position of the neutral surface for the laminar, purely azimuthal flow is plotted for comparison. For slight counter-rotation and fixed a , the mean neutral surface is increasingly pushed towards the outer cylinder with increasing Ta

Figure 4.22: Contour plots of the dimensionless angular velocity $\tilde{\omega}(\tilde{r}, \tilde{z})$ with indicated neutral surface (black line) for $Ta = 1.03 \cdot 10^5$. (a) $a = 0.2$ and three vortex pairs. (b) $a = 0.4$ and four vortex pairs with a non-square aspect ratio and a highly distorted neutral surface. (c) $a = 0.6$, four vortex pairs which do not penetrate the whole gap. See also figure 4.23 for an enlargement.

due to enhanced turbulence. On the other hand, with increasing a the Coriolis force pushes the neutral surface more and more towards the inner cylinder. Once the neutral surface reaches the laminar and purely azimuthal flow value, the flow is stabilized.

The curves $\bar{r}_N(a)$ for different Ta can also cross. At a constant rotation ratio, some of the lower Ta curves have a neutral surface which is further away from the inner cylinder than for some of the higher Ta ones. This is due to changes in the number and strengths of vortex pairs in the flow, which happen earlier for lower Ta . With a further increase of a the trend reverses again, since respectively smaller values of a stabilize the flow at already lower Ta . In an inertial reference frame, this simply means that there is no longer a radial velocity which can push the neutral surface outwards, so $\bar{r}_N$ falls back to the laminar, purely azimuthal flow value $\bar{r}_{N,pa} = (\eta)/(1 - \eta)(\sqrt{((1 + a)/(\eta^2 - \mu))} - 1)$.

Figure 4.23: Enlarged portions of the contour plots of $\bar{\omega}(\tilde{r}, \tilde{z})$ for $Ta = 1.03 \cdot 10^5$ with the neutral surface indicated as a black line. (a) $a = 0.2$, normal state with three vortices. (b) $a = 0.4$, the distorted vortices are strong enough to fully penetrate the gap. (c) $a = 0.6$, the distorted vortices cannot penetrate the whole gap due to the stabilizing effects beyond the neutral line. Moreover, with increasing a the distance between the vortex centres shrinks.

Figure 4.24: Contour plots of $\bar{\omega}(\tilde{r}, \tilde{z})$ with neutral surface indicated for $Ta = 3.90 \cdot 10^6$ and a -values (a-c) of 0.2, 0.4, and 0.6. All of the simulations show three vortex pairs.

Figure 4.25: (a) The neutral surface location $\bar{r}_N$ of the time, azimuthally, and axially averaged angular velocity versus rotation ratio a for various Ta . The neutral radius $\bar{r}_N$ is moving inwards with increasing counter-rotation (increasing a), but is pushed back towards the outer cylinder for increasing Ta at given a . The black line shows the position of neutral radius $\bar{r}_{N,pa}$ in laminar, purely azimuthal flow. (b) $\bar{r}_N - \bar{r}_{N,pa}$ is shown, quantifying the pushing of the neutral line towards the outer cylinder with increasing Ta .

4.7 Summary and conclusions

An extensive direct numerical simulation (DNS) exploration of the parameter space of a Taylor-Couette (TC) system at Taylor numbers in the range of $10^4 < Ta < 10^8$ is presented. First the code was validated versus existing numerical and experimental data. After this, the transition from the laminar but still purely azimuthal regime to the Taylor vortex state was analysed. The regime where these vortices dominate the flow was studied in detail, revealing scaling laws between the Taylor number Ta , the angular velocity flux $Nu_\omega - 1$, and the wind Reynolds number Re_w . These scaling laws ceased to be valid when the Taylor number was increased beyond $Ta \approx 3 \cdot 10^6$. At this driving strength the coherence structures become so small that they lose importance and are no longer the dominating feature of the flow.

Then the effect of the outer-cylinder rotation on these scaling laws was analysed. If both cylinders are co-rotating, the scaling laws were (slightly) modified, but for counter-rotating cylinders no significant differences could be seen. After the shrinking of the coherent structure and loss of their importance the value for optimal transport a_{opt} shifted towards counter-rotation. This drift is expected to continue at higher Taylor numbers and will be the subject of future DNS investigations.

Next, the behaviour of local flow variables was studied. Analysing the profiles $\omega(r)$ sheds light on the two transport mechanisms, convective and diffusive, cf. the two contributions in equation (4.4). The optimal transport of ω could be linked to a balance between the Coriolis force and the inertial terms in the equations of motion.

This balance is best achieved when the bulk profile is flattest, and is broken with increasing counter-rotation. This leads to the appearance of a neutral line and of “stabilizing” bursts.

The outer boundary layer of the ω -profile is much thicker than the inner boundary layer. The quality of the approximation of the ω -profiles by three straight lines was found to improve with increasing Ta , as the (bulk-)turbulence becomes stronger. But although the bulk is turbulent, the boundary layers are still of Prandtl-Blasius type. TC flow only reaches the ultimate state, if the boundary layers also undergo a shear instability and become turbulent, too. The present analysis showed that this transition is expected to happen in the Ta range between 10^8 and 10^9 , which is just outside the range of the present DNS. It will be analyzed in future work.

Our ambition is to further extend the Ta range in our DNS of TC flow in order to allow a one-to-one comparison between experiments and simulations in the ultimate regime of TC turbulence and to explore the physics of this ultimate regime, in particular to understand the transition to this regime, and the bulk–boundary layer interaction in that regime. This ultimate regime in TC flow has recently been observed and analysed in the experiments by Huisman *et al.* [45] and van Gils *et al.* [77], as well as in Rayleigh-Bénard (RB) experiments of He *et al.* [76]. As the mechanical driving in TC is more efficient than heating in RB convection, it is easier to reach the ultimate regime in TC experiments than in RB experiments. Therefore we expect also numerically to reach the ultimate regime earlier in TC flow than in RB flow.

5

Optimal Taylor-Couette: radius ratio dependence *

Whenever a theory appears to you as the only possible one, take this as a sign that you have neither understood the theory, nor the problem which it was intended to solve.

—KARL POPPER

Taylor–Couette flow with independently rotating inner (i) and outer (o) cylinders is explored numerically and experimentally to determine the effects of the radius ratio η on the system response. Numerical simulations reach Reynolds numbers of up to $Re_i = 9.5 \cdot 10^3$ and $Re_o = 5 \cdot 10^3$, corresponding to Taylor numbers of up to $Ta = 10^8$ for four different radius ratios $\eta = r_i/r_o$ between 0.5 and 0.909. The experiments, performed in the Twente Turbulent Taylor–Couette (T^3C) setup, reach Reynolds numbers of up to $Re_i = 2 \cdot 10^6$ and $Re_o = 1.5 \cdot 10^6$, corresponding to $Ta = 5 \cdot 10^{12}$ for $\eta = 0.714–0.909$. Effective scaling laws for the torque $J^\omega(Ta)$ are found, which for sufficiently large driving Ta are independent of the radius ratio η . As previously reported for $\eta = 0.714$, optimum transport at a non-zero Rossby number $Ro = r_i|\omega_i - \omega_o|/[2(r_o - r_i)\omega_o]$ is found in both experiments and numerics. Ro_{opt} is found to depend on the radius ratio and the driving of the system. At a driving in the range between $Ta \sim 3 \cdot 10^8$ and $Ta \sim 10^{10}$, Ro_{opt} saturates to an asymptotic η -dependent value. Theoretical predictions for the asymptotic value of Ro_{opt} are compared to the experimental results, and found to differ notably. Furthermore, the local angular velocity profiles from experiments and numerics are compared, and a link between a flat bulk profile and optimum transport for all radius ratios is reported.

*Based on: R. Ostilla-Mónico, S. G. Huisman, T. J. G. Jannink, D. P. M. van Gils, R. Verzicco, S. Grossmann, C. Sun and D. Lohse, *Optimal Taylor-Couette flow: radius ratio dependence*, J. Fluid Mech., **747**, 1–29 (2014). Simulations performed by Ostilla-Mónico under the supervision of Verzicco. Experiments performed by Jannink and Huisman under the supervision of van Gils and Sun. Discussion and writing done by everyone.

5.1 Introduction

Recent experiments [42–44, 80] and simulations [112, 123, 156] of Taylor–Couette (TC) flow have shown that at fixed Ta an optimal angular momentum transport is obtained at non-zero a_{opt} , and that the location of this maximum a_{opt} varies with Ta . However, both experiments and simulations have been restricted to two radius ratios, namely $\eta = 0.5$ and $\eta = 0.714$. The same radius ratios were also used for studies carried out on scaling laws of the torque and the “wind” of turbulence at highly turbulent Taylor numbers [41, 43–45, 80]. Up to now, it is not clear how the radius ratio affects the scaling laws of the system response or the recently found phenomena of optimal transport as a function of Ta .

Two suggestions were made to account for the radius ratio dependence of optimal transport. Van Gils *et al.* [43] wondered whether the optimal transport in general lies in or at least close to the Voronoi boundary (meaning a line of equal distance) of the Esser-Grossmann stability lines [152] in the (Re_o, Re_i) phase space as it does for $\eta = 0.714$. However, this bisector value does not give the correct optimal transport for $\eta = 0.5$ [80, 156]. Therefore, Brauckmann & Eckhardt [156] proposed a dynamic extension of the Esser-Grossmann instability theory. This model correctly gives the observed optimal transport (within experimental error bars) between $\eta = 0.5$ and $\eta = 0.714$ for three experimental datasets [31, 43, 44] and one numerical dataset [156], but it is not clear how it performs outside the η -range $[0.5, 0.714]$.

In this chapter, we study the following questions. How does the radius ratio η affect the flow? How are the scaling laws of the angular momentum transport affected? What is the role of the geometric parameter called the pseudo-Prandtl number σ introduced in EGL2007? Can the effect of the radius ratio be interpreted as a kind of non-Oberbeck-Boussinesq effect, analogous to this effect in Rayleigh–Bénard flow? Finally, are the predictions and insights of van Gils *et al.* [43], Ostilla *et al.* [112] and Brauckmann & Eckhardt [156] on the optimal transport also valid for other values of η ?

In order to answer these questions, both direct numerical simulations (DNS) and experiments have been undertaken. Numerical simulations, with periodic axial boundary conditions, have been performed using the finite-difference code previously used in [112]. In these simulations, three more values of η have been investigated: one in which the gap is larger ($\eta = 0.5$), and two in which the gap is smaller ($\eta = 0.833$ and 0.909). With the previous simulations from Ref. [112] at $\eta = 0.714$, a total of four radius ratios has been analysed.

In both experiments and numerics, only one aspect ratio Γ has been studied for every radius ratio. Since the work of Benjamin & Mullin [157] it has been known that multiple flow states with a different amount of vortex pairs can coexist in TC flow for the same non-dimensional flow parameters. However, with increased driving, the bifurcations become less important and many branches do not survive. Indeed, Lewis

& Swinney [41] found that for pure inner cylinder rotation only one branch with 8 vortices (for $\Gamma = 11.4$ and $\eta = 0.714$) remains when Re_i is increased above $2 \cdot 10^4$. As the Reynolds numbers reached in the experiments greatly exceed this value we do not expect to see the effect of multiple states in the current experimental results.

For the numerical simulations, axially periodic boundary conditions have been taken. Brauckmann & Eckhardt [123] already studied the effect of the axial periodicity length on the system, and found that, for a fixed vortical wavelength, the number of vortices does not affect the overall flow behaviour. It was also found that, in analogy to experiments, the effect of vortical wavelength, and hence of multiple states, becomes smaller with increased driving. However, changes between states can still be seen for small drivings.

Figure 5.1 shows the $(Ta, 1/Ro)$ parameter space explored in the simulations for the four selected values of the radius ratio η . A higher density of points has been used in places where the global response (Nu_ω, Re_w) of the flow shows more variation with the control parameters Ta and $1/Ro$. A fixed aspect ratio of $\Gamma = 2\pi$ has been taken for all simulations, and axially periodic boundary conditions have been used. These simulations have the same upper bounds of Ta (or Re_i) as those of Ref. [112].

In addition to these simulations, experiments have been performed with the Twente Turbulent Taylor-Couette (T³C) facility, with which we achieve larger Ta numbers. Details of the set-up are given in Ref. [42]. Once again, four values of η have been investigated, but, owing to experimental constraints, we have been limited to investigate only smaller gap widths, i.e. values $\eta \geq 0.714$, namely $\eta = 0.716, 0.769, 0.833$ and 0.909 . The experimentally explored parameter space are shown in figure 5.2. In the experimental set-up, $\eta = 0.716$ is used, while in the simulations $\eta = 0.714$ is used. These two are directly compared to each other in the present chapter. When referring to either numerical data, or numerical and experimental data in combination, $\eta = 0.714$ is used, while when referring to exclusively experimental data, $\eta = 0.716$ is used.

The structure of the chapter is as follows. In §2 and §3, we start by describing the numerical code and the experimental setup, respectively. In §4, the global response of the system, quantified by the non-dimensionalized angular velocity current Nu_ω is analysed. To understand the global response, we analyse the local data that can be obtained from the DNS in §5. Angular velocity profiles in the bulk and in the boundary layers are analysed and related to the global angular velocity optimal transport. We finish in §6 with a discussion of the results and an outlook for further investigations.

5.2 Numerical method

In this section, the used numerical method is explained in some detail. The rotating frame in which the Navier-Stokes equations are solved and the employed non-

Figure 5.1: Control parameter phase space which was numerically explored in this chapter in the $(Ta, 1/Ro)$ representation: (a) $\eta = 0.5$, (b) $\eta = 0.714$, (c) $\eta = 0.833$, (d) $\eta = 0.909$. $\Gamma = 2\pi$ was fixed, and axial periodicity was employed. The grey-shaded area signals boundary conditions for which the angular momentum $L = r^2\omega$ of the outer cylinder (L_o) is larger than the angular momentum of the inner cylinder (L_i). This causes the flow to have an overall transport of angular momentum towards the inner cylinder. In this region, the Rayleigh stability criterion applies, which states that if $dL^2/dr > 0$ the flow is linearly stable to axisymmetric perturbations [84].

Figure 5.2: Control parameter phase space which was explored in experiments in the (Ta, Ro^{-1}) representation: (a) $\eta = 0.716$, (b) $\eta = 0.769$, (c) $\eta = 0.833$, (d) $\eta = 0.909$.

dimensionalizations are introduced in the first section. This is followed by detailing the spatial resolution checks which have been performed.

5.2.1 Code description

The employed code is a finite-difference code, which solves the Navier-Stokes equations in cylindrical coordinates. A second-order spatial discretization is used, and the equations are advanced in time by a fractional time integration method. This code is based on the so-called Verzicco code, whose numerical algorithms are detailed in Ref. [103]. A combination of MPI and OpenMP directives are used to achieve large-scale parallelization. This code has been extensively used for Rayleigh-Bénard flow; for recent simulations see Refs. [108, 150]. In the context of TC flow, Ostilla-Mónico *et al.* [112] have already validated the code for $\eta = 0.714$.

The flow was simulated in a rotating frame, which was chosen to rotate with $\Omega = \omega_o e_z$. This was done in order to simplify the boundary conditions. In that frame the outer cylinder is stationary for any value of a , while the inner cylinder has an azimuthal velocity of $u_\theta(r = r_i) = r_i(\omega_i^\ell - \omega_o^\ell)$, where the superscript ℓ denotes variables in the laboratory frame, while no superscript denotes variables in the rotating frame. We then choose the inner cylinder rotation rate in the rotating frame as the characteristic velocity of the system $U \equiv |u_\theta(r_i)| = r_i|\omega_i - \omega_o|$ and the characteristic length scale d to non-dimensionalize the equations and boundary conditions.

Using this non-dimensionalization, the inner cylinder velocity boundary condition simplifies to $\hat{u}_\theta(r = r_i) = \text{sgn}(\omega_i - \omega_o)$. In this chapter, $\omega_i - \omega_o$ is always positive. Thus, in this rotating frame the flow geometry is simplified to a pure inner cylinder rotation with the boundary condition $\hat{u}_\theta(r_i) = 1$. The outer cylinder's effect on the flow is felt as a Coriolis force in this rotating frame of reference. The Navier-Stokes equations then read

$$\frac{\partial \hat{\mathbf{u}}}{\partial \hat{t}} + \hat{\mathbf{u}} \cdot \hat{\nabla} \hat{\mathbf{u}} = -\hat{\nabla} \hat{p} + \frac{f(\eta)}{Ta^{1/2}} \hat{\nabla}^2 \hat{\mathbf{u}} - Ro^{-1} e_z \times \hat{\mathbf{u}}, \quad (5.1)$$

where Ro was defined previously in equation (4.23), and $f(\eta)$ in equation (4.24).

It is useful to continue the non-dimensionalization by defining the normalized radius $\tilde{r} = (r - r_i)/d$ and the normalized height $\tilde{z} = z/d$. We define the time- and azimuthally averaged velocity field as

$$\hat{\mathbf{u}}(r, z) = \langle \hat{\mathbf{u}}(\theta, r, z, t) \rangle_{\theta, t}, \quad (5.2)$$

where $\langle \phi(x_1, x_2, \dots, x_n) \rangle_{x_i}$ indicates averaging of the field ϕ with respect to x_i .

To quantify the torque in the system, we will use the pseudo-Nusselt number Nu_ω , which is defined as $Nu_\omega = T/T_{pa}$, where T is the torque at the cylinders, and T_{pa} the torque at the purely azimuthal, laminar state (cf. §4.1 for further details). This quantity Nu_ω can be seen as an angular velocity “Nusselt” number.

When Nu_ω is calculated numerically, the values will depend on the radial position, owing to finite time averaging. We can define Δ_J to quantify this radial dependence as

$$\Delta_J = \frac{\max(Nu_\omega(r)) - \min(Nu_\omega(r))}{\langle Nu_\omega(r) \rangle_r} \quad (5.3)$$

which analytically equals zero but will deviate when calculated numerically.

The convective dissipation per unit mass can be calculated either from its definition as a volume average of the local dissipation rate for an incompressible fluid,

$$\varepsilon_u = \varepsilon_u^\ell = \frac{\nu}{2} \left\langle (\partial_i^\ell u_j + \partial_j^\ell u_i)^2 \right\rangle_{V,t}, \quad (5.4)$$

or a global balance can be used. The exact relationship (EGL 2007)

$$\varepsilon_u^\ell - \varepsilon_{u,0}^\ell = \frac{\nu^3}{d^4} \sigma^{-2} Ta (Nu_\omega - 1), \quad (5.5)$$

links the volume-averaged dissipation to the global driving Ta and response Nu_ω .

This link can be and has been used for code validation and for checking spatial resolution adequateness. The volume-averaged dissipation can be calculated from both equations (5.4) and (5.5) and later checked for sufficient agreement. We define the quantity Δ_ε as the relative difference between the two ways of numerically calculating the dissipation, namely either via Nu_ω with (equation (5.5)) or directly from the velocity gradients, (equation (5.4)),

$$\Delta_\varepsilon = \frac{\nu^3 d^{-4} \sigma^{-2} Ta (Nu_\omega - 1) + \varepsilon_{u,0}^\ell - \frac{\nu}{2} \left\langle (\partial_i^\ell u_j + \partial_j^\ell u_i)^2 \right\rangle_{V,t}}{\frac{\nu}{2} \left\langle (\partial_i^\ell u_j + \partial_j^\ell u_i)^2 \right\rangle_{V,t}}. \quad (5.6)$$

The quantity Δ_ε is equal to 0 analytically, but will deviate when calculated numerically. The deviation of Δ_J and Δ_ε from zero is an indication of the adequateness of the resolution.

We would like to emphasize that the requirement for $\Delta_J < 0.01$ is much stricter than torque balance, which can simply be expressed as $Nu_\omega(r_i) = Nu_\omega(r_o)$. As analysed in Ref. [112], a value of less than 1% for Δ_J and about 1% for Δ_ε is linked to grid adequateness at the Taylor number simulated. To ensure convergence in time, the time averages of the Nusselt number and the energy dissipation calculated locally (cf. equation (5.4)) were also checked to converge in time within 1%.

5.2.2 Resolution checks

Spatial resolution checks were performed in two ways. First, as mentioned previously, the values of Δ_J and Δ_ε were checked. As an additional check, simulations

at selected values of Ta were performed at a higher resolution. As the explored parameter space is large, these checks were performed only for the highest value of Ta simulated for the grid size. A lower driving of the flow for the same grid size is expected to have a smaller error due to spatial discretization, as spatial discretization errors increase with increased Re , and thus increased Ta .

Concerning the temporal resolution there are numerical and physical constraint. The former requires a time step small enough to keep the integration scheme stable, and this is achieved by using an adaptive time step based on a Courant–Frederichs–Lewy (*CFL*) criterion. The third-order Runge–Kutta time-marching algorithm allows for a *CFL* of up to $\sqrt{3}$, but this can be reduced due to the implicit factorization of the viscous terms. For safety, the maximum *CFL* has been taken as 1.4. From the physical point of view, the time step size must also be small enough to properly describe the fast dynamics of the smallest flow scale, which is the Kolmogorov scale. Although the time step size should be determined by the most restrictive among the two criteria above, our experience suggests that, as long as the *CFL* number criterion is satisfied, which guarantees numerical stability, the results become insensitive to the time step size and all the flow scales are adequately described temporally. Direct confirmation of this statement can be found in Ref. [112].

The results for $\eta = 0.5, 0.833$ and 0.909 are presented in table 5.1. Uniform discretization was used in the azimuthal and axial directions. In the radial direction, points were clustered near the walls by using hyperbolic tangent-type clustering, or a clipped Chebychev type clustering for higher values of Ta . A table including the results for the spatial resolution tests at $\eta = 0.714$ can be found in Ref. [112].

5.3 Experimental set-up

The Twente Turbulent Taylor-Couette (T^3C) apparatus has been built to obtain high Ta numbers. It has been described in detail in Refs. [42, 77]. The inner cylinder with outside radius $r_i = 0.200$ m consists of three sections. The total height of those axially stacked sections is $L = 0.927$ m. We measure the torque only on the middle section of the inner cylinder, which has a height of $L_{\text{mid}} = 0.536$ m, to reduce the effect of the torque losses at the end-plates in our measurements. This approach has already been validated in Ref. [77]. The transparent outer cylinder is made of acrylic and has an inside radius of $r_o = 0.2794$ m. We vary the radius ratio by reducing the diameter of the outer cylinder by adding a ‘filler’ that is fixed to the outer cylinder and sits between the inner and the outer cylinder, effectively reducing r_o while keeping r_i fixed. We have 3 fillers giving us 4 possible outer radii: $r_o = 0.2794$ m (without any filler), 0.26 m, 0.24 m, and 0.22 m, giving experimental access to $\eta = 0.716, 0.769, 0.833$ and 0.909 , respectively. Note that by reducing the outer radius, we not only change η , but also change $\Gamma = L/(r_o - r_i)$ from $\Gamma(\eta = 0.716) = 11.68$ to

η	Ta	$N_\theta \times N_r \times N_z$	Nu_ω	$100\Delta_J$	$100\Delta_\epsilon$	Case
0.5	$2.5 \cdot 10^5$	$100 \times 100 \times 100$	2.03372	0.30	1.11	R
0.5	$2.5 \cdot 10^5$	$150 \times 150 \times 150$	2.03648	0.76	0.89	E
0.5	$7.5 \cdot 10^5$	$150 \times 150 \times 150$	2.56183	0.47	0.92	R
0.5	$7.5 \cdot 10^5$	$256 \times 256 \times 256$	2.55673	0.74	0.31	E
0.5	$1 \cdot 10^7$	$300 \times 300 \times 300$	4.23128	0.33	1.07	R
0.5	$1 \cdot 10^7$	$400 \times 400 \times 400$	4.22574	0.97	1.06	E
0.5	$2.5 \cdot 10^7$	$350 \times 350 \times 350$	5.07899	0.85	1.12	R
0.5	$2.5 \cdot 10^7$	$512 \times 512 \times 512$	5.08193	0.87	1.98	E
0.5	$5 \cdot 10^7$	$768 \times 512 \times 1536$	6.08284	0.45	1.56	R
0.5	$1 \cdot 10^8$	$768 \times 512 \times 1536$	7.48561	1.46	0.88	R
0.833	$2.5 \cdot 10^5$	$180 \times 120 \times 120$	2.72293	0.21	0.76	R
0.833	$2.5 \cdot 10^5$	$300 \times 180 \times 180$	2.72452	0.29	0.29	E
0.833	$1 \cdot 10^7$	$384 \times 264 \times 264$	7.07487	0.29	0.61	R
0.833	$1 \cdot 10^7$	$512 \times 384 \times 384$	7.17245	0.13	1.16	E
0.833	$2.5 \cdot 10^7$	$512 \times 384 \times 384$	8.62497	0.71	1.05	R
0.833	$2.5 \cdot 10^7$	$768 \times 576 \times 576$	8.51678	0.90	1.26	E
0.833	$5 \cdot 10^7$	$512 \times 384 \times 384$	9.68437	0.26	2.92	R
0.833	$1 \cdot 10^8$	$768 \times 576 \times 576$	11.4536	0.89	2.29	R
0.909	$2.5 \cdot 10^5$	$180 \times 120 \times 120$	2.31902	0.16	0.91	R
0.909	$2.5 \cdot 10^5$	$300 \times 200 \times 200$	2.30810	0.07	0.17	E
0.909	$2 \cdot 10^6$	$256 \times 180 \times 180$	3.76826	0.55	0.49	R
0.909	$2 \cdot 10^6$	$384 \times 256 \times 256$	3.77532	0.39	0.21	E
0.909	$2.5 \cdot 10^7$	$384 \times 256 \times 256$	7.83190	0.43	3.15	R
0.909	$2.5 \cdot 10^7$	$450 \times 320 \times 320$	7.86819	0.81	2.07	E
0.909	$5 \cdot 10^7$	$2305 \times 400 \times 1536$	9.74268	0.46	1.02	R
0.909	$1 \cdot 10^8$	$2305 \times 400 \times 1536$	11.3373	0.57	1.06	R

Table 5.1: Resolution tests for $\Gamma = 2\pi$ and $\eta = 0.5, 0.833$ and 0.909 . The first column displays the radius ratio, the second column displays the Taylor number, the third column displays the resolution employed, the fourth column the calculated Nu_ω , the fifth column and sixth columns the relative discrepancies Δ_J and Δ_ϵ , and the last column the 'case': (R)esolved and (E)rror reference. Δ_ϵ is positive, and exceeds the 1% threshold reported in Ref. [112] for some cases at the largest η , but even so resolution appears to be sufficient as variations of Nu_ω are small.

$$\Gamma(\eta = 0.909) = 46.35.$$

For high Ta the heating up of the system becomes apparent and it has to be actively cooled in order to keep the temperature constant. We cool the working fluid (water) from the top and bottom end plates and maintain a constant temperature within $\pm 0.5K$ through both the spatial extent and the time run of the experiment. The set-up has been constructed in such a way that we can rotate the two cylinders independently while keeping the set-up cooled.

As said before, we measure the torque on the middle inner cylinder. We do this by measuring the torque that is transferred from the axis to the cylinder by using a load cell that is inside the aforementioned cylinder. Torque measurements are performed using a fixed procedure. The inner cylinder is spun up to its maximum rotational frequency of 20 Hz and kept there for several minutes. Then the system is brought to rest. The cylinders are then brought to their initial rotational velocities (with the chosen $1/Ro$), corresponding to a velocity for which the torque is accurate enough; generally of order 2–3 Hz. We then slowly increase both velocities over 3–6 hours to their final velocities while maintaining $1/Ro$ fixed during the entire experiment. During this velocity ramp we continuously acquire the torque of this quasi-stationary state. The calibration of the system is done in a similar way; first we apply the maximum load on the system, going back to zero load, and then gradually adding weight while recording the torque. These procedures ensure that hysteresis effects are kept to a minimum, and that the system is always brought to the same state before measuring. More details about the set-up can be found in Ref. [42].

Local velocity measurements are done by laser Doppler anemometry (LDA). We measure the azimuthal velocity component by focusing two beams in the radial–azimuthal plane. We correct for curvature effects of the outer cylinder by using a ray tracer, see Ref. [158]. The velocities are measured at mid-height ($z = L/2$) unless specified otherwise. For every measurement position, we measured for long enough such as to have a statistically stationary result, for which about $\sim 10^5$ samples were required for every data point. This ensured a statistical convergence of $< 1\%$.

5.4 Global response: Torque

In this section, the global response of the TC system for the four simulated radius ratios is presented. This is done by measuring the scaling law(s) of the non-dimensional torque Nu_ω as function(s) of Ta . The transition between different types of local scaling laws in different Ta ranges is investigated, and related to previous simulations [112] and experiments [77].

5.4.1 Pure inner cylinder rotation

The global response of the system is quantified by Nu_ω . By definition, for purely azimuthal laminar flow, $Nu_\omega = 1$. Once the flow is driven stronger than a certain critical Ta , large azimuthal roll structures appear, which enhance angular transport through a large scale wind.

Figure 5.3 shows the response of the system for increasing Ta in case of pure inner cylinder rotation for four values of η . Experimental and numerical results are shown in the same panels, covering different ranges, and thus complementary to, but consistent with each other. Numerical results for $\eta = 0.714$ from Ref. [112] have been added to both panels.

As has already been noticed in Ref. [112] for $\eta = 0.714$, a change in the local scaling law relating Ta to Nu_ω occurs at around $Ta \approx 3 \cdot 10^6$. We can interpret these changes in the same way as Ref. [112] and relate the transition in the Ta - Nu_ω local scaling law to the break-up of coherent structures. A large jump in Nu_ω can be seen around $Ta \approx 5 \cdot 10^6$ for $\eta = 0.909$, which corresponds to a change in the number of rolls. The system goes from two pairs of rolls to three pairs of rolls.

It is also worth mentioning that the exponent in the local scaling laws in the regime before the transition depends on the radius ratio. This can be seen in the compensated plot, and explains the curve crossings that we see in the graphs.

For experiments (solid lines of figure 5.3), a different local scaling law can be seen. In this case the experiments are performed at much higher Ta than the simulations. The scaling $Nu_\omega \sim Ta^{0.38}$ can be related to the so-called “ultimate” regime, a regime where the boundary layers have become completely turbulent [78, 79, 82]. As indicated for the case at $\eta = 0.714$ we expect that for increasing Ta also the simulations become turbulent enough to reach this scaling law (cf. Ref. [112]). In this regime, the local scaling law relating Ta and Nu_ω has no dependence on η and thus is universal.

In the experiments with large Ta , the value of η corresponding to the smallest gap, i.e. $\eta = 0.909$, has the highest angular velocity transport (Nu_ω) at a given Ta . This can be phrased in terms of the pseudo-Prandtl number σ , introduced in EGL2007. As a smaller gap means a smaller σ , we thus find a decrease of Nu_ω with increasing σ , for the drivings explored in experiments, similarly as predicted [154] and found [159] for $Nu(Pr)$ in RB convection for $Pr > 1$.

5.4.2 Rossby-number dependence

In this subsection, the effect of outer cylinder rotation on angular velocity transport will be studied. Previous experimental and numerical work at $\eta = 0.714$ [43, 44, 112, 123] revealed the existence of an optimum transport where, for a given Ta , the transport of momentum is highest at a Rossby number Ro_{opt}^{-1} , that depends on Ta and

Figure 5.3: The global system response for pure inner cylinder rotation as function of the driving Ta . (a) Plot of $Nu_\omega - 1$ versus Ta for both simulations (points on the left of the graph) and experiments (lines on the right of the graph). Numerical data from Ref. [112] for $\eta = 0.714$ have been added to these figures. (b) Plot of the compensated Nusselt $((Nu_\omega - 1)/Ta^{1/3})$ versus Ta , with added lines with scaling law $Ta^{0.20}$ and $Ta^{0.39}$ to guide the eye. (c,d) Zoom-in of the experimental data.

saturates around $Ta \sim 10^{10}$. In this subsection, this work will be extended to the other values of η .

Figure 5.4 shows the results of the numerical exploration of the Ro^{-1} parameter space between $Ta = 4 \cdot 10^4$ and $Ta = 2.5 \cdot 10^7$. The shape of the $Nu_\omega = Nu_\omega(Ro^{-1})$ curves and the position of Ro_{opt}^{-1} depend very strongly on η in the Ta range studied in numerics. For the largest gap (i.e., $\eta = 0.5$), the optimum can be seen to be in the counter-rotating range (i.e., $Ro^{-1} < 0$) as long as Ta is high enough. On the other hand, for the smallest gap (i.e., $\eta = 0.909$), the optimum is at co-rotation (i.e., $Ro^{-1} > 0$) in the whole region studied. The other values of studied η reveal an intermediate behaviour. Optimum transport is located for co-rotation at lower values of Ta and slowly moves towards counter-rotation. For all values of η , when the driving is increased, Ro_{opt}^{-1} tends to shift to more negative values.

For two values of Ta ($Ta = 4 \cdot 10^6$ and $Ta = 10^7$) for a radius ratio $\eta = 0.5$, two distinct peaks can be seen in the $Nu_\omega(Ro^{-1})$ curve. This can be understood by looking at the flow topology. For $Ro^{-1} = 0$, three distinct rolls can be seen. However, when decreasing Ro^{-1} , the rolls begin to break up. Some remnants of large-scale structures can be seen, but these are weaker than in the $Ro^{-1} = 0$ case. Having a large-scale roll helps the transport of angular momentum, leading to the peak in Nu_ω at $Ro^{-1} = 0$. Further increasing the driving causes the rolls to also break up for $Ro^{-1} = 0$, and eliminates the anomalous peak.

The shift seen in the numerics may or may not continue with increasing Ta . The experiments conducted explore a parameter space of $10^{10} < Ta < 10^{13}$ and thus serve to explore the shift at higher driving. Figure 5.5 presents the obtained results. Figure 5.5(a) shows Nu_ω versus Ta for all measurements. Figure 5.5(b) shows the exponent γ , obtained from least-square linear fit of the log-log plots. Across the η and Ro^{-1} ranges studied, the average exponent is $\gamma \approx 0.39$. This value is used in figure 5.6 to compensate Nu_ω . The horizontality of all data points reflects the good scaling and the universality of this ultimate scaling behaviour $Nu_\omega \propto Ta^{0.39}$.

To determine the optimal rotation ratio for the experimental data, a Ta -averaged compensated Nusselt number $\langle Nu_\omega / Ta^{0.39} \rangle_{Ta}$ was used. This is defined as

$$\langle Nu_\omega / Ta^{0.39} \rangle_{Ta} = \frac{1}{Ta_{max} - Ta_{co}} \int_{Ta_{co}}^{Ta_{max}} Nu_\omega / Ta^{0.39} dTa, \quad (5.7)$$

where Ta_{max} is the maximum value of Ta for every (η, Ro^{-1}) dataset, and Ta_{co} is a cut-off Ta number used for the larger η ($Ta_{co} = 2 \cdot 10^{11}$ for $\eta = 0.833$, and $Ta_{co} = 3 \cdot 10^{10}$ for $\eta = 0.909$) to exclude the initial part of the $Nu_\omega / Ta^{0.39}$ data points, which seem to have a different scaling for some of the values of Ro^{-1} explored. For the smaller values of η , we have $Ta_{co} = Ta_{min}$, the minimum value of Ta for every (η, Ro^{-1}) dataset. An error bar on this average is estimated as one standard deviation of the data from the computed average.

Figure 5.4: Plots of $Nu_{\omega} - 1$ versus Ro^{-1} for the four values of η studied numerically, (a) $\eta = 0.5$, (b) $\eta = 0.714$, (c) $\eta = 0.833$, (d) $\eta = 0.909$. The shape of the curve and the position of the maximum depend very strongly on both Ta and η .

Figure 5.5: (a) Plot of Nu_ω versus Ta for all values of η and Ro^{-1} studied in experiments. (b) The exponent γ of the scaling law $Nu_\omega \propto Ta^\gamma$ for various Ro^{-1} , obtained by a least-square linear fit in log-log space. The average value of γ for each η is represented by the dashed lines, while the solid line represents the average value of $\gamma = 0.39$ for all η , which will be used for compensating Nu_ω .

Figure 5.6: Plots of $Nu_\omega/Ta^{0.39}$ versus Ta for all explored values of Ro^{-1} for the four studied values of η : (a) $\eta = 0.714$, (b) $\eta = 0.769$, (c), $\eta = 0.833$, (d) $\eta = 0.909$. The scaling law $Nu_\omega \sim Ta^{0.39}$ is seen to approximately hold throughout the whole parameter space explored. No trends can be appreciated which would lead us to expect further shift of the optimum with increased driving. Dashed lines, indicating the cut-off regions used for determining $\langle Nu_\omega/Ta^{0.39} \rangle_{Ta}$ have been plotted for $\eta = 0.833$ and $\eta = 0.909$.

Figure 5.7(a,b) shows $\langle Nu_\omega / Ta^{0.39} \rangle_{Ta}$ as a function of Ro^{-1} or alternatively of a for the four values of η considered in experiments. The increased driving changes the characteristics of the flow. This is reflected in the very different shapes of the Ro^{-1} dependence of Nu_ω when comparing figures 5.4 and 5.7, and in the shift of Ro_{opt}^{-1} .

To summarize these effects, figure 5.8 presents both the 95% peak width ΔRo_{max}^{-1} and the position of optimal transport Ro_{opt}^{-1} determined as the realization with the maximum torque as a function of Ta and η obtained from numerics as well as the asymptotic value from experiments. The peak width ΔRo_{max}^{-1} is defined as

$$\Delta Ro_{max}^{-1} = \frac{\int_{Ro_{-0.95}^{-1}}^{Ro_{0.95}^{-1}} Nu_\omega(Ro^{-1}) dRo^{-1}}{\max(Nu_\omega - 1)} \quad (5.8)$$

where $Ro_{-0.95}^{-1}$ and $Ro_{0.95}^{-1}$ are the values of Ro^{-1} for which Nu_ω is 95% of the peak value.

The 95% peak width can be seen to vary with driving, reflecting what is seen in figure 5.4. The shape of the Ro^{-1} - Nu_ω curve is highly dependent on both η and Ta . Also, Ro_{opt}^{-1} shows a very large variation across the Ta range studied in numerics. The shift of the Ro_{opt}^{-1} with Ta is expected to continue until it reaches the value found in the experiments. This can be seen in figure 5.4 for $\eta = 0.5$ to $\eta = 0.833$. For $\eta = 0.909$, the trend seems to change for the last point. However, this is due to the very large and flat peak of the $Nu_\omega(Ro^{-1})$ curve— this can also be seen in the left panel and in figure 5.4(d).

One may also ask the question: has the value of Ro_{opt}^{-1} already saturated in our experiments? Figure 5.6 shows the trend for Nu_ω for increasing Ta . This trend does not seem to vary much for different values of Ro^{-1} . Therefore, we expect the value of Ro_{opt}^{-1} to have already reached saturation in our experiments.

We can compare these new experimental results to the available results from the literature, the speculation made in Ref. [77] and the prediction made in Ref. [156] for the dependence of the saturated a_{opt} on η . This is shown in figure 5.9. Error bars are either taken from the literature, or computed by fitting a quadratic polynomial to $Nu_\omega(a)$ for three values of a below the maximum, i.e. $a < a_{opt}$, three values of a after the maximum, i.e. $a > a_{opt}$ and the six values of a , and comparing the different results obtained. Both dependences are shown to deviate substantially from the experimental results obtained in the present work. Even if the speculation from Ref. [77] appears to be better for this η range, for previous experimental data at $\eta = 0.5$, it is clearly difference from the experimentally measured value for optimal transport by Ref. [80].

This section has shown that the radius ratio has a very strong effect on the global response and especially on optimal transport. Significantly increased transport for co-rotation has been found at the lowest drivings based on the DNS results. This finding was already reported in [112] for $\eta = 0.714$, but the transport increase was

Figure 5.7: Plots of $\langle Nu_\omega / Ta^{0.39} \rangle_{Ta}$ versus either (a) Ro^{-1} or (b) a (bottom) at the cut-off region highlighted in figure 5.6 for the values of η studied experimentally. Insets containing a zoom-in around the optimum have been added for clarity. Error bars indicate one standard deviation from the mean value, and are too small to be seen for most data points. There is a strong η -dependence of the curve $Nu_\omega / Ta^{0.39}$ versus Ro^{-1} , even at the largest drivings studied in experiments. Optimal transport is located at $Ro_{opt}^{-1} = -0.20$ for $\eta = 0.716$, $Ro_{opt}^{-1} = -0.15$ for $\eta = 0.769$, $Ro_{opt}^{-1} = -0.10$ for $\eta = 0.833$ and $Ro_{opt}^{-1} = -0.05$ for $\eta = 0.909$, corresponding to $a \approx 0.33 - 0.35$ for all values of η . In the bottom panel, the maximum of the graph is less pronounced, i.e. it becomes more flat with increasing η . In the limit $\eta \rightarrow 1$, a does not tend to a finite limit, while Ro^{-1} does. This result highlights the advantage of using Ro^{-1} instead of a as a control parameter.

Figure 5.8: (a) The 95% peak width $\Delta Ro_{\max}^{-1}$ versus Ta for the four values of η analyzed in numerics. The peak width can be seen to vary with driving, and for smaller gaps is larger for larger values of Ta . (b) Plot of Ro_{opt}^{-1} versus Ta for the same four values of η . The location of the optimal transport has a very strong dependence on the driving, especially for the largest values of η . As driving increases beyond the numerically studied range and overlaps with experiments, Ro_{opt}^{-1} should tend to the experimentally found values, represented as dashed lines in the figure. The asymptote for $\eta = 0.5$ is obtained from Ref. [80]. The trend appears to be less clear for $\eta = 0.909$, but this might be understandable from the peak width at the highest driving Ta .

Figure 5.9: State-of-the-art data for $a_{opt}(\eta)$, from both experiments [31, 44, 80] and numerics [123]. The speculation of van Gils *et al.* [77] and the prediction of Brauckmann & Eckhardt [156] are plotted as lines on the graph. The new experimental results deviate up substantially from both predictions, even when taking into account error bars.

marginal. For $\eta = 0.833$, and especially for $\eta = 0.909$ the transport can be increased up to three times. The shift of Ro_{opt}^{-1} has also been seen to be much bigger and to happen in a much slower way for smaller gaps. The reason for this will be studied in §5.5, using the local data obtained from experiments and numerics.

5.5 Local results

In this section, the local angular velocity profiles will be analysed. Angular velocity is the transported quantity in TC flow and shows the interplay between the bulk, where the transport is convection-dominated, and the boundary layers, where the transport is diffusion dominated. Numerical profiles and experimental profiles obtained from LDA will be shown. The angular velocity gradient in the bulk will be analysed and connected to the optimal transport. In addition, the boundary layers will be analysed and compared to the results from the analytical formula from EGL 2007 for the BL thickness ratio in the non-ultimate regime.

5.5.1 Angular velocity profiles

Angular velocity ω profiles obtained from numerics are shown in figure 5.10. Results are presented for four values of η and selected values of Ro^{-1} at $Ta = 2.5 \cdot 10^7$ (and $Ta = 2.39 \cdot 10^7$ for $\eta = 0.714$). Experimental data obtained by using LDA are shown in figure 5.11 for three values of η : (a) $\eta = 0.714$ for $Re_i - Re_o = 10^6$, (b) $\eta = 0.833$ for $Ta = 5 \cdot 10^{11}$, and (c) $\eta = 0.909$ for $Ta = 1.1 \cdot 10^{11}$.

The different radius ratios affect the angular velocity profiles on both boundary layers, as the two boundary layers are more asymmetric for the wide gaps; and they affect the bulk, as the bulk angular velocity is smaller for wide gaps. These effects will be analysed in the next sections.

5.5.2 Angular velocity profiles in the bulk

We now analyse the properties of the angular velocity profiles in the bulk. We find that the slope of the profiles in the bulk is controlled mainly by Ro^{-1} and less so by Ta . This can be understood as follows. The Taylor number Ta acts through the viscous term, dominant in the boundary layers, while Ro^{-1} acts through the Coriolis force, present in the whole domain. These results extend the finding from Ref. [112] to other values of η .

To further quantify the effect of Ro^{-1} on the bulk profiles, we calculate the gradient of $\langle \tilde{\omega} \rangle_z$. For the DNS data, this is done by numerically fitting a tangent line to the profile at the inflection point using the two neighbouring points on both sides (at a distance of 0.01 – 0.02 r-units); such a fit is shown in figure 5.12(a).

Figure 5.10: Azimuthally, axially and temporally averaged angular velocity $\langle \bar{\omega} \rangle_z$ versus radius $\tilde{r}$ for (a) $\eta = 0.5$, (b) $\eta = 0.714$, (c) $\eta = 0.833$ and (d) $\eta = 0.909$. Data are for $Ta = 2.5 \cdot 10^7$ ($Ta = 2.39 \cdot 10^7$ for $\eta = 0.714$) and selected values of Ro^{-1} . For smaller η , the ω bulk profiles differ more from a straight line, and have, on average, a smaller value.

Figure 5.11: Angular velocity profiles obtained by LDA for either (a) $\eta = 0.716$ at $Re_i - Re_o = 10^6$, (b) $\eta = 0.833$ at $Ta = 5 \cdot 10^{11}$ and (c) $\eta = 0.909$ at $Ta = 1.1 \cdot 10^{11}$, to explore different dependencies in parameter space. Data is taken at a fixed axial height (i.e. the cylinder mid-height, $z = L/2$), but as the Taylor number Ta is much larger than in the numerics, the axial dependence is much weaker.

Figure 5.12: An example of the two fitting procedures for the bulk angular velocity gradient and for boundary layer thicknesses done on the DNS data is shown here. Both panels show the θ , z , and t averaged azimuthal velocity and angular velocity for $\eta = 0.5$, $Ta = 1 \cdot 10^7$, and pure inner cylinder rotation. In (a) a line is fitted to the bulk of the angular velocity to obtain the bulk gradient. The dashed lines indicate the EGL approximation. (b) shows the three-lines-fit to the whole profile to obtain the width of its boundary layers, used in §5.5.3. Both bulk fits are done at the inflection point, but for different variables ($\bar{\omega}$ or $\bar{u}_\theta$), which gives slightly different slopes (and intersection points).

As the spatial resolution of the LDA data is more limited, the fit is done differently. A linear regression to the ω profile between $0.2 < \tilde{r} < 0.8$ is carried out. The larger range of $\tilde{r}$ is chosen in experiments because: (i) the boundary layers are small enough due to the high Ta that they are outside of the fitting range, and (ii) the fluctuations of the data are much higher in experiments, especially for the LDA of the narrow gaps ($\eta = 0.833$ and $\eta = 0.909$). From this regression, we calculate $\langle \bar{\omega} \rangle_z$, and an error taken from the covariance matrix of the fit.

Figure 5.13 shows four panels, each containing the angular velocity gradient in the bulk from the numerical simulations and experiments for a given value of η . We first notice that the angular velocity gradients from experiment and numerics are in excellent agreement. Next the connection between a flat angular velocity profile and optimal transport for the highest drivings explored in the experiments can now be seen for other values of η and not just for $\eta = 0.714$ as reported previously in Ref. [43]. Once $Ro^{-1} < Ro_{opt}^{-1}$, the large-scale balance analysed in Ref. [112] breaks down, and a “neutral” surface which reduces the transport appears in the flow.

In simulations, because of resolution requirements, we are unable to drive the flow strongly enough to see a totally flat bulk profile. Also, the influence of the large-scale structures causes a small discrepancy between the flattest profile and the value of Ro_{opt}^{-1} measured from Nu_ω . This is expected to slowly disappear with increasing

Ta .

In Ref. [112], a linear extrapolation of the bulk angular velocity gradient was made to give an estimate for the case when this profile would become horizontal, i.e., $d\langle\bar{\omega}\rangle_z/d\tilde{r} = 0$, and thus give an estimate of Ro_{opt}^{-1} . For $\eta = 0.714$ this estimate agreed with the numerical result within error bars. Here, we extend this analysis for the other values of η and, as we shall see, successfully.

As in Ref. [112], an almost linear relationship between Ro^{-1} and $d\langle\bar{\omega}\rangle_z/d\tilde{r}$ can be seen. This linear relationship is extrapolated and plotted in each panel. This extrapolation gives an estimate for $Ro_{opt}^{-1}(Ta \rightarrow \infty)$, which we can compare to the experimentally determined $Ro_{opt}^{-1}(Ta \rightarrow \infty)$. For $\eta = 0.833$, $Ro^{-1}(Ta \rightarrow \infty) \approx -0.12$, corresponding to $a \approx 0.38$, is obtained; and for $\eta = 0.909$, $Ro^{-1}(Ta \rightarrow \infty) \approx -0.05$, corresponding to $a \approx 0.31$, is obtained. These values are (within error bars) also obtained for Ro_{opt}^{-1} at the large Ta investigated in experiments, namely $Ro_{opt}^{-1} = -0.10$ and -0.05 , respectively.

For $\eta = 0.5$, $Ro_{opt}^{-1}(Ta \rightarrow \infty) \approx -0.33$ is obtained, corresponding to $a \approx 0.2$. This value is consistent with the numerical results in [156], which report $a_{opt} \approx 0.2$. However, care must be taken, as fitting straight lines to the ω profiles gives higher residuals for $\eta = 0.5$ as the profiles deviate from straight lines (cf. top-left panel of figure 5.10). A fit to the “quarter-Couette” profile derived from upper bound theory [61] is much more appropriate for $\eta = 0.5$ at the strongest drivings achieved in experiments [80]. This is because the flow feels much more the effect of the curvature at small η . At the other end of the scale, the linear relationship works best for the smallest gaps, i.e. $\eta = 0.909$ (cf. figure 5.10(d)) where curvature plays a smaller effect.

To further elaborate the link between η , flat ω profiles and Ro^{-1} , data for the smallest gap $\eta = 0.95$ from Ref. [31] has been digitized, and $d\langle\bar{\omega}\rangle_z/d\tilde{r}$ has been determined for it. These data correspond to a driving of $Ta \sim 10^8 - 10^9$. Figure 5.14 shows $d\langle\bar{\omega}\rangle_z/d\tilde{r}$ against Ro^{-1} for Wendt’s data and also for the current experimental data. The flattest profile can be seen to occur for increasing (in absolute value) Ro^{-1} for larger gaps, similar to the shift of Ro^{-1} . For $\eta = 0.95$, almost no curvature is felt by the flow and flat profiles can be seen for $-0.05 < Ro^{-1} < 0$. However, adding a Coriolis force (in the form of Ro^{-1}), a large ω gradient is sustained in the bulk. This corroborates the balance between Ro^{-1} and the bulk ω gradients proposed in Ref. [112].

5.5.3 Angular velocity profiles in the boundary layers in the classical turbulent regime

As the driving is increased, the transport is enhanced. To accommodate for this, the boundary layers (BLs) become thinner and therefore the ω slopes ($\partial_r \omega$) become steeper. Owing to the geometry of the TC system an intrinsic asymmetry in the BL

Figure 5.13: Bulk angular velocity gradient $d\langle\bar{\omega}\rangle_z/d\tilde{r}$ against Ro^{-1} for the four values of η explored in simulation: (a) $\eta = 0.5$, (b) $\eta = 0.714$, (c) $\eta = 0.833$, and (d) $\eta = 0.909$. Data from experiments obtained by LDA are also plotted for the three values of η for which it was experimentally measured (green circles). For all values of η except $\eta = 0.5$, for co-rotation and slight counter-rotation there is once again an almost linear relationship between Ro^{-1} and $d\langle\bar{\omega}\rangle_z/d\tilde{r}$. A black straight line is added to extrapolate this relationship in order to estimate Ro_{opt}^{-1} . A plateau, in which the radial gradient of $\langle\bar{\omega}\rangle_z$ is small can be seen around optimal transport, indicating a large convective transport of angular velocity.

Figure 5.14: Bulk angular velocity gradient $d\langle\bar{\omega}\rangle_z/d\tilde{r}$ against Ro^{-1} for the three values of η explored in experiments, and for $\eta = 0.95$ (digitized from Ref. [31]). The error bars of Wendt’s data are larger due to the quality of the digitization. As seen previously, the flattest profile occurs around weak counter-rotation, for all values of η including $\eta = 0.95$.

layer widths is present. More precisely, the exact relationship $\partial_r\langle\omega\rangle|_o = \eta^3\partial_r\langle\omega\rangle|_i$ holds for the slopes of the boundary layers, due to the r independence of J^ω , cf. EGL 2007 and equation (4.4).

An analysis of the boundary layers was not possible in the present experiments because the present LDA measurements have insufficient spatial resolution to resolve the flow in the near-wall region. Therefore, only DNS results will be analysed here. In simulations the driving is not as large as in experiments, and as a consequence the shear in the BLs is expected to be not large enough to cause a shear instability. This means that the BLs are expected to be of Prandtl–Blasius (i.e. laminar) type, even if the bulk is turbulent. On the other hand, in the experiments both boundary layers and bulk are turbulent, i.e. the system is in the “ultimate regime”.

Using the DNS data, we can compare the ratio of the numerically obtained boundary layer widths with the analytical formula for this ratio obtained by EGL 2007 for laminar boundary layers, namely

$$\frac{\lambda_\omega^o}{\lambda_\omega^i} = \eta^{-3} \frac{|\omega_o - \omega_{bulk}|}{|\omega_i - \omega_{bulk}|}, \quad (5.9)$$

where the value of ω_{bulk} is some appropriate value in the range from which the angular velocity at the inflection point of the profile might be chosen, i.e., the point at which the linear bulk profile fit was done to obtain λ_ω^o and λ_ω^i . The value ω_{bulk} is

taken from the numerics, and may bias the estimate.

To calculate the boundary layer thicknesses, the profile of the mean azimuthal velocity $\langle \bar{u}_\theta \rangle_z$ is approximated by three straight lines, one for each boundary layer and one for the bulk. For the boundary layers the slope of the fit is calculated by fitting (by least-mean-squares) a line through the first two computational grid points. For the bulk, first the line is forced to pass through the grid point that is numerically closest to the inflection point of the profile. Then its slope is taken from a least-mean-squares fit using two grid points on both sides of this inflection point. The respective boundary layer line will cross this bulk line at a point that then defines the thickness of that boundary layer.

The results obtained for $\lambda_\omega^o/\lambda_\omega^i$ both from equation (5.9) and directly from the simulations are shown in figure 5.13. Results are presented for the four values of η and only for the highest value of Ta achieved in the simulations. The boundary layer asymmetry for counter-rotating cylinders (i.e., $Ro^{-1} < 0$) grows with larger gaps. This is to be expected, as the η^{-3} term is much larger (≈ 8) for the largest gap as compared to the smallest gap (≈ 1.3). This is consistent with the η and thus σ restriction in EGL 2007 to a range of smaller gap widths.

As noted already in Ref. [112] we find that the fit is not satisfactory for co-rotation (i.e., $Ro^{-1} > 0$) at the lowest values of η , but is satisfactory for counter-rotation (ie. $Ro^{-1} < 0$). In EGL 2007, equation (5.9) is obtained by approximating the profile by three straight lines, two for the BLs and a constant ω line for the bulk. Therefore, we expect the approximation to hold best when the bulk has a flat gradient. For co-rotating cylinders and strongly counter-rotating cylinders, the bulk has a steep gradient (see figure 5.13), but characteristically different shapes. The only free parameter in equation (5.9) is ω_{bulk} , which is chosen to be ω at the point of inflection. Owing to the different shapes of the ω profiles, this choice seems more correct for counter-rotating cylinders, as there is a clear inflection point in the profile. On the other hand, for co-rotating cylinders, the profile appears to be more convex-like, and there the choice of ω_{bulk} as the inflection point induces errors in the approximation (cf. figure 5.12(a)). For $\eta = 0.5$, the error from the constant- ω approximation is even more pronounced, and the formula fails.

For co-rotation the boundary layers are approximately of the same size, and the ratio $\lambda_\omega^o/\lambda_\omega^i$ is very close to 1. If one inverts equation (5.9) by approximating this ratio by 1, an estimate of what the angular velocity will be in the bulk due to the boundary layer slope asymmetry is obtained

$$\omega_{bulk} = \frac{\eta^3}{1 + \eta^3}, \quad (5.10)$$

corresponding to

Figure 5.15: Plots of $\lambda_{\omega}^o/\lambda_{\omega}^i$ from simulations (dots) and from equation (5.9) (solid lines) versus Ro^{-1} for the four values of η studied numerically: (a) $\eta = 0.5$, (b) 0.714, (c) 0.833 and (d) 0.909 at $Ta = 2.5 \cdot 10^7$. The numerical results and the estimate from (5.9) match very well for larger values of η and especially for counter-rotating cylinders ($1/Ro < 0$). The vertical black dashed line indicates Ro_{opt}^{-1} , while the vertical dash-dot line indicates the 90% peak width on the counter-rotation side. $\lambda_{\omega}^o/\lambda_{\omega}^i$ begins to deviate from 1 between the two lines, as counter rotation increases. The asymmetry between the boundary layers can be seen to be larger for smaller values of η , which is expected as equation (5.9) contains the explicit factor η^{-3} .

Figure 5.16: ω_{bulk} as a function of η , taken from both eq. 5.10 and from the crosspoints of the ω -profiles in figure 5.10. The trend is the same in both data sets. A smaller value of η decreases the value of the bulk angular velocity.

$$\omega_{bulk}^{\ell} = \frac{-\omega_o^{\ell} + \eta^3 \omega_i^{\ell}}{1 + \eta^3}, \quad (5.11)$$

in the laboratory frame. This expression gives an estimate for ω_{bulk} when the profile is flattest, and has been represented graphically in figure 5.10. Indeed, one can take this estimate (e.g. ≈ 0.27 for $\eta = 0.714$) and compare it with figures 5.10 and 5.11. We note that the value of ω in the bulk for the flattest profile in the numerics (at $Ro^{-1} \approx Ro_{opt}^{-1}(Ta)$) lies around ω_{bulk} . We can also note that the profiles for $Ro^{-1} > Ro_{opt}^{-1}$ approximately cross each other at the same point, and this point has a value of $\omega \approx \omega_{bulk}$. This effect can only be seen in the numerics, as these approximations break down once the boundary layers become turbulent. The crossing points of the curves are taken as an estimate for ω_{bulk} , and this is represented against 5.10 in figure 5.16.

To understand why the boundary layers are of approximately the same thickness despite the different initial slopes at the cylinders, one has to go back to §4.4. The angular velocity current has a diffusive part and a convective part. Per definition in the boundary layer the diffusion dominates, and in the bulk the convection does. Thus the boundary layer ceases when convection becomes significant. But convection is controlled by the wind. Thus in essence the boundary layer size is controlled by the wind and not immediately by the initial slope at the wall. Owing to continuity, if the rolls penetrate the whole domain the wind may be expected to be the same close to

the inner and close to the outer cylinder. This suggests that the flow organizes itself in such a way that the boundary layer extensions (or widths) might be similar, even if the initial slopes at the walls are different.

What happens for counter-rotation, or more precisely when $Ro^{-1} < Ro_{opt}^{-1}$? For Ro^{-1} below the optimum Ro_{opt}^{-1} , a so-called neutral surface will be present in the flow, which separates the Rayleigh-stable and -unstable areas. The wind drastically changes in the Rayleigh-stable areas [112], leading to very different wind velocities close to the outer and inner cylinder, respectively. The wind at the outer cylinder will be weaker, as the rolls cannot fully penetrate the Rayleigh-stable domain. This means that the outer cylinder boundary layer will extend deeper into the flow, in accordance with is seen in figure 5.15.

5.6 Summary and conclusions

Experiments and direct numerical simulations (DNS) were analysed to explore the effects of the radius ratio η on turbulent Taylor-Couette flow. Numerical results corresponding to Taylor numbers in the range of $10^4 < Ta < 10^8$ alongside with experiments in a Taylor number range of $10^{10} < Ta < 10^{12}$ were presented for four values of the radius ratio η .

First the influence of the radius ratio on the global scaling law $Nu_\omega \sim Ta^\gamma$ was studied. The local scaling exponent γ describing the response of the torque caused by a Taylor number increase, is barely modified by varying the radius ratio η . Indeed, in experiments a universal exponent of $\gamma \approx 0.39$ is obtained, independent of radius ratio and outer cylinder rotation. For the numerical simulations at lower Ta similar universal behaviour can be observed. The transition associated to the vanishing of coherent structures can also be appreciated at $Ta \sim 10^6$ for all values of η . Before this transition local exponents of $\gamma \approx 0.33$ are seen and after the transition these decrease to about $\gamma \approx 0.2$.

The radius ratio does play a very important role in optimal transport. At smaller gaps, i.e., for larger η , at the lower end of the Ta range a very large increase in transport for co-rotating cylinders can be seen. The shift towards the asymptotic optimal transport happens in a much slower way for small gaps, but this shift is seen for all studied radius ratios. For the largest gap ($\eta = 0.5$), optimal transport for pure inner cylinder rotation at the lowest drivings is obtained. The shift towards the asymptotic value happens suddenly, as two peaks can be seen in the Nu_ω versus Ro^{-1} curve, and one of the peaks becomes larger than the other one as driving increases. This might point in the direction of different phenomena and transitions in the flow topology happening at larger gaps. Finally, the asymptotic values of Ro_{opt}^{-1} obtained in experiments were compared to the speculation of van Gils *et al.* [77] and the prediction of Brauckmann & Eckhardt [156]. Both of the models were found to deviate from

Radius ratio, η	Extrapolated Ro_{opt}^{-1}/a_{opt}	Measured Ro_{opt}^{-1}/a_{opt}
0.5	-0.33/0.20	-/-
0.714	-0.20/0.33	-0.20/0.33
0.769	-/-	-0.20/0.36
0.833	-0.12/0.41	-0.10/0.37
0.909	-0.05/0.34	-0.05/0.34

Table 5.2: Summary of values obtained for Ro_{opt}^{-1} and a_{opt} through both the extrapolation of $d\langle\bar{\omega}\rangle_z/d\tilde{r}$ (§5.5.2) and direct measurement of the torque (§5.4.2) for the various values of η studied in this chapter.

experimental and numerical results.

When looking at the local results, as in Ref. [112] we can link the optimal transport in the smallest gaps to a balance between Coriolis forces and the inertia terms in the equations of motion. The flattest profiles in the bulk are linked to optimal transport in experiments. With the numerics the extrapolation presented in Ref. [112] for predicting optimal transport was extended to other radius ratios. It is found to work well for all selected η except for $\eta = 0.5$. At this η , i.e., for the largest gap considered here, the most obvious problem is that the ω profiles strongly feel the effect of curvature difference at the inner and outer cylinders, and a straight-line fit (i.e. constant ω) to the bulk is not appropriate. There may be additional reasons for this discrepancy and optimal transport in large gap TC flow requires more investigation. A summary of the results for determining Ro_{opt}^{-1} using both the experimentally measured torque maxima from §5.4.2 and the numerical extrapolation from §5.5.2 are presented in table 5.2.

Finally, the boundary layers have been analysed. The outer boundary layer is found to be much thicker than the inner boundary layer when $Ro^{-1} < Ro_{opt}^{-1}$. We attribute this to the appearance of Rayleigh-stable zones in the flow. This prevents the turbulent Taylor vortices from covering the full domain between the cylinders. As the boundary layer size is essentially determined by the wind, if the rolls penetrate the whole domain (which is the case for $Ro^{-1} > Ro_{opt}^{-1}$), the two boundary layers are approximately of the same size. If the rolls do not penetrate the full domain, the outer boundary layer will be much larger than the inner boundary layer, in accordance with the smaller initial slope of $\omega(r)$ at the cylinder walls.

In this work, simulations and experiments have been performed on a range of radius ratios between $0.5 \leq \eta \leq 0.909$. Insights for the small gaps seem to be consistent with what was discussed in Ref. [112]. However, for $\eta = 0.5$ the phenomenon of optimal transport appears to be quite different. Therefore, our ambition is to extend the DNS towards values of η smaller than 0.5 to improve the understanding of that regime.

6

Boundary layer dynamics at the transition between the classical and the ultimate regimes

*

The fact that we live at the bottom of a deep gravity well, on the surface of a gas covered planet going around a nuclear fireball 90 million miles away and think this to be normal is obviously some indication of how skewed our perspective tends to be.

—DOUGLAS ADAMS

Direct numerical simulations of turbulent Taylor-Couette flow are performed up to inner cylinder Reynolds numbers of $Re_i = 10^5$ for a radius ratio of $\eta = r_i/r_o = 0.714$ between the inner and outer cylinder. With increasing Re_i , the flow undergoes transitions between three different regimes: (i) a flow dominated by large coherent structures, (ii) an intermediate transitional regime, and (iii) a flow with developed turbulence. In the first regime the large-scale rolls completely drive the meridional flow while in the second one the coherent structures recover only on average. The presence of a mean flow allows for the coexistence of laminar and turbulent boundary layer dynamics. In the third regime the mean flow effects fade away and the flow becomes dominated by plumes. The effect of the local driving on the azimuthal and angular velocity profiles is quantified, in particular we show when and where those profiles develop.

*Based on: R. Ostilla-Mónico, E. P. van der Poel, R. Verzicco, S. Grossmann and D. Lohse, *Boundary layer dynamics at the transition between the classical and the ultimate regime of Taylor-Couette flow*, Phys. Fluids., **26**, 015114 (2014).

6.1 Introduction

Taylor-Couette flow (TC), the flow between two independently rotating coaxial cylinders, and Rayleigh-Bénard flow (RB), the thermally driven flow in a fluid cell heated from below and cooled from above, are twin flows [160], with analogous global balances [161]. The temperature difference between the plates in RB flow corresponds to the different rotation rates of the inner and outer cylinders in TC flow (hereafter referred to as drivings of the flow).

Kraichnan [74] postulated that RB flow would reach an asymptotic or “ultimate” state if the driving is sufficiently large. In this ultimate regime, turbulence is fully developed throughout both bulk and boundary layers. In that regime, heat transport, which for small driving is limited by the laminar boundary layers, has no explicit, unit-wise dependence on viscosity. The ultimate-regime scaling laws relating heat transfer and flow driving are expected to be extendible to arbitrarily large Rayleigh numbers, like those present in both geo- and astrophysical flows.

Indeed, such transitions in the scaling laws expected in this ultimate state have recently been found in experiments for heat transport and wind in RB flow [75, 76] and angular velocity transport and wind in TC flow [41, 45, 77], confirming the existence of the ultimate regime, 50 years after Kraichnan’s prediction.

RB and TC flow are expected to show the transition to the ultimate regime when the boundary layers are sheared strongly enough so that they undergo a shear instability and become turbulent [77–79]. A universal characteristic of turbulent boundary layers is the so-called law of the wall [9, 162, 163]. For wall distances much larger than the internal length scale and much smaller than the outer length scale, the mean velocity profile has a logarithmic dependence on the distance to the wall. This region has been found in many different types of flows, both experimentally [164] and numerically [94]. For further details on the empirical evidence for the universality of this region, we refer the reader to recent reviews by Smith, McKeon & Marusic [11] and Jiménez [12].

As in other flows, in TC and RB turbulent boundary layers are expected to produce this characteristic logarithmic signature, not only in the mean velocity, but, in RB, also in the mean temperature profile [79]. These log-layers, which extend significantly into the flow, have been experimentally detected [54] and measured [82] in TC. Later, in a theoretical investigation, Grossmann *et al.* [165] pointed out that in addition to the axial velocity, it should be the (properly shifted) angular velocity which should be closest to a log-behaviour, rather than the azimuthal velocity, but both should show curvature corrections. In RB flow, logarithmic mean temperature profiles were measured [166] and theoretically accounted for [79] beyond the onset of the ultimate regime, which suggested that indeed, in the ultimate regime, the boundary layers are fully turbulent [166]. However, unexpectedly, log-layers (in the bulk) were also found below the ultimate transition in the global scaling laws [166].

In view of all these findings we readdress the mechanism of log-layer formation since at those low drivings of the flow prior to the BL transition the shear alone is not large enough to generate turbulent boundary layers. Based on direct numerical simulations and on local (i.e., z dependent) velocity profiles we will develop a novel viewpoint of the laminar-turbulence transitions, linking them not only to either bulk or boundary layer transitions as has been hitherto done, but also to the interaction between bulk and boundary layers. This viewpoint will lead to the following picture: for low driving, the boundary layers are of Prandtl–Blasius type (PB) [167]. For the strongest driving, in the ultimate regime, the boundary layers are turbulent and coherence is lost in the bulk. Both these regimes were predicted within the unifying theory of Eckhardt, Grossmann and Lohse [161]. In between these two regions we will identify yet another regime which we denote as transitional regime, where PB-type and turbulent boundary layers coexist. In this transitional regime, the local scaling laws do not yet show the characteristic increase of transport which is seen in the ultimate regime.

Direct numerical simulations (DNS) of the Taylor-Couette system have been performed by numerically integrating the Navier–Stokes equations using a second-order finite-difference code [103]. Simulations give access to the complete flow field, and this allows for the analysis of the different flow regimes and for an identification of the mechanisms which lead to the transitions. The flow is driven by the rotation rate difference of the inner and outer cylinders, which can be expressed in non-dimensional form as the Taylor number $Ta = \frac{1}{4}\sigma(r_o - r_i)^2(r_o + r_i)^2\omega_i^2\nu^{-2}$. Here r_o and r_i are the radii and ω_o and ω_i the angular velocities of the outer and inner cylinder, respectively. Ta is the analogue of the Rayleigh number in RB while $\sigma = [(r_o + r_i)/(2\sqrt{r_i r_o})]^4$ can be considered as a geometrical Prandtl number, which gives the relationship between the “wind” (i.e. u_r and u_z) boundary layer, and the angular velocity boundary layer [161]. The response of the system is the torque required to drive the cylinders. It can be non-dimensionalized as a pseudo–Nusselt number [161] $Nu_\omega = T/T_{pa}$ where T is the torque and T_{pa} the torque required to drive the cylinders when the flow is purely azimuthal. We also define the non-dimensional radius $\tilde{r}$ to be $\tilde{r} = (r - r_i)/(r_o - r_i)$ and the non-dimensional height $\tilde{z}$ to be $\tilde{z} = z/(r_o - r_i)$.

For the present simulations, the geometry of the system will have a fixed radius ratio of $\eta = r_i/r_o = 0.714$, and periodic boundary conditions in the axial direction. For the smaller Ta , data will be taken from Ref. [112] (§4). These originate from a simulation of the full domain, with $\Gamma = L/(r_o - r_i) = 2\pi$, where L is the axial domain length. For the largest Ta , a reduced geometry has been used: This is done in the spirit of Brauckmann and Eckhardt [123], where it is shown that (i) one pair of vortices in the axial direction gives the same first order statistics as three pairs and (ii) forcing the system to have a rotational symmetry of order 6 does not affect the mean flow statistics. Accordingly, for the largest Ta , i.e. $Ta > 10^8$, the aspect

ratio has been taken as $\Gamma = 2\pi/3$ and a rotational symmetry of order 6 is imposed on the system. This reduces the computational requirements by a factor of 20 and allows us to perform the largest Ta -range simulations. The outer cylinder will be kept at rest, and only the inner cylinder will drive the flow. A uniform grid is used in the azimuthal and axial directions, while a Chebychev-type clustering near the cylinders is used in the radial direction. For temporal convergence, two criteria must be satisfied: simulations are run until the difference between the time-averaged torque of inner and outer cylinder is less than 1%, and the average between these two values is taken for Nu_ω . The simulations are then run for at least 40 large eddy turnover times ($d/(r_i\omega_i)$).

6.2 Results

Figure 6.1 shows the relationship between Ta and $Nu_\omega - 1$ including existing experimental [41, 168] and numerical data [112, 123] and those of the present study. Grid resolutions between $N_\theta \times N_r \times N_z = 256 \times 640 \times 512$ and $N_\theta \times N_r \times N_z = 512 \times 800 \times 1024$ were used, where N_i is the number of grid points in the i -th direction. We refer the reader to the end of the chapter for further details of all the numerical simulations not covered in Ref. [112]. After the onset of Taylor vortices at $Ta \approx 10^4$ and up to $Ta \approx 3 \cdot 10^6$, a laminar regime with a scaling law of $(Nu_\omega - 1) \sim Ta^{1/3}$ is found. With increasing Ta , time dependence of the flow sets in, and the large-scale coherent structures break up into smaller structures. The vortex topology changes from large scale rolls in the center of the flow to hairpin vortices near the boundary layer [112]. The transition to the ultimate regime appears around $Ta \approx 3 \cdot 10^8$, when the exponent of the scaling law $(Nu_\omega - 1) \sim Ta^\alpha$ grows to values $\alpha > 1/3$. As mentioned previously, this has been linked to a transition towards the turbulent state of the boundary layer [77–79].

However, the situation is more complicated. Between the laminar and ultimate regimes there is a transitional regime in which a mixture of both laminar and turbulent boundary layers exists. Analysing the simulated flow shows that the large-scale wind generated by the coherent vortex pairs interacts with the angular velocity boundary layers, and induces regions where an axial pressure gradient is present. This pressure gradient is either favourable, and the flow is accelerated, or adverse, and the flow is decelerated. In the favourable pressure gradient case, the boundary layer tends to remain laminar even for intense shear rates owing to the stabilizing action of the pressure gradient. We wish to emphasize that this pressure gradient comes from the wind boundary layer, i.e. that one of the axial velocity, and acts on the ω -boundary layer. Only when Ta is large enough such that the large scale vortices are weakened and eventually fade away, can the boundary layers be turbulent all over the axial extent, giving rise to the ultimate regime. This can be understood because the shear

Figure 6.1: Compensated Nu_w vs Ta for $\eta = 0.714$. The data are from experiments (blue dots: Ref. [41], black dots: Ref. [168]) and numerical simulations (black circles: Ref. [123], red squares: older works (Ref. [112]) and present work). The three dark yellow diamonds are also numerical simulations from present work and correspond to the cases shown in figure 6.2. The thick black vertical lines indicate the transition between the regimes.

Reynolds number Re_s due to the wind scales as $Re_s \sim \sqrt{Re_w} \sim \sqrt{Re_i} \sim Ta^{1/4}$ [112] where Re_w is a “wind”-Reynolds number, while the strength of the driving scales as $Re_i \sim Ta^{1/2}$. Therefore, in TC flow, with increasing Ta , the direct driving will eventually dominate the wind shear, and the plume growing regions will extend.

Figure 6.2 shows three contour plots of the instantaneous azimuthal velocity u_θ in a meridional plane (i.e., a constant θ cut) for three values of Ta , in the laminar (a), transitional (b), and ultimate regime (c). Figure 6.2(a) shows a stationary flow field in the laminar regime. The axial structure in the azimuthal velocity is produced by the Taylor vortices. Figure 6.2(b) shows some plumes in the flow which are not ejected from the entire boundary layer but from preferential positions where the pressure gradient is adverse and separation of the boundary layer is favoured. Finally, figure 6.2(c) shows a vanishing mean flow that, due to the negligible wall pressure gradient in the axial direction, allows plumes to be ejected from all over the boundary layer.

The weakening of the coherence of the wind for increasing Ta can be seen in figure 6.3, showing two contour plots of azimuthal- and time-averaged azimuthal velocity $\langle u_\theta \rangle_{t,\theta}$ for two values of Ta . Figure 6.3(a) is for $Ta = 4.77 \cdot 10^7$, in the transitional regime. The time averaging reveals the underlying large-scale structures; plumes are ejected from the inner cylinder boundary layer at the preferential positions where the pressure gradient is adverse. These plumes travel to the outer cylinder and impact the BL at preferential positions. Similar dynamics occurs at the outer cylinder,

Figure 6.2: Contour plots of the instantaneous azimuthal velocity field u_θ for (a) $Ta = 7.04 \cdot 10^5$ (b) $Ta = 4.77 \cdot 10^7$, and (c) $Ta = 4.63 \cdot 10^9$. In (a) the flow is laminar, while in (b), the flow is in the transition region, where laminar regions in the boundary layer coexist with turbulent, plume ejection regions. Panel (c) shows fully turbulent boundary layers in which plumes have no preferential ejection regions.

where plumes are ejected from preferential positions, leading to plume impacts on the inner cylinder at preferential positions. In between the ejection and impact spots there are “quiet” regions, where the wind shears the BL, but its dynamics remains laminar. Figure 6.3(b) is for $Ta = 4.63 \cdot 10^9$ and corresponds to a flow field in the ultimate regime. Some structure is still present, but its strength is negligible, and plumes are ejected from all over the surface.

Two contour plots of azimuthally- and time-averaged r.m.s. of the velocity fluctuations $u'_\theta = (\langle u_\theta^2 \rangle_{t,\theta} - \langle u_\theta \rangle_{t,\theta}^2)^{1/2}$ for two values of Ta can be seen in figure 6.4. Figure 6.4(a) corresponds to the transitional regime ($Ta = 4.77 \cdot 10^7$) where the fluctuations occur in very localized regions which can be associated to the plume ejection spots. Outside this region, the fluctuations, especially in the boundary layer, are much less. The ultimate regime ($Ta = 4.63 \cdot 10^9$) is shown in figure 6.4(b) In this regime the fluctuations are not localized and are present in the entire boundary layer.

To define the BL regions more precisely, we use the following criteria: “wind sheared” regions are defined to correspond to an axial coordinate where $\max_r |\langle u_z \rangle_{t,\theta}|$ has a local minimum. Respectively, “plume-impacting/ejecting” regions correspond to regions where $\max_r |\langle u_z \rangle_{t,\theta}|$ has a local maximum. To distinguish between plume-impacting and plume-ejecting regions, the sign of u_r at the mid-gap has also been measured. When u_r is positive, i.e. directed towards the outer cylinder, plumes are ejected from the inner cylinder and impact on the outer cylinder and vice versa.

Following these criteria, figure 6.5 presents the non-dimensional local azimuthal

Figure 6.3: Contour plots of the azimuthally- and time averaged azimuthal velocity field $\langle u_\theta \rangle_{t,\theta}$ for (a) $Ta = 4.77 \cdot 10^7$, and (b) $Ta = 4.63 \cdot 10^9$. In (a) there are some remnants of large-scale structures, indicating preferential ranges for plume ejection. In (b) this large scale structure has been largely washed away by the wind, but a small structure of about 5% variation still exists.

Figure 6.4: Contour plots of the azimuthal- and time-averaged r.m.s. of the velocity fluctuations u'_θ for (a) $Ta = 4.77 \cdot 10^7$ and (b) $Ta = 4.63 \cdot 10^9$. The colour scales used is the same for both plots. In (a), the large fluctuations in the boundary layer are very localized in the plume ejection spots. In (b), the large-scale structure in the r.m.s. of the velocity fluctuation seen in (a) is almost gone and fluctuations (slightly smaller in absolute value) are distributed over the axial extent of the boundary layer.

velocity $u_{loc}^+ \equiv (u_i - \langle u_\theta \rangle_{t,\theta})/u_*$ and the non-dimensional local angular velocity $\omega_{loc}^+ \equiv (\omega_i - \langle \omega \rangle_{t,\theta})/\omega_*$ plotted against r^+ for the inner cylinder boundary layers for two values of Ta , one in the transitional, the other in the ultimate regime. For the inner cylinder, we define u_* , the friction velocity, as $u_\tau = (-v \langle \partial_r u_\theta(r_i) \rangle_{t,z,\theta})^{1/2}$, with ∂_r the derivative normal to the wall and r^+ is the non-dimensional wall distance $r^+ = (r - r_i)u_\tau/\nu$ in wall units. Conversely, for the outer cylinder, the frictional velocity is now defined as $u_\tau = (-v \langle \partial_r u_\theta(r_o) \rangle_{t,z,\theta})^{1/2}$, and the non-dimensional wall distance is defined as $r^+ = (r_o - r)u_\tau/\nu$. Finally, the frictional angular velocity ω_τ is defined as $\omega_\tau = u_{\tau,i,o}/r_{i,o}$ for the inner and outer cylinders.

In figure 6.5(a,c), which is in the transitional regime ($Ta = 4.77 \cdot 10^7$), a logarithmic region appears only for the ejection regions, while for figure 6.5(b,d) which shows the boundary layers in the ultimate regime ($Ta = 2.15 \cdot 10^9$), two log-layer ranges appear, though with different slopes, for both the wind-sheared and the plume ejection regions, but not for the plume impact region. In the ultimate regime plumes are ejected from a larger portion of the domain and thus the distinction between these two regions becomes less sharp. While for the ejection regions, a logarithmic fit seems to be better for the azimuthal velocity, for the wind-sheared regions, logarithmic fits are better for angular velocity.

Straight lines presenting a logarithmic Prandtl-von Kármán-type law of the wall (i.e. $u^+ = (1/\kappa) \log r^+ + B$ and $\omega^+ = (1/\kappa_\omega) \log r^+ + B_\omega$) were fitted through the data in the log-layer regime. This regime is expected to begin at $r^+ \approx 30$, see Ref. [15], but begins at an even lower r^+ because we do not account for the axial dependence of u_τ and ω_τ when non-dimensionalizing velocities and distances. In the ejection regions, these fits give coefficients of $\kappa = 0.82$ ($Ta = 4.77 \cdot 10^7$) and $\kappa = 0.85$ ($Ta = 2.15 \cdot 10^9$). For the wind-sheared region, the ω -fit gives coefficients of $\kappa_\omega = 0.51$ and $B_\omega = 4.9$ at $Ta = 2.15 \cdot 10^9$. In the ejection region, κ shows at most a weak dependence on Ta . Also, these coefficients deviate significantly from the classical von Kármán constant $\kappa = 0.4$. This is not surprising since the value $\kappa = 0.4$ was obtained for a zero pressure gradient boundary layer. Consistently here we obtain values closer to $\kappa = 0.4$ in the wind-sheared region where the wall pressure gradient is zero on average and switches from favorable to adverse.

What distinguishes the ultimate regime from the transitional regime physically? The transition from the laminar Taylor vortex regime (which includes the wavy Taylor vortex) to the transitional regime, is sharp and can be associated with the onset of time dependence in the instantaneous Nusselt number, which happens when the wavy Taylor vortices transition to the modulated wavy Taylor vortices (MWTV) regime [37]. The MWTV regime is only stable for a small range of Ta , and further increasing Ta causes the MWTV to become turbulent Taylor vortices as the breakup of coherence takes place, and the system transitions to chaos [37, 112]. On the other hand, the transition from the transitional regime to the ultimate regime resembles more a saturation

Figure 6.5: Local inner cylinder boundary layer profiles of u_{loc}^+ vs. r^+ for (a) $Ta = 4.77 \cdot 10^7$ and (b) $Ta = 2.15 \cdot 10^9$. (c,d): Local inner cylinder boundary layer profiles of ω_{loc}^+ vs. r^+ for (c) $Ta = 4.77 \cdot 10^7$ and (d) $Ta = 2.15 \cdot 10^9$. In (a) and (c) a clear distinction between the three regions can be seen. A log-layer appears for the ejection regions at around $10 < r^+ < 50$, but not in the other regions. This log-layer was fitted with a straight line (in the log-lin plot) with $\kappa = 0.82$ and $B = 3.4$ for u_{loc}^+ and $\kappa_\omega = 0.65$ and $B_\omega = 2.9$ for ω_{loc}^+ . In (b) and (d), this log-layer appears also for the wind-sheared region, and now extends more into the flow, up to $r^+ < 300$. The difference in profiles between the three regions is smaller. In the ultimate regime, the ejection log-layer was fitted with a straight line of coefficients $\kappa = 0.85$ and $B = 6.5$ for u_{loc}^+ and $\kappa_\omega = 0.61$ and $B_\omega = 4.90$ for ω_{loc}^+ , while the wind-sheared log-layer can be fitted with a straight line of coefficients $\kappa_\omega = 0.64$ and $B_\omega = 6.3$ for u_{loc}^+ and $\kappa_\omega = 0.51$ and $B_\omega = 4.9$ for ω_{loc}^+ .

process. As the plume-ejection region grows, the relative portion of the BL which is turbulent grows, and in consequence the angular velocity transport increases. In the ultimate regime, this region cannot grow significantly any more, and the BLs behave on average like turbulent BLs instead of Prandtl–Blasius laminar type BLs.

This can be observed in figure 6.6, where the azimuthal and angular velocity in wall units, i.e. $u^+(r^+) = (u_i - \langle u_\theta \rangle_{t,\theta,z})/u_\tau$ and $\omega^+(r^+) = (\omega_i - \langle \omega_\theta \rangle_{t,\theta,z})/\omega_\tau$ are plotted against r^+ for the inner and outer cylinder BLs for three values of Ta . In the transitional regime, the azimuthal and angular velocity profiles lose the logarithmic behaviour when averaged axially. On the other hand, in the ultimate regime, the log-behaviour is maintained after averaging.

As also observed in experiments [82], a logarithmic dependence can be seen for $u^+(r^+)$ and for $\omega^+(r^+)$ when r^+ is in the range $50 < r^+ < 600$. Strictly speaking, not both can show a logarithmic profile, as they are related by: $\omega^+(r^+) = u^+(r^+) + (u_\theta[r - r_i])/(ru_\tau)$. As expected from the theory [165], the match is better for the angular velocity profiles and not the azimuthal velocity profiles. The value of κ and κ_ω may depend on the strength of the flow driving, i.e. Ta . We can quantify the dependence of κ and κ_ω on the driving by fitting logarithmic curves to $u^+(r^+)$ and $\omega^+(r^+)$, the time- and axially-averaged non-dimensional velocity profiles, in the range of $50 < r^+ < 600$ (following Huisman *et al.* [82]) for various Ta . For the small Ta cases where $r^+ = 600$ is farther away from the cylinder walls than the mid-gap, the mid-gap is taken as upper limit for the fit. The results are shown in figure 6.7(a,b) for Ta between 10^9 and 10^{10} . For data below $Ta \approx 10^9$, truncating the upper value of r^+ at the mid-gap makes the r^+ range used for the fit span less than half a decade. Experimental data from Huisman *et al.* [82] for the Ta range between $Ta = 10^9$ and $Ta = 6.2 \cdot 10^{12}$ are also plotted in figure 6.7(a).

Both experimental and numerical data show the same trend for the lower drivings, i.e. for the lower values of Ta , the obtained values for κ and κ_ω deviate from the classical 0.41 von Kármán constant. This might be due to driving which is still not large enough and/or to the remnants of the large-scale structures, which is itself a consequence of this insufficient driving. Indeed, in experiments up to $Ta = 6.2 \cdot 10^{12}$ is achieved, and a convergence of $\kappa(Ta)$ to a value of 0.40 can be seen, see figure 6.7(a). However, when looking at the values of κ itself, deviations between experimental and numerical data can be seen. Discrepancies are probably caused by the axial dependence of κ , as the experimental values of κ are taken at a fixed height at the mid-cylinder, while the numerical values originate from an axially averaged azimuthal velocity.

To further quantify this statement, the axial dependence of κ and κ_ω for $Ta = 4.2 \cdot 10^9$ is shown in figure 6.7(c,d). Large variations across the axial direction of the cylinder can be seen, which are smaller in the case of κ_ω . In figure 6.7(c) the experimental value of κ at $Ta = 3.9 \cdot 10^9$ is plotted as a dashed line. This value is

Figure 6.6: (a,b) Axially averaged inner-cylinder (a) and outer-cylinder (b) boundary layer profiles u^+ versus r^+ for three values of Ta . (c,d) Axially averaged inner cylinder (c) and outer cylinder (d) boundary layer profiles ω^+ versus r^+ for three values of Ta . Profiles are plotted from the cylinder walls to mid-gap. No log-layers can be seen for the smallest value of Ta , while for the other two, in the ultimate regime, a logarithmic signature appears, which extends for increasing values of Ta . This characterizes the transition of the whole flow to the ultimate regime. The logarithmic signatures can be better seen when looking at the ω^+ profiles rather than the u^+ profiles, consistent with the theory by Grossmann *et al.* [165]. Straight line fits for the log-region of the highest Ta profile have been added, with coefficients $\kappa = 0.76$ and $B = 8.3$ (u^+), and $\kappa_\omega = 0.56$ and $B_\omega = 6.6$ (ω^+) for the inner cylinder (a,c) and $\kappa = 1.07$ and $B = 11.3$ (u^+) and $\kappa_\omega = 0.63$ and $B_\omega = 8.5$ (ω^+) for the outer cylinder (b,d). To highlight the differences across the profiles, a small subpanel to the right of each panel shows the three curves compensated by the logarithmic fit of the profile with the highest Ta , i.e. $u^+ / u_{fit}^+(Ta = 10^{10})$ and $\omega^+ / \omega_{fit}^+(Ta = 10^{10})$.

measured at a fixed height, which can correspond to any value of the axial coordinate in the numerical domain.

As a confirmation of our statement that the large scale structures are washed away, we quantify the reduction in axial dependence of the angular and azimuthal velocity profiles for increasing driving strength by defining the normalized velocity increment $\Delta_U = (\max_z(\bar{u}_\theta(r_a)) - \min_z(\bar{u}_\theta(r_a))) / \langle \bar{u}_\theta(r_a) \rangle_z$. The meaning of Δ_U is that the larger this increment, the stronger the axial dependence. Δ_U versus Ta is shown in figure 6.8. As expected, the axial dependence strongly decreases in the Ta -range of the transition to the ultimate regime. After the transition to the ultimate regime, Δ_U fluctuates between 0.1 to 0.15, though the strength of the large scale wind continuously diminishes with increasing Ta . Indeed, some degree of axial dependence remains, even at the highest drivings. This result is remarkable, as even at $Ta = 10^{10}$, corresponding to Reynolds numbers of 10^5 , an effect of the initial roll state seems to remain.

The residual axial dependence of the profiles can be understood by looking at the local Nusselt number $Nu_\omega(r, z) = r^3 \langle \langle u_r \omega \rangle_{\theta, t} - v \partial_r \langle \omega \rangle_{\theta, t} \rangle / T_{pa}$. Figure 6.9 shows azimuthally cut contour plots of $Nu_\omega(r, z)$ for two values of Ta , i.e. figure 6.9(a) in the transitional ($Ta = 4.77 \cdot 10^7$) and figure 6.9(b) in the ultimate regimes ($Ta = 1.00 \cdot 10^{10}$). At the highest drivings, a very strong axial dependence can still be seen in Nu_ω . In the bulk, Nu_ω is (apart from the non-dimensionalization), the correlation between u_r and ω . The axial dependence in Nu_ω is two orders of magnitude larger than for the average values, and negative Nu_ω of even 3000 can be seen for the largest driving. This means that even if the azimuthal velocity loses most of its axial dependence, structures can be seen in u_r (and in consequence in u_z) up to Reynolds numbers of about 10^5 , which in turn is causing the residual axial dependence of the profiles $u_\theta(r)$ and $\omega(r)$.

Up to now, we have focused on the loss of axial dependence, and have omitted from our analysis the azimuthal structure of the flow. This is justified by the fact that the flow is statistically homogeneous in the θ -direction, so it does not play an important role in the transitions. As a confirmation of this, figure 6.10 shows contour plots taken at a constant radius $\tilde{r} = \tilde{r}_{cut}$ of the instantaneous azimuthal velocity field u_θ for two values of $Ta = 4.77 \cdot 10^7$ (transitional regime) and $Ta = 4.63 \cdot 10^9$ (ultimate regime) both in the BLs ($\tilde{r}_{cut} = 3.2 \cdot 10^{-2}$ for $Ta = 4.77 \cdot 10^7$ and $\tilde{r}_{cut} = 1.1 \cdot 10^{-2}$ for $Ta = 4.63 \cdot 10^9$) and in the bulk ($\tilde{r}_{cut} = 0.5$ for both Ta).

Figure 6.10(c,d) shows the formation of ω -plumes in the BLs. These were previously interpreted to be herring-bone streaks by Dong [48]. The axial structure of the flow present in the transitional regime can be appreciated for figure 6.10(a,c), and it can be seen to disappear when looking at figure 6.10(b,d). However, the flow shows no clear azimuthal structure in any of the panels, and there is no indication of a flow transition if one looks at the azimuthal structure of the flow.

Figure 6.7: (a,b) Dependence of the log-layer fit coefficient κ (a) and κ_ω (b) on the driving Ta for both cylinders, based on the present DNSs (N), averaged over time and in axial direction. Experimental data points (E) measured for a fixed height from Huisman *et al.* [82] are also shown in (a). For the inner cylinder, the inverse-slope coefficient κ tends to $\kappa = 0.41$ with increased driving. For the outer cylinder, this trend is not yet visible in the numerics, probably due to insufficient driving. (c,d) Axial dependence of κ (c) and κ_ω (d) for $Ta = 4.2 \cdot 10^9$ for both inner and outer cylinder. Large fluctuations across the axial direction of the cylinder can be seen, which are smaller for the case of κ_ω . On (d), the dashed line shows κ at the experimental value at $Ta = 3.9 \cdot 10^9$ (closest to the numerical data) but this value of κ was measured at a fixed height, i.e. the cylinder's mid-height.

Figure 6.8: Normalized velocity increment $\Delta_U = (\max_z(\bar{u}_\theta(r_a)) - \min_z(u_\theta(r_a)))/\langle \bar{u}_\theta(r_a) \rangle_z$ against Ta . There is a transition in the axial dependence, dropping to about 10% at $Ta = 3 \cdot 10^8$, corresponding to the transition to the ultimate regime. Remarkably, an axial dependence of about 10% is still present at $Ta = 10^{10}$, which corresponds to a Reynolds number of about 10^5 .

Figure 6.9: Contour plots of the time- and azimuthally averaged local Nusselt number $Nu_\omega(r, z)$ for (a) $Ta = 4.77 \cdot 10^7$ ($Nu_\omega = 8.68$) and (b) $Ta = 1.00 \cdot 10^{10}$ ($Nu_\omega = 51.5$). Fluctuations two orders of magnitude higher than the axially averaged values of Nu_ω can be seen even at the largest drivings, indicating that remnants of the initial roll state are still present.

Figure 6.10: Contour plots of the instantaneous azimuthal velocity u_θ for two Taylor numbers at constant radius cuts (at $\tilde{r} = \tilde{r}_{cut}$) either in BL or bulk. (a): $Ta = 4.77 \cdot 10^7$, $\tilde{r}_{cut} = 0.5$ (transition regime, bulk), (b): $Ta = 4.63 \cdot 10^9$, $\tilde{r}_{cut} = 0.5$ (ultimate regime, bulk), (c): $Ta = 4.77 \cdot 10^7$, $\tilde{r}_{cut} = 3.2 \cdot 10^{-2}$ (transition regime, BL) and (d): $Ta = 4.63 \cdot 10^9$, $\tilde{r}_{cut} = 1.1 \cdot 10^{-2}$ (ultimate regime, BL).

6.3 Summary and conclusions

In conclusion, the logarithmic azimuthal velocity profile of the ultimate regime in TC is triggered by plume ejection that in turn is regulated by the relative strength of the axial and radial mean flow with respect to the azimuthal one. This is not a finite-size effect (i.e., triggered by the upper and lower plates), as simulations in this study are done in a periodic domain. The transition to the ultimate regime in TC flow is determined by the growth of the plume ejection regions until they cover the majority of the cylinder surface. This occurs when the coherent wind is sufficiently weak so that the axial pressure gradient at the wall becomes negligible, when compared to the shear from the inner cylinder, resulting in plume emission from the complete axial extent of the boundary layer.

In the ultimate regime two logarithmic boundary layers (with curvature corrections, see Ref. [165] for the angular velocities form, one at each cylinder, whose (inverse) slopes are Ta -dependent. Also, the azimuthal velocity profiles can still reasonably well be fitted by a log-law, as done in experiments [82], though strictly speaking not both u^+ and ω^+ can follow log-laws because they differ by an additive constant. If u^+ is fitted with a log-law nonetheless, the (inverse) slopes κ are also Ta -dependent and thus differs from the classical von Kármán constant. In the limit of very large Ta , the (inverse) slopes κ seem to tend to the universal von Kármán constant $\kappa = 0.4$. Surprisingly, in this regime some dependence on the initial roll state can still be observed up to the highest driving achieved in these simulations.

Re_i	Ta	Nu_ω	$N_\theta \times N_r \times N_z$	η_K/d	$Re_{\tau,i}$
$1.19 \cdot 10^4$	$2.15 \cdot 10^8$	13.105	$256 \times 640 \times 512$	$4.46 \cdot 10^{-3}$	$2.3 \cdot 10^2$
$1.74 \cdot 10^4$	$4.62 \cdot 10^8$	16.940	$256 \times 640 \times 512$	$3.46 \cdot 10^{-3}$	$3.2 \cdot 10^2$
$2.53 \cdot 10^4$	$9.75 \cdot 10^8$	22.081	$256 \times 640 \times 512$	$2.68 \cdot 10^{-3}$	$4.7 \cdot 10^2$
$3.76 \cdot 10^4$	$2.15 \cdot 10^9$	29.860	$256 \times 640 \times 512$	$2.04 \cdot 10^{-3}$	$6.7 \cdot 10^2$
$5.52 \cdot 10^4$	$4.63 \cdot 10^9$	38.544	$384 \times 640 \times 768$	$1.58 \cdot 10^{-3}$	$9.2 \cdot 10^2$
$8.10 \cdot 10^4$	$1.00 \cdot 10^{10}$	51.554	$512 \times 800 \times 1024$	$1.21 \cdot 10^{-3}$	$1.3 \cdot 10^3$

Table 6.1: This table presents the numerical results which are new to this chapter. For the other data points see Ref. [112]. The first two columns show the driving, expressed as either Re_i or Ta . The third column shows the non-dimensionalized torque, Nu_ω . The fourth column shows the amount of grid points used in azimuthal (N_θ), radial (N_r) and axial direction (N_z). The last two columns show additional details on the grid resolution and the Kolmogorov scale η_K in the bulk as well as the frictional Reynolds number Re_τ in the boundary layers. For TC we obtain η_K from the exact dissipation relationships as $\eta_K/d = (\sigma^{-2}Ta(Nu_\omega - 1))^{-1/4}$ (see Ref. [161] for a full derivation), and the frictional Reynolds number at the inner cylinder as $Re_{\tau,i} = u_{\tau,i}d/(2\nu)$. All these results are for the “reduced” geometry with $\Gamma = 2\pi/3$ and $n_{sym} = 6$.

Exploring the phase diagram of fully turbulent Taylor-Couette flow *

A good traveler has no fixed plans and is not intent upon arriving. A good artist lets his intuition lead him wherever it wants. A good scientist has freed himself of concepts and keeps his mind open to what is.

—LÄOZĬ

Direct numerical simulations of Taylor-Couette flow were performed. Shear Reynolds numbers of up to $3 \cdot 10^5$, corresponding to Taylor numbers of $Ta = 4.6 \cdot 10^{10}$, were reached. Effective scaling laws for the torque are presented. The transition to the ultimate regime, in which asymptotic scaling laws (with logarithmic corrections) for the torque are expected to hold up to arbitrarily high driving, is analysed for different radius ratios, different aspect ratios and different rotation ratios. It is shown that the transition is approximately independent of the aspect- and rotation-ratios, but depends significantly on the radius-ratio. We furthermore calculate the local angular velocity profiles and visualize different flow regimes that depend both on the shearing of the flow, and the Coriolis force originating from the outer cylinder rotation. Two main regimes are distinguished, based on the magnitude of the Coriolis force, namely the co-rotating and weakly counter-rotating regime dominated by Rayleigh-unstable regions, and the strongly counter-rotating regime where a mixture of Rayleigh-stable and Rayleigh-unstable regions exist. Furthermore, an analogy between radius-ratio and outer-cylinder rotation is revealed, namely that smaller gaps behave like a wider gap with co-rotating cylinders, and that wider gaps behave like smaller gaps with weakly counter-rotating cylinders. Finally, the effect of the aspect ratio on the effective torque versus Taylor number scaling is analysed and it is shown that different branches of the torque-versus-Taylor relationships associated to differ-

*Based on: R. Ostilla-Mónico, E. P. van der Poel, R. Verzicco, S. Grossmann and D. Lohse, *Exploring the phase diagram of fully turbulent Taylor-Couette flow*, J. Fluid Mech., **761**, 1–26 (2014).

ent aspect ratios are found to cross within 15% of the Reynolds number associated to the transition to the ultimate regime. The chapter culminates in phase diagram in the inner versus outer Reynolds number parameter space and in the Taylor vs inverse Rossby number parameter space, which can be seen as the extension of the Andereck *et al.* [37] phase diagram towards the ultimate regime.

7.1 Introduction

Taylor-Couette flow has for long been used as a model system in fluid dynamics. Recently, a mathematically exact analogy between TC and Rayleigh-Bénard flow (RB), i.e. the convective flow between two parallel plates heated from below and cooled from above was found by Eckhardt, Grossmann and Lohse (2007) [25] (here referred to as EGL07). Within this context, TC can be viewed as a convective flow, driven by the shear between both cylinders where angular velocity is transported from the inner to the outer cylinder. As explained by Grossmann *et al.* [165], as long as the driving of the system is small, the transport is limited by the laminar boundary layers. But if the driving becomes strong enough the boundary layers become turbulent and the system enters the so-called “ultimate” regime. The study of the transition to this regime, expected to be also present in RB, has attracted recent interest, as most applications of TC and RB in geo- and astro-physics are expected to be in this ultimate regime.

For RB flow, the transition to an ultimate regime was first qualitatively predicted by Kraichnan [74], and later quantitatively by Grossmann & Lohse [70, 78, 154] and then experimentally found by several groups [75, 76, 169, 170]. It lies outside the present reach of DNS. The analogous boundary layer transition to an ultimate regime in TC flow was first found in the experiments by Lathrop and coworkers [39, 40], and analysed more precisely in Lewis & Swinney [41], even though earlier work by Wendt [31] already showed some transition in the torque scaling around the same Reynolds number. The transition was not related to the transition to the ultimate regime until later [43–45, 165]. In DNS, it was observed for the first time by Ostilla-Mónico *et al.* [113].

In TC flow this transition is easier to achieve as the mechanical driving is more efficient than the thermal one, and thus the frictional Reynolds numbers in the boundary layer are much larger. By using the analogy between both systems, better understanding of the transition in TC can thus also lead to new insight in RB, where it is more elusive. Ostilla-Mónico *et al.* [113] numerically studied the transitions in TC for pure inner cylinder rotation for a radius ratio of $\eta = r_i/r_o = 0.714$, where r_o and r_i are the outer and inner radii respectively, and an aspect ratio $\Gamma = L/(r_o - r_i) = 2\pi/3$, where L is the axial period in the DNS. In that study, the flow transitions and boundary layer dynamics were revealed in the range of Taylor numbers Ta between 10^4 and 10^{10} ,

where the Taylor number is defined as:

$$Ta = \frac{1}{4} \sigma d^2 (r_o + r_i)^2 (\omega_i - \omega_o)^2 \nu^{-2}, \quad (7.1)$$

with ω_o and ω_i the angular velocities of the outer and inner cylinder, respectively, $d = r_o - r_i$ the gap width, and ν the kinematic viscosity of the fluid. $\sigma = [(r_o + r_i)/(2\sqrt{r_i r_o})]^4$ can be considered as a geometric quasi-Prandtl number (EGL07).

We now describe the series of events when increasing Ta . For small enough Ta , the flow is in the purely azimuthal, laminar, state. When the system is driven beyond a critical driving, one passes the onset of instability and the purely azimuthal, laminar, flow disappears and large-scale Taylor rolls form. Further increasing of the driving breaks up these rolls, causing the onset of time-dependence as the system transitions from the stationary Taylor vortex regime to the modulated Taylor vortex regime and finally the breakup of these into chaotic turbulent Taylor vortices. These changes of the flow are reflected in transitions of the local scaling laws for the torque versus driving, i.e. versus Taylor number Ta . All this has been studied extensively and summarized e.g. in Refs. [39–41, 171]. The mentioned breakup of the rolls leads to the existence of a transitional regime, where the large-scale coherent structures still can be identified when looking at the time-averaged quantities. Looking at the details of the flow, a mixture of turbulent and laminar boundary layers is present.

In this transitional regime, hairpin vortices, which, in the context of RB, can also be viewed as plumes, are ejected from both inner and outer cylinders and contribute to large-scale bulk structures. These structures in turn cause an axial pressure gradient, which couples back to the boundary layers, causing plumes to be ejected there. But this only happens from preferential spots in the boundary layers. Once the driving is strong enough, the large-scale structures slowly vanish, and the plumes no longer feel an axial pressure gradient. The boundary layers now become fully turbulent and the flow transitions to the “ultimate” regime. As the flow enters the ultimate regime, and the boundary layer become turbulent, a logarithmic signature in the angular velocity boundary layers is expected, which indeed has been found experimentally [82] and numerically [113].

In the ultimate regime, an effective scaling relation between the Nusselt number Nu_ω , i.e. the non-dimensional torque $Nu_\omega = T/T_{pa}$ where T is the torque, and T_{pa} the torque in the purely azimuthal state, and the Taylor number Ta is expected, with an effective scaling exponent which exceeds that for the laminar-type boundary layer case [172], for which $\alpha = 1/3$. That is, in the ultimate regime, we expect an effective scaling law $Nu_\omega \sim Ta^\alpha$ with $\alpha > 1/3$. In fact, for that regime, the relation law $Nu_\omega \sim Ta^{1/2}$ with logarithmic corrections was suggested, [74, 78, 128]. The logarithmic corrections are quite large, and lead to an effective scaling law with $\alpha \approx 0.38$ for $Ta \sim 10^{11}$ [77, 78]. We note that this scaling law is analog to the scaling of the friction factor with Reynolds number in fully turbulent pipes [8].

For the largest drivings, remnants of the larger rolls, which can be seen as a large scale wind, are still observed at even the largest Reynolds numbers studied numerically [113], and experimentally, even up to $Re \sim 10^6$ [173]. In Ref. [113], it was shown that the remnants of the large scale structures played a crucial role in the transition to the ultimate regime. However, large scale structures are not present in the whole parameter space of TC. Andereck *et al.* [37] showed how rich a variety of different states exists at low Reynolds number when the outer cylinder is also rotated. Brauckmann & Eckhardt [156] reported that the strength of the large scale wind was most pronounced at the position of optimal transport. However, if the outer cylinder is counter-rotated past the position of optimal transport, bursts arise from the outer cylinder. The flow is very different outside and inside the neutral surface, which separates Rayleigh-stable from Rayleigh-unstable regions of the gap, changing completely the dynamics of the system. The Taylor vortices no longer penetrate the whole gap, extending thus the unstable region effectively somewhat outside the neutral surface of laminar type flow [112].

The geometry of the system can be expected to play an important role in determining the strength of the large scale wind, and how the transition takes place. In the context of understanding the radius-ratio dependence of the transition to the ultimate regime, Merbold *et al.* [80] reported a higher transitional Reynolds numbers for $\eta = 0.5$ than what was seen for $\eta = 0.714$ by Ostilla-Mónico *et al.* [113] and for $\eta = 0.909$ by Ravelet *et al.* [53]. Also the aspect-ratio plays a role. Although different vortical states were known to coexist at low Reynolds number [157], it was previously thought that if the driving was sufficiently large, only one branch of the torque versus Taylor number curve would survive [41]. Brauckmann & Eckhardt [123] found that the difference in the global response between different vortical states becomes smaller with increasing Reynolds number. Recently, Martinez-Arias *et al.* [174] reported on the existence of different vortical states associated to different global torques at a given Taylor number for $\eta = 0.909$, and that there is a crossing between those torque-versus- Ta curves around the transition to the ultimate regime. Furthermore, Huisman *et al.* [173] showed that different vortical states survive up to Reynolds number of 10^6 , corresponding to Taylor numbers of order 10^{12} . By combining measurements of global torque and local velocity, Huisman *et al.* [173] found that the optimal transport is connected to the existence of the large-scale coherent structures at high Taylor numbers.

Therefore, some questions arise which we want to address in the present chapter: How does the transition in the boundary layers take place across the full parameter space of TC? Is the vanishing of the large-scale wind a necessary and/or a sufficient condition for the boundary layer transition? Why does the transition occur later for $\eta = 0.5$ than for larger values of η ? And finally, what is the effect of the vortical wavelength and why do different branches of the torque versus Taylor number scaling

curves cross near the transition to the ultimate regime?

7.2 Explored parameter space

7.2.1 Control parameters

To answer these questions, direct numerical simulations (DNS) of TC have been performed across all dimensions of the parameter space, not only adding outer cylinder rotation, but also varying both geometrical parameters η and Γ . To do this, the rotating-frame formulation of Ostilla-Mónico *et al.* [112] (cf. §4) was used. TC was formulated in a frame rotating with the outer cylinder, such that it looks like a system in which only the inner cylinder is rotating, but with a Coriolis force term, which represents the original presence of the outer cylinder rotation. The shear driving of the system is non-dimensionally expressed as a Taylor number, introduced previously:

$$Ta = \frac{1}{4} \sigma d^2 (r_o + r_i)^2 (\omega_i - \omega_o)^2 \nu^{-2}, \quad (7.2)$$

Ta is the analog to the Rayleigh number in RB, as elaborated in EGL07. The outer cylinder rotation reflects in a Coriolis force, characterized by a Rossby number $Ro = |\omega_o - \omega_i| r_i / (2\omega_o d)$. The Rossby number or rather Ro^{-1} is the parameter which appears in the equations of motion for the fluid:

$$\frac{\partial \tilde{\mathbf{u}}}{\partial \tilde{t}} + \tilde{\mathbf{u}} \cdot \tilde{\nabla} \tilde{\mathbf{u}} = -\tilde{\nabla} \tilde{p} + \frac{f(\eta)}{Ta^{1/2}} \tilde{\nabla}^2 \tilde{\mathbf{u}} - Ro^{-1} \mathbf{e}_z \times \tilde{\mathbf{u}}, \quad (7.3)$$

where $f(\eta) = \frac{1}{4} \sigma ((1 + \eta)/\eta)^2$, a geometrical parameter. The Rossby number is related to the frequency ratio $\mu = \frac{\omega_o}{\omega_i}$ via

$$Ro^{-1} = \text{sgn}(\omega_o) \left| \frac{\mu}{\mu - 1} \right| \frac{2(1 - \eta)}{\eta}. \quad (7.4)$$

Thus fixed Ro^{-1} means fixed μ and vice versa. $Ro^{-1} > 0$ describes co-rotation or $\omega_o > 0$, while $Ro^{-1} < 0$ means counter-rotation. The radius ratio η is presented by the geometrical amplitude factor $2(1 - \eta)/\eta$, being small for small gap ($\eta \rightarrow 1$) and large for large gap ($\eta \rightarrow 0$). A resting outer cylinder is described by $Ro^{-1} = 0$.

There are also other ways of choosing the control parameters. Classically, they have been expressed as two non-dimensional Reynolds numbers corresponding to the inner and outer cylinders: $Re_{i,o} = u_{i,o}^\theta \cdot d / \nu$, where $u_{i,o}^\theta$ are the azimuthal velocities of the inner and outer cylinders. The classical flow control parameters (Re_i, Re_o) can be transformed to the (Ta, Ro^{-1}) parameter space by:

$$Ta = f(\eta)^2 |Re_i - \eta Re_o|^2, \quad (7.5)$$

and

$$Ro^{-1} = \frac{2(1-\eta)Re_o}{|\eta Re_o - Re_i|}. \quad (7.6)$$

Vice-versa, we have

$$Re_i = \frac{Ta^{1/2}}{f(\eta)} \left(1 + \frac{\eta Ro^{-1}}{2(1-\eta)} \right), \quad (7.7)$$

and

$$Re_o = \frac{Ro^{-1}Ta^{1/2}}{2f(\eta)(1-\eta)}. \quad (7.8)$$

The shear driving of the flow can also be expressed as a shear Reynolds number $Re_s = \sqrt{Ta/\sigma}$.

7.2.2 Numerical scheme

A second-order finite-difference code was used with fractional time integration. The code was parallelized using hybrid OpenMP and MPI-slab decomposition. Simulations were run on local clusters and on the supercomputer CURIE (Thin nodes) using a maximum of 8192 cores. Details about the code can be found in Verzicco & Orlandi [103] and in Ostilla-Mónico *et al.* [112]. The explored parameter space from previous work [112, 113] was extended through further simulations. Figure 7.1 shows the parameter space explored in this chapter. Circles show simulations of a “full” geometry, i.e. a complete cylinder and with $\Gamma = 2\pi$. Following the work of Brauckmann & Eckhardt [123], the simulations with the largest Ta were performed on “reduced” geometries to reduce computational costs, and these are indicated as squares in the plots. The idea is that instead of simulating the whole cylinder, a cylinder wedge with rotational symmetry of order n_{sym} is considered. The aspect ratio was also reduced to $\Gamma = 2\pi/3$, accommodating a single vortex pair with the wavelength $\lambda_z = 2\pi/3 = 2.09$. The vortical wavelength remains the same, although there is a single vortex instead of the three vortex pairs having also the wavelength $\lambda_z = 2\pi/3$. Other vortical wavelengths were also simulated using reduced geometries for $\eta = 0.909$. We note that the aspect ratio Γ is a geometrical control parameter, but λ_z is a response of the system, which depends both on Γ and on the amount of vortex pairs which fit in the system. They are related by $\lambda_z = \Gamma/n$, where n is the amount of vortex pairs which fit in the system. For all simulations axially periodic boundary conditions were used. Its consequences on the vortex wavelength are analyzed in §7.5. Further details on the numerical resolution can be found in table 7.1 at the end of the chapter.

Figure 7.1: (a): Explored (Re_i, Re_o) parameter space for $\eta = 0.714$, $\lambda_z = 1.04$. (b): Same as left panel, but now in the $(Ta, 1/Ro)$ parameter space. In both panels the solid line indicates pure inner cylinder rotation, the dot-dash line indicates the Rayleigh stability criterion, while the dashed line indicates the asymptotic position of optimum transport in experiments, i.e. $Ro^{-1} = -0.20$. The Rayleigh-stability line lies outside the top right panel, at $Ro^{-1} = 0.83$. These lines divide the parameter space into the Rayleigh-stable zone, the co-rotating or weakly counter-rotating regime (CWCR) and the strongly counter-rotating regime (SCR). (c): explored (Ta, η) parameter space for $Ro^{-1} = 0$, $\lambda_z = 1.04$. (d): explored (Ta, λ_z) parameter space for $Ro^{-1} = 0$, $\eta = 0.909$. The dashed line indicates the cross-point of branches with different λ_z in Martinez-Arias *et al.* [174]. For (c,d), the same colour coding is maintained throughout the chapter. On all panels, circles indicate simulations of the “full” geometry, with three vortex pairs, while squares indicate simulations of “reduced” geometries with forced rotational symmetry and one vortex pair.

7.2.3 Explored parameter space

Figure 7.1(a,b) shows the parameter space explored for $\eta = 0.714$ in both (Re_i, Re_o) and $(Ta, 1/Ro)$ to study the effects of outer cylinder rotation. For $\eta = 0.714$, reduced geometries simulate one sixth of the cylinder, i.e. $n_{sym} = 6$ as used in Ostilla-Mónico *et al.* [113]. The chosen values of Ro^{-1} include a co-rotating outer cylinder ($Ro^{-1} = 0.20$), a weakly counter-rotating outer cylinder ($Ro^{-1} = -0.13$), counter-rotation near the asymptotic position of optimal transport, Ro_{opt}^{-1} ($Ro^{-1} = -0.22$), and two values of Ro^{-1} in the strongly counter-rotating regime ($Ro^{-1} = -0.30$ and $Ro^{-1} = -0.40$). No simulations were run in the Rayleigh-stable regime (i.e. when $r_o^2 \omega_o > r_i^2 \omega_i$) as no evidence of turbulence was found in that regime up to $Ta \sim 10^{10}$ in Ostilla-Mónico *et al.* [175].

In addition, to study the effects of geometry, i.e. both the radius ratio η , and the vortical wavelength λ_z (controlled through the aspect ratio Γ), additional simulations were performed. Figure 7.1(c) shows that two additional radius ratios were simulated up to $Ta = 4 \cdot 10^{10}$, one with a larger gap ($\eta = 0.5$) and one with a smaller gap ($\eta = 0.909$). For $\eta = 0.5$, one third of the cylinder ($n_{sym} = 3$) was simulated for Ta larger than 10^8 . This value of n_{sym} for $\eta = 0.5$ was shown not to affect the values of the torque obtained in the simulations in Brauckmann & Eckhardt [156]. For $\eta = 0.909$, one twentieth of the geometry ($n_{sym} = 20$) was used.

Figure 7.1(d) shows the simulations with varying vortical wavelength λ_z done for $\eta = 0.909$ and pure inner cylinder rotation. $\eta = 0.909$ was chosen as we expect the effects of the coherent structures, and thus of λ_z , to be stronger for larger η (see later sections §7.4 and §7.5 for an explanation). The values of Ta simulated are around the range where Martinez-Arias *et al.* [174] have experimentally observed the crossing of different branches in $Nu_\omega(Ta)$ and also coincides with the onset of the “ultimate” regime.

7.2.4 Non-dimensionalization

The following non-dimensionalizations will be used: as the flow is simulated in a rotating frame, the outer cylinder is stationary, and the system has an unique velocity scale, equal to $U \equiv r_i(\omega_i^\ell - \omega_o^\ell)$ in the laboratory frame. All velocities are non-dimensionalized using U , i.e. $\tilde{\mathbf{u}} = \mathbf{u}/U$. The gap width d is the characteristic length scale, and thus used for normalizing distances.

We define the normalized (non-dimensional) distance from the inner cylinder $\tilde{r} = (r - r_i)/d$ and the normalized height $\tilde{z} = z/d$. We furthermore define the time- and azimuthally-averaged velocity fields as:

$$\bar{\tilde{\mathbf{u}}}(r, z) = \langle \tilde{\mathbf{u}}(\theta, r, z, t) \rangle_{\theta, t}, \quad (7.9)$$

where $\langle \phi(x_1, x_2, \dots, x_n) \rangle_{x_i}$ indicates averaging of the field ϕ with respect to x_i . As men-

tioned previously, the torque is non-dimensionalized as an angular velocity “Nusselt” number (EGL07), defined as $Nu_\omega = T/T_{pa}$, where T_{pa} is the torque in the purely azimuthal flow. The torque is calculated from the radial derivative of $\langle \tilde{\tilde{\omega}} \rangle_z$ at the inner and outer cylinders. The simulations are run in time until the respective values are equal within 1%. The torque is then taken as the average value of the inner and outer cylinder torques. Therefore, the error due to finite time statistics is smaller than 1%. From here on, for convenience we will drop the overhead tilde on all non-dimensionalized variables.

7.2.5 Structure of chapter

The organization of the chapter is as follows. In §7.3, we analyze the effect of rotating the outer cylinder. This is followed by §7.4, where we study the influence of η , and notice an analogy between the effects of smaller η and larger Ro^{-1} . In §7.5, we consider the effects of the last parameter, the vortical wavelength λ_z . We finish in §7.6 with a summary of the results and an outlook for future work.

7.3 The inverse Rossby number dependence

In this section we will study the effect of the Coriolis force (Ro^{-1}), originating from the rotation of the outer cylinder, on the scaling of $Nu_\omega(Ta)$ with Ta and, more specifically, the effect of Ro^{-1} on the transition to the ultimate regime. Depending on the value of Ro^{-1} , two distinct regimes will be identified: First a co- and weakly counter-rotating Ro^{-1} range, denoted from here on as CWCR regime, and second the strongly counter-rotating Ro^{-1} range, denoted from here on as SCR regime. The CWCR regime is found when the outer cylinder either is at rest, co-rotates with the inner cylinder, or slowly counter-rotates. The counter-rotation must be slow enough such that no Rayleigh-stable zones are generated in the bulk of the flow. In this CWCR regime the Coriolis force is balanced through the bulk gradient of ω . This can be derived from a large scale balance in the θ -component of the velocity in equation (7.3). In summary, the non-linear term $u_r(\partial_r u_\theta + u_r u_\theta / r)$ and the Coriolis force term $-u_r Ro^{-1}$ balance each other out on average (cf. Ref. [112] for the full derivation). This results in a linear relationship between Ro^{-1} and $\partial_r \langle \tilde{\tilde{\omega}} \rangle_z$ [176].

Taylor-Couette flow can be considered as being in the SCR regime, if the outer cylinder strongly counter-rotates and generates a Coriolis force which exceeds what the ω -gradient can balance. The threshold value of Ro^{-1} corresponds to the flattest ω profile. This also is the value of Ro^{-1} , for which $Nu_\omega(Ro^{-1})$ is found to be largest [77, 112], denoted henceforth as Ro_{opt}^{-1} . In this regime the turbulent plumes originating from the inner cylinder are not strong enough to overcome the stabilizing effect of the outer cylinder, and the flow is divided into two regions, a Rayleigh-stable region

Figure 7.2: Left panel: the non-dimensional torque $Nu_\omega - 1$ versus the driving, i. e., the Taylor number Ta , for $\eta = 0.714$ and six values of Ro^{-1} . Right panel: the compensated Nusselt $(Nu_\omega - 1)/Ta^{1/3}$ versus the driving Ta for the same six values of Ro^{-1} . The effective scaling law of $Nu_\omega \sim Ta^{0.38}$ is reached for all Ro^{-1} at the highest drivings Ta beyond about 10^9 . However, the behavior in the classical regime Ta less than 10^9 depends heavily on Ro^{-1} . Before the onset of the ultimate regime, we observe a transitional Ta -regime ranging from about 10^6 to about 10^8 associated to the breakup of coherent structures for co-rotating and weakly counter-rotating cylinders ($-0.13 \leq Ro^{-1} \leq 0.2$). For more positive values of Ro^{-1} this regime can be seen earlier, and is persistent for a larger Ta -range. For the strongly counter-rotating cases ($Ro^{-1} \leq -0.22$), an effective local scaling exponent with $\alpha > 1/3$ is seen in the classical regime. This can be related to the interplay between Rayleigh-stable and unstable regions.

in the outer gap region, which plumes do not reach, and a Rayleigh-unstable region in the inner parts of the gap. For given Coriolis force, the relative sizes of these spatial regions depend on Ta , as for a stronger driving (i.e. larger Ta), the turbulence originating from the inner cylinder “pushes” these zones more towards the outer cylinder. This may lead to switching between vortical states and jumps in global quantities as seen in Ostilla-Mónico *et al.* [112]. The boundary between both regimes is at Ro_{opt}^{-1} . Of course, Ro_{opt}^{-1} depends on Ta too, due to effect of viscosity in the Coriolis force balance [112], and only saturates to $Ro_{opt}^{-1}(Ta \rightarrow \infty) = -0.20$ for sufficiently high drivings of $Ta \sim 5 \cdot 10^8$ and more (cf. both panels of figure 7.2 and Ref. [123]).

Figure 7.2 shows both $Nu_\omega - 1$ and the compensated Nusselt number $(Nu_\omega - 1)/Ta^{1/3}$ versus Ta for $\eta = 0.714$ and the six values of Ro^{-1} studied. For the largest drivings (i.e. $Ta > 10^9$) all values of Ro^{-1} reach the effective scaling law $Nu_\omega \sim Ta^{0.38}$ (with a different amplitude), similar to what was reported in the experiments by van Gils *et al.* [43]. However, very different behaviour can be seen for $Ta < 10^9$, i.e. before the onset of the ultimate regime.

In the CWCR regime ($Ro^{-1} \geq Ro_{opt}^{-1} = -0.20$), the Coriolis force is reflected in the flow structure through the bulk gradient of ω , making it either flatter as in

the case of weak-counter rotation, or steeper, as in the case of co-rotation (if the driving is sufficiently large). A consequence of the angular velocity gradient in the bulk is that large scale structures can be weakened or even completely disappear in the CWCR regime. These changes in ω -gradient strongly affect the capability of plumes to “coordinate” and form a large-scale wind, which in turn leads to an earlier (or later) onset of the sharp decrease in the local exponent α in the scaling law $(Nu_\omega - 1) \sim Ta^\alpha$ associated with the breakdown of coherence, and the onset of time dependence in Nu_ω [113].

For the case of co-rotating cylinders ($Ro^{-1} = 0.20$), this happens when the system enters the so-called “wavelet” regime, characterized by moving waves in the boundary regions between a pair of Taylor vortices [37, 171]. These waves move with different speeds, and as a consequence this regime is not stationary in any reference frame. This regime only persists for a small range of Ta , and eventually all remnants of Taylor vortices vanish. Axial dependence of the flow structure is almost completely lost, even at Ta as low as $Ta \approx 5 \cdot 10^7$. Unlike the case of $Ro^{-1} = 0$ studied by Ostilla-Mónico *et al.* [113], in this transitional regime the large-scale rolls already completely vanished, but for $Ro^{-1} = 0.20$ this does not immediately lead to the transition to the ultimate regime. After its sharp decrease, α does not exceed $1/3$. Instead, at a driving strength around $Ta \approx 10^7$ (coinciding with the disappearance of the structures), the local effective scaling exponent α has increased to $\alpha \approx 1/3$, and then stops growing. Only if Ta increases further and the shear in the boundary layers grows past a threshold, a shear-instability takes place, and the system transitions to the ultimate regime.

For the case of counter-rotating cylinders, (i.e. $Ro^{-1} < 0$), α can locally grow beyond $\alpha = 1/3$ in the classical regime. This is unexpected, as values of α larger than one third have been associated to the transition to turbulence of the boundary layers in the context of both Rayleigh-Bénard convection [76], and TC with a stationary outer cylinder [113]. However, in this case, the shear in the boundary layers is too low so the boundary layers still remain laminar.

For counter-rotating cylinders, a wide range of flow configurations is available in the low- Ta regime [37]. We can relate local steps in α to the switching between such flow configurations. The interplay between Rayleigh-stable and -unstable regions can also play a role. Larger drivings cause the Rayleigh-unstable region to grow, and thus to increase the transport. These two effects lead to larger increases in the non-dimensional torque than what is expected for pure inner cylinder rotation, and explain the large values of α seen.

To further illustrate the effect of the Coriolis force on the large-scale structures, figure 7.3 presents a contour plot of $\bar{\omega}$ in the CWCR regime $Ro^{-1} = 0.20$, around the optimum $Ro^{-1} = -0.22 \approx Ro_{opt}^{-1}$ and in the SCR regime $Ro^{-1} = -0.40$. Figure 7.4 shows the axially-averaged angular velocity profiles $\langle \bar{\omega} \rangle_z$ for $\eta = 0.714$ and the six

Figure 7.3: Contour plots of the azimuthally- and time averaged angular velocity field $\bar{\omega}$ for $Ta = 10^{10}$, $\eta = 0.714$ and three values of Ro^{-1} . (a) corresponds to $Ro^{-1} = 0.2$ (CWCR regime) and shows no traces of axial dependence. Plumes detach rapidly into the bulk, mix there strongly, and thus cannot form large-scale structures. (b) corresponds to $Ro^{-1} = -0.22$ ($\approx Ro_{opt}^{-1}$). The reduced plume mixing enables the formation of large-scale structures, and a strong signature of them can be seen in the averaged angular velocity field. (c) corresponds to $Ro^{-1} = -0.40$ (SCR regime) and also shows some signatures of large-scale structures. However, these do not fully penetrate the gap but stop at the border to the Rayleigh-stable zones near the outer cylinder.

values of Ro^{-1} simulated here. The large-scale structures cannot be seen in figure 7.3(a), which corresponds to $Ro^{-1} = 0.20$ (co-rotating cylinders), but they are pronounced for 7.3(b,c) ($Ro^{-1} = -0.22$ and $Ro^{-1} = -0.40$). As shown in figure 7.4, in the CWC regime, the bulk sustains a large $\bar{\omega}_z$ gradient, and to accommodate for this, there is smaller $\bar{\omega}_z$ jump across the boundary layers. Plumes ejected from both cylinders can now mix easier when entering the bulk. As a consequence, the large-scale structures, which essentially consist of unmixed plumes, break up easier and thus do that for lower values of Ta . For this reason they have completely vanished in the left panel of figure 7.3.

If we now decrease Ro^{-1} , the profile becomes flatter. The effect of this is visible in figure 7.3(b) showing $\bar{\omega}$ for $Ro^{-1} = -0.22$. It can be seen from figure 7.4 that this value of Ro^{-1} corresponds to the flattest ω -profile available, and it is also the closest to the experimental optimum transport $Ro_{opt}^{-1}(Ta \rightarrow \infty) = -0.20$. A very marked signature of the large-scale structure on $\bar{\omega}$ can be seen. This is because a very flat $\bar{\omega}$ profile will sustain a large $\bar{\omega}$ jump across the boundary layer, and thus plumes detach less violently into the bulk, thus stabilizing the large-scale structures. Therefore, we can relate the flatness of the $\bar{\omega}$ -profile to the strength of the large-scale circulation, and this in turn can be related to the optimum in $Nu_\omega(Ro^{-1})$. As mentioned in Brauckmann & Eckhardt [156], optimum transport coincides with the strongest mean circulation. Plumes travel faster from one cylinder to the other when the large-scale circulation is strongest, and thus more angular momentum is transferred. We also highlight that the signature of the large-scale structures on the mean azimuthal flow remains even in the ultimate regime, and is also seen in experiment at $Ta \sim 10^{12}$ [173]. Thus in general the vanishing of the rolls appears to be independent from the transition to the ultimate regime. Only in the special case of pure inner cylinder rotation these two effects coincidentally occur at the same Ta .

In the right panel of figure 7.3, we can see that once the Coriolis force is sufficiently large, the vortices cannot fully penetrate the domain. Near the outer cylinder, the flow is predominantly Rayleigh-stable. Rayleigh-stable zones are well mixed, as transport here happens through intermittent turbulent bursts, instead of convective transport by plumes and vortices [156]. Thus, in Rayleigh-stable regions, no rolls can be seen in the averaged fields. The effect of the neutral surface can also be observed in the averaged $\bar{\omega}$ profiles (cf. figure 7.4). The two simulated cases in the SCR regime, ($Ro^{-1} = -0.30$ and $Ro^{-1} = -0.40$) show an outer cylinder boundary layer which with more and more negative Ro^{-1} extends deeper into the flow, and the distinction from the bulk is blurred away.

To further disentangle the effect of axial dependence and the transition to the ultimate regime we show the loss of axial dependence characterized by a special spread measure Δ_U as a function of the driving Ta in figure 7.5. Δ_U is defined as $\Delta_U = (\max_z(\bar{u}_\theta(r_a, z)) - \min_z(\bar{u}_\theta(r_a, z))) / \langle \bar{u}_\theta(r_a, z) \rangle_z$, with r_a (the mid-gap radius) defined

Figure 7.4: Azimuthally, axially and time-averaged and non-dimensionalized angular velocity profiles $\langle \bar{\omega} \rangle_z$ for $\eta = 0.714$, $Ta = 10^{10}$, and six values of Ro^{-1} . For co- and weakly counter-rotating cylinders, we see that the bulk $\bar{\omega}$ profiles become flatter as Ro^{-1} becomes more negative. Thus, the angular velocity difference, which the plumes encounter when detaching from the BL and entering the bulk, is larger for more negative Ro^{-1} . (b) shows a zoom-in in the bulk region of the left panel. LDA data from experiments from van Gils *et al.* [77], for which $Re_i - Re_o = 10^6$ have been superimposed. Note the good agreement between both datasets for values of Ro^{-1} in the CWCR regime, while discrepancies exist for values of Ro^{-1} around the optimum and in the SCR regime. This is attributed to the axial dependence of the profiles, which exists in this regime (see figure 7.3) as experimental data is measured at fixed height, while numerical data are axially averaged.

as $r_a = r_i + d/2$, i.e. the arithmetic mean of the inner and outer cylinder radii. When measuring the axial spread, the velocity is averaged in time, and azimuthally, as the flow is homogeneous in the azimuthal direction. As stated previously, for co-rotating cylinders the axial dependence disappears for low drivings corresponding to those in the transitional regime, and associated to the appearance of the “wavelet” states. For counter-rotating cylinders, a sharp jump in Δ_U can be noticed. This is due to Δ_U being measured at the mid-cylinder $\tilde{r} = \tilde{r}_a$. For low drivings, $\tilde{r}_a$ is located in the Rayleigh-stable zones, and the flow is mixed better. As the driving increases, turbulence from the inner cylinder pushes the neutral surface, which divides the stable and unstable zones further towards the outer cylinder. As a consequence of this pushing, $\tilde{r}_a$ is no longer in the Rayleigh-stable zone, but instead in the Rayleigh-unstable zone. This zone is dominated by large-scale structures. This makes the axial dependence increase and provides more evidence that the vanishing of the Taylor-rolls is only coincidental with the transition to the ultimate regime for pure inner cylinder rotation.

As mentioned previously, the value of Ro_{opt}^{-1} , and thus of the border between the CWCR and the SCR regimes depends on Ta . This is summarized in figure 7.6, which shows the approximate division between the different flow regimes explored in this chapter in both the (Ta, Ro^{-1}) and the (Re_i, Re_o) parameter spaces, both for $\eta = 0.714$. Ro_{opt}^{-1} , and thus the division between the regimes can be seen to saturate for $Ta \sim 5 \cdot 10^8$, when driving is large enough, and the mean $\bar{\omega}(r)$ profile at Ro_{opt}^{-1} is completely flat.

Finally, to further justify the division of the flow into the CWCR and the SCR regimes with decreasing inverse Rossby number Ro^{-1} , we can quantify the distribution of Rayleigh-stable and unstable zones as a function of Ro^{-1} . This is done by looking at the PDF of $\tilde{r}_N$, i.e. the collection of points outlining the neutral surface $\tilde{r}_N = \tilde{r}_N(t, \theta, z)$. This is, the border between Rayleigh-stable outer gap range and the Rayleigh-unstable inner gap parts, and given as the points for which $\omega(t, \theta, z, \tilde{r}_N) = 0$ in the laboratory (non-rotating) frame. For counter-rotating cylinders, the neutral surface defines the instantaneous border between Rayleigh-stable and Rayleigh-unstable zones. For co-rotating cylinders, the neutral line does not exist, and the whole flow is either Rayleigh-stable or Rayleigh-unstable. In principle, the neutral surface might be fragmented, and thus the position of $\tilde{r}_n$ multivalued. However, this is usually not the case. When taking the ensemble, all values are considered, as this does not change the PDFs significantly.

Figure 7.7 shows the PDFs of $\tilde{r}_N$ calculated for the four negative values of Ro^{-1} at the largest driving simulated here. The difference between the two regimes can clearly be noticed. In the CWCR regime and near the optimum, the border between the zones is located very closely to the outer cylinder, which means that almost all the domain is Rayleigh-unstable and dominated by plumes or rolls. In the SCR regime, the border between the zones is pushed closer towards the inner cylinder,

Figure 7.5: The axial velocity spread measure Δ_U versus Ta for the four values of Ro^{-1} in the CWCR regime. The dashed line indicates the approximate value of Ta where the flow transitions to the ultimate regime for all values of Ro^{-1} , which was previously associated with the vanishing of the large-scale structures. For co-rotating cylinders ($Ro^{-1} = 0.20$), at Ta as low as $Ta \approx 10^7$ no axial dependence is seen, well before the transition. For counter-rotating cylinders a sharp increase of the axial velocity spreading measure Δ_U can be seen, which then slowly decreases with increasing Ta . The sharp increase in Δ_U can be associated to the growth of the Rayleigh-unstable zones. For low Ta , the mid-gap is in a Rayleigh-stable zone mixed by bursts, while for large Ta , the mid-gap is in a Rayleigh-unstable zone, dominated by rolls leading to a strong height dependence. The large axial spreads explain the discrepancies when comparing (axially averaged) DNS data to experimental data measured at a fixed height.

Figure 7.6: Transition between different regimes in the (Ta, Ro^{-1}) (top-left) and (Re_i, Re_o) (top-right and bottom) parameter spaces for $\eta = 0.714$. The hollow circles indicate the location of optimal transport, and serve as an indication of the movement of the division between CWCRC (blueish and reddish) and SCR (greenish) regimes with Ta . In both DNS and experiments, Ro_{opt}^{-1} reaches an asymptotic value for $Ta > 5 \cdot 10^8$. For larger η (smaller gap), this separation line moves towards smaller Ro^{-1} . For $Ta \lesssim 10^7$, we have the rich variety of different states of Andereck *et al.* [37], not detailed in this diagram. This region appears explicitly in the top-left panel as “lam TRs” and “lam TRs at IC”, but is not shown in the other two due to the axes used. Abbreviations: boundary layer (BL), Taylor rolls (TR), ultimate regime (UR), and inner cylinder (IC).

Figure 7.7: DNS results for the PDFs of the radial position $\tilde{r}_N$ of the neutral surface, at the border between Rayleigh-stable and Rayleigh-unstable regions, for the four simulated negative values of Ro^{-1} (i.e. for counter-rotating cylinders) for $Ta = 10^{10}$. For $Ro^{-1} = -0.14$ (CWCR) and $Ro^{-1} = -0.22$ (close to the optimum), the PDFs show that the destabilizing action of the inner cylinder causes the Rayleigh-stable regions to be confined only very closely to the outer cylinder. For $Ro^{-1} = -0.22$ we can begin to see a limited amount of Rayleigh-stable zones in the whole domain, as -0.22 is slightly more negative than Ro_{opt}^{-1} . For $Ro^{-1} = -0.30$ (SCR), the stabilization due to the Coriolis force increases, and the border between the regions can be anywhere in the gap, indicating a mixture of stable and unstable zones everywhere in the gap. Finally, for $Ro^{-1} = -0.40$ (also SCR), the border between both zones never gets close to the outer cylinder. For this case, the portion of the gap width with $\tilde{r} > 0.84$ is always Rayleigh-stable.

and Rayleigh-stable zones appear all over the gap. For the most negative simulated value of Ro^{-1} , i.e. $Ro^{-1} = -0.40$, the areas near the outer cylinder are permanently Rayleigh-stable, and transport occurs in intermittent bursts which mix this zone well. This causes the partial disappearance of axial dependence seen in the right panel of figure 7.3.

7.4 The effect of radius ratio or the η -dependence

In the previous section we showed that for $\eta = 0.714$ the transition to the ultimate regime and the vanishing of the rolls only (incidentally) co-occur at the same Ta for pure inner cylinder rotation. Flatter bulk ω -profiles result in stronger large-scale structures, and steeper bulk ω -profiles result in weaker large-scale structures which vanish at $Ta \sim 10^6$. Now we will show that we can modify the $\bar{\omega}(r)$ profile in the bulk not only by varying the Coriolis force, but also by changing the radius ratio

η (or the gap width). In this section, we will thus analyse the influence of η , to understand whether the co-occurrence of the vanishing large scales and the boundary layer transition observed for pure inner cylinder rotation is just a coincidence seen in the case $\eta = 0.714$.

Figure 7.8 shows both the Nusselt number and the compensated Nusselt number plotted as a function of Ta for the three values of η simulated. As seen by Ostilla-Mónico *et al.* [176] for $\eta = 0.714$ (and now also for $\eta = 0.909$), the flow undergoes a structural transition at around $Ta \approx 3 \cdot 10^6$, where the local exponent α of the effective scaling law $Nu \sim Ta^\alpha$ rapidly decreases. This is associated with the breakdown of coherence in the flow and the onset of time-dependence in the Nusselt number. For $\eta = 0.714$ and $\eta = 0.909$, the effective exponent α begins to increase again after this breakdown. We can say that the flow transitions to the ultimate regime once $\alpha > 1/3$, and this happens at about $Ta \approx 3 \cdot 10^8$. This Ta value coincides with the experimentally observed value for the transition to the ultimate regime for $\eta = 0.909$ by Ravelet *et al.* [53].

For $\eta = 0.5$ a different behaviour can be seen. After the breakdown of coherence, the transitional regime with $\alpha \approx 1/3$ goes on for three decades in Ta , up to $Ta \approx 10^{10}$ (last three data points of the panel). An increase in α only happens for the last three data points, with $Ta > 10^{10}$. This might be the beginning of the transition to the ultimate regime, observed at about that value of Ta in the experiments by Merbold *et al.* [80]. We emphasize that the behavior of the $Nu_\omega(Ta)$ curve for $\eta = 0.5$ is very similar to the one seen for $\eta = 0.714$ and $Ro^{-1} = 0.20$ (cf. figure 7.2), while the $Nu_\omega(Ta)$ curve for $\eta = 0.909$ is similar to the one for $Ro^{-1} = -0.14$ and $\eta = 0.714$.

We thus can draw an analogy between the effects of varying η and those of changing Ro^{-1} . The larger the gap or the smaller η is, the more the flow feels the curvature. This is reflected in an asymmetry between inner and outer cylinder, since the inner cylinder curvature becomes increasingly stronger relative to the outer cylinder curvature. Also the exact relationship $\eta^{-3} \partial_r \langle \omega \rangle|_o = \partial_r \langle \omega \rangle|_i$ (cf. Ref. [77]) must hold in both boundary layers due to the r -independence of the angular velocity current $J^\omega = r^3 (\langle u_r \omega \rangle_{z,\theta,t} - \nu \partial_r \langle \omega \rangle_{z,\theta,t})$ (EGL07). For $\eta = 0.5$ we have $\eta^{-3} = 8$ and the ω -slope at the inner cylinder is eight-fold steeper than the outer cylinder ω -slope. Thus the inner-outer asymmetry is expected to become much more dominant for $\eta = 0.5$ in comparison to $\eta = 0.714$ ($\eta^{-3} = 2.75$) as well as $\eta = 0.909$ ($\eta^{-3} = 1.331$), for which it is hardly visible any more.

While the inner and outer cylinder boundary layers extend into the bulk equally for pure inner cylinder rotation [176], the *jump* of ω in the boundary layers is much larger in the inner cylinder as compared to the outer cylinder due to the different slopes and equal extents. Therefore, the plumes are highly asymmetric, and smaller drivings Ta break up the “plume conveyor belts”, which form the large-scale structures seen in the time-averaged azimuthal velocity. On top of this plume asymmetry,

Figure 7.8: (a): The nondimensional torque $Nu_\omega - 1$ versus driving Ta for pure inner cylinder rotation $Ro^{-1} = 0$ and three values of η . (b): the compensated Nusselt $(Nu_\omega - 1)/Ta^{1/3}$ versus driving strength Ta for the same three values of η . The asymptotic effective scaling laws of the ultimate regime are reached for all values of η at large enough drivings. For $\eta = 0.909$ jumps in $Nu_\omega(Ta)$ can be seen for the highest drivings (around $Ta \sim 10^{10}$). These jumps cause the exponent of the local scaling laws to be around 0.44, and are caused by changes in the large scale structures.

originating from the boundary layers, a larger curvature has an effect on the bulk. The underlying $\bar{\omega}(r)$ profile is less flat, and thus the drop in angular velocity inside the bulk is the larger the smaller the value of η is.

Both effects can be appreciated in figure 7.9, which shows contour plots of the azimuthally- and time-averaged angular velocity $\bar{\omega}$ at $Ta = 10^{10}$ for the three simulated values of η . This also explains figure 7.10(a), where the now also axially averaged angular velocity $\langle \bar{\omega} \rangle_z$ is shown for the same three values of η . For comparison, figure 7.10(b) shows three profiles of $\langle \bar{\omega} \rangle_z$ in the CWCR regime for $\eta = 0.714$.

The analogy between the effect of η and the effect of Ro^{-1} on $\omega(\tilde{r})$ is also demonstrated in figure 7.10. The rolls are weak for $\eta = 0.5$, as they are weak for co-rotating cylinders, and the rolls are strongest for $\eta = 0.909$ and for $Ro^{-1} \approx Ro_{opt}^{-1}$. This also explains why, for large enough Ta , Nu_ω is highest at a given Ta for the largest η . However, the analogy is not perfect. For pure inner cylinder rotation, i.e. for $Ro^{-1} = 0$, the wide variety of flow states seen by Andereck and coworkers [37, 171] is greatly reduced. The system essentially goes from Taylor vortex flow to modulated Taylor vortex flow to finally turbulent Taylor vortex flow. It does not undergo transitions to different states (such as e.g. the “wavelet” state), and thus the rolls do not vanish for the lower drivings at which this happens in co-rotating cylinders. This can be seen in figure 7.11, which shows the measure Δ_U for the axial velocity spread as function of Ta . With increased driving, the rolls progressively lose importance until Ta reaches a value of $Ta \approx 3 \cdot 10^8$. However, the effect of η , and thus of the cylinder wall curvature on the ω profiles can be clearly noticed in the

Figure 7.9: Contour plots of the azimuthally- and time averaged angular velocity field $\bar{\omega}$ for $Ta = 10^{10}$ and $Ro^{-1} = 0$ and three values of η : (a) $\eta = 0.5$, (b) $\eta = 0.714$ and (c) $\eta = 0.909$. The colour scale has been shifted in order to account for the different bulk angular velocities at different η . Almost no axial dependence can be noticed for $\eta = 0.5$, while it is still very marked for $\eta = 0.909$

Figure 7.10: (a): Axially averaged angular velocity profiles $\langle \bar{\omega} \rangle_z$ for $\eta = 0.5$, $\eta = 0.714$, and $\eta = 0.909$ at moderate driving $Ta = 10^{10}$ and $Ro^{-1} = 0$. Solid lines are DNS data, while squares and triangles correspond to LDA data from experiments ($Ta = 1.51 \cdot 10^{12}$ for $\eta = 0.714$ and $Ta = 1.1 \cdot 10^{11}$ for $\eta = 0.909$, from Ref. [176]). A larger decrease of ω across the bulk can be seen for $\eta = 0.5$. The angular velocity in the bulk also deviates more from $\omega = 0.5$, the expected value in the limit case of $\eta \rightarrow 1$ (plane Couette flow). (b): Axially averaged angular velocity profiles $\langle \bar{\omega} \rangle_z$ for $\eta = 0.714$, and three values of Ro^{-1} in the CWCR regime. The analogy between the effects of η and Ro^{-1} on $\bar{\omega}(\tilde{r})$ can be clearly seen.

Figure 7.11: The measure Δ_U for the axial velocity spread versus Ta for the three values of η simulated. A decrease in axial dependence can be seen for all values of η around $Ta \approx 10^8$, unlike what was seen for varying Ro^{-1} , where the Ta at which the decrease of axial dependence took place is Ro^{-1} dependent. However, the residual axial spread at the largest drivings increases with increasing η , as we would expect from the analogy between decreasing Ro^{-1} and increasing η .

residual axial dependence and behaves as expected from the analogy. The behaviour of the transition to the ultimate regime and associated sub-regimes is summarized in figure 7.12, which is analogous to figure 7.6, but now for the (Ta, η) parameter space explored.

Finally, one may ask the question of why the onset of the ultimate regime happens at a much higher Ta for $\eta = 0.5$ than for the two other values of η studied. For $\eta = 0.714$, the transition seems to set in for the same value of Ta independently of Ro^{-1} . A factor ten increase in shear in the boundary layers is required for the boundary layer instability to occur and the ultimate regime to set in. Convex curvature is known to produce a stabilizing effect on boundary layers [177, 178], and this will have a more significant effect on the inner cylinder for $\eta = 0.5$ than for the larger η . On the other hand we might expect that the destabilizing effect of concave curvature [179, 180] would also play a role in accelerating the transition. Due to the boundary layer asymmetry however, the outer boundary layer is much more “quiet”, and has less fluctuations. This also delays the transition, and can be seen in figure 7.13, which shows the rms-fluctuations of the angular velocity $\omega' = \langle \langle \omega^2 \rangle_{t,\theta} - \bar{\omega}^2 \rangle_z^{1/2}$, for $Ta = 10^9$ and the three values of η simulated. The levels of fluctuations at the outer cylinder are significantly reduced for $\eta = 0.5$ when compared to the other values of η . Finally, the large gradient of angular velocity sustained in the bulk will also reduce

Figure 7.12: Transition between different regimes in the (Ta, η) parameter space for pure inner cylinder rotation $Ro^{-1} = 0$. The transition to the ultimate regime occurs at a higher Ta for smaller η (wider gap), while the vanishing of the large scale structures occurs at around the same Ta for $0.5 < \eta < 0.714$. Remains of the Taylor rolls can only be seen for large η , i.e. smaller gap. Abbreviations: boundary layer (BL), Taylor rolls (TR), ultimate regime (UR).

the shear in the outer cylinder, as the bulk angular velocity is smaller for $\eta = 0.5$. Thus, a combination of reduced fluctuations, stabilizing effect due to curvature at the inner cylinder, and reduced shear due to bulk angular velocity gradients is causing the delayed transition.

7.5 Dependence on number and size of rolls

Finally, we will quantify how the torque depends on the number and the size of the rolls, i.e. the vortical wavelength. The wavelength of a roll λ_z is restricted to the values $\lambda_z = \Gamma/n$, where n is a strictly positive integer. For all simulations in this chapter, $n = 1$, and thus $\lambda_z = \Gamma$. This is not necessarily always the case, n is a response of the system, and if Γ is large enough, i.e. the system can accommodate more than one vortex pair, n can take several values depending on how the final state of the system is reached. Brauckmann & Eckhardt [123] showed that for $\eta = 0.714$, the “optimal” vortex wavelength, i.e. the vortex wavelength λ_z which corresponds to a maximum Nu_ω , increased when comparing $Nu_\omega(\lambda_z)$ for two Taylor numbers, one in the Taylor vortex regime and another in the turbulent Taylor vortex regime. For the higher Ta , the dependence of Nu_ω on λ_z was quite weak. Martinez-Arias *et al.* [174] showed that for $\eta = 0.909$, different branches in the $Nu_\omega(Ta)$ relationship, associated to distinct vortical states cross around $Re_i = 1.3 \cdot 10^4$. This corresponds to a driving of $Ta = 1.8 \cdot 10^8$, around the value at which the transition to the ultimate regime occurs for $\eta = 0.909$. The large-scale circulation could still be seen to play a role in deter-

Figure 7.13: Root mean square (rms) profiles of the angular velocity fluctuations, $\omega'(\tilde{r})$, at $Ta = 10^9$ and $Ro^{-1} = 0$ for the three η values simulated here. The boundary layer asymmetry causes the fluctuations to be strongly reduced at the outer cylinder for $\eta = 0.5$, as compared to those at the inner cylinder.

mining the system response after the transition to the ultimate regime. Furthermore, large scale patterns were observed in Ostilla-Mónico *et al.* [113] when looking at the $\langle \tilde{\omega} \tilde{u}_r \rangle$ correlation at $Ta \sim 10^{10}$, even though they are absent when looking only at $\tilde{\omega}$.

Figure 7.14 shows the compensated torque Nu_ω as function of Ta for the four values of the vortical wavelength studied. Experimental data from Ref. [174] and DNS data from Ref. [176] is also plotted. It is worth noting that experimental data will have some end-plate effects, even if the aspect ratio Γ of the experiments is larger than 30, while the DNSs have periodic axial boundary conditions. Even so, very similar behaviour can be seen. The transition to the asymptotic scaling laws of the ultimate regime seem to occur around the same value of Ta , but are less pronounced the smaller the vortical wavelength is.

The change in behaviour of the $Nu_\omega(Ta)$ curves can be associated to the change of behaviour of the wind-sheared regions in the ultimate regime. As seen in Ostilla-Monico *et al.* [113], plume ejection is suppressed outside the ultimate regime in regions of the flow, the so called “wind-sheared” regions due to the sweeping by the large scale rolls. This reduction in plume ejection results in a reduced transport of angular velocity (torque). A similar reduction in the torque caused by a mean flow was also seen when forcing the flow with an axial pressure gradient by Manna & Vacca [181]. Vortices with a smaller wavelength have smaller wind-sheared regions and thus result in a larger Nu_ω , if this suppression is taking place. After the transition to the ultimate regime, the suppression ceases, and these regions become active

Figure 7.14: Compensated torque Nu_ω versus driving strength Ta for $\eta = 0.909$ and the three different vortical wavelengths. Experimental data from the T^3C apparatus ($\Gamma = 46.35$, number of rolls not determined, cf. Ref. [176]) and from Ref. [174] (denoted MPCM14, $\lambda_z = 2$ corresponds to 30 rolls and $\lambda_z = 3$ corresponds to 18 rolls) are also plotted. Axial boundary conditions are different in experiments and DNSs. Experiments have end-plates, while DNS are axially periodic and thus end effects are absent. Both in experiment and in numerics, different branches associated to different states cross at $Ta \approx 2 \cdot 10^8$, shown as a vertical dashed line in the graph. This value of Ta corresponds to the transition to the ultimate regime for radius ratio $\eta = 0.909$.

ejectors of plumes, leading to increased transport.

The difference between $\lambda_z = 2.09$, $\lambda_z = 3.0$ and the $\lambda_z = 4.0$ is very small for $Ta = 10^9$, of the order of 5%, but for the $\lambda_z = 1.5$ branch the difference is almost 15%. Only at $Ta = 10^{10}$, when the distinction between wind-sheared and ejection regions is completely blurred away, and the whole inner cylinder can emit plumes (or hairpin vortices), $Nu_\omega(Ta)$ loses its Γ dependence, within the error bars of the numerics. This sudden transition of wind-sheared regions to ejection regions causes the jump we see in the $Nu_\omega(Ta)$ curve at around $Ta = 5 \cdot 10^9$ for $\eta = 0.909$.

Note that for the largest drivings axially periodic boundary conditions have been used, with only one vortex pair. This does not prevent the creation of two pairs of vortices with wavelength $\lambda_z = 1.5$ by a breakup of one pair of vortices of $\lambda_z = 3.0$ in a domain, which has $\Gamma = 3.0$. And indeed this is seen to happen for the lower drivings both in DNS and experiment. On the other hand, this axial periodicity affects the stability of one pair of vortices of wavelength $\lambda_z = 1.5$ in a domain of $\Gamma = 1.5$. Therefore, vortices with $\lambda_z = 1.5$ might be an artefact due to the numerical constraints, and not be stable if a system with large Γ at large Ta is considered. States with $\lambda_z < 2$ are not reported by Martinez-Arias *et al.* [174].

Even if we do not expect a quantitative agreement of the present DNS results with those of Ref. [123] and experimental data by Huisman *et al.* [173], as we simulate a different η , the results reported in this section even do not agree qualitatively. Brauckmann & Eckhardt [123] see a maximum in torque for $\lambda_z = 1.93$ in the turbulent Taylor vortex regime ($Ta \sim 10^7$), while in the present simulations for $\eta = 0.909$ at the same Ta , this maximum is clearly at $\lambda_z = 1.5$, and not near $\lambda_z = 2.09$. In the experiments of Ref. [174], states with λ_z smaller than one are not reported, and a direct comparison cannot be made.

We also note that the relationship between larger vortices and larger torque in the ultimate regime is the inverse of what was recently reported by Huisman *et al.* [173], who found multiple states, with different λ_z in highly turbulent TC flow. For different states they found that the torque differs less than 5%, although they note that this might be due to the fact the torque is only measured on part of the inner cylinder, not on the entire inner cylinder. Furthermore their results are for $Ro^{-1} \neq 0$, for higher Ta , and for different η , as compared to the current research.

7.6 Summary and conclusions

Numerical simulations of turbulent Taylor-Couette flow in the range $10^4 < Ta < 4.6 \cdot 10^{10}$ were performed to explore the transition of TC flow to the (fully turbulent) ultimate regime. The four dimensions of the parameter space were explored, including the dependence of the transition on the radius ratio η , the vortex wavelength λ_z and Coriolis force Ro^{-1} or rotation ratio μ .

First, the effect of the outer cylinder rotation, in the equations of motion in the frame co-rotating with the outer cylinder, present as a Coriolis force, was analysed for $\eta = 0.714$. Depending on the value of Ro^{-1} two regimes were identified, (i) the co-rotating and weakly counter-rotating cylinder regime (CWCR) and (ii) the strongly counter-rotating cylinder regime (SCR), both with their respective sub-regime. Our findings of that chapter culminate in the phase diagram fig. 7.6, in the (Ta, Ro^{-1}) regime (figure 7.6(a) and in the (Re_i, Re_o) regime (figures 7.6(b & c). The transition to the ultimate regime could be observed for all values of Ro^{-1} around $Ta \sim 3 \cdot 10^8$. However, for these two regimes a rather different behaviour in the scaling laws $Nu_\omega(Ta)$ was found before the transition. We also found very different flow structures in the respective ultimate regimes in accordance with the description by Brauckmann & Eckhardt [156]. An explanation why the Coriolis force, proportional to Ro^{-1} stabilizes the large-scale structures was illustrated; the large-scale structures were found to *not* vanish at the transition to the ultimate regime for $Ro^{-1} = -0.22 \approx Ro_{opt}^{-1}$, unlike what was seen in Ref. [113] for stationary outer cylinder.

After this, the transition was analysed for various gap widths, namely for $\eta = 0.5, 0.714$, and 0.909 without Coriolis forces, i.e., for $Ro^{-1} = 0$. The transition was found to occur at about the same Ta for $\eta = 0.714$ and 0.909 . However, the transition was considerably delayed to $Ta \approx 10^{10}$ for $\eta = 0.5$, due to the combined effects of stabilizing curvature of the inner cylinder, and the reduced shear as well as smaller fluctuations in the vicinity of the outer cylinder. An analogy between the effect of Ro^{-1} in the CWCR regime and the effect of η on the large scale rolls was described: Decreasing η was found to have the same effect as adding a positive Ro^{-1} –corresponding to co-rotating cylinders–, while increasing η behaved like (weakly) counter-rotating the outer cylinder.

Finally, as the large-scale structures were found to be strongest for $\eta = 0.909$, the effect of varying the vortical wavelength was analysed for this value of η . As in Martinez-Arias *et al.* [174], different branches of the $Nu_\omega(Ta)$ curve were found to cross around the transition to the ultimate regime. Before this transition, the influence of the vortical wavelength (and thus of the aspect ratio) on Nu_ω was quite noticeable. After the ultimate range transition, this effect decreased drastically. The results of our DNS agree qualitatively with those in the experiments in Ref. [174] for $\eta = 0.909$ even though the axial boundary conditions are different. However, they are qualitatively different from those reported for $\eta = 0.714$ in Refs. [123, 173].

In this work, the vortical wavelength by using periodic boundary conditions was fixed. Some of these states might not be accessible in experiment or might be a product of the periodic boundary conditions. Studying the coexistence of different states for large Γ , like done in Refs. [173, 174] with DNS requires a large amount of computational resources for high Ta . Switches between two and three vortex pairs

were seen at lower Ta for $\eta = 0.909$ [176]. Switching between states might also occur at high Ta , although they are not captured in the DNS presented in this work. In the future, additional DNS for $\eta = 0.909$ with large Γ at high Ta should be run to improve the understanding of the switching between different states.

Our ambition also is to further understand why the transition is delayed at $\eta = 0.5$, but also the curvature effects on the ω -profiles in the boundary layers along the ideas of Grossmann *et al.* [165]. Curvature effects at $\eta = 0.714$ and $\eta = 0.909$ are too small to be appreciated, and the flow for $\eta = 0.5$ is still in the transition to the “ultimate” regime. Thus, higher Ta simulations for $\eta = 0.5$ will provide further understanding on how curvature makes the boundary layers of TC flow different from those of channel and pipe flow.

Table 7.1: Summary of the numerical results for $\eta = 0.714$ which are new to this chapter. for the other data points see Ref. [113] (§6). The columns show the driving, Ta , the outer cylinder rotation as either a Coriolis force Ro^{-1} or a rotation frequency ratio $\mu = \omega_o/\omega_i$, the non-dimensionalized torque, Nu_ω , the amount of grid points used in azimuthal (N_θ), radial (N_r) and axial direction (N_z). All these simulations use a rotational symmetry order six in the azimuthal direction, and are for $\Gamma = 2.09$.

Ta	Ro^{-1}	μ	Nu_ω	$N_\theta \times N_r \times N_z$
$2.15 \cdot 10^8$	0.20	0.2	11.48	$256 \times 640 \times 512$
$2.15 \cdot 10^8$	-0.13	-0.2	13.43	$256 \times 640 \times 512$
$2.15 \cdot 10^8$	-0.22	-0.4	12.85	$256 \times 640 \times 512$
$2.15 \cdot 10^8$	-0.30	-0.6	11.13	$256 \times 640 \times 512$
$2.15 \cdot 10^8$	-0.40	-1.0	8.565	$256 \times 640 \times 512$
$4.64 \cdot 10^8$	0.20	0.2	14.21	$256 \times 640 \times 512$
$4.64 \cdot 10^8$	-0.13	-0.2	17.20	$256 \times 640 \times 512$
$4.64 \cdot 10^8$	-0.22	-0.4	17.77	$256 \times 640 \times 512$
$4.64 \cdot 10^8$	-0.30	-0.6	15.81	$256 \times 640 \times 512$
$4.64 \cdot 10^8$	-0.40	-1.0	11.36	$256 \times 640 \times 512$
$1.00 \cdot 10^9$	0.20	0.2	18.57	$256 \times 640 \times 512$
$1.00 \cdot 10^9$	-0.13	-0.2	23.10	$256 \times 640 \times 512$
$1.00 \cdot 10^9$	-0.22	-0.4	23.18	$256 \times 640 \times 512$
$1.00 \cdot 10^9$	-0.30	-0.6	19.85	$256 \times 640 \times 512$
$1.00 \cdot 10^9$	-0.40	-1.0	14.73	$256 \times 640 \times 512$
$2.15 \cdot 10^9$	0.20	0.2	24.96	$256 \times 640 \times 512$
$2.15 \cdot 10^9$	-0.13	-0.2	31.26	$256 \times 640 \times 512$
$2.15 \cdot 10^9$	-0.22	-0.4	31.41	$256 \times 640 \times 512$
$2.15 \cdot 10^9$	-0.30	-0.6	27.46	$256 \times 640 \times 512$
$2.15 \cdot 10^9$	-0.40	-1.0	20.15	$256 \times 640 \times 512$
$4.64 \cdot 10^9$	0.20	0.2	32.51	$384 \times 640 \times 768$
$4.64 \cdot 10^9$	-0.13	-0.2	41.44	$384 \times 640 \times 768$
$4.64 \cdot 10^9$	-0.22	-0.4	41.13	$384 \times 640 \times 768$
$4.64 \cdot 10^9$	-0.30	-0.6	36.39	$384 \times 640 \times 768$
$4.64 \cdot 10^9$	-0.40	-1.0	26.01	$384 \times 640 \times 768$
$1.00 \cdot 10^{10}$	0.20	0.2	41.01	$512 \times 800 \times 1024$
$1.00 \cdot 10^{10}$	-0.13	-0.2	57.50	$512 \times 800 \times 1024$
$1.00 \cdot 10^{10}$	-0.22	-0.4	58.61	$512 \times 800 \times 1024$
$1.00 \cdot 10^{10}$	-0.30	-0.6	49.98	$512 \times 800 \times 1024$
$1.00 \cdot 10^{10}$	-0.40	-1.0	34.42	$512 \times 800 \times 1024$
$2.15 \cdot 10^{10}$	0	0	66.57	$768 \times 1024 \times 1536$
$4.64 \cdot 10^{10}$	0	0	94.77	$768 \times 1200 \times 2048$

Table 7.2: Summary of the numerical results for the various geometries at $Ro^{-1} = 0$, i.e., for resting outer cylinder and $\Gamma = 2.09$ with one roll (thus $\lambda_z = 2.09$), which are new to this chapter. For the other data points see Ref. [112] and Ref. [176] (§§4 & 5). The columns show the driving Ta , the radius ratio η , the non-dimensionalized torque Nu_ω , the amount of grid points used in azimuthal (N_θ), radial (N_r) and axial direction (N_z).

Ta	η	Nu_ω	$N_\theta \times N_r \times N_z$
$2.15 \cdot 10^8$	0.5	9.33	$384 \times 512 \times 768$
$4.64 \cdot 10^8$	0.5	11.9	$384 \times 701 \times 768$
$1.00 \cdot 10^9$	0.5	14.9	$512 \times 768 \times 768$
$2.15 \cdot 10^9$	0.5	18.8	$768 \times 768 \times 1024$
$4.64 \cdot 10^9$	0.5	24.1	$768 \times 768 \times 1024$
$1.00 \cdot 10^{10}$	0.5	31.3	$1024 \times 1024 \times 1536$
$2.15 \cdot 10^{10}$	0.5	40.9	$1024 \times 1024 \times 1536$
$4.64 \cdot 10^{10}$	0.5	53.9	$1024 \times 1024 \times 2048$
$2.76 \cdot 10^7$	0.909	12.8	$256 \times 512 \times 480$
$5.26 \cdot 10^7$	0.909	16.8	$256 \times 512 \times 480$
$1.00 \cdot 10^9$	0.909	22.6	$512 \times 768 \times 768$
$2.15 \cdot 10^9$	0.909	31.3	$512 \times 768 \times 768$
$4.64 \cdot 10^9$	0.909	43.6	$1024 \times 768 \times 768$
$1.00 \cdot 10^{10}$	0.909	67.2	$1024 \times 1024 \times 1024$
$2.15 \cdot 10^{10}$	0.909	99.3	$1536 \times 1536 \times 1024$
$4.64 \cdot 10^{10}$	0.909	138	$2048 \times 1536 \times 1024$

Table 7.3: Summary of the numerical results for the various geometries at $Ro^{-1} = 0$ i. e., for resting outer cylinder, $\eta = 0.909$ and various Γ which are new to this chapter. For the other data points see Ref. [112] and Ref. [176] (§§4 & 5). The columns show the driving Ta , the aspect ratio Γ , the vortical wavelength λ_z , the non-dimensionalized torque Nu_ω , the amount of grid points used in azimuthal (N_θ), radial (N_r) and axial direction (N_z). We note that the $Ta = 1.11 \cdot 10^{10}$, $\Gamma = 4$, $\eta = 0.909$ simulation has $n_{sym} = 10$.

Ta	Γ	λ_z	Nu_ω	$N_\theta \times N_r \times N_z$
$1.00 \cdot 10^6$	1.50	1.50	4.31	$256 \times 512 \times 480$
$1.00 \cdot 10^7$	1.50	1.50	7.46	$256 \times 512 \times 480$
$2.76 \cdot 10^7$	1.50	1.50	9.15	$256 \times 512 \times 480$
$5.26 \cdot 10^7$	1.50	1.50	9.91	$256 \times 512 \times 480$
$1.00 \cdot 10^8$	1.50	1.50	10.9	$256 \times 512 \times 480$
$2.15 \cdot 10^8$	1.50	1.50	12.6	$256 \times 512 \times 480$
$4.64 \cdot 10^8$	1.50	1.50	15.6	$256 \times 512 \times 480$
$1.00 \cdot 10^9$	1.50	1.50	20.8	$512 \times 512 \times 480$
$1.00 \cdot 10^6$	3.00	3.00	3.60	$256 \times 512 \times 480$
$1.00 \cdot 10^7$	3.00	3.00	6.10	$256 \times 512 \times 480$
$2.76 \cdot 10^7$	3.00	3.00	7.50	$256 \times 512 \times 480$
$5.26 \cdot 10^7$	3.00	3.00	8.58	$256 \times 512 \times 480$
$1.00 \cdot 10^8$	3.00	3.00	10.3	$256 \times 512 \times 480$
$2.15 \cdot 10^8$	3.00	3.00	12.8	$256 \times 512 \times 480$
$4.64 \cdot 10^8$	3.00	3.00	17.4	$256 \times 512 \times 480$
$1.00 \cdot 10^9$	3.00	3.00	23.1	$512 \times 512 \times 720$
$1.11 \cdot 10^{10}$	3.00	3.00	68.9	$1024 \times 1024 \times 3072$
$2.76 \cdot 10^7$	4.00	4.00	7.12	$256 \times 512 \times 480$
$5.26 \cdot 10^7$	4.00	4.00	8.40	$256 \times 512 \times 480$
$1.00 \cdot 10^8$	4.00	4.00	10.3	$256 \times 512 \times 480$
$2.15 \cdot 10^8$	4.00	4.00	13.5	$256 \times 512 \times 480$
$4.64 \cdot 10^8$	4.00	4.00	18.1	$256 \times 512 \times 480$
$1.00 \cdot 10^9$	4.00	4.00	24.0	$512 \times 512 \times 720$
$1.11 \cdot 10^{10}$	4.00	4.00	69.8	$2048 \times 1024 \times 4096$

8

The effect of Taylor rolls on highly turbulent Taylor-Couette flow *

Half of science is asking the right questions

—ROGER BACON

Direct numerical simulations of the Taylor-Couette problem have been performed. Reynolds numbers of up to $3 \cdot 10^5$, corresponding to frictional Reynolds numbers of $Re_\tau \approx 4000$ were reached. The gap between the cylinders was kept small by fixing the radius ratio to $\eta = r_i/r_o = 0.909$. Small gap TC was found to be dominated by large scale structures, which are permanent in time and known as Taylor rolls (TRs). TRs are attached to the boundary layer, and are active, i.e they transport angular velocity through Reynolds stresses. Evidence for the existence of logarithmic velocity fluctuations, and of an overlap layer where the velocity fluctuations collapse in outer units was also found. An externally imposed axial flow of comparable strength as the wind of the TRs was found to convect them without any weakening effect.

*Based on: R. Ostilla-Mónico, R. Verzicco, S. Grossmann and D. Lohse, *The effect of Taylor rolls on highly turbulent Taylor-Couette flow*, under review (arXiv:1501.03012).

8.1 Introduction

Direct numerical simulations (DNS) are a very versatile tool for the study of turbulence, and have led to deep insight on its nature in the last 30 years. The seminal work of Kim, Moin and Moser [94] began a line of simulations of wall-bounded flows which is very active in the present day. Kim, Moin and Moser studied a particular kind of wall-bounded flow, namely pressure-driven (Poiseuille) flow bounded by two parallel plates, from now on referred to as channel flow. After this seminal work, which achieved $Re_\tau = 180$, simulations of wall-bounded flows focused mainly on channel flows, even if some attention was given to zero-pressure-gradient boundary layers (ZPGBL) and pipes. Present day simulations achieve much higher frictional Reynolds numbers: $Re_\tau = 4600$ for channels [95, 96], $Re_\tau = 2000$ for ZPGBL [182] and $Re_\tau \approx 1100$ for pipes [183]. DNS has allowed for detailed study of the near-wall energy cascade and for correlations between different canonical flows to be established. This has led to further understanding of the attached-eddy model of the logarithmic layer by Townsend [9], further developed by Perry and coworkers [10, 83]. We refer the reader to Ref. [12] for a recent review of advances in this field.

DNS of shear driven flow between two parallel plates, i.e. plane Couette (PC) flow has been even more challenging. This is due to the extremely large and wide structures present in the turbulent flow, seen both experimentally and numerically [184]. DNS of PC requires much larger computational boxes than that of channels, a factor 10 in both spanwise and streamwise directions. This meant that only this year $Re_\tau = 550$ was achieved by Avsarkisov *et al.* [185] and more recently $Re_\tau \approx 1000$ by Pirozoli *et al.* [186]. An alternative to having two independently moving parallel plates is having two independently rotating coaxial cylinders. Such system is known as cylindrical Couette flow or Taylor-Couette (TC) flow. TC is a closed system, which makes experimental realizations easier to construct. Furthermore, TC does not require the large computational boxes of PC, which makes higher Re_τ easier to achieve in DNS.

However, TC flow is fundamentally different from PC flow. Unlike PC (and all the flows mentioned previously) TC is linearly unstable if $(d(r^2\omega)/dr)^2 < 0$, where r is the radial coordinate and ω the angular momentum [84]. Faisst & Eckhardt [187] explored the transition from PC to TC, and found that only for radius ratios $\eta = r_i/r_o > 0.99$, where r_i and r_o are the inner and outer cylinder radii respectively, the sub-critical PC instabilities overcame the super-critical TC instabilities.

Therefore, adding a minute curvature to PC changes the flow properties drastically. Bradshaw [188] was the first to note "the surprisingly large effect exerted on shear-flow turbulence by curvature of the streamlines in the plane of the mean shear", when studying boundary layers over curved surfaces. For curved channels Hunt & Joubert [189] and Hoffmann *et al.* [180] found that adding a very weak destabilizing (concave) curvature $d/R \sim 0.01$, where d is the boundary layer thickness and R the

radius of curvature was found to produce significant differences in the mean and fluctuation velocity profiles. This is because destabilizing curvature adds a new mode of instability, which is reflected as Taylor-Görtler vortices in the flow and a change in the flow dynamics. We note that pipes are not a member of this group. The natural curvature of pipes is in planes perpendicular to the mean shear, so its effects are milder.

Both channel flow and TC have been used as a playground for investigating drag reduction through deformable bubbles or riblets [133, 190–193]. From the previous discussion, we would expect significant differences between both systems, but instead, very similar results were obtained. In both cases drag was reduced, and both the bubbles and riblets modified the boundary layer. This may be taken as an indication that as similar drag reduction mechanisms can be seen in both systems, the boundary layers in both systems may also be similar. Further evidence for such a universal behaviour was provided by recent experiments: Huisman *et al.* [82] measured the mean velocity profiles in TC and obtained a von Kármán constant $\kappa \approx 0.4$ for the highest Reynolds numbers achieved (with $Re_\tau \approx 30000$), in line with that seen in experimental pipes [194]. Outside the boundary layer, things change. The linear instability in TC causes the formation of large-scale structures, which fill the entire channel and are permanent in time. These are known as Taylor rolls, after the seminal work by G. I. Taylor [33]. Taylor rolls have been observed in both experiments and simulations, even up to $Re = 10^6$ [173]. If the Reynolds number is large enough, Taylor rolls play very little effect in determining the torque required to drive the cylinders [129]. However, their signature in the velocity field is still apparent, clearly appearing in the mean fields. An analogous structure which fills the entire domain and is persistent in time was also found in curved channels with fully developed turbulence: large scale Taylor-Görtler vortices [189].

Many fundamental differences seem to exist between pipe & channel flows and TC flow. In this chapter we wish to further characterize the TC system in general, and in particular the effect of the Taylor rolls and of curvature. To do this, four high Reynolds number DNSs of TC flow were performed. All simulations were done for pure inner cylinder rotation, i.e. $\omega_o = 0$, and with an axial periodicity aspect ratio $\Gamma = L_z/d = 2\pi/3$, where L_z is the axial periodicity length. With this Γ , the system fits a single Taylor roll pair of wavelength $\lambda_{TR} = 2\pi/3$. We focus on $\eta = 0.909$ which has a mild enough curvature to make it linearly unstable, but small enough not to be dominated by streamline topology. Simulations at shear Reynolds numbers of $Re_s = dr_i\omega_i/\nu = 10^5$, $Re_s = 2 \cdot 10^5$ and $Re_s = 3 \cdot 10^5$, were performed, where d is the gap-width $d = r_o - r_i$, r_i and r_o are the inner and outer cylinder radius respectively, ω_i and ω_o are the inner and outer cylinder angular velocity respectively, and ν is the kinematic viscosity of the fluid. This results in a frictional Reynolds numbers at the inner cylinder between $Re_{\tau,i} \sim 1000$ and $Re_{\tau,i} \sim 4000$. For TC, we define the inner

cylinder frictional Reynolds number $Re_{\tau,i}$ as $Re_{\tau,i} = u_{\tau,i}d/(2\nu)$, where the frictional velocity is $u_{\tau,i} = \sqrt{\tau_{w,i}/\rho}$, with ρ the fluid density and τ_w the stress at the inner cylinder wall. The frictional Reynolds number at the outer cylinder is obtained from $Re_{\tau,o} \approx \eta Re_{\tau,i}$.

To probe the stability of the rolls to a transversal velocity, an additional simulation at $\eta = 0.909$ and $Re_s = 10^5$ with an imposed axial (spanwise) pressure gradient was performed. The pressure gradient resulted in an average axial velocity $U_w \approx \frac{1}{10}U_i$, where U_i is the inner cylinder velocity $U_i = r_i\omega_i$. This U_w is of the order of magnitude of the characteristic velocity of the Taylor rolls, but small enough so that the system dynamics is not dominated by this secondary flow.

The chapter is organized as follows. §8.2 presents the numerical setup, and details the simulations. §8.3 presents the results and discussions. The findings are summarized in §8.4 and an outlook for future work is given.

8.2 Numerical details

The DNS were performed using a second order centred finite difference scheme with fractional-time stepping [103]. This scheme has been used and validated extensively in the context of TC (cf. comparison to experiments in Refs. [113, 129] (§6 & §7)). In order to perform the simulations at high Reynolds numbers, following Brauckmann & Eckhardt [123] “small” computational boxes were used. Instead of simulating the full azimuthal extent of the cylinder, a cylindrical wedge was simulated by imposing a rotational symmetry $n_{sym} = 20$ for $\eta = 0.909$. This gives an azimuthal extent at midgap of $1.05\pi d$. The axial box size was also minimal, Γ was set to $\Gamma = 2\pi/3$ for all simulations, meaning that a single Taylor roll pair could fit in to the domain. Simulating more than one roll was shown to be unnecessary for producing accurate results for the torque by Brauckmann & Eckhardt [123]. Furthermore, very sharp dropoffs of the axial spectra can be seen by Dong [48] for wavelengths larger than one roll.

Full details of the numerical simulations are presented in table 8.1. To demonstrate that our small boxes are sufficient, we refer the reader to §8.3.4, which show that the boxes are large enough for the autocorrelations to change sign in both axial and azimuthal direction. As previously mentioned, the azimuthal (streamwise) correlations decay much faster than those of PC flow, and thus smaller boxes can be used. The axial and azimuthal spectra shown in §8.3.5 demonstrate that the mesh is sufficient to capture the small scales. We mention that $\Delta z^+ \approx 5$ is a marginally resolved case, while for the azimuthal direction, coarser resolutions can be used without loss of accuracy.

The simulations were run for about 100 large eddy turnover times based on $\tilde{t} = d/U_i$, after transient behaviour had died out. In time units based on the frictional

Case	Re_s	N_θ	N_r	N_z	Δx^+	Δz^+	Nu_ω	$Re_{\tau,i}$	Colour
R1	$1 \cdot 10^5$	1024	1024	2048	9.1	2.7	69.5 ± 0.2	1410	Blue
R2	$2 \cdot 10^5$	1536	1536	3072	11.4	3.4	126 ± 2.1	2660	Black
R3	$3 \cdot 10^5$	1536	1536	3072	16.8	5.1	171 ± 2.5	3920	Red
AF	$1 \cdot 10^5$	1024	1024	2048	9.1	2.7	66.2 ± 0.5	1390	Green

Table 8.1: Details of the numerical simulations. The first column is the name with which the simulation will be refereed to in the chapter. The second column is Re_s , the shear Reynolds number. The third to fifth columns represent the amount of points in the azimuthal, radial and axial directions, while the sixth and seventh columns show the resolution in inner wall-units at the mid-gap, $x = \frac{1}{2}(r_i + r_o)\Delta\theta$. The eighth column is the non-dimensional torque Nu_ω , and the ninth column is $Re_{\tau,i} = u_{\tau,i}d/(2\nu)$, the frictional Reynolds number at the inner cylinder. $Re_{\tau,o}$ can be obtained from $Re_{\tau,o} = \eta Re_{\tau,i}$. The AF case has an imposed axial flow with a mean velocity $U_w = \frac{1}{10}r_i\omega_i$. The final column shows the color code used for that simulation for figures 8.3, 8.5, 8.6 and 8.8-8.10.

velocity u_τ and half the gap width, this is between 3 and 4 turnover times. This might seem small when comparing to channels, but it is sufficient for TC flow with pure inner cylinder rotation. The characteristic timescale for TC appears to be $\mathcal{O}(d/U_i)$. Dong [49] already showed that the decay of temporal autocorrelations in TC to happen in $\tilde{t} \approx 3$. This fast time scale is further quantified in the figure 8.1(a), which presents the instantaneous torque, non-dimensionalized as a quasi-Nusselt number $Nu_\omega = T/T_{pa}$ for the R1 simulation, where T is the torque and T_{pa} is the torque in the purely azimuthal, laminar state, and in figure 8.1(b) which shows the instantaneous azimuthal velocity at two points for the R2 case.

As we will see in more detail in §8.3.5, the Reynolds stresses which transport angular velocity are mainly localized in low-wavelength eddies, which have velocities of order $\mathcal{O}(U_i)$. This is not the case in channel flow, where the largest eddies are inactive, i.e. they do not transport Reynolds-stresses [9, 195]. Reynolds-stresses are transported by boundary layer detachments, or streaks, with velocity scales of $\mathcal{O}(u_\tau)$, which then naturally leads to the time scales $\mathcal{O}(d/u_\tau)$ for shear transport. However, for the large scales to saturate in energy, much larger times are needed. Large scales are “fed” by detachments from the boundary layers with velocity $\mathcal{O}(u_\tau)$. This means that the transients in TC flow have time scale $\mathcal{O}(d/u_\tau)$.

Figure 8.1: (a): Instantaneous Nu_ω at inner (IC) and outer (OC) cylinders for the R1 simulation. (b): Instantaneous azimuthal velocity near the inner cylinder (blue) and in the mid-gap (black) for the R2 simulation. The temporal origin is arbitrary. Fluctuations are on timescales of the order $\mathcal{O}(d/U_i)$.

8.3 Results

8.3.1 Visualization of the effect of pressure gradient

As mentioned previously, we have attempted to weaken the rolls by the addition of an axial pressure gradient, such that a mean axial flow $U_w \approx 0.1r_i\omega_i$ was sustained (case AF in table 8.1). Figure 8.2 shows a pseudocolour plot of the instantaneous azimuthal velocity for the R2 case, and the AF case in two separate instances in time. The large scale patterns caused by the presence of an underlying Taylor-roll can be appreciated in all the panels. This roll is stationary in time in the case of the R2 simulation. For the AF case, the mean axial velocity satisfies $u_\tau \ll U_w \ll U_i$, but it does not prevent the formation of the Taylor rolls. Instead, it is slowly convected upwards in the computational domain, reappearing on the other side due to the axial periodicity. The formation of Taylor rolls appears to be inevitable, and this happens by the merging of the boundary layer detachments- the hairpin vortices.

8.3.2 Mean velocity profiles

Figure 8.3(a,b) shows the mean azimuthal velocity profiles at the inner and outer cylinder, in wall units for all cases. Inner cylinder wall units are defined using $u_{\tau,i}$ as a velocity scale and $\delta_{v,i} = \nu/u_{\tau,i}$ as a length scale. The mean azimuthal profile is defined as a velocity difference, i.e. $U^+ = (U_i - \langle u_\theta \rangle_{t,z,\theta})/u_{\tau,i}$, where $\langle \phi \rangle_{x_i}$ denotes the variable ϕ averaged with respect to x_i . r^+ is the distance to the inner cylinder $r^+ = (r - r_i)/\delta_{v,i}$. Outer cylinder wall units are defined using $\delta_{v,o}$ and $u_{\tau,o}$, and $r^+ = (r_o - r)/\delta_{v,o}$. Inner and outer cylinder wall units use different scales. The

Figure 8.2: Visualization of the instantaneous azimuthal velocity u_θ for an azimuthal cut for (a,b): AF case at times separated by $\tilde{t} = 10$, (c) R2 case. The large scale pattern can be clearly seen all panels. The roll in (a) and (b) is slightly bent upwards due to the presence of an axial mean flow.

outer cylinder branch is always slightly above than the inner cylinder branch, because $u_{\tau,i} \approx \eta u_{\tau,o}$ must hold, and thus U^+ at the mid-gap is larger in outer cylinder wall units.

Very large deviations from the classic law-of-the-wall with von Kármán constants of $\kappa = 0.4$ and $B = 5.2$ are seen. This can be also appreciated in figure 8.3(c,d), which show the logarithmic diagnostic function

$$\Xi^+ = r^+ dU^+ / dr^+. \quad (8.1)$$

If U^+ is logarithmic, Ξ^+ should be horizontal and equal to the inverse of κ . For channels, deviations from the universal von Kármán law were proposed by Jiménez and Moser [196] based on the overlap arguments of Afzal and Yajnik [197] to have the following shape:

$$\Xi^+ = \kappa^{-1} + \alpha y/h + \beta Re_\tau^{-1/2}, \quad (8.2)$$

where h is the channel half-gap, and $\alpha = 1$ and $\beta = 150$ are obtained from fits to the data. Bernardini *et al.* [95] found the data up to $Re_\tau = 4000$ to fit well to this equation, with slightly modified constants. For PC, Pirozoli *et al.* [186] were not able to quantify the deviations of Ξ^+ in a systematic way.

In the case of TC, fitting an equation similar to equation (8.2) to our DNS data in the logarithmic region does not lead to a good description of the data, as the effect of

Figure 8.3: (a,b): Mean azimuthal (streamwise) velocity profile at the inner (a) and outer (b) cylinders in wall units. Dashed lines are $u^+ = r^+$ and $u^+ = 2.5 \log(r^+) + 5.2$. (c,d): Logarithmic diagnostic function Ξ^+ at the inner (c) and outer (d) cylinders. The dashed straight line corresponds to $\Xi^+ = 2.5$ on both panels. Colours are as in table 8.1.

Figure 8.4: (a): mean streamwise component for several canonical flows. (b): Ξ^+ for the same flows. Legend: R2 case (black line), ZPG-BL at $Re_\tau \approx 2000$ from [182] (red), channel from [198] at $Re_\tau \approx 2000$ (blue), pipe flow data at $Re_\tau \approx 1100$ from [183] (dark green) and PC data at $Re_\tau \approx 1000$ from [186] (yellowish-brown).

curvature and of the Taylor rolls is of extreme importance. A tilted S-like behaviour in Ξ^+ around $r^+ \approx 100$ can be appreciated for the R2 and R3 cases, which is very similar to the one seen in PC by Pirozoli *et al.* [186]. When comparing the cases, there is still a strong dependence on Re_τ of the inner layer. The peak in Ξ^+ at $r^+ \approx 10$ is far from universal across our four simulations, and is probably due to the presence of the large scale structures. Our DNS data for TC at $\eta = 0.909$ are clearly insufficient to conclude anything about the deviations of Ξ^+ .

A “bending” of the profiles away from the logarithmic law is seen in the bulk for all simulations. This is especially significant from $r^+ > 500$. For $\eta = 0.909$, this corresponds to one tenth of the gap, or only 1% curvature. Even if 1% curvature is small, it is consistent with the notion that curvature plays a surprisingly large effect, and thus we may take $r^+ \approx 0.1Re_\tau$ as an upper bound for the log-layer in TC for $\eta = 0.909$. This means that the effective Re_τ of the simulations decreases substantially, and the possible logarithmic regions extend less into the bulk. Higher Re_τ are needed in TC to see the same logarithmic profiles as in channels or pipes. This is shown in more detail in figure 8.4, which compares the streamwise mean profile and Ξ^+ in several canonical flows. TC shows very large variability in Ξ^+ when compared to the other canonical flows.

If one takes the start of the logarithmic layer to be approximately at $r^+ = 3Re_\tau^{1/2}$ [199], and curvature effects to be dominant at $r^+ \approx 0.1Re_\tau$, the range of validity of the assumptions made to obtain a logarithmic layer is reduced to $150 < r^+ < 400$, a clearly insufficient separation of scales to see a well-developed profile. Therefore, we cannot conclude anything about the mean azimuthal profile, even if we may speculate based on experimental results (cf. Ref. [82]) that further DNS at larger Re_τ will lead to the development of a region with properties consistent with those of other canonical

Figure 8.5: Mean fluctuation profiles for velocities and pressure at the inner cylinder in wall units. Colours are as in table 8.1.

flows.

8.3.3 Velocity and pressure fluctuations

Figure 8.5 shows the velocity and pressure fluctuations at the inner cylinder. The fluctuations, i.e. the root mean squared (rms) of a field ϕ is computed as $\phi' = [\langle \phi^2 \rangle_{\theta,t} - \langle \phi \rangle_{\theta,t}^2]^{1/2}$. The axial average and the subtraction operations must be computed in a precise order, due to the signature of the static large scale structures on the mean fields.

The velocity fluctuations are generally lower than those seen for channels and PC flow at comparable Re_τ . Except for the AF case, there is no presence of a second peak, and the level of fluctuations slightly increases with Re_τ , as expected. Velocity fluctuations are also shown in outer units in figure 8.6. A reasonable collapse can be seen in an “overlap” layer. This layer is defined between $\tilde{r} < 0.2$, i.e. $r^+ < 0.4Re_\tau$, and $r^+ > 100$. Here, u'_θ and u'_z show a small region which is consistent with logarithmic behaviour. This opens the door for the possibility of a region where overlap

Figure 8.6: Fluctuations of the azimuthal (a) and axial (b) velocity components for R1, R2 and R3 cases in outer units. An overlap region, where the fluctuations collapse can be seen in both panels. Colours are as in table 8.1

arguments are valid in TC. However, DNS with larger Re_τ is needed to sufficiently decouple the scales, and obtain a convincing overlap layer with logarithmic profiles in the fluctuations.

The addition of an axial flow greatly increases the level of fluctuations, especially those of the radial and axial velocities. A second peak in u'_θ for high r^+ can be seen to form, and the maximum values of u'_r and u'_z becomes much larger, especially in the bulk. These results are probably due to a combined effect of the axial flow and the large scale rolls, and can be discounted.

It seems that the Taylor rolls dampen fluctuations by fixing the position of the plume ejection regions. This is further quantified in figure 8.7 which compares the streamwise fluctuations for the R1 and AF cases with those of several canonical flows at around $Re_\tau \approx 1000$. Except for PC and the R1 case, a remarkable agreement in the value of u' at the peak of $r^+ \approx 12$ is obtained, including the data from the AF case. The R1 case has a lower amount of fluctuations than the rest of the cases, again due to the constraints imposed by the fixed Taylor rolls.

Finally, we note that clear logarithmic behaviour is also seen in the overlap layer region for p' in figure 8.5, and a collapse in the overlap region also happens between the R2 and R3 cases.

8.3.4 Axial and azimuthal autocorrelations

Figure 8.8 shows the autocorrelations for the azimuthal ($R_{\theta\theta}$) and radial (R_{rr}) velocities at the mid-gap. As mentioned previously, the azimuthal extent of the domain is sufficient, as it is large enough for the autocorrelation to change sign. This happens between 0.5 and 1 gap widths for all cases. This fast decorrelation is consistent with the finding that small computational boxes can obtain accurate results for the torque

Figure 8.7: Streamwise velocity fluctuations for several canonical flows. Legend: R1 case (blue line), AF case (dark green), ZPG-BL at $Re_\tau \approx 1300$ from Ref. [182] (purple), channel from Ref. [198] at $Re_\tau \approx 950$ (red), pipe flow data at $Re_\tau \approx 1100$ from Ref. [183] (black) and PC data at $Re_\tau \approx 1000$ from Ref. [186] (brownish yellow).

in TC, and is in contrast to PC, where azimuthal decorrelation takes place in a length scale which is one order of magnitude larger, which leads to the necessity of large computational boxes [184, 200]. The large scale Taylor roll has a clear signature on the axial autocorrelations $R_{\theta\theta}$ and especially R_{rr} . The presence of a roll can be inferred for all cases. Again we see that the axial pressure gradient does not weaken the roll, it only convects the whole structure.

8.3.5 Velocity power spectra

Taylor rolls contain a significant amount of energy, both in the bulk and in the boundary layer. Figure 8.9 shows the azimuthal and axial spectra near the wall, for $r^+ \approx 12$. The low-wavelength signature of the Taylor rolls is apparent in the axial spectra. This signature is not only present for $\Phi_{\theta\theta}$ and Φ_{zz} , but also for the radial velocity spectra Φ_{rr} for all cases. All cases show a maximum in the cospectra $\Phi_{\theta r}$ corresponding to $k_\theta = 0$, $k_z = 2\pi/\lambda_{TR}$, i.e. the wavelength of axisymmetric Taylor rolls. The rolls dominate the convective transport of angular velocity through the Reynolds stresses in the boundary layer. This means that Taylor rolls are not inactive in the sense of Townsend [9] and Hoyas & Jiménez [195], they actively transport angular velocity. Unlike the large scale structures in channel flow and PC flow, the rolls may be considered to be “attached” to the wall.

It might seem strange that such a large wavelength appears with the radial (wall normal) velocity near the wall, considering the impermeability condition. In channels and PC, the large scale structures reflect on the streamwise velocity near the wall.

Figure 8.8: Autocorrelation functions for the azimuthal and radial velocities in the azimuthal and axial directions at the mid-plane $\tilde{r} = \tilde{r}_a$. The azimuthal extent of the computational boxes can be seen to be enough. The signature of the Taylor rolls can be seen on the axial autocorrelations for all cases. Colours are as in table 8.1

Figure 8.9: Azimuthal (left) and axial (right) spectra for all three velocity components for all simulations at $r^+ \approx 13$. Colours are as in table 8.1.

In TC, a perturbation of locally larger azimuthal (streamwise) velocity is linearly unstable, and will gain radial velocity. Therefore, one might think of the rolls as causing a large wavelength pattern in the azimuthal velocity near the walls, which causes a locally larger formation of hairpin vortices which have radial velocity, and this reflect back on the spectra. This is due to the fundamental difference between PC and TC, the linear instability.

As they form, the hairpin vortices, just like thermal plumes in thermal convection [25] tend to attract each other, and to merge, forming larger structures. The merging of the vortices is very apparent if the system is started from zero velocity initial conditions. The merging of the hairpin vortices forms the Taylor rolls in the bulk. Rolls dominate the spectra at the mid-gap, shown in figure 8.10. A prominent saw-tooth behaviour in the axial spectra can be seen for the lowest wavenumbers. The spectra are consistent with the experiments by Lewis & Swinney [41], and neither show -1 nor $-5/3$ scaling. This is unlike the case of pipe flow [10], for which a -1 scaling was found, and also unlike the case of curved channel flow, which show the expected -1 scaling laws for the streamwise energy spectra in the streamwise direction [189]. Also, the axial velocity spectra in the azimuthal direction do not show a $-5/3$ exponent which could be expected from Ref. [10].

As noticed already by Brauckmann & Eckhardt [123], the large-scale structures are responsible in TC for the transport of angular velocity, while the fluctuations transport on average very little, even though their instantaneous transport can be orders of magnitude higher than the mean transport [45]. Unlike in channel flow, and curved channel flow, the Reynolds stresses in TC are maximal at the mid-gap to satisfy the conservation of angular velocity current (torque). In PC, large scale structures also form, but these are also inactive, i.e. they do not transport shear stresses [185]. Therefore, it seems that small gap TC flow is the only flow examined up to date which involves transport by Reynolds stresses of a conserved quantity, in this case angular velocity, by large scale structures which are attached to the wall.

8.4 Summary and conclusions

A series of large DNS simulations of small gap Taylor-Couette flow were conducted, reaching frictional Reynolds numbers of $Re_\tau \approx 4000$. Large scale structures, known as Taylor rolls form, and to play an active role in the system dynamics. Taylor rolls are present the bulk, and also are “attached” to the cylinders. They play an active role, by transporting angular velocity through Reynolds stresses. Even deep inside in the boundary layers, for $r^+ \approx 12$, their signature is significant in the $\phi_{\theta r}$ cospecra. Adding an axial pressure gradient was found to convect them slowly, but not to weaken them significantly. Their signature was still present in axial spectra in the boundary layer, and in the axial correlations of radial velocity.

Figure 8.10: Azimuthal and axial spectra for all three velocity components for all simulations at mid-gap. Colours are as in table 8.1.

The mean velocity profile of TC has significant differences from the Prandtl-von Kármán log-layer profile $u^+ = \kappa^{-1} \ln y^+ + B$ with $\kappa \approx 0.4$ and $B \approx 5.2$. Large deviations for the logarithmic law are found even for 1% curvature, consistent with the notion that curvature effects are orders of magnitude larger than expected by mere dimensional analysis [188]. Fluctuation profiles were also analysed, and the presence of rolls was found to reduce the level of fluctuations. When adding a mean axial flow, the azimuthal velocity fluctuations presented coincide with those seen in other canonical flows at comparable Re_τ . The velocity fluctuations were found to collapse in outer units, indicating the presence of an overlap region. Based on these findings, we expect that simulations at higher Re_τ will provide enough decoupling of scales for a large enough overlap region to form, and logarithmic profiles to form in both the mean azimuthal velocity, and the axial and azimuthal fluctuations.

Further investigations should also look at the effect of strongly counter-rotating cylinders, which have a Rayleigh-stable region, and how this region affects the large scale structures. Furthermore, in Ostilla-Mónico *et al.* [129], it was found that large gaps, i.e. small η , severely weakens the rolls. Additional simulations at η around 0.5 provide more data on how the large scales interact with the boundary layers.

Effects of the computational domain size on the numerical simulations *

De omnibus dubitandum est: Everything must be doubted

—RENÉ DESCARTES

In search for the cheapest but still reliable numerical simulation, a systematic study on the effect of the computational domain (“box”) size on direct numerical simulations of Taylor-Couette flow was performed. Four boxes, with varying azimuthal and axial extents were used. The radius ratio between the inner cylinder and the outer cylinder was fixed to $\eta = r_i/r_o = 0.909$. The outer cylinder was kept stationary, while the inner rotated at a Reynolds number $Re_i = 10^5$. Profiles of mean and fluctuation velocities are compared, as well as autocorrelations and velocity spectra. The smallest box is found to accurately reproduce the torque and mean azimuthal velocity profiles of larger boxes, while having smaller values of the fluctuations than the larger boxes. The axial extent of the box directly reflects on the Taylor-rolls and plays a crucial role on the correlations and spectra. The azimuthal extent is found to play a minor role in the simulations, as the boxes are large enough. For all boxes studied, the spectra does not reach a box independent maximum.

*Based on: R. Ostilla-Mónico, R. Verzicco and D. Lohse, *Effects of the computational domain size on DNS of Taylor-Couette turbulence*, under review (arXiv:1411.3826).

9.1 Introduction

Taylor-Couette (TC), the flow between two coaxial and independently rotating cylinders, is commonly used as a basic model and paradigmatic system for shear flows in very diverse topics, for example probing the stability of astrophysical flows [85] or more direct applications such as bubbly drag reduction [191, 193]. TC is very accessible experimentally, as it is a closed system, it has a high number of symmetries and a simple geometry. Experiments of TC have been conducted up to $Re \sim \mathcal{O}(10^6)$ [39, 40, 42, 43]. These large Re allow for the study of the “ultimate” regime, in which the flow is fully turbulent both in the boundary layers and in the bulk. It is expected that the scaling laws which hold in this regime can be extrapolated to arbitrarily large Reynolds numbers, such as those present in geo- and astro-physics [74, 78].

The flow is sheared by the angular velocity difference between the two cylinders. The driving of the cylinders can be expressed non-dimensionally with two Reynolds numbers: $Re_i = r_i d \omega_i / \nu$ for the inner cylinder and $Re_o = r_o d \omega_o / \nu$ for the outer cylinder, where r_i and r_o are the inner and outer cylinder radius, ω_i and ω_o the inner and outer cylinder angular velocity, d is the gap width, $d = r_o - r_i$ and ν the kinematic viscosity of the fluid. The shear driving of the flow can then be expressed as a shear Reynolds number $Re_s = r_i |\omega_i - \omega_o| d / \nu = |Re_i - \eta Re_o|$, where η is the radius ratio $\eta = r_i / r_o$.

Direct numerical simulations (DNS) of TC have received increasing level of sophistication in the last years. The first attempt at high Re was done using “large” computational boxes by Dong [48, 49]. Dong used a periodic aspect ratio of $\Gamma = 2\pi$, where $\Gamma = L/d$, with L the axial periodicity length, for $\eta = 0.5$, reaching $Re_s = 8000$. Ostilla-Mónico *et al.* [112, 176] also achieved $Re_s \approx 8000$ for $\eta = 0.714$ using large boxes with $\Gamma = 2\pi$. A breakthrough was achieved by Brauckmann & Eckhardt [123, 156], who showed that simulation boxes could be heavily reduced in two ways, while still obtaining accurate data for the torque. For simulation boxes with $\Gamma = 2\pi$, three Taylor roll pairs fits in the system. However, only one pair of rolls was sufficient to calculate the torque, so Γ could be reduced to $\Gamma = 2$ ($\approx 2\pi/3$). Secondly, simulating the full azimuthal extent of the cylinder is also not necessary to obtain an accurate result for the torque. A cylindrical wedge, with a rotational symmetry can be imposed, and, for $\eta = 0.714$, only a ninth [123] of the cylinder was necessary. The use of these “small” boxes reduces the computational requirements by a factor thirty or more, and made later DNS deep inside the ultimate regime by Ostilla-Mónico *et al.* [113, 129] possible, who achieved $Re_s \sim \mathcal{O}(10^5)$. Recently, Shi *et al.* [201] also achieved $Re_s \sim \mathcal{O}(10^5)$ in TC in the quasi-Keplerian regime using reduced boxes.

The torque at the inner and outer cylinders are a first order, integral statistic. The finiteness of the computational domain however may play a role for other statistics, both higher order statistics, or one-, two- and many-point statistics. This is the case for example in the channel flow simulations of Lozano-Durán *et al.* [96], where even

if the stream-wise mean velocity profile is well reproduced with small boxes, accurate results for velocity and pressure fluctuations (root mean squared) require larger boxes.

Before we continue, it is worth noting that there are two main difference between channels (and pipes) and TC. In TC, a natural constraint on the azimuthal (stream-wise) extent of the domain always exists, i.e. the full cylinder, while one could think of channels and pipes extending infinitely. Second, the axial (spanwise) periodicity length in TC fixes the size of the Taylor rolls. The effect of the box-size in the axial direction is not purely numerical. The use of axially periodic conditions in axially “small” computational domains, where only one pair of rolls fit in means that the axial periodicity length sets the wavelength of the Taylor vortices. This wavelength is directly analogous to the wavelength of the Taylor vortices seen in experiments. In large aspect ratio experiments, the physical parameter is the axial length which is given by distance between the end-plates. No preferential wavelength is set for the rolls, but experimental evidence [173, 174] suggests that vortex pairs tend to have wavelengths λ_{TR} between $2 \leq \lambda_{TR}/d \leq 4$. These rolls survive up to very large Reynolds numbers, having been observed experimentally up to $Re_s \sim \mathcal{O}(10^6)$ by Huisman *et al.* [173], and, in the corresponding parameter regimes all DNS simulations up to now. Taylor rolls have no direct analog in pipes and channel flow between two parallel plates, and can still play a large role in the DNS for various physical quantities, even though the torque has been shown to become independent of Γ in the range $2 \leq \Gamma \leq 4$ at about $Re_s \sim 3 \cdot 10^4$ [129].

In this chapter we attempt to answer the question: how does the size of the computational domain affect other statistics of TC DNSs and in particular higher order moments? To do so, we performed a series of DNSs of TC using a second-order finite difference code, with fractional time-stepping detailed in Ref. [103]. This code has been used for all our previous DNS of TC, and has been extensively validated against experiments [112, 113, 129].

The radius ratio was fixed to $\eta = r_i/r_o = 0.909$, the inner cylinder was rotated at $Re_i = 10^5$, while the outer cylinder was kept stationary, i.e. $Re_o = 0$. This resulted in a total shear driving of $Re_s = 10^5$. With the chosen parameters, the simulations are in the fully turbulent (ultimate) regime, and still have a strong large-scale axial circulation. A small-gap geometry was chosen, as it has the strongest large-scale rolls [129]. Therefore, we do not expect the effect of η to be significant on the results. By fixing the outer cylinder, we ensure that the flow is fully Rayleigh unstable, i.e. $d|\omega r^2|/dr < 0$ everywhere [84]. TC flow in the Rayleigh-stable regime may present other kinds of turbulent structures, Taylor columns instead of Taylor vortices [202], and this limits the applicability of the present work.

For the parameters mentioned previously, four simulations were conducted with computational boxes of varying sizes. Details of the geometry and resolutions used

are available in table 9.1. The adequacy of the mesh can be further checked in the spectra shown in later sections. A variable time-step with a maximum CFL of 1.4 was used, which does not reduce the accuracy of the results [112]. Simulations were started from zero-velocity initial conditions on a coarse grid, at low Reynolds numbers. Once these simulations had reached a statistically stationary state, the Reynolds number was increased, and the grid was refined, until the desired Reynolds number and grid refinement was reached. After a sufficiently long time to let transient behaviour die out, simulations were run between 30 large eddy turnover times (defined as $d/(r_i\omega_i)$) in the case of the $\Gamma 4N10$ run, or up to 70 for the $\Gamma 2N20$ run. This run-time is smaller than time required to overcome the transient, which is between 100 to 150 time units. For a more detailed analysis of the time-scales in TC flow, we refer the reader to Ostilla-Mónico *et al.* [130] (§8).

Throughout this chapter, the following conventions will be used: $\langle \phi \rangle_{x_i}$ denotes the average of a quantity ϕ with respect to the independent variable x_i . The torque T is non-dimensionalized as a pseudo-Nusselt number [25] $Nu_\omega = T/T_{pa}$ where T_{pa} is the torque in the purely azimuthal and laminar state. We also define $\tilde{r}$, the normalized radius as $\tilde{r} = (r - r_i)/d$, the normalized height as $\tilde{z} = z/d$, and the normalized azimuthal distance at the mid-gap as $\tilde{x} = (r_o + r_i)\theta/(2d)$.

Normalizations with respect to “wall” variables are denoted with a plus superscript, i.e. ϕ^+ . Wall variables are first averaged azimuthally, axially and temporally. Then, the frictional velocity at the corresponding cylinder $u_\tau = \sqrt{\tau_w/\rho}$ is computed, and used as velocity scale, where τ_w is the mean friction at the corresponding cylinder, and ρ is the fluid density. As length scale to non-dimensionalize the viscous length $\delta_v = \nu/u_\tau$ is used as usual. In these wall variables, we denote the distance to the cylinder(s) with r^+ . For the inner cylinder wall variables, this is defined as $r^+ = (r - r_i)/\delta_{v,i}$, while for the outer cylinder wall variables $r^+ = (r_o - r)/\delta_{v,o}$.

9.2 Results and analysis

Table 9.1 also shows that Nu_ω is the same within the statistical temporal error due to the necessarily limited time averaging for all simulations. This finding, even if expected from previous research, is still remarkable, considering the large-scale flow patterns, i.e. the Taylor rolls, which are still present in the flow. These patterns can be appreciated from figure 9.1. Structures are emitted from the boundary layers. These can be thought of as hairpin vortices, or as plumes when speaking in the language common for thermal convection. Plumes tend to attract each other, and merge together forming regions with large angular velocity transport, i.e. with a strong correlation between u_θ and u_r [130] (cf. also §8). These regions have a very large positive or negative angular velocity transport, which can be orders of magnitude larger than the mean [113]. “Neutral”-transport regions, in the core of the Taylor rolls, lie be-

Case	$n_{sym} / \tilde{x}$	Γ	$N_\theta \times N_r \times N_z$	Nu_ω	Line Style
$\Gamma 2N20$	20 / 3.3	2.09	$1024 \times 1024 \times 2048$	69.5 ± 0.2	Solid light blue
$\Gamma 2N10$	10 / 6.6	2.09	$2048 \times 1024 \times 2048$	69.4 ± 0.4	Solid black
$\Gamma 3N20$	20 / 3.3	3.00	$1024 \times 1024 \times 3072$	69.6 ± 0.2	Dashed dark green
$\Gamma 4N10$	10 / 6.6	4.00	$2048 \times 1024 \times 4096$	69.8 ± 1.6	Dash-dot dark red

Table 9.1: Details of the numerical simulations. The first column is the name with which the simulation will be refereed to in the chapter. The second column shows n_{sym} , the order of the rotational symmetry imposed on the system and the resulting azimuthal extent of the domain at the mid-gap, $\tilde{x} = 2\pi(r_i + r_o)/(2dn_{sym})$. The third column gives Γ , the axial periodicity aspect ratio. The fourth column represents the amount of points in the azimuthal, radial and axial directions used for the simulations. For the uniformly discretized directions, the resolutions used correspond to approximately $r_o\Delta\theta^+ \approx 9$ and $\Delta z^+ \approx 3$ in inner cylinder wall units. For the radial direction, points are clustered near the wall using a clipped Chebychev arrangement, with the resolution varying from $\Delta r^+ \approx 0.3$ near the cylinders to $\Delta r^+ \approx 4.2$ at the mid-gap. The fifth column shows the non-dimensional torque Nu_ω . The last column indicates the line shapes used for Figs. 9.4-9.8. The frictional Reynolds number at the inner cylinder $Re_{\tau,i} = u_{\tau,i}d/(2\nu)$, is $Re_{\tau,i} = 1410$ for all cases. The outer cylinder friction number, $Re_{\tau,o}$, is $Re_{\tau,o} \approx \eta Re_{\tau,i} = 1280$.

tween them, and in these regions u_r and u_z are small. For a more detailed analysis of the dynamics of these regions, see Refs. [113, 130].

Taylor-rolls are stationary in time. This is displayed by figure 9.2, which shows a pseudocolour plot of the azimuthally- and temporally- averaged azimuthal velocity $\langle u_\theta \rangle_{\theta,t}$. For all panels, a single vortex pair is present, which fills up the whole computational domain. This large-scale structure has little to no effect on the total angular velocity transport. The simulations were ran for more than 30 large eddy turnover times (defined as $\tilde{t} = d/r_i\omega_i$), and this did not significantly modify the position of the roll. Remarkably, for the $\Gamma 4N10$ case, only one roll with wavelength $\lambda_{TR} = 4$ fills the domain, instead of two “square” rolls with $\lambda_{TR} = 2$. This is consistent with the findings that the preferred λ_{TR} increases with Re , and that large Re simulations and experiments tend to find rectangular Taylor rolls with $\lambda_{TR} > 2$, which however cannot be sustained at lower drivings [129] (cf. the experiments of Huisman *et al.* [173] at $Re \sim \mathcal{O}(10^6)$ who found Taylor rolls with $\lambda_{TR} > 3$).

Long wavelength patterns, present in the axial direction, can also form in the azimuthal direction. Figure 9.3 shows a pseudocolour plot of the instantaneous velocity u_θ at the mid-gap $\tilde{r} = 0.5$ for the $\Gamma 2N20$ and $\Gamma 2N10$ cases. The axial signature of the Taylor rolls can be clearly appreciated in the panels. On top of this signature, additional azimuthal patterns can be seen. On both panels of figure 9.3 structures similar to wavy Taylor vortices can be seen, which we will refer to as “wavyness” of the roll. We note that unlike Taylor rolls, long wavelength azimuthal structures are

Figure 9.1: Pseudocolor of the instantaneous velocity in the (a) azimuthal, (b) radial and (c) axial direction for an azimuthal cut of $\Gamma 2N20$. The presence of large scale rolls, which make the flow inhomogeneous, is apparent here.

Figure 9.2: Azimuthally- and time- averaged azimuthal velocity $\bar{u}_\theta$ for the cases (a) $\Gamma 2N20$, (b) $\Gamma 3N20$ and (c) $\Gamma 4N10$. As Γ increases, the Taylor roll pair also grows, filling the entire box. In these panels (and also in Figs. 9.1 & 9.3) there is a slight preference for blue (low velocity) regions in the plots, because the mean azimuthal velocity in the bulk is slightly below 0.5 due to the inherent asymmetry between both cylinders.

not stationary in time, as they are convected with the mean flow velocity and do not show up on temporal averages.

We will now quantify the effect of these patterns and thus of the computational box size on the flow. We start with one-point statistics, in particular with the mean azimuthal velocity. Figure 9.4(a) shows the azimuthally-, temporally- and axially- averaged azimuthal velocity $\langle u_\theta \rangle_z$ at the inner cylinder in inner cylinder wall units, while figure 9.4(b) shows the so-called diagnostic function $\Xi^+ = du^+/d(\log r^+)$, which is the local slope of a lin-log plot of the profile.

Only little variation between all cases can be seen between these panels. In the right panel, a slight decrease of the intercept of the logarithmic profile is observed for the $\Gamma 3N20$ and $\Gamma 4N10$ simulations. This is caused by the remnant axial dependence of u_θ , which increases with increasing Γ . On the other hand, there are no appreciable differences in the boundary layers between the $\Gamma 2N20$ and $\Gamma 2N10$ simulations indicating that the azimuthal extent of the box is enough to capture the mean profiles in the boundary layer.

Figure 9.5 shows the velocity fluctuation profiles at the inner cylinder. The root mean square (rms) of a field ϕ is computed as $\phi' = \langle \langle \phi^2 \rangle_{\theta,t} - \langle \phi \rangle_{\theta,t}^2 \rangle_z$. Note that the order of the axial average and subtraction operations is crucial, due to the remnant and significant axial dependence in the mean velocity fields, (cf. figure 9.2). Axially

Figure 9.3: Instantaneous azimuthal velocity u_θ at the mid-gap $\tilde{r}_a$ for the cases (a) $\Gamma 2N20$ and (b) $\Gamma 2N10$, as a function of the azimuthal variable θ . Increasing n_{sym} allows larger wavelength structures to fit in the simulation. These structures azimuthally modulate the Taylor rolls, and allow for more and stronger fluctuations in the Taylor roll cores.

Figure 9.4: (a): Axially-, azimuthally- and time-averaged azimuthal velocity in inner cylinder wall units. The dashed lines indicate $u^+ = r^+$ and $u^+ = 2.5 \log r^+ + 5.2$. (b): Diagnostic function $\Xi^+ = du^+/d(\log r^+)$. Very little to no dependence on the box-size can be seen near the boundary layers, while differences can be appreciated from $r^+ > 500$, i.e. in the bulk. In the bulk, u_θ , has a strong axial dependence, and this is probably causing the discrepancies. For both panels, symbols are as in table 9.1.

averaging before subtracting results in rms values which are considerably higher, but originate simply from the Taylor rolls and have nothing to do with the underlying statistics.

While the box-size can be seen to play a small effect on the u_θ^+ profile, it is critical for other averages else. In general, increased box sizes lead to increased fluctuations, in line with what is seen in channels for small boxes [96]. The effect of increasing the azimuthal size of the domain on the fluctuations can be appreciated by comparing the $\Gamma 2N10$ and the $\Gamma 2N20$ cases. There is a clear increase in the u_r^+ and u_z^+ fluctuations in the bulk. This is a direct reflection of the increased mobility of the rolls. The axial extent of the domain can also be seen to affect the velocity fluctuation profiles. A larger axial domain, i.e. increasing Γ , again leads to larger fluctuations in the boundary layer, due to the increased mobility of the plumes, and the increased axial dependence of the velocity.

Finally, the pressure fluctuations become smaller for larger domains. This is probably due to the pressure playing a damping role on the velocity fluctuations in the smaller domains. Again, this is in line with what is seen in channels [96]. The geometric (computational) constraint on the flow inside the boundary layers is enforced through these pressure fluctuations. The sharp increase of velocity fluctuations in the bulk seen for the $\Gamma 2N10$ case also results in a sharp increase of pressure fluctuations in the bulk.

Figure 9.6 shows the two-point autocorrelation function for all velocity fields in both axial and azimuthal direction. Two main effects of the box size can be seen on the flow. The azimuthal extent of the box plays a negligible role in the decorrelation

Figure 9.5: Velocity and pressure fluctuations near the inner cylinder in inner cylinder wall units for all simulations. Symbols are as in table 9.1.

in all panels. If one compares the $\Gamma 2N10$ case to the $\Gamma 2N20$ cases, both show a region of anti-correlation at around $\tilde{r}\theta \approx 1.6$. This suggests that $n_{sym} = 20$ is large enough to capture the largest structures in the azimuthal extent, with wavelength equal to the small-box size $\tilde{r}\theta \approx 3.3$. This can be confirmed by visual inspection of figure 9.3, where the right panel seems to consist of two copies of the left panel. We also note that the azimuthal decorrelation lengths are an order of magnitude smaller than those seen in plane Couette flow [184]; which have allowed TC simulations to reach higher Re_τ with heavily reduced computational costs. Current state-of-the art plane Couette simulations “only” reach $Re_\tau = 550$ while requiring 2.2 billion points [185].

On the other hand, the axial extent plays an important role in both axial and azimuthal correlations. As expected, the axial autocorrelations are dominated by the effect of Taylor rolls. This is especially true in the case of the radial velocity autocorrelation R_{rr} . The axial velocity autocorrelation R_{zz} remains relatively unaffected, as axial velocities in the mid-gap are very small (cf. Fig. 9.3). Additionally, increasing Γ allows for a faster drop of $R_{\theta\theta}$ in the azimuthal direction, but not for faster drops in R_{rr} and R_{zz} . This is due to the larger cores of the Taylor rolls, which result in more mixing. In these regions, the radial and axial velocities are small, so R_{rr} and R_{zz} are dominated by the strongly correlated regions seen in figure 9.1.

We now turn to the velocity spectra. Figure 9.7 shows the premultiplied velocity spectra in the inner cylinder boundary layer ($r^+ \approx 12$), while figure 9.8 shows the spectra at the mid-gap ($\tilde{r} = 0.5$). In both figures, the size of the computational box can be seen to play a negligible role for the spectra at the small scales, while, as seen in Ref. [130] the large scales contain a very significant amount of energy both deep inside the boundary layer and in the bulk.

Inside the boundary layer there are two main energy containing scales- that of the plumes at high k , and that of the Taylor rolls, at low k . In the mid-gap, the plumes have merged with each other, and the main energy content can be found only in the large scales. It is apparent from both figures, that the maximum of the spectra is not converged to their very large box-size value. The box is not large enough to contain all large scales which are energetic. Even then, good collapse for the small scales can be seen. The largest scales contain energy from all three velocity components, and not only for the azimuthal and spanwise components, which is the case in channels [195] and plane Couette [185]. All mid-gap spectra neither display clearly the inertial range Kolmogorov scaling with $-5/3$ scaling, nor the -1 scaling for the $E_{\theta\theta}$ in the θ -direction predicted by Perry & Chong [10]. This is consistent with the experimental findings of Lewis & Swinney [41] and those of Huisman *et al.* [81].

Saw-tooth patterns, indicating preferred even or odd modes can be seen for the axial spectra for all cases. Increasing Γ shifts the maxima in k_z accordingly, to accommodate for the different size of the Taylor-roll. We note that inside the boundary layer, simulations with larger Γ have a smaller energy peak at the Taylor-roll wave-

Figure 9.6: Two-point autocorrelation functions at the mid-gap ($\tilde{r} = 0.5$) for all simulations. Panels on the left column are autocorrelations in the azimuthal direction (θ), while those on the right are autocorrelations in the axial directions. The three rows are for each velocity component: azimuthal (top), radial (middle) and axial (bottom). Symbols are as in table 9.1.

length k_{TR} . This means that the Taylor-roll has a larger effect on the plume generation and its aspect ratio is smaller. Using a small domain may artificially strengthen the roll and lead to an increased correlation inside the boundary layer. Further increases in Γ will accommodate for more rolls. This was seen to produce a very sharp dropoff in the spectra for $k_z < k_{TR}$ by Dong [48]. The effect of end-plates has also been omitted, and this will require attention in the future.

Remarkably, sawtooth behaviour at low frequencies is also present in the azimuthal direction if the azimuthal extent of the domain is increased. A preference for even modes over odd modes is indicated, and this again suggests the fundamental wavelength of the system is $\tilde{r}\theta \approx 3.3$, i.e. half the domain. This is further evidence that the statistics for the $\Gamma 2N10$ case can be accurately reproduced with the $\Gamma 2N20$ case.

However, the spectra are not saturated, i.e. a clear maximum is not seen. This means that all the energy containing scales are not captured, and thus extending the azimuthal extent of the simulation seems to be a necessity to fully capture these scales, which do not affect the flow significantly, as $n_{sym} = 20$ boxes are able to reproduce the main features. It could also be the case that the spectra will not saturate even if $n_{sym} = 1$, and thus all large scales would be energy containing due to the natural finiteness of the azimuthal extent. This point seems to require further study.

9.3 Summary and conclusions

In summary, a systematic study of the effect of the computational box size on TC DNS was performed. From previous studies [123], it was already known that small boxes can obtain accurate results for the non-dimensional torque. Furthermore, similar to what was found by Lozano-Durán *et al.* [96] for DNS of channel flow, small boxes also have accurate mean azimuthal (streamwise) velocity profiles in the boundary layers. Larger boxes are needed in order to obtain box-independent results for fluctuation values, two-point autocorrelations and low-wavelength spectra. The artificial truncation of the spectra by using a reduced box does not bring about significant changes in its structure at low wavelengths- even if the most energetic scales are not accounted for. Azimuthally small boxes show a reduced level of fluctuations, even if the fundamental wavelengths are already captured.

In the axial direction, things are different. The size of the underlying Taylor roll dominates the autocorrelations, especially for the radial velocity. As we mentioned previously, the effect of a the computational box-size in the axial direction is not purely numerical, the wavelength of the Taylor vortex is a physical parameter. In experimental and natural realizations of TC, the wavelength of the Taylor rolls is determined by the axial constraints, i.e. end plates or periodicity. It appears from recent work [129, 174] that for larger Γ domains, which can accommodate more than

Figure 9.7: Premultiplied azimuthal (left column) and axial (right column) spectra for the azimuthal (top), radial (middle) and axial (bottom) velocities at $r_i^+ \approx 12$. Symbols are as in table 9.1.

Figure 9.8: Premultiplied azimuthal (left column) and axial (right column) spectra for the azimuthal (top), radial (middle) and axial (bottom) velocities at the mid-gap $\tilde{r} = 0.5$. Symbols are as in table 9.1. No significant effect of the box size is seen for the small wavelength (large k) scales, while for the large scales, the azimuthal size of the domain appears to play a role in the axial structure of the rolls. The “wavy” modulations seen in figure 9.3 are clearly reflected in the spectra.

a pair of rolls, the preferred wavelength of Taylor rolls increases with increasing Re . This effect cannot be captured in small box simulations, but can lead to bifurcations at high Reynolds numbers [173]. Future work should consider simulating large $\Gamma \sim \mathcal{O}(10)$ to check the window of coexistence of different states at high Re , and how these affect the flow.

In the near future, a study of larger TC boxes seems mandatory, as to determine the minimal box size for accurate statistics of the velocity and pressure fluctuations and higher order moments. It can be the case that the axial extent is too small and thus non-physical, and that the azimuthal extent is not to capture all wavelengths. It is also unclear where the increased level of fluctuations of the $\Gamma 2N10$ case comes from. These increased fluctuations may provide a way for the system to overcome the energy barrier, and to switch between vortical states, i.e. from two vortex pairs with $\lambda_{TR} = 2$ to one vortex pair with $\lambda_{TR} = 4$. Another reason for doubting their physicality comes from the correlations, as the $\Gamma = 2$ cases show unusually large decorrelation lengths in the azimuthal direction, and this might cause the formation of said patterns.

We also point out that the size of this minimal box is larger than those required for channels. The largest box in this manuscript, for which it is not clear yet whether its statistics are box-independent, has relative dimensions of 8 half-gap lengths in the axial (spanwise) direction and 4.2π in the azimuthal (streamwise) direction, while accurate statistics were obtained for a box of size π half-gaps in the spanwise direction and 2π half-gaps in the streamwise direction in Lozano-Durán *et al.* [96].

Turbulence decay towards the linearly-stable regime *

If we are uncritical we shall always find what we want: we shall look for, and find, confirmations, and we shall look away from, and not see, whatever might be dangerous to our pet theories.

—KARL POPPER

*Taylor-Couette (TC) flow is used to probe the hydrodynamical stability of astrophysical accretion disks. Experimental data on the subcritical stability of TC are in conflict about the existence of turbulence (cf. Ji et al. *Nature*, **444**, 343-346 (2006) and Paoletti et al. , *Astron. & Astroph.*, **547**, A64 (2012)), with discrepancies attributed to end-plate effects. In this chapter we numerically simulate TC flow with axially periodic boundary conditions to explore the existence of sub-critical transitions to turbulence when no end-plates are present. We start the simulations with a fully turbulent state in the unstable regime and enter the linearly stable regime by suddenly starting a (stabilizing) outer cylinder rotation. The shear Reynolds number of the turbulent initial state is up to $Re_s \lesssim 10^5$ and the radius ratio is $\eta = 0.714$. The stabilization causes the system to behave as a damped oscillator and correspondingly the turbulence decays. The evolution of the torque and turbulent kinetic energy is analysed and the periodicity and damping of the oscillations are quantified and explained as a function of shear Reynolds number. Though the initially turbulent flow state decays, surprisingly, the system is found to absorb energy during this decay.*

*Based on: R. Ostilla-Mónico, R. Verzicco, S. Grossmann and D. Lohse, *Turbulence decay towards the linearly-stable regime of Taylor-Couette flow*, *J. Fluid Mech.*, **768**, R3 (2014).

10.1 Introduction

Quasars are the most luminous objects in the Universe. They are thought to consist of super-massive black holes in the centres of galaxies, which accrete matter and emit radiation, thereby transforming mass into energy with an efficiency between 5-40% [85]. In order for orbiting material to move radially inwards into the central object in the so-called accretion disk, it must lose its angular momentum. Molecular viscosity alone is not enough to account for this loss, so some sort of turbulent viscosity, causing the enhanced transport of angular momentum, has been conjectured [203], as otherwise gravitationally bound objects, such as quasars, stars or planets, would not exist. This implies that the flow of the material must be turbulent, but the origin of turbulence in some types of accretion disks is currently disputed. Accretion disks can either be “hot”, and therefore highly ionized and electrically conducting, or “cold”, and therefore extremely poorly ionized and poorly electrically conducting. Hot disks are found around quasars and active galactic nuclei, and thought to undergo a magneto-rotational instability (MRI) to become turbulent [204, 205]. In contrast, cold disks are found in protoplanetary systems, and are thought to have “MRI-dead” regions where turbulence due the MRI cannot exist as the magnetic Reynolds number is too low [206, 207]. The places where planets eventually form and reside coincide with these MRI-dead regions, so additional transport of angular momentum must take place. To account for this transport another mechanism to generate turbulence has been proposed: hydrodynamical (HD) non-linear instabilities.

Taylor-Couette flow (TC) is the flow between two coaxial cylinders which rotate independently. It is used as a model for probing the HD stability of accretion disks [85]. Accretion disks have velocity profiles that are linearly stable, but just as in pipe and channel flows, the shear Reynolds numbers are so large that non-linear instabilities may play a role in the formation of turbulence. The question whether TC flow does indeed undergo a non-linear transition [86, 87] to turbulence at large shear Reynolds numbers has been probed experimentally by several authors with conflicting results [43, 44, 77, 88–93]. The discrepancies between the different experiments have been attributed to the end plates of the TC devices, i.e. the solid boundaries which axially confine TC flow. These cause secondary flows, known as Ekman circulation, which propagate into the flow and influence global stability. Avila [93], based on his numerical simulations with end plates imitating the Ji *et al.* [90] and the Paoletti *et al.* [92] experiments, concluded that both of these experiments were inadequate to study the non-linear transition due to the presence of end plates. However, these simulations are at shear Reynolds numbers of the order of $Re_s \sim 10^4$, while Ji *et al.* [90] mention the counter-intuitive result that Re_s must exceed a certain threshold before the flow relaminarizes.

Indeed, confinement of the flow plays a very important role in this transition, but it is unavoidable in experiments. Direct numerical simulations (DNS) do not

have these limitations, as they can be run with periodic boundary conditions to completely prevent end-plate effects. Following this idea, Lesur & Longaretti [208] simulated rotating plane Couette flow with perturbed initial conditions, going up to shear Reynolds numbers of order of 10^4 ; however, due to computational limitations, the quasi-Keplerian regime velocity (i.e., the regime where the two velocity boundary conditions are related by Kepler's third law) was not reached.

In this chapter, we adopt a similar approach, but now for Taylor-Couette flow and, in particular, reaching larger shear Reynolds numbers $Re_s \sim 10^5$. We proceed as follows: We start from a turbulent flow field, corresponding to a pure inner cylinder rotation in the laboratory frame, and then switch on stabilizing outer cylinder rotation, wondering whether the turbulence is sustained. We find that it is not: the turbulence decays, and the torque decreases down to a value corresponding to purely azimuthal, laminar flow. This means that not only the TC system is linearly stable in this geometry, but even an initially turbulent flow decays towards the linearly stable regime.

10.2 Numerical details

The DNS were performed using a second-order finite difference code with fractional time-stepping [103]. This code has already been validated and used extensively in the context of TC flow [112, 113]. The turbulent initial conditions are taken from Ref. [113] (cf. §6). The radius ratio is $\eta = r_i/r_o = 0.714$, where r_i and r_o are the inner and outer radii, respectively, and the spatial period aspect ratio is $\Gamma = L/d = 2.094$, where L is the axial periodicity length and d the gap width $d = r_o - r_i$. A rotational symmetry of order 6 was forced on the system to reduce computational costs while not affecting the results [113, 123]. With this rotational symmetry, the system has an azimuthal periodic length of $L_\theta/d = 3.141$ at the mid-gap.

The simulations are performed in a frame of reference co-rotating with the outer cylinder. In the rotating frame, the inner cylinder has an azimuthal velocity $U = r_i(\omega_i - \omega_o)$, with ω_i and ω_o the inner and outer cylinder angular velocities, respectively, while the outer cylinder is now stationary. The resulting shear drives the flow and can be expressed non-dimensionally as a shear Reynolds number. $Re_s = dU/\nu$, where ν is the kinematic viscosity of the fluid. We note that in the astrophysical context, (e.g. Ref. [89]), often an extra factor $2/(1 + \eta)$ is used in this definition. In this rotating frame, the outer cylinder rotation in the lab frame manifests itself as a Coriolis force $Ro^{-1}(\mathbf{e}_z \times \mathbf{u})$, where the Rossby number is defined as $Ro = U/(2\omega_o d) = (\eta[1 - \mu])/(2\mu[1 - \eta])$, with $\mathbf{e}_z$ the unit vector in the axial direction and $\mu = \omega_o/\omega_i$ the angular velocity ratio (in the stationary lab system). Note that Ro can be either negative or positive. We define the non-dimensional radius as $\tilde{r} = (r - r_i)/d$, the non-dimensional axial height as $\tilde{z} = z/d$ and the non-dimensional

Figure 10.1: Initial conditions and simulations ran in the (Re_i, Re_o) phase space (left panel) and the $(Re_s, 1/Ro)$ phase space (right panel). The shaded region indicates the zone for which angular velocity is transported outwards $\omega_o \leq \omega_i$, and angular momentum is transported inwards $r_i^2 \omega_i \leq r_o^2 \omega_o$, i.e. the flow is Rayleigh-stable and can be formulated in the rotating frame as $\omega_i > \omega_o$. The dashed line indicates quasi-Keplerian boundary conditions. The arrows indicate the movement of the initial conditions from pure inner cylinder rotation to the Rayleigh-stable states by adding a Coriolis force.

time as $\tilde{t} = tU/d$, where d/U is the large-eddy turnover time of the initial state.

Six simulations were run: Three with fixed $Re_s = 8.10 \cdot 10^3$, but different strengths of the Coriolis force, namely: a) $Ro^{-1} = 0.83$, corresponding to a system on the Rayleigh stability line, equivalent to $r_i^2 \omega_i = r_o^2 \omega_o$, or $\mu = \eta^2 \approx 0.51$ in the lab frame of reference; b) $Ro^{-1} = 1.22$, corresponding to $\mu = \eta^{3/2} \approx 0.60$ or a system in the quasi-Keplerian regime, i.e. the regime in which the angular velocities at the two cylinders are related by Kepler's third law $r_i^3 \omega_o^2 = r_o^3 \omega_i^2$; and c) $Ro^{-1} = 2.50$, an even stronger stabilization, equivalent to $\mu \approx 0.76$. The other three simulations were done by fixing $Ro^{-1} = 1.22$ and using the values $Re_s = 8.10 \cdot 10^3$, $Re_s = 2.52 \cdot 10^4$, and $Re_s = 8.10 \cdot 10^4$. Grids of $N_\theta \times N_r \times N_z = 256 \times 640 \times 512$ were used for all simulations, except for $Re_s = 8.10 \cdot 10^4$, where a grid of $N_\theta \times N_r \times N_z = 512 \times 800 \times 1024$ is required to guarantee sufficient resolution. The adequacy of these meshes for the simulated Re_s was demonstrated in Ref. [113] (§6). Figure 10.1 shows an overview of how the initial conditions from the Rayleigh-unstable regime were moved across the phase space by adding the Coriolis force.

Simulating in a rotating frame adds a new perspective to the problem. In this frame the (non-dimensional) Navier-Stokes equations read:

$$\frac{\partial \hat{\mathbf{u}}}{\partial \hat{t}} + \hat{\mathbf{u}} \cdot \hat{\nabla} \hat{\mathbf{u}} = -\hat{\nabla} \hat{p} + \frac{f(\eta)}{Ta^{1/2}} \hat{\nabla}^2 \hat{\mathbf{u}} - Ro^{-1} \mathbf{e}_z \times \hat{\mathbf{u}}. \quad (10.1)$$

Strong enough outer cylinder co-rotation in the lab-frame implies a large stabilizing Coriolis force Ro^{-1} in the rotating frame, leading to the suppression of the turbu-

Figure 10.2: Three contour plots of the instantaneous angular velocity field ω for $Re_s = 8.10 \cdot 10^3$, $Ro^{-1} = 1.22$, separated by 3.1 non-dimensional time units. Arrows represent the underlying axial and radial velocities. The first snapshot shows the initial field, and in the other two snapshots (at $\tilde{t} = 3.1$ and $\tilde{t} = 6.2$) reversals of the Taylor rolls can be seen: The radial velocity in the gap at mid-height ($\tilde{z} \approx 1.1$) changes from inwards in the first snapshot, to outwards in the second, and back again to inwards in the last.

lence as the large-scale balance in the bulk between the angular velocity gradient and the Coriolis force is broken [112]. This argument for stabilization can be compared to that of Taylor-Proudman's theorem [209, 210] which when applied to rotating Rayleigh-Bénard flow implies that the flow is stable even at large thermal driving, as long as the background rotation is large enough [34].

10.3 Results and analysis

Figure 10.2 shows three instantaneous snapshots of the angular velocity ω for $Re_s = 8.10 \cdot 10^3$, at times $\tilde{t} = 0$, $\tilde{t} = 3.1$ and $\tilde{t} = 6.2$ after the stabilizing Coriolis force corresponding to the quasi-Keplerian regime, $Ro^{-1} = 1.22$, was added. The sequence of figures shows the reversal of the Taylor vortices, which occurs with a period of $\tilde{t} \approx 6.2$. If the system is simulated for a large enough time, here $\tilde{t} \sim 200$, the Taylor vortices progressively fade away. The system behaves approximately as a damped oscillator.

To quantify these reversals, we define the quasi-Nusselt number [25] $Nu_\omega(r, t)$ as $Nu_\omega(r, t) = r^3 (\langle u_r \omega \rangle_{\theta, z} - \nu \partial_r \langle \omega \rangle_{\theta, z}) / T_{pa}$, where T_{pa} is the torque required to drive the system in the purely azimuthal and laminar case. Here $Nu_\omega(r, z)$ represents the

Figure 10.3: Time series of $Nu_\omega - 1$ at mid-gap for $Re_s = 8.10 \cdot 10^3$ and $Ro^{-1} = 1.22$, (quasi-Keplerian). Large oscillations of $Nu_\omega - 1$ from positive to negative can be seen, which, when averaged, show that there is no net turbulent transport of angular velocity. The fluctuations decay with time. For a long enough time $\tilde{t} \approx 3000$, they stop and $Nu_\omega - 1$ finally stabilizes at 0.

transport of angular velocity. For a purely azimuthal flow with no turbulence, $Nu_\omega = 1$, and therefore $Nu_\omega - 1$ represents the additional transport of angular velocity due to turbulence.

Figure 10.3 shows a time series of the axially averaged $Nu_\omega - 1$ at the mid-radius $\tilde{r} = 0.5$, for $Re_s = 8.10 \cdot 10^3$ and $Ro^{-1} = 1.22$. Here, $Nu_\omega - 1$ can be seen to oscillate between large positive and large negative values, with an average of approximately zero, i.e. no net turbulent transport. For large enough times ($\tilde{t} > 3000$), the oscillations damp out and the flow becomes purely azimuthal and laminar, $Nu_\omega = 1$. In the other simulations, the period of oscillation depends strongly on Ro^{-1} (decreasing with increasing Ro^{-1}) but only weakly on Re_s . Oscillations can be seen even if a non-stabilizing co-rotation is added, but the system does not return to a purely azimuthal state with $Nu_\omega = 1$. Instead, after some oscillations in which Nu_ω decreases, a new turbulent state with $Nu_\omega > 1$ is reached.

Figure 10.4 shows the axially averaged $Nu_\omega - 1$, but now measured directly at the inner cylinder (probing the inner boundary layer (BL)) for different values of Ro^{-1} and Re_s . Unlike the previous mid-gap case, no oscillations can be seen, reflecting that the oscillations occur only in the bulk of the flow, where the Taylor vortices are, but that the BLs are unaffected. Of course, the decay towards the purely azimuthal, laminar state $Nu_\omega = 1$ is also observed in the BLs (figure 10.4(a,b)), independent of Ro^{-1} and Re_s .

Finally, we quantify the decay of turbulence. To do so, we define the turbulent kinetic energy substitute of the flow as $K(t) = \frac{1}{2}(u_r^2 + u_z^2)_{r,z,\theta}$. Note that this is not the

Figure 10.4: (a): Semi-log time series of $Nu_\omega - 1$ at the inner cylinder ($\tilde{r} = 0$, z -averaged) for $Re_s = 8.10 \cdot 10^3$ and three values of Ro^{-1} , namely $Ro^{-1} = 0.83$, (Rayleigh-stability line, red dashed line), $Ro^{-1} = 1.22$ (quasi-Keplerian, black line), and $Ro^{-1} = 2.50$ (blue dash-dot line). (b): semi-log time series of $Nu_\omega - 1$ at the inner cylinder for $Ro^{-1} = 1.22$ and three values of Re_s (black line: $Re_s = 8.10 \cdot 10^3$, blue dashed line: $Re_s = 2.52 \cdot 10^4$, and red dash-dot line: $Re_s = 8.10 \cdot 10^4$). Two decay modes can be seen in this figure, with the crossovers marked with arrows. In both (a) and (b), fluctuations are absent, unlike at the mid-gap, but $Nu_\omega - 1$ decays. For large enough time, $Nu_\omega - 1$ drops to zero, taking more time to do so for higher Re_s or smaller Ro^{-1} . Zoom-ins of the first few large eddy turnover times are shown as insets (linear scale) for both figures. Transients, corresponding to an increase of Nu_ω can be seen during the initial stages.

Figure 10.5: (a): Time series of $K(\tilde{t})/K_0$ for $Re_s = 8.10 \cdot 10^3$ and three values of Ro^{-1} , $Ro^{-1} = 0.83$, (Rayleigh-stability line), $Ro^{-1} = 1.22$, (quasi-Keplerian) and $Ro^{-1} = 2.50$. (b): Time series of K/K_0 for $Ro^{-1} = 1.22$ and three values of Re_s . After the initial transient, an exponential decay of the turbulent kinetic energy can be seen, even at the largest Reynolds numbers.

full turbulent energy, as for that also fluctuations in u_θ contribute, but, by definition, K has the desirable property of being zero in the purely azimuthal state. Figure 10.5(a) shows the decay of K for $Re_s = 8.10 \cdot 10^3$ and three values of Ro^{-1} , while figure 10.5(b) shows the decay for $Ro^{-1} = 1.22$ and three values of Re_s . In the simulations, $K(\tilde{t})$ shows fluctuations with a short time scale, similar to that of Nu_ω in the mid-gap, so for clarity only the peak of $K(\tilde{t})$ in every cycle is represented.

As expected from the previous results on $Nu_\omega(\tilde{t}) - 1$, $K(\tilde{t})$ decays to zero when given enough time. The only mechanism for energy dissipation is viscosity, so it is not surprising that the decay time is proportional to the energy content, thus leading to an approximately exponential decay, $K(\tilde{t}) = K_0 \exp(-\tilde{t}/\tilde{\tau})$ where $\tilde{\tau}$ is the characteristic (dimensionless) decay time and $K_0 = K(\tilde{t} = 0)$. Fits to the data in figure 10.5(b) were performed to obtain estimates for the dependence $\tilde{\tau}(Re_s)$, which is shown in figure 10.6(a). The absolute value of $\tilde{\tau}$ is large (around 70 large eddy turnover times for $Re_s \sim 10^4$ and $Ro^{-1} = 1.22$) and in the range studied shows a power law

$$\tilde{\tau} \sim \sqrt{Re_s}. \quad (10.2)$$

This scaling can be understood by realizing that the decay of the turbulence is first dominated by the energy dissipation rate ε_{BL} in the BLs, which scales as [70] $\varepsilon_{BL} \sim \nu^3 d^{-4} Re_s^{5/2}$. This immediately implies the scaling relation (10.2) for the typical decay time $\tau = \tilde{\tau} d/U \sim U^2/\varepsilon_{BL}$. Only for very large times or for very small Reynolds numbers from the very beginning, i.e., when the thickness of the BLs is comparable to the gap thickness d and the Taylor rolls have died out or are absent, does $\varepsilon_{BL} \sim \nu^3 d^{-4} Re_s^2$ (cf. Ref. [154]) and thus $\tau \sim Re_s$. This second regime can be seen

in figures 10.4(b) and 10.6(b).

We highlight that even though $K(\tilde{r})$ decays, the total energy of the system *increases*, as it absorbs more energy through the boundaries than it dissipates viscously. This can be seen in figure 10.6b, which shows a time series of the non-dimensionalized energy input into the system $\dot{E} = (T_i\omega_i + T_o\omega_o)/(\rho\Omega\epsilon_{pa})$, where $T_{i,o}$ is the torque applied at the inner (outer) cylinder, ρ and Ω are the fluid density and volume, respectively, and ϵ_{pa} is the viscous dissipation in the purely azimuthal and laminar state, and a time series of the non-dimensionalized energy dissipation rate inside the system, $\tilde{\epsilon}_v = \frac{1}{2}v\langle[\partial_i u_j + \partial_j u_i]^2\rangle_{r,\theta,z}/\epsilon_{pa}$. When transitioning from turbulence to the purely azimuthal $u_\theta(r)$ -profile, the system absorbs energy, while at the same time the initial turbulence decays. The purely azimuthal $u_\theta(r)$ -profile has more energy than its turbulent counterpart. Indeed, this is the reason for the Rayleigh instability, now suppressed by the Coriolis force.

Again, two time scales for the decay can be seen in the inset of figure 10.6(b). After the initial transient, turbulence in the bulk decays rapidly, due to the efficient angular momentum transport by the Taylor rolls. After about 200-300 large eddy turnover times this decay mechanism is exhausted, as then the (turbulent) Taylor rolls in the bulk are so weak that they can no longer transport angular momentum. Then a second decay mechanism takes over, whose onset is marked by an arrow in the inset of figure 10.6(b). In this regime, characterized by purely viscous dissipation $\sim v^3 L^{-4} Re_s^2$ in the whole gap, the profile $u_\theta(r)$ returns to the laminar form and Nu_ω thus becomes unity. It is important to note that Nu_ω must return to unity, because in the statistically stationary state Nu_ω must be independent of r .

10.4 Summary and conclusions

In summary, consistent with what had been found in the experiments of Ji *et al.* [90] and Schartman *et al.* [91], no turbulence, and, therefore, no turbulent transport of angular momentum can be seen in the Rayleigh-stable regime for the control parameter range studied once the effect of axial boundaries is mitigated. An initially turbulent state is stabilized by the co-rotation of the outer cylinder and the transport of angular momentum in the bulk ceases; the turbulence decays in all our simulations, for Re_s up to $8.10 \cdot 10^4$.

We point out, however, that even though the system is stable, the decay times for turbulence are long, ranging between hundreds and thousands of large eddy turnover times. If one extrapolates the scaling relation (10.2) to shear Reynolds numbers attainable in cold accretion disks, which are of the order of 10^{14} , decay times for the turbulence become very large. This is done ignoring possible crossovers to an ultimate regime in which the energy dissipation rate may be dominated by the bulk dissipation [70, 78] $\epsilon_{bulk} \sim v^3 d^{-4} Re_s^3$. In that regime $\tilde{\tau}$ would become independent

Figure 10.6: (a) Log-log plot of the turbulence decay time $\tilde{\tau}$ vs Re_s for $Ro^{-1} = 1.22$ (quasi-Keplerian) and the three values of Re_s simulated. The power law $\tilde{\tau} = 0.79\sqrt{Re_s}$ fits the data points. (b): Semi-logarithmic plot of the time series of both the energy input through the boundaries $\dot{E} - 1$ and the viscous dissipation of energy in the system $\epsilon_\nu - 1$ for $Re_s = 2.49 \cdot 10^4$ and $Ro^{-1} = 1.22$. The laminar dissipation, i.e. 1 with the non-dimensionalization chosen, is subtracted from both variables. An inset with the ratio between both energies is shown. The arrow marks the transition between the two decay mechanisms. For quite a long time, the system still absorbs more energy than it dissipates; equilibrium is reached only for $\tilde{t}$ of order beyond 1000, due to the small decay rate.

of Re_s . Indeed, for astrophysical disks without inner- and outer- boundary layers, we expect this to be the case. Viscous dissipation would be dominated by the bulk dynamics, and 10.2 would not be applicable.

The decay times observed in the simulations are in discrepancy with experiments by Borrero-Echeverry *et al.* [211] and Edlund & Ji [212]. Edlund & Ji [212] use a similar set-up to that of Ji *et al.* [90], but now switch on and off a pump to perturb the quiescent TC flow with quasi-Keplerian boundary conditions. The decay time of the perturbation is much smaller than those observed in the DNS. The discrepancy in time scale can probably be attributed to increased bulk dissipation at higher Re and to end-plate effects, which increase the dissipation of the system. On the other hand, Borrero-Echeverry *et al.* [211] showed that for pure outer cylinder rotation, perturbations decay stochastically (i.e. are memoryless). The decay times grow super-exponentially, and thus are much larger than those ones seen in the DNS. Turbulent dynamics can be seen before the relaminarization, and a mechanism by which turbulence regenerates is present. We point out that for pure outer cylinder rotation, TC flow transports angular velocity inwards, and this completely changes the dynamics of the problem, thus causing a different behaviour in the decay.

It is also unclear whether the extreme values of η present in astrophysical accretion disks will have an effect on the decay times. The Rossby number has a strong

η -dependence, and for $\eta \rightarrow 0$ it diverges. Other types of mechanism might dominate the angular momentum transport in cold accretion disks. These are axially stratified in density, and thus other types of instability such as subcritical baroclinic instabilities [213–215] can be present. Finally, transport associated to self-gravity [216] may also play a role. We refer the reader to Ref. [207] for a more detailed discussion on these issues.

In conclusion, in our TC simulations without end plates, in the parameter range studied up to shear Reynolds numbers of 10^5 , the flow was seen to not only remain laminar in the presence of small perturbations, but also when starting from an initially turbulent state the turbulence decays and the system returns to its laminar state. All of this happens while the system absorbs energy. Our findings extend the results from Avila [93], who also attributed the turbulence and increased angular momentum transport found in some of the experiments [44, 89, 92] which probe the linearly stable regime of TC flow to the effects of end-plate confinement.

Initial conditions are taken from the Rayleigh-unstable region, and thus are dominated by toroidal Taylor vortices with a strong axial dependence. These are chosen as they are as general as possible. When the stabilizing rotation is switched on, the Taylor–Proudman theorem implies that these vortices cannot be sustained, and thus vanish. It is worth noting that if different boundary conditions are used, such as radial inflow [217] or differentially heated cylinders [218], Taylor columns which have no axial dependence may develop. However, with the basic TC boundary conditions, these structures do not develop, as the system might be too constrained. This result is in agreement with theoretical work by Rincon *et al.* [219], who reach a similar conclusion for plane Couette flow, as the non-linear mechanisms appear to be fundamentally differently beyond the Rayleigh stability line. An analysis based on linear transient growth such as done by Marezke *et al.* [220], which also predicts columnar structures as the optimal initial conditions, may thus not be valid for determining the stability of TC flow at high Reynolds numbers.

11

Conclusions and Outlook

Whereof one cannot speak, thereof one must be silent

—LUDWIG WITTGENSTEIN

The search for an objective and incontestable truth permeates Western thought since the time of Plato's myth of the cavern. According to the epistemological perspectives developed during the Enlightenment, science has since widely been regarded as the systematic search for these eternal truths through reason and experiment. The existence of these truths was heavily scrutinized by the Berkeley and the German idealists. Following the failure of the logical positivists to completely objectify science, Kuhn, Popper and other philosophers of science in the 20th century attempted to redefine what science exactly is. According to Kuhn [221], science endlessly accumulates knowledge, but most of the time this does not lead to eternal truth, but instead extends the scope and precision of working models, having periodic paradigm-shifts in which new models are developed. This thesis takes place in the normal, puzzle-solving stage of science. We are trying to extend a central paradigm, i.e. fluid mechanics, to cope with an increasing number of problems, i.e. incompressible wall-bounded turbulent flows, and Taylor-Couette flow in particular. In this last chapter, we will summarize and reflect on the knowledge we have accumulated in all the other chapters. We will analyse which puzzles have been solved, helping us achieve a higher understanding of first Taylor-Couette flow, and then turbulence in general.

The study of TC flow is not a static field. During the years the research of this thesis took place, the explored parameter space which was shown in the introduction has been greatly expanded. This was due not only by the simulations (and experiments) from this thesis, but also with other experimental and numerical work [45, 80, 82, 118, 123, 156, 173]. In figure 11.1, we present an updated parameter space, which can be contrasted with figure 1.4 in the introduction.

The thesis starts with the development of the numerical code for the simulations. In §2 the high Reynolds number code for simulations of wall-bounded turbulence

Figure 11.1: State-of-the-art of the parameter space exploration of TC flow at the end of the thesis, including all data from the previous figure, as well as the new data from [45, 80, 82, 118, 123, 156, 173] and all previous chapters, which also correspond to Refs. [112, 113, 129, 175, 176].

is detailed. The second-order finite-difference scheme by Verzicco & Orlandi [103] was modified by integrating explicitly the viscous in the wall-parallel directions. A pencil-type decomposition, with the pencils aligned in the wall-normal direction, was developed. This synergy between algorithm and data parallelization allowed for the total amount of communication to be minimal, and this made the code scale up to sixty-four thousand processors, allowing the largest simulations in the thesis (as those in §8) to be performed. In §3 we detailed a further addition to the DNS code in order to efficiently simulate scalar fields. Sharp fronts, present in the scalar field and absent from the momentum fields [142–148] impose that scalars require finer meshes than momentum fields to be adequately resolved, even if momentum and the scalar have similar diffusivities. A multiple resolutions strategy was developed to decouple the fields. The method requires at least a tri-cubic Hermitian interpolation of velocity on to the refined mesh, and the use of sub-timestepping of the scalar field to keep the integration stable. This strategy was applied to Rayleigh-Bénard convection, allowing for considerable speed-ups of the simulations, up to a factor seven.

In §§4–7 we concerned ourselves with the exploration and characterization of the TC flow parameter space, in an attempt to understand the phenomena of optimal transport, and the transition to the ultimate regime. In §4, we started by validating the code against existing experimental and numerical data. A new formulation of TC, based on a rotating frame which co-rotates with the outer cylinder was developed. Using this formulation, we followed by examining the phenomena of optimal transport, i.e., a maximum in the $Nu_\omega(Ro^{-1})$ relationship for a given Ta , for $\eta = 0.714$ and moderate drivings $10^4 < Ta < 10^8$. Similarly to experiments [43, 44], an optimum transport at a non-zero rotation rate was found. The position of the maximum was found to shift with Ta , and an explanation for this shift was elucidated. This explanation is consistent with the experimental result that the optimal a becomes approximately independent of Ta for large enough drivings. Using the rotating frame, a link between the local angular velocity profiles and the global transport quantities was found.

We continued in §5 by exploring three other values of η at moderate driving. The DNS were complimented by a set of experiments in the T³C provided by Sander Huisman and coworkers. Radius ratios between $\eta = 0.5$ and $\eta = 0.909$ were used, reaching up to $Ta \sim \mathcal{O}(10^8)$ in the DNS and $Ta \sim \mathcal{O}(10^{12})$ in the experiments. This new data allowed us to test some of the insights of §4, while developing our understanding of the effect of η on the flow. Effective scaling laws for the torque were found, which for sufficiently large driving Ta are independent of the radius ratio η , and correspond to those expected in the ultimate regime. Optimum transport at a non-zero rotation rates was also found in both experiments and numerics. When comparing the predictions of Refs. [77, 156] to the new experimental data, discrep-

ancies were found. In addition, the link between the local angular velocity profiles with the optimum transport was further corroborated.

In §6 we continued by exploring in detail the transition to the ultimate regime for $\eta = 0.714$ and pure inner-cylinder rotation. When increasing the driving past the onset of instability, three regimes were found. First, the flow was dominated by large coherent structures. Further increasing the driving leads to the breakup of these structures, and to a transitional regime where the coherent structures are only recovered on average. The presence of an underlying structure allows for the coexistence of laminar and turbulent boundary layer dynamics. The flow was found to only transition to the ultimate regime when the mean flow effects fade away, and the flow becomes completely dominated by plumes. Once this happens, the global scaling laws relating torque to driving become those expected in the ultimate regime. Logarithmic boundary layer profiles for the angular velocity were found to develop in the ultimate regime, according to the predictions of Ref. [165].

In §7, we concluded the exploratory part of the thesis by mapping out the transition to the ultimate regime across the four dimensions of the parameter space, i.e. the driving Ta required to reach the ultimate regime as a function of aspect ratio Γ , the radius ratio η and the rotation ratio Ro^{-1} . We showed that the transition is approximately independent of Γ and Ro^{-1} , but depends significantly on η . We furthermore calculated the local angular velocity profiles and visualized different flow regimes, depending on both the shearing of the flow, and the Coriolis force originating from the outer cylinder rotation. Two main regimes were distinguished, based on the magnitude of the Coriolis force, namely the co-rotating and weakly counter-rotating regime dominated by Rayleigh-unstable regions, and the strongly counter-rotating regime where a mixture of Rayleigh-stable and Rayleigh-unstable regions exist. Furthermore, an analogy between radius-ratio and outer-cylinder rotation was revealed, namely that smaller gaps behave like a wider gap with co-rotating cylinders, and that wider gaps behave like smaller gaps with weakly counter-rotating cylinders. Finally, the effect of the aspect ratio on the effective torque versus Taylor number scaling is analysed and it was shown that different branches of the $Nu_\omega(Ta)$ relationships associated to different Γ cross within 15% of the Taylor number associated to the transition to the ultimate regime. The chapter culminates in a phase diagram, which can be seen as the extension of the Andereck *et al.* [37] phase diagram towards the ultimate regime.

In §8, a detailed look at the effect of the Taylor-rolls was performed. Taylor-rolls were found to dominate the dynamics of small-gap TC, even at the highest Reynolds numbers simulated. Imposing a weak axial flow was found to convect the rolls, but not to weaken them. By closely looking at the spectra, the rolls were found to be attached to the wall, and to actively transport angular velocity through Reynolds stresses. Very important differences in the behaviour of TC when comparing with

other canonical flows were found. With increasing Reynolds number, an overlap region was found. This region is not very pronounced for the largest simulations in §8, but further increasing the Reynolds numbers of the DNS would open the door for observing logarithmic profiles which are consistent with the overlap-layer assumptions.

In §9, we revisited the accuracy of small-box simulations by running larger computational boxes. The dependence on the axial periodicity aspect ratio Γ , which directly fixes the wavelength of the Taylor rolls, was found to be relevant for some statistics. The velocity fluctuations and the energy spectra were clearly affected, but the small scale statistics remained unchanged. One big disadvantage of DNS of TC flow is that Γ must be kept small so that the computational cost remains reasonable for high Reynolds numbers. By keeping Γ small, we completely fix the wavelength of the underlying single Taylor roll. In §9, it was shown that even if the Nusselt number loses its Γ dependence at high Re , other statistics are affected. As the axial periodicity length fixes the wavelength of the vortices, the possibility of transitions between different vortical states is lost. Therefore, it probably is worth running a DNS with aspect ratio of at least $\Gamma = 12$ to check the preferred wavelength of the Taylor rolls, and to check the possibility of coexistence of different vortical states.

Sub-critical TC flow was explored in §10. Experiments were in conflict on whether TC flow has a sub-critical transition to turbulence. Discrepancies were attributed to end-plate effects by Avila [93]. Using DNS, end-plates were avoided completely, and using the rotating-frame formulation to instantaneously add a strong rotation, without causing any singularities at the cylinders, the possibility of exploring the transition opened up. Starting from a fully turbulent state, taken at $Re_s \sim \mathcal{O}(10^4 - 10^5)$, a stabilizing rotation was added. Turbulence was found to decay for all Rayleigh-stable simulated rotations (i.e. those for which $L_o > L_i$ and $\omega_i > \omega_o$), and for all Reynolds numbers. The decay time τ was found to increase with increasing Reynolds following an approximate scaling relationship of $\tau \sim Re_s^{1/2}$. However, the answer remains slightly unsatisfactory. Even if consistent with the large scale balance of §4.6, it still seems strange that absolutely no turbulence can exist in the Rayleigh-stable regime with $\omega_i > \omega_o$. The parameter range of $\omega_i < \omega_o$ was explored experimentally by Borrero-Echeverry *et al.* [211], and they presented a different kind of transition of sub-critical TC flow to turbulence. The rotating frame formulation is not valid in this case, as $\omega_i < \omega_o$. Therefore, open questions in this regime remain, and more simulations in this regime would open the door to further understanding of sub-critical TC flow.

The behaviour of large gap TC has also raised some questions. The transition to the ultimate regime was found to be heavily delayed for $\eta = 0.5$, while Taylor vortices were seen to vanish from the flow. Also, the wall-velocity profiles for $\eta = 0.714$ look very different from the classical law-of-the-wall with von Kármán constants of $\kappa \approx 0.4$ and $B \approx 5.2$. Indeed, $\kappa_\omega \approx 0.56$ was found in §6. Understanding how (or

even, if) overlap layers form in TC flow with curvature is an open question. Further study to understand this regime is necessary.

The behaviour, or even, the mere existence of optimal transport is also a big open question for large gap TC flow. We are very far from a full understanding of the effect of curvature, and especially of the limit $\eta \rightarrow 0$. More simulations, and also more experiments, would surely help. The main caveat of this regime is that, from the experience gathered in performing the simulations in this thesis, transients become longer with decreasing η . The asymmetry between cylinders grows very rapidly, and the viscous wall unit at the inner cylinder becomes very small even at moderate Reynolds numbers ($Re \sim \mathcal{O}(10^4)$). This forces the use of very fine grids, especially in the axial direction, and as a result, these simulations will probably have larger computational requirements than naively expected.

Also, all simulations in this thesis have been performed with axially periodic boundary conditions. End-plates, which axially confine the flow are required in experiments, even if they are absent from some of the applications of TC flow. In the high Reynolds, linearly unstable regime of TC, which has been the focus of this thesis, experimental and numerical data are found to be in very good agreement despite the use of axial periodicity (cf. figures 6.1 and 7.4(b)). Even so, it would probably be worth studying a single, or a couple of cases with end-plates, to check the validity of the assumptions made. These simulations will have very high computational cost, as higher resolutions, and non-uniform meshes in two dimensions will be required. This will have a very strong impact on the Poisson equation solver of the code, which becomes prohibitively costly for even moderate Reynolds numbers of $Re \sim \mathcal{O}(10^4)$.

Finally, further work can also be performed in developing the code detailed in §2 to scale up to today's systems, which consist of up to 10^6 cores. Porting the code for use in massively parallel GPUs is also a challenge, as no single efficient implementation of an incompressible Navier-Stokes solver exists for these systems, and the Exascale supercomputers coming up soon are based on these paradigms. The code has been open-sourced for the benefit of the community, under the name of AFiD, and this will allow it to be further tested and improved not only by its authors, but by everybody who works in the Rayleigh-Bénard and Taylor-Couette fields.

Of course, we need not limit ourselves to single phase “platonic” TC flow. Many more variations of TC flow exist, and can be explored. Some possibilities have already been explored experimentally, such as rough inner and outer cylinders [222, 223], or the addition of a second phase, such as bubbles [191, 193], and can be further explored numerically. Other variations which are hard to implement in experiments can be tried numerically due to the inherent flexibility of simulations, like mixed free-slip/no-slip cylinders. The possibilities seem endless. There are still many more puzzles to solve not only in TC flow, but also in the entire Universe.

The gods had condemned Sisyphus to ceaselessly rolling a rock to the top of a mountain, whence the stone would fall back of its own weight. They had thought with some reason that there is no more dreadful punishment than futile and hopeless labour. [...] One always finds one's burden again. But Sisyphus teaches the higher fidelity that negates the gods and raises rocks. He too concludes that all is well. [...] The struggle towards the heights is enough to fill a man's heart. One must imagine Sisyphus happy.

—ALBERT CAMUS

References

- [1] Various, “Turbulence - Wikipedia, the free encyclopedia”, (2014), URL <http://en.wikipedia.org/wiki/Turbulence>.
- [2] L. F. Richardson, “The supply of energy from and to atmospheric eddies”, Proc. R. Soc. London Ser. A **97**, 354–373 (1920).
- [3] A. N. Kolmogorov, “The local structure of turbulence in incompressible viscous fluid for very large Reynolds numbers”, Dokl. Akad. Nauk SSSR **30**, 9–13 (1941), reprinted in Proc. R. Soc. Lond. A **434**, 9-13 (1991).
- [4] A. N. Kolmogorov, “On degeneration (decay) of isotropic turbulence in incompressible viscous liquid”, Dokl. Akad. Nauk SSSR **31**, 538–540 (1941).
- [5] A. N. Kolmogorov, “Dissipation of energy in locally isotropic turbulence”, Dokl. Akad. Nauk SSSR **32**, 16–18 (1941), reprinted in Proc. R. Soc. Lond. A **434**, 15-17 (1991).
- [6] K. R. Sreenivasan, “On the universality of the Kolomogorov constant”, Phys. Fluids **7**, 2778–2784 (1995).
- [7] S. B. Pope, *Turbulent Flow* (Cambridge University Press, Cambridge) (2000).
- [8] L. Prandtl, “Neuere Ergebnisse der Turbulenzforschung”, Zeit. Vereines Deutscher Ingenieure **77**, 105–114 (1933).
- [9] A. A. Townsend, *The structure of turbulent shear flow* (Cambridge University Press, Cambridge, UK) (1976).
- [10] A. E. Perry, S. Henbest, and M. S. Chong, “A theoretical and experimental study of wall turbulence”, J. Fluid Mech. **165**, 163–199 (1986).
- [11] A. J. Smits, B. J. McKeon, and I. Marusic, “High-Reynolds number wall turbulence”, Ann. Rev. Fluid. Mech. **43**, 353–375 (2011).
- [12] J. Jimenez, “Cascades in wall-bounded turbulence”, Ann. Rev. Fluid. Mech. **44**, 27–45 (2012).
- [13] J. Pedlosky, *Geophysical fluid dynamics* (Springer-Verlag, Berlin, Heidelberg) (1987).
- [14] J. Anderson, *Computational Fluid Dynamics* (McGraw-Hill Science/Engineering/Math, New York City) (1995).

- [15] L. D. Landau and E. M. Lifshitz, *Fluid Mechanics* (Pergamon Press, Oxford) (1987).
- [16] S. A. Orszag and G. S. Patterson, “Numerical simulation of three-dimensional homogeneous isotropic turbulence”, *Phys. Rev. Lett* **28**, 76–79 (1972).
- [17] Y. Kaneda, T. Ishihara, M. Yokokawa, K. Itakura, and A. Uno, “Energy dissipation rate and energy spectrum in high resolution direct numerical simulations of turbulence in a periodic box”, *Phys. Fluids* **15**, L21 (2003).
- [18] T. Ishihara, Y. Kaneda, K. Morishita, M. Yokokawa, and A. Uno, “Energy spectra of higher Reynolds number turbulence by the DNS with up to 12288^3 grid points”, *Bulletin of the American Physical Society* **59** (2014).
- [19] I. Bermejo-Moreno, J. Bodart, J. Larsson, B. M. Barney, J. W. Nichols, and S. Jones, “Solving the compressible Navier-Stokes equations on up to 1.97 million cores and 4.1 trillion grid points”, *SC’13 Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis* **62** (2014).
- [20] D. Haim and L. Pismen, “Performance of a photochemical reactor in the regime of Taylor-Görtler vortical flow”, *Chem. Eng. Sci.* **49**, 1119–1129 (1994).
- [21] S. T. Wereley and R. M. Lueptow, “Inertial particle motion in a Taylor-Couette rotating filter”, *Phys. Fluids* **11** (1999).
- [22] R. M. Lueptow and J. J. Peterson, “Blood processing system for cell washing”, (1984).
- [23] B. Kong, J. V. Shanks, and R. D. Vigil, “Enhanced algal growth rate in a Taylor vortex reactor”, *Biotechnol. Bioengin.* **110**, 2140–2149 (2013).
- [24] B. Kong and R. D. Vigil, “Light-limited continuous culture of *Chlorella vulgaris* in a Taylor vortex reactor”, *Environ. Prog. & Sustain. Energy* **32**, 884–890 (2013).
- [25] B. Eckhardt, S. Grossmann, and D. Lohse, “Torque scaling in turbulent Taylor-Couette flow between independently rotating cylinders”, *J. Fluid Mech.* **581**, 221–250 (2007).
- [26] M. Couette, *Études sur le frottement des liquides* (Gauthier-Villars et fils) (1890).

- [27] G. G. Stokes, “On the theories of the internal friction of fluids in motion and of the equilibrium and motion of elastic solids”, Trans. Cambridge Phil. Soc. VIII **3**, 287–319 (1849).
- [28] M. Magules, “Über die Bestimmung des Reibungs- und Gleitcoefficienten aus ebenen Bewegungen einer Flüssigkeit”, Wien. Ber. (2nd ser) **83**, 588–602 (1881).
- [29] A. Mallock, “Determination of the viscosity of water”, Proc. R. Soc. Lond. **45**, 126–132 (1888).
- [30] A. Mallock, “Experiments on fluid viscosity”, Phil. Trans. R. Soc. Lond. A **187**, 41–56 (1896).
- [31] F. Wendt, “Turbulente Strömungen zwischen zwei rotierenden Zylindern”, Ingenieurs-Archiv **4**, 577–595 (1933).
- [32] G. I. Taylor, “Experiments on the motion of solid bodies in rotating fluids”, Proc. R. Soc. Lond. A **104**, 213–218 (1923).
- [33] G. I. Taylor, “Fluid friction between rotating cylinders”, Proc. R. Soc. London A **157**, 546–564 (1936).
- [34] S. Chandrasekhar, *Hydrodynamic and Hydromagnetic Stability* (Dover, New York) (1981).
- [35] D. Coles, “Transition in circular Couette flow”, J. Fluid Mech. **21**, 385–425 (1965).
- [36] G. P. Smith and A. A. Townsend, “Turbulent Couette flow between concentric cylinders at large Taylor numbers”, J. Fluid Mech. **123**, 187–217 (1982).
- [37] C. D. Andereck, S. S. Liu, and H. L. Swinney, “Flow regimes in a circular Couette system with independently rotating cylinders”, J. Fluid Mech. **164**, 155–183 (1986).
- [38] P. Tong, W. I. Goldburg, J. S. Huang, and T. A. Witten, “Anisotropy in turbulent drag reduction”, Phys. Rev. Lett. **65**, 2780–2783 (1990).
- [39] D. P. Lathrop, J. Fineberg, and H. S. Swinney, “Turbulent flow between concentric rotating cylinders at large Reynolds numbers”, Phys. Rev. Lett. **68**, 1515–1518 (1992).
- [40] D. P. Lathrop, J. Fineberg, and H. S. Swinney, “Transition to shear-driven turbulence in Couette-Taylor flow”, Phys. Rev. A **46**, 6390–6405 (1992).

- [41] G. S. Lewis and H. L. Swinney, “Velocity structure functions, scaling, and transitions in high-Reynolds-number Couette-Taylor flow”, *Phys. Rev. E* **59**, 5457–5467 (1999).
- [42] D. P. M. van Gils, G. W. Bruggert, D. P. Lathrop, C. Sun, and D. Lohse, “The Twente turbulent Taylor-Couette (T^3C) facility: strongly turbulent (multi-phase) flow between independently rotating cylinders”, *Rev. Sci. Instr.* **82**, 025105 (2011).
- [43] D. P. M. van Gils, S. G. Huisman, G. W. Bruggert, C. Sun, and D. Lohse, “Torque scaling in turbulent Taylor-Couette flow with co- and counter-rotating cylinders”, *Phys. Rev. Lett.* **106**, 024502 (2011).
- [44] M. S. Paoletti and D. P. Lathrop, “Angular momentum transport in turbulent flow between independently rotating cylinders”, *Phys. Rev. Lett.* **106**, 024501 (2011).
- [45] S. G. Huisman, D. P. M. van Gils, S. Grossmann, C. Sun, and D. Lohse, “Ultimate turbulent Taylor-Couette flow”, *Phys. Rev. Lett.* **108**, 024501 (2012).
- [46] H. Fasel and O. Booz, “Numerical investigation of supercritical Taylor-vortex flow for a wide gap”, *J. Fluid Mech.* **138**, 21–52 (1984).
- [47] K. Coughlin and P. S. Marcus, “Turbulent Bursts in Couette-Taylor Flow”, *Phys. Rev. Lett.* **77**, 2214–2217 (1996).
- [48] S. Dong, “Direct numerical simulation of turbulent Taylor-Couette flow”, *J. Fluid Mech.* **587**, 373–393 (2007).
- [49] S. Dong, “Turbulent flow between counter-rotating concentric cylinders: a direct numerical simulation study”, *J. Fluid Mech.* **615**, 371–399 (2008).
- [50] D. Pirro and M. Quadrio, “Direct numerical simulation of turbulent Taylor-Couette flow”, *Eur. J. Mech. B-Fluids* **27**, 552 (2008).
- [51] T. Gebhardt and S. Grossmann, “The Taylor-Couette eigenvalue problem with independently rotating cylinders”, *Z. Phys. B* **90**, 475–490 (1993).
- [52] R. Donnelly, “Taylor-Couette flow: the early days”, *Physics Today* 32–39 (1991).
- [53] F. Ravelet, R. Delfos, and J. Westerweel, “Influence of global rotation and Reynolds number on the large-scale features of a turbulent Taylor-Couette flow”, *Phys. Fluids* **22**, 055103 (2010).

- [54] R. van Hout and J. Katz, “Measurements of mean flow and turbulence characteristics in high-Reynolds number counter-rotating Taylor-Couette flow”, *Phys. Fluids* **23**, 105102 (2011).
- [55] M. Bilson and K. Bremhorst, “Direct numerical simulation of turbulent Taylor-Couette flow”, *J. Fluid Mech.* **579**, 227 (2007).
- [56] W. Q. He, M. Tanahashi, and T. Miyauchi, “Direct numerical simulation of turbulent Taylor-Couette flow with high Reynolds number”, in *Advances in turbulence XI* (Porto, Portugal) (2007).
- [57] H. L. Swinney and J. P. Gollub, *Hydrodynamic instabilities and the transition to turbulence*, volume 45 (Topics in Applied Physics) (Springer-Verlag, Berlin) (1981).
- [58] G. Pfister and I. Rehberg, “Space dependent order parameter in circular Couette flow transitions”, *Phys. Lett.* **83**, 19–22 (1981).
- [59] G. Pfister, H. Schmidt, K. A. Cliffe, and T. Mullin, “Bifurcation phenomena in Taylor-Couette flow in a very short annulus”, *J. Fluid Mech.* **191**, 1–18 (1988).
- [60] P. Drazin and W. H. Reid, *Hydrodynamic stability* (Cambridge University Press, Cambridge) (1981).
- [61] F. H. Busse, “The stability of finite amplitude cellular convection and its relation to an extremum principle”, *J. Fluid Mech.* **30**, 625–649 (1967).
- [62] E. N. Lorenz, “Deterministic nonperiodic flow”, *J. Atmos. Sci* **20**, 130–141 (1963).
- [63] G. Ahlers, “Low temperature studies of the Rayleigh-Bénard instability and turbulence”, *Phys. Rev. Lett.* **33**, 1185–1188 (1974).
- [64] R. P. Behringer, “Rayleigh-Bénard convection and turbulence in liquid-helium”, *Rev. Mod. Phys.* **57**, 657 – 687 (1985).
- [65] M. A. Dominguez-Lerma, D. S. Cannell, and G. Ahlers, “Eckhaus boundary and wavenumber selection in rotating Couette-Taylor flow”, *Phys. Rev. A* **34**, 4956 (1986).
- [66] S. H. Strogatz, *Nonlinear dynamics and chaos* (Perseus Press, Reading) (1994).
- [67] M. C. Cross and P. C. Hohenberg, “Pattern formation outside of equilibrium”, *Rev. Mod. Phys.* **65**, 851 (1993).

- [68] E. Bodenschatz, W. Pesch, and G. Ahlers, “Recent developments in Rayleigh-Bénard convection”, *Ann. Rev. Fluid Mech.* **32**, 709–778 (2000).
- [69] E. D. Siggia, “High Rayleigh number convection”, *Annu. Rev. Fluid Mech.* **26**, 137–168 (1994).
- [70] S. Grossmann and D. Lohse, “Scaling in thermal convection: A unifying view”, *J. Fluid. Mech.* **407**, 27–56 (2000).
- [71] L. P. Kadanoff, “Turbulent heat flow: Structures and scaling”, *Phys. Today* **54**, 34–39 (2001).
- [72] G. Ahlers, S. Grossmann, and D. Lohse, “Heat transfer and large scale dynamics in turbulent Rayleigh-Bénard convection”, *Rev. Mod. Phys.* **81**, 503 (2009).
- [73] D. Lohse and K.-Q. Xia, “Small-scale properties of turbulent Rayleigh-Bénard convection”, *Ann. Rev. Fluid Mech.* **42**, 335–364 (2010).
- [74] R. H. Kraichnan, “Turbulent thermal convection at arbitrary Prandtl number”, *Phys. Fluids* **5**, 1374–1389 (1962).
- [75] P. E. Roche, G. Gauthier, R. Kaiser, and J. Salort, “On the triggering of the ultimate regime of convection”, *New J. Phys.* **12**, 085014 (2010).
- [76] X. He, D. Funfschilling, H. Nobach, E. Bodenschatz, and G. Ahlers, “Transition to the ultimate state of turbulent Rayleigh-Bénard convection”, *Phys. Rev. Lett.* **108**, 024502 (2012).
- [77] D. P. M. van Gils, S. G. Huisman, S. Grossmann, C. Sun, and D. Lohse, “Optimal Taylor-Couette turbulence”, *J. Fluid Mech.* **706**, 118–149 (2012).
- [78] S. Grossmann and D. Lohse, “Multiple scaling in the ultimate regime of thermal convection”, *Phys. Fluids* **23**, 045108 (2011).
- [79] S. Grossmann and D. Lohse, “Logarithmic temperature profiles in the ultimate regime of thermal convection”, *Phys. Fluids* **24**, 125103 (2012).
- [80] S. Merbold, H. Brauckmann, and C. Egbers, “Torque measurements and numerical determination in differentially rotating wide gap Taylor-Couette flow”, *Phys. Rev. E* **87** (2013).
- [81] S. G. Huisman, C. Sun, and D. Lohse, “Statistics of turbulent fluctuations in counter-rotating Taylor-Couette flows”, *Phys. Rev. E* **88**, 063001 (2013).

- [82] S. G. Huisman, S. Scharnowski, C. Cierpka, C. Kähler, D. Lohse, and C. Sun, “Logarithmic boundary layers in strong Taylor-Couette turbulence”, *Phys. Rev. Lett.* **110**, 264501 (2013).
- [83] A. E. Perry and M. S. Chong, “On the mechanism of wall turbulence”, *J. Fluid Mech.* **119**, 173–217 (1982).
- [84] L. Rayleigh, “On the dynamics of revolving fluids”, *Proc. Roy. Soc. Lond. A* **93**, 148–154 (1917).
- [85] H. Ji and S. A. Balbus, “Angular momentum transport in astrophysics and in the lab”, *Physics Today* **66**, 27–33 (2013).
- [86] L. N. Trefethen, A. E. Trefethen, S. C. Reddy, and T. A. Driscoll, “Hydrodynamic stability without eigenvalues”, *Science* **261**, 578–584 (1993).
- [87] S. Grossmann, “The onset of shear flow turbulence”, *Rev. Mod. Phys.* **72**, 603–618 (2000).
- [88] D. Richard, “Instabilités hydrodynamiques dans les écoulements en rotation différentielle”, Ph.D. thesis, Univ. Paris 7 (2001).
- [89] B. Dubrulle, O. Dauchot, F. Daviaud, P. Y. Longaretti, D. Richard, and J. P. Zahn, “Stability and turbulent transport in Taylor–Couette flow from analysis of experimental data”, *Phys. Fluids* **17**, 095103 (2005).
- [90] H. Ji, M. Burin, E. Schartman, and J. Goodman, “Hydrodynamic turbulence cannot transport angular momentum effectively in astrophysical disks”, *Nature* **444**, 343–346 (2006).
- [91] E. Schartman, H. Ji, M. J. Burin, and J. Goodman, “Stability of quasi-Keplerian shear flow in a laboratory experiment”, *Astron. & Astrophys* **543** (2012).
- [92] M. S. Paoletti, D. P. M. van Gils, B. Dubrulle, C. Sun, D. Lohse, and D. P. Lathrop, “Angular momentum transport and turbulence in laboratory models of Keplerian flows”, *Astron. & Astrophys* **547**, A64 (2012).
- [93] M. Avila, “Stability and angular-momentum transport of fluid flows between corotating cylinders”, *Phys. Rev. Lett.* **108**, 124501 (2012).
- [94] H. T. Kim, P. Moin, and R. Moser, “Turbulence statistics in fully developed channel flow at low Reynolds number”, *J. Fluid Mech.* **177**, 133–160 (1987).
- [95] M. Bernardini, S. Pirozzoli, and P. Orlandi, “Velocity statistics in turbulent channel flow up to $Re_\tau = 4000$ ”, *J. Fluid Mech.* **742**, 171–191 (2014).

- [96] A. Lozano-Durán and J. Jiménez, “Effect of the computational domain on direct simulations of turbulent channels up to $Re_\tau = 4200$ ”, *Phys. Fluids* **26**, 011702 (2014).
- [97] E. A. Fadlun, R. Verzicco, P. Orlandi, and J. Mohd-Yusof, “Combined immersed-boundary finite-difference methods for three-dimensional complex flow simulations”, *J. Comput. Phys.* **161**, 35–60 (2000).
- [98] S. Ghosal, “An analysis of numerical errors in large-eddy simulations of turbulence”, *J. Comput. Phys.* **125**, 187–206 (1996).
- [99] A. G. Kravchenko and P. Moin, “On the effect of numerical errors in large eddy simulations of turbulent flows”, *J. Comput. Phys.* **131**, 310–322 (1997).
- [100] M. M. Rai and P. Moin, “Direct numerical simulation of turbulent flow using finite-difference schemes”, *J. Comput. Phys.* **96**, 15–53 (1991).
- [101] M. Duponcheel, P. Orlandi, and G. Winckelmans, “Time-reversibility of the Euler equations as a benchmark for energy conserving schemes”, *J. Comput. Phys.* **227**, 8736–8752 (2008).
- [102] A. W. Vreman and J. G. M. Kuerten, “Comparison of direct numerical simulation databases of turbulent channel flow at $Re_\tau = 180$ ”, *Phys. Fluids* **26**, 015102 (2014).
- [103] R. Verzicco and P. Orlandi, “A finite-difference scheme for three-dimensional incompressible flow in cylindrical coordinates”, *J. Comput. Phys.* **123**, 402–413 (1996).
- [104] G. Stringano and R. Verzicco, “Mean flow structure in thermal convection in a cylindrical cell of aspect-ratio one half”, *J. Fluid Mech.* **548**, 1–16 (2006).
- [105] R. Verzicco and R. Camussi, “Numerical experiments on strongly turbulent thermal convection in a slender cylindrical cell”, *J. Fluid Mech.* **477**, 19–49 (2003).
- [106] R. Verzicco and K. R. Sreenivasan, “A comparison of turbulent thermal convection between conditions of constant temperature and constant heat flux”, *J. Fluid Mech.* **595**, 203–219 (2008).
- [107] R. P. J. Kunnen, H. J. H. Clercx, B. J. Geurts, L. J. A. Bokhoven, R. A. D. Akkermans, and R. Verzicco, “A numerical and experimental investigation of structure function scaling in turbulent Rayleigh-Bénard convection”, *Phys. Rev. E* **77**, 016302 (2008).

- [108] R. J. A. M. Stevens, D. Lohse, and R. Verzicco, “Prandtl and Rayleigh number dependence of heat transport in high Rayleigh number thermal convection”, *J. Fluid Mech.* **688**, 31–43 (2011).
- [109] E. P. van der Poel, R. J. A. M. Stevens, and D. Lohse, “Comparison between two- and three-dimensional Rayleigh-Bénard convection”, *J. Fluid Mech.* **736**, 177–194 (2013).
- [110] E. P. van der Poel, R. Ostilla-Mónico, R. Verzicco, and D. Lohse, “Effect of velocity boundary conditions on the heat transfer and flow topology in two-dimensional rayleigh-bénard convection”, *Phys. Rev. E* **90**, 013017 (2014).
- [111] R. J. A. M. Stevens, R. Verzicco, and D. Lohse, “Sidewall effects in Rayleigh-Bénard convection”, *J. Fluid Mech.* **741**, 1–27 (2014).
- [112] R. Ostilla, R. J. A. M. Stevens, S. Grossmann, R. Verzicco, and D. Lohse, “Optimal Taylor-Couette flow: direct numerical simulations”, *J. Fluid Mech.* **719**, 14–46 (2013).
- [113] R. Ostilla-Monico, E. P. van der Poel, R. Verzicco, S. Grossmann, and D. Lohse, “Boundary layer dynamics at the transition between the classical and the ultimate regime of Taylor-Couette flow”, *Phys. Fluids* **26** (2014).
- [114] R. Courant, K. Friedrichs, and H. Lewy, “Über die partiellen Differenzengleichungen der mathematischen Physik”, *Mathematische Annalen* **100**, 32–74 (1928).
- [115] E. P. van der Poel, R. Verzicco, S. Grossmann, and D. Lohse, “Plume emission statistics in turbulent Rayleigh-Bénard paper”, in review (2015).
- [116] D. Lohse and F. Toschi, “The ultimate state of thermal convection”, *Phys. Rev. Lett.* **90**, 034502 (2003).
- [117] E. Calzavarini, C. R. Doering, J. D. Gibbon, D. Lohse, A. Tanabe, and F. Toschi, “Exponentially growing solutions of homogeneous Rayleigh-Bénard flow”, *Phys. Rev. E* **73**, R035301 (2006).
- [118] L. E. Schmidt, E. Calzavarini, D. Lohse, F. Toschi, and R. Verzicco, “Axially homogeneous Rayleigh-Bénard convection in a cylindrical cell”, *J. Fluid Mech.* **691**, 52–68 (2012).
- [119] E. Calzavarini, F. Toschi, and R. Tripiccone, “Evidences of Bolgiano-Obukhov scaling in three-dimensional Rayleigh-Bénard convection”, *Phys. Rev. E* **66**, 016304 (2002).

- [120] A. Parodi, J. von Hardenberg, G. Passoni, A. Provenzale, and E. A. Spiegel, “Clustering of plumes in turbulent convection”, *Phys. Rev. Lett.* **92**, 194503 (2004).
- [121] J. Schumacher, “Lagrangian dispersion and heat transport in convective turbulence”, *Phys. Rev. Lett.* **100**, 134502 (2008).
- [122] B. Gayen, G. O. Hughes, and R. W. Griffiths, “Completing the mechanical energy pathways in turbulent Rayleigh-Bénard convection”, *Phys. Rev. Lett.* **111**, 124301 (2013).
- [123] H. Brauckmann and B. Eckhardt, “Direct Numerical Simulations of Local and Global Torque in Taylor-Couette Flow up to $Re=30.000$ ”, *J. Fluid Mech.* **718**, 398–427 (2013).
- [124] R. Ostilla-Monico, E. P. van der Poel, Y. Yang, D. Lohse, and R. Verzicco, “A multiple resolutions strategy for direct numerical simulation of scalar turbulence”, in preparation, see Chapter 2 (2015).
- [125] M. J. Lee, B. D. Oh, and Y. B. Kim, “Canonical fractional-step methods and consistent boundary conditions for the incompressible navier-stokes equations”, *J. Comput. Phys.* **168**, 73–100 (2001).
- [126] N. Li and S. Laizet, “2DECOMP&FFT – a highly scalable 2D decomposition library and FFT interface”, in *Cray User Group 2010 Conference* (Edinburgh) (2010).
- [127] B. I. Shraiman and E. D. Siggia, “Heat transport in high-Rayleigh number convection”, *Phys. Rev. A* **42**, 3650–3653 (1990).
- [128] E. A. Spiegel, “Convection in stars”, *Ann. Rev. Astron. Astrophys.* **9**, 323–352 (1971).
- [129] R. Ostilla-Monico, E. P. van der Poel, R. Verzicco, S. Grossmann, and D. Lohse, “Exploring the phase diagram of fully turbulent Taylor-Couette flow”, *J. Fluid Mech.* **761**, 1–26 (2014).
- [130] R. Ostilla-Mónico, R. Verzicco, S. Grossmann, and D. Lohse, “The effect of Taylor rolls on highly turbulent Taylor-Couette flow”, in preparation, see Chapter 7 (2014).
- [131] Y. Yang, E. P. van der Poel, R. Ostilla-Monico, C. Sun, R. Verzicco, S. Grossmann, and D. Lohse, “Salinity transfer in bounded double diffusive convection”, submitted to *J. Fluid Mech.* (2014).

- [132] V. Spandan, R. Ostilla-Monico, and R. Verzicco, “Direct numerical simulations of two-phase Taylor-Couette turbulence”, in preparation (2015).
- [133] X. Zhu, R. Ostilla-Mónico, R. Verzicco, and D. Lohse, “Direct numerical simulations of Taylor-Couette flow with grooved cylinders”, in preparation (2015).
- [134] S. Arya, *Air Pollution Meteorology and Dispersion* (Oxford University Press) (1990).
- [135] W. R. Peltier and C. P. Caulfield, “Mixing efficiency in stratified shear flows”, *Ann. Rev. Fluid Mech.* **35**, 135–167 (2003).
- [136] F. P. Incropera, “Convection heat transfer in electronic equipment cooling”, *Journal of Heat Transfer* **110**, 1097–1111 (1988).
- [137] P. F. Linden, “The fluid mechanics of natural ventilation”, *Ann. Rev. Fluid Mech.* **31**, 201–238 (1999).
- [138] J. S. Turner, “Multicomponent convection”, *Annu. Rev. Fluid Mech.* **17**, 11–44 (1985).
- [139] R. W. Schmitt, “Double diffusion in oceanography”, *Annu. Rev. Fluid Mech.* **26**, 255–285 (1994).
- [140] O. Shishkina, R. J. A. M. Stevens, S. Grossmann, and D. Lohse, “Boundary layer structure in turbulent thermal convection and its consequences for the required numerical resolution”, *New J. Phys.* **12**, 075022 (2010).
- [141] J. Kim and P. Moin, “Application of a fractional-step method to incompressible Navier-Stokes equations”, *J. Comp. Phys.* **59**, 308–323 (1985).
- [142] B. I. Shraiman and E. D. Siggia, “Scalar turbulence”, *Nature* **405**, 639–646 (2000).
- [143] K. R. Sreenivasan, “On local isotropy of passive scalars in turbulent shear flows”, *Proc. R. Soc. London A* **434**, 165 (1991).
- [144] A. Pumir, “A numerical study of the mixing of a passive scalar in three dimensions in the presence of a mean gradient”, *Phys. Fluids* **6** (1994).
- [145] A. Pumir, “Structure of the three-point correlation function of a passive scalar in the presence of a mean gradient”, *Phys. Rev. E* **57** (1998).
- [146] A. Celani, A. Lanotte, A. Mazzino, and M. Vergassola, “Universality and saturation of intermittency in passive scalar turbulence”, *Phys. Rev. Lett.* **84**, 2385 (2000).

- [147] Z. Warhaft, “Passive scalars in turbulent flows”, *Ann. Rev. Fluid Mech.* **32**, 203–240 (2000).
- [148] F. Moisy, H. Willaime, J. Andersen, and P. Tabeling, “Passive scalar intermittency in low temperature helium flows”, *Phys. Rev. Lett.* **86**, 4827–4830 (2001).
- [149] J. Burgers, “A mathematical model illustrating the theory of turbulence”, *Adv. Appl. Mech.* **1**, 171–199 (1948).
- [150] R. J. A. M. Stevens, R. Verzicco, and D. Lohse, “Radial boundary layer structure and Nusselt number in Rayleigh-Bénard convection”, *J. Fluid Mech.* **643**, 495–507 (2010).
- [151] M. A. T. van Hinsberg, J. H. M. ten Thije Boonkkamp, F. Toschi, and H. J. H. Clercx, “Optimal interpolation schemes for particle tracking in turbulence”, *Phys. Rev. E* **87**, 043307 (2013).
- [152] A. Esser and S. Grossmann, “Analytic expression for Taylor-Couette stability boundary”, *Phys. Fluids* **8**, 1814–1819 (1996).
- [153] K. Sugiyama, E. Calzavarini, S. Grossmann, and D. Lohse, “Non-Oberbeck-Boussinesq effects in Rayleigh-Bénard convection: beyond boundary-layer theory”, *Europhys. Lett.* **80**, 34002 (2007).
- [154] S. Grossmann and D. Lohse, “Thermal convection for large Prandtl number”, *Phys. Rev. Lett.* **86**, 3316–3319 (2001).
- [155] J. Jeong and F. Hussain, “On the identification of a vortex”, *J. Fluid Mech.* **285**, 69–94 (1995).
- [156] H. J. Brauckmann and B. Eckhardt, “Intermittent boundary layers and torque maxima in Taylor-Couette flow”, *Phys. Rev. E* **87**, 033004 (2013).
- [157] T. B. Benjamin, “Bifurcation phenomena in steady flows of a viscous liquid”, *Proc. R. Soc. London A* **359**, 1–43 (1978).
- [158] S. G. Huisman, D. P. M. van Gils, and C. Sun, “Applying laser Doppler anemometry inside a Taylor-Couette geometry using a ray-tracer to correct for curvature effects”, *European J. Mech. B/Fluids* **36**, 115–119 (2012).
- [159] K.-Q. Xia, S. Lam, and S. Q. Zhou, “Heat-flux measurement in high-Prandtl-number turbulent Rayleigh-Bénard convection”, *Phys. Rev. Lett.* **88**, 064501 (2002).

- [160] F. H. Busse, “Viewpoint: The twins of turbulence research”, *Physics* **5**, 4 (2012).
- [161] B. Eckhardt, T. Schneider, B. Hof, and J. Westerweel, “Turbulence transition in pipe flow”, *Annu. Rev. Fluid Mech.* **39**, 447–468 (2007).
- [162] L. Prandtl, “Bericht über Untersuchungen zur ausgebildeten Turbulenz”, *Z. Angew. Math. Mech.* **5**, 136–139 (1925).
- [163] T. von Karman, *Mechanische Ähnlichkeit und Turbulenz* (Proc. Third Int. Congr. Applied Mechanics, Stockholm, Sweden, 1930, Weidmannsche, Berlin) (1930).
- [164] T. Wei and W. Willmarth, “Reynolds-number effects effects on the structure of a turbulent channel flow”, *J. Fluid Mech.* **204**, 57–95 (1989).
- [165] S. Grossmann, D. Lohse, and C. Sun, “Velocity profiles in strongly turbulent Taylor-Couette flow”, *Phys. Fluids* **26**, 025114 (2014).
- [166] G. Ahlers, E. Bodenschatz, D. Funfschilling, S. Grossmann, X. He, D. Lohse, R. Stevens, and R. Verzicco, “Logarithmic temperature profiles in turbulent Rayleigh-Bénard convection”, *Phys. Rev. Lett.* **109**, 114501 (2012).
- [167] H. Schlichting, *Boundary layer theory*, 7th edition (McGraw Hill book company) (1979).
- [168] R. C. A. van der Veen, S. G. Huisman, and C. Sun, “Twente Turbulent Taylor-Couette data”, Personal communication (2013).
- [169] X. He, D. Funfschilling, E. Bodenschatz, and G. Ahlers, “Heat transport by turbulent Rayleigh-Bénard convection for $Pr = 0.8$ and $4 \times 10^{11} \lesssim Ra \lesssim 2 \times 10^{14}$: ultimate-state transition for aspect ratio $\Gamma = 1.00$ ”, *New J. Phys.* **14**, 063030 (2012).
- [170] G. Ahlers, X. He, D. Funfschilling, and E. Bodenschatz, “Heat transport by turbulent Rayleigh-Bénard convection for $Pr = 0.8$ and $3 \times 10^{12} \lesssim Ra \lesssim 10^{15}$: aspect ratio $\Gamma = 0.50$ ”, *New J. Phys.* **14**, 103012 (2012).
- [171] C. D. Andereck, R. Dickman, and H. L. Swinney, “New flows in a circular Couette system with corotating cylinders”, *Phys. Fluids* **26** (1983).
- [172] M. V. R. Malkus, “The heat transport and spectrum of thermal turbulence”, *Proc. R. Soc. London A* **225**, 196–212 (1954).
- [173] S. G. Huisman, R. C. A. van der Veen, C. Sun, and D. Lohse, “Multiple states in ultimate Taylor-Couette turbulence”, *Nature Comm.* **5** (2014).

- [174] B. Martinez-Arias, J. Peixinho, O. Crumeyrolle, and I. Mutabazi, “Effect of the number of vortices on the torque scaling in Taylor-Couette flow”, *J. Fluid Mech.* **748**, 756–767 (2014).
- [175] R. Ostilla-Monico, R. Verzicco, S. Grossmann, and D. Lohse, “Turbulence decay towards the linearly-stable regime of Taylor-Couette flow”, *J. Fluid Mech.* **747**, 1–29 (2014).
- [176] R. Ostilla-Monico, S. G. Huisman, T. Jannink, D. P. M. van Gils, R. Verzicco, S. Grossmann, C. Sun, and D. Lohse, “Optimal Taylor-Couette flow: Radius ratio dependence”, *J. Fluid Mech.* **747**, 1–29 (2014).
- [177] H. Görtler, “Über den Einfluss der Wandkrümmung auf die Entstehung der Turbulenz”, *Z. Angew. Math. Mech.* **20**, 138–147 (1940).
- [178] K. C. Muck, P. H. Hoffmann, and P. Bradshaw, “The effect of convex surface curvature on turbulent boundary layers”, *J. Fluid Mech.* **161**, 347–369 (1985).
- [179] H. Görtler, “Über eine dreidimensionale Instabilität laminarer Grenzschichten an konkaven Wänden”, *Z. Angew. Math. Mech.* **21**, 250–252 (1940).
- [180] P. H. Hoffmann, K. C. Muck, and P. Bradshaw, “The effect of concave surface curvature on turbulent boundary layers”, *J. Fluid Mech.* **161**, 371–403 (1985).
- [181] M. Manna and A. Vacca, “Torque reduction in Taylor–Couette flows subject to an axial pressure gradient”, *J. Fluid Mech.* **639**, 373–401 (2009).
- [182] J. A. Sillero, J. Jiménez, and R. D. Moser, “One-point statistics for turbulent wall-bounded flows at Reynolds numbers up to $\delta^+ \approx 2000$ ”, *Phys. Fluids* **25**, 105102 (2013).
- [183] X. Wu and P. Moin, “A direct numerical simulation study on the mean velocity characteristics in turbulent pipe flow”, *J. Fluid Mech.* **608**, 81–112 (2008).
- [184] K. H. Bech, N. Tillmark, P. H. Alfredsson, and H. I. Andersson, “An investigation of turbulent plane Couette flow at low Reynolds numbers”, *J. Fluid Mech.* **286**, 291–325 (1995).
- [185] V. Avsarkisov, S. Hoyas, M. Oberlack, and J. P. García-Galache, “Turbulent plane Couette flow at moderately high Reynolds number”, *J. Fluid Mech.* **751**, R1–8 (2014).
- [186] S. Pirozzoli, M. Bernardini, and P. Orlandi, “Turbulence statistics in Couette flow at high Reynolds number”, *J. Fluid Mech.* **758**, 327–343 (2014).

- [187] H. Faisst and B. Eckhardt, “Transition from the Couette-Taylor system to the plane Couette system”, *Phys. Rev. E* **61**, 7227–7230 (2000).
- [188] P. Bradshaw, “The effects of streamline curvature on turbulent flow”, Technical Report, AGARDograph no. 169 (1973).
- [189] I. A. Hunt and P. N. Joubert, “Effects of small streamline curvature on turbulent duct flow”, *J. Fluid Mech.* **91**, 633–659 (1979).
- [190] H. Choi, P. Moin, and J. Kim, “Direct numerical simulation of turbulent flow over riblets”, *J. Fluid Mech.* **255**, 503–539 (1993).
- [191] T. H. van den Berg, S. Luther, D. P. Lathrop, and D. Lohse, “Drag reduction in bubbly Taylor-Couette turbulence”, *Phys. Rev. Lett* **94**, 044501 (2005).
- [192] J. Lu, A. Fernández, and G. Tryggvason, “The effect of bubbles on the wall drag in a turbulent channel flow”, *Phys. Fluids* **17**, 095102 (2005).
- [193] D. P. M. van Gils, D. Narez-Guzman, C. Sun, and D. Lohse, “The importance of bubble deformability for strong drag reduction in bubbly turbulent Taylor-Couette flow”, *J. Fluid Mech.* **722**, 317–347 (2013).
- [194] S. C. C. Bailey, M. Vallikivi, M. Hultmark, and A. J. Smits, “Estimating the value of von Kármán’s constant in turbulent pipe flow”, *J. Fluid Mech.* **749**, 79–98 (2014).
- [195] S. Hoyas and J. Jiménez, “Scaling of the velocity fluctuations in turbulent channels up to $Re_\tau = 2003$ ”, *Phys. Fluids* **18**, 011702 (2006).
- [196] J. Jiménez and R. Moser, “What are we learning from simulating wall turbulence?”, *Phil. Trans. R. Soc. Lond. A* **365**, 715–732 (2007).
- [197] N. Afzal and K. Yajnik, “Analysis of turbulent pipe and channel flows at moderately large Reynolds number”, *J. Fluid Mech.* **61**, 23–31 (1973).
- [198] J. C. del Álamo and J. Jiménez, “Spectra of the very large anisotropic scales in turbulent channels”, *Phys. Fluids* **15**, L41–L44 (2003).
- [199] I. Marusic, J. P. Monty, M. Hultmark, and A. J. Smits, “On the logarithmic region in wall turbulence”, *J. Fluid. Mech.* **716**, R3 (2013).
- [200] T. Tsukahara, H. Kawamura, and K. Shingai, “DNS of turbulent Couette flow with emphasis on the large-scale structure in the core region”, *J. Turb.* **7** (2006).

- [201] L. Shi, M. Rampp, B. Hof, and M. Avila, “A hybrid MPI-OpenMP parallel implementation for pseudospectral simulations with application to Taylor-Couette flow”, *Computers & Fluids* **106**, 1–11 (2015).
- [202] L. S. Tuckerman, “Taylor vortices versus Taylor columns”, *J. Fluid Mech.* **750**, 1–4 (2014).
- [203] N. I. Shakura and R. A. Sunyaev, “Black holes in binary systems. Observational appearance.”, *Astron. & Astrophys* **24**, 337–355 (1973).
- [204] E. P. Velikhov, “Stability of an ideally conducting liquid flowing between rotating cylinders in a magnetic field”, *Zhur. Eksptl'. i Teoret. Fiz.* **36** (1959).
- [205] S. A. Balbus and J. Hawley, “A powerful local shear instability in weakly magnetized disks”, *Astrophys. J.* **376**, 214–233 (1991).
- [206] C. F. Gammie, “Layered accretion in T-Tauri disks”, *Astrophys. J.* **457**, 355–362 (1996).
- [207] P. J. Armitage, “Dynamics of protoplanetary disks”, *Ann. Rev. Astron. Astrophys.* **49**, 195–236 (2011).
- [208] G. Lesur and P. Y. Longaretti, “On the relevance of subcritical hydrodynamic turbulence to accretion disk transport”, *Astron. & Astrophys* **444**, 25–44 (2005).
- [209] G. I. Taylor, “Motion of solids in fluids when the flow is not irrotational”, *Proc. R. Soc. Lond. A* **93**, 92–113 (1917).
- [210] J. Proudman, “On the motion of solids in a liquid possessing vorticity”, *Proc. R. Soc. Lond. A* **92**, 408–424 (1916).
- [211] D. Borrero-Echeverry, M. F. Schatz, and R. Tagg, “Transient turbulence in Taylor-Couette flow”, *Phys. Rev. E* **81**, 025301 (2010).
- [212] E. M. Edlund and H. Ji, “Nonlinear stability of laboratory quasi-Keplerian flows”, *Phys. Rev. E* **89** (2014).
- [213] E. M. Withjack and C. F. Chen, “An experimental study of Couette instability of stratified fluids”, *J. Fluid Mech.* **66**, 725–737 (1974).
- [214] H. Klahr and P. Bodenheimer, “Turbulence in accretion disks: Vorticity generation and angular momentum transport via the global baroclinic instability”, *The Astrophysical Journal* **508**, 869–892 (2003).

- [215] M. R. Petersen, K. Julien, and G. R. Stewart, “Baroclinic vorticity production in protoplanetary disks. I. Vortex formation”, *The Astrophysical Journal* **658**, 1236–1251 (2007).
- [216] B. Paczynski and G. Bisnovatyi-Kogan, “A model of a thin accretion disk around a black hole”, *Acta Astronomica* **31**, 283–291 (1981).
- [217] B. Gallet, C. R. Doering, and E. A. Spiegel, “Destabilizing Taylor–Couette flow with suction”, *Phys. Fluids* **22** (2010).
- [218] J. M. Lopez, F. Marques, and M. Avila, “The Boussinesq approximation in rapidly rotating flows”, *J. Fluid Mech.* **737**, 56–77 (2013).
- [219] F. Rincon, G. I. Ogilvie, and C. Cossu, “On self-sustaining processes in Rayleigh-stable rotating plane couette flows and subcritical transitions to turbulence in accretion disks”, *Astron. & Astrophys* **463**, 817–832 (2008).
- [220] S. Maretzke, B. Hof, and M. Avila, “Transient growth in linearly stable Taylor–Couette flows”, *J. Fluid Mech* **742**, 254–290 (2014).
- [221] T. Kuhn, *The structure of scientific revolutions: 50th Anniversary Edition*. (University of Chicago Press, Chicago, IL.) (2012).
- [222] O. Cadot, Y. Couder, A. Daerr, S. Douady, and A. Tsinober, “Energy injection in closed turbulent flows: Stirring through boundary layers versus inertial stirring”, *Phys. Rev. E* **56**, 427–433 (1997).
- [223] T. H. van den Berg, C. Doering, D. Lohse, and D. Lathrop, “Smooth and rough boundaries in turbulent Taylor–Couette flow”, *Phys. Rev. E* **68**, 036307 (2003).
- [224] J. R. R. Tolkien, *The Lord of the Rings* (George Allen & Unwin, United Kingdom) (1954).

Summary

A systematic study of Taylor-Couette flow, the flow in a fluid layer between two coaxial and independently rotating cylinders was performed with direct numerical simulations. First, the numerical strategies are detailed. A highly parallelized code, specific to the high Reynolds number regime was developed. The code was found to have adequate performance up to 64K cores on a supercomputer. A multiple resolutions in space and time strategy, to simulate scalar fields in turbulence was also developed, and implemented into the code. The use of this multiple resolutions strategy was found to be advantageous not only when the diffusivity of the scalar is smaller than the momentum diffusivity, but also for similar diffusivities, reducing computational costs by up to a factor sixty. Scalar fields were found to sustain more pronounced gradients, due to the absence of pressure from the diffusion-advection equations. Using the code, the four dimensions of the TC flow parameter space were explored. Effective scaling laws relating the shear driving (expressed as a Taylor number Ta) with the torque (expressed as a quasi-Nusselt number Nu_ω) and other system responses were revealed. An “optimal” rotation ratio, for which Nu_ω is maximum at a given Ta , previously seen in experiments, was also found in the simulations. The optimal rotation rate was found to depend on the shear, and on the radius ratio between inner and outer cylinder. A formulation based on a rotating frame, where the outer cylinder rotation is reflected as a Coriolis force with magnitude Ro on the flow was developed, and using this rotating frame an explanation for the moving optimum was elucidated. The radius ratio dependence of the optimum was also analysed, both experimentally and numerically. Experimental results yielded a value for the optimal rotation ratio in conflict with existing theoretical predictions. The Ta dependence of TC flow was analysed in detail, and three regimes were found: (i) a flow dominated by large coherent structures, called Taylor vortices, which completely drive the meridional flows (ii) an intermediate transitional regime, where the presence of a mean flow allows for the coexistence of laminar and turbulent dynamics and (iii) a flow with fully developed turbulence, where some remnants of the Taylor vortices, called turbulent Taylor vortices were found to remain. In this third regime, also known as the “ultimate”-regime, scaling laws relating Nu_ω and Ta with a one-to-one correspondence to Rayleigh-Bénard flow, the flow in a fluid layer heated from below and cooled from above were found. These scaling laws are expected to hold up to arbitrarily large drivings. The transition to this ultimate regime was explored throughout the four dimensions of the parameter space, and found to depend significantly on η , but not on Ro and Γ , the aspect ratio of the system. Based on the magnitude of the Coriolis force, two main regimes were distinguished: (i)

co-rotating and weakly counter-rotating cylinders, dominated by Rayleigh-unstable dynamics and (ii) strongly counter-rotating cylinders, where a mixture of stable and unstable zones exist. An analogy between outer-cylinder rotation and η was found, namely that smaller gaps behave like a wider gap with co-rotating cylinders, and that wider gaps behave like smaller gaps with weakly counter-rotating cylinders. A phase diagram, which can be seen as an extension of the Andereck *et al.* diagram towards the ultimate regime was presented. A detailed analysis of the Turbulent Taylor vortices was performed, where the vortices were found to be attached to the boundary layer, and to dominate the transport of angular velocity. Furthermore, they could be weakened by curvature, but imposing a weak axial flow was found to convect them, without weakening them. Evidence for the development of an overlap layer, consistent with the logarithmic behaviour in the boundary layers was found. The effect of the computational boxes on the statistics was also examined. The axial extent of the domain was found to play a role, due to the underlying strong axial dependence of TC flow, while the effect of the azimuthal extent was milder. Finally, the stability of sub-critical TC flow was examined in the rotating frame, in an attempt to probe the hydrodynamic stability of astrophysical accretion disks. For this, simulations were initialized in the super-critical parameter range, and a sudden, stabilizing rotation was added. The system behaved like a damped oscillator, and turbulence was found to decay, regardless of Reynolds number. The system was found to absorb energy while the turbulence decayed.

Samenvatting

Deze systematische studie van Taylor-Couette (TC) stroming, de stroming in een vloeistoflaag tussen twee coaxiale en onafhankelijk draaiende cilinders, is uitgevoerd met directe numerieke simulaties. Als eerst wordt de numerieke methode uitgelegd. Deze methode bestaat uit een uiterst parallele code die specifiek voor hoge Reynolds getallen is ontwikkeld. Het is aangetoond dat de code adequaat tenminste 64K cores tegelijk op een supercomputer kan gebruiken. Om scalaire velden in turbulentie te simuleren is er een strategie ontwikkeld die gebruik maakt van verschillende resoluties in ruimte en tijd voor het snelheid en scalaire veld. Het gebruiken van deze strategie is niet alleen voordelig wanneer de diffusiviteit van de scalar kleiner is dan die van de snelheid; ook wanneer deze gelijk zijn, is een reductie van de computationele intensiviteit met een factor zestig te realiseren. Dit komt omdat zelfs dan scalaire velden hogere gradienten hebben door de afwezigheid van de druk term in de diffusie-advectie vergelijking. Met behulp van de code zijn de vier dimensies van de TC parameter ruimte onderzocht. Effectieve schalingswetten met betrekking tot de schering (uitgedrukt als een Taylor getal Ta) met het koppel (uitgedrukt als een quasi-Nusselt getal Nu_ω) en andere systeem reacties zijn geopenbaard. Een 'optimale' omloopsnelheid waarvoor Nu_ω maximaal is bij een gegeven Ta , eerder gezien in experimenten, is ook gevonden in de simulaties. De optimale rotatiesnelheid bleek afhankelijk van de schering en de straal verhouding tussen binnenste en buitenste cilinder η . Een formulering gebaseerd op een roterend referentiestelsel, waarbij de rotatie van de buitenste cilinder wordt uitgedrukt als een Corioliskracht met grootte Ro is ontwikkeld. Deze formulering is gebruikt om een verklaring voor het verplaatsende optimum te vinden. De afhankelijkheid het optimum op de straalverhouding is ook geanalyseerd, zowel numeriek als experimenteel. Experimentele resultaten leverden waarden voor de optimale rotatie verhouding op die in conflict zijn met bestaande theoretische voorspellingen. De Ta afhankelijkheid van TC stroming is in detail geanalyseerd en daarbij zijn drie regimes gevonden: (i) een stroming gedomineerd door grote coherente structuren, genaamd Taylor-vortices, die volledig de meridionale stromingen aandrijven (ii) een overgangsgebied, waar de aanwezigheid van een tijdsgemiddelde stroming zorgt voor naast elkaar bestaande laminaire en turbulente dynamiek en (iii) een stroming met volledig ontwikkelde turbulentie, waarin de Taylor-vortices voortduren en turbulent zijn. In dit derde regime, ook bekend als het "ultieme"- regime, is de schalingswet van Nu_ω as functie van Ta n-op-n overeenkomend met die van Rayleigh-Bénard convectie; de stroming in een van onder verwarmde en van boven gekoelde vloeistoflaag. Deze schalingswetten zullen naar verwachting tot willekeurig grote aandrijving valide blijven. De over-

gang naar dit asymptotische regime is onderzocht in de vier dimensionale controle-parameterruimte en is substantieel afhankelijk van η , maar niet van Ro en Γ . Hier is Γ de verhouding tussen de breedte en hoogte van het systeem. Op basis van de grootte van de Coriolis kracht, kunnen er twee regimes onderscheiden worden. Namelijk, een regime met (i) co-roterende en zwak tegengesteld draaiende cilinders, gedomineerd door Rayleigh-instabiele dynamiek en een regime met (ii) sterk tegengestelde draaiende cilinders, waarin een mengsel van stabiele en instabiele zones bestaat. Er is een analogie tussen de rotatie van de buitenste cilinder en η gevonden. Systemen met kleinere openingen gedragen zich als die met een bredere opening maar dan met meedraaiende cilinders, en die met grotere openingen gedragen zich als die met kleinere openingen maar dan met zwak tegendraaiende cilinders. Een uitbreiding van het TC fasediagram van Andereck *et al.* naar het ultieme regime is gepresenteerd. Een gedetailleerde analyse van de turbulente Taylor-vortices is uitgevoerd, waaruit bleek dat de wervels aan de grensslaag vast zitten en ook het transport van de hoeksnelheid domineren. Bovendien kunnen zij worden verzwakt door kromming, maar een zwakke axiale drukgradient bleek ze alleen maar te verplaatsen en niet af te zwakken. Bewijs voor de ontwikkeling van een overlappende laag, in overeenstemming met het logaritmische gedrag in klassieke grenslagen, is gevonden. De invloed van het computationele domein op de statistieken is eveneens bestudeerd. De axiale lengte van het domein bleek een significante rol te spelen via de onderliggende sterke afhankelijkheid van de statistiek op de axiale stroming. Het effect van de azimuthale lengte is veel kleiner gebleken. Tenslotte is de stabiliteit van sub-kritische TC stroming bestudeerd in het roterende referentiestelsel in een poging om de hydrodynamische stabiliteit van astrofysische accretieschijven te verduidelijken. Hiervoor zijn simulaties genitaliseerd met velden uit een super-kritische TC simulatie, die plotseling aan een stabiliserende rotatie zijn onderworpen. Deze systemen gedroegen zich als een gedempte oscillator en de turbulentie bleek te vervallen, ongeacht de hoogte van het Reynolds getal. Ook bleek het systeem energie te absorberen terwijl de turbulente zwakker werd.

Publications

Published in refereed journals:

- R. Ostilla, R. J. A. M. Stevens, S. Grossmann, R. Verzicco, D. Lohse. *Optimal Taylor-Couette flow: direct numerical simulations*. Journal of Fluid Mechanics, **719**, 14-46. (2013) – see §4.
- R. Ostilla-Mónico, E. P. van der Poel, R. Verzicco, S. Grossmann, D. Lohse. *Boundary layer dynamics at the transition between the classical and the ultimate regime of Taylor-Couette flow*. Physics of Fluids, **26**, 015114, (2014) – see §6.
- R. Ostilla-Mónico, S. G. Huisman, T. J. G. Jannink, D. P. M. van Gils, R. Verzicco, S. Grossmann, C. Sun, D. Lohse. *Optimal Taylor-Couette flow: radius ratio dependence*. Journal of Fluid Mechanics, **747**, 1-29, (2014) – see §5.
- R. Ostilla-Mónico, R. Verzicco, S. Grossmann, D. Lohse. *Turbulence decay towards the linearly-stable regime of Taylor-Couette flow*. Journal of Fluid Mechanics, **748**, R3. (2014) – see §10.
- R. Ostilla-Mónico, E. P. van der Poel, R. Verzicco, S. Grossmann, D. Lohse. *Exploring the phase diagram of fully turbulent Taylor-Couette flow*. Journal of Fluid Mechanics, **761**, 1-26. (2014) – see §7.
- E. P. van der Poel, R. Ostilla-Mónico, R. Verzicco, D. Lohse. *Effect of velocity boundary conditions on the heat transfer and flow topology in two-dimensional Rayleigh-Bénard convection*. Physical Review E, **90**, 013017. (2014)
- Y. Yang, R. Ostilla-Mónico, J. Z. Wu, P. Orlandi. *Inertial waves and mean velocity profiles in a rotating pipe and a circular annulus with axial flow*. Physical Review E, *in press*.

Under review:

- E. P. van der Poel, R. Ostilla-Mónico, J. Donners, R. Verzicco. *A pencil distributed finite difference code for strongly turbulent wall-bounded flows*. – see §2.
- R. Ostilla-Mónico, R. Verzicco, D. Lohse. *Effects of the computational domain size on DNS of Taylor-Couette turbulence*. – see §8.

- R. Ostilla-Mónico, R. Verzicco, S. Grossmann, D. Lohse. *The effect of Taylor rolls on Taylor-Couette flow.* – see §9.
- R. Ostilla-Mónico, E. P. van der Poel, R. P. J. Kunnen, R. Verzicco, D. Lohse. *Geostrophic convective turbulence: The effect of boundary layers.*
- Y. Yang, E. P. van der Poel, R. Ostilla-Mónico, C. Sun, R. Verzicco, S. Grossmann, D. Lohse. *Salinity transfer in bounded double diffusive convection.*

Acknowledgements

In my understanding, the acknowledgements is the part which is most telling about the person behind the thesis, and also the most read by everyone. On the personal side, I do not wish to simply write a list of names, where everyone would barely mentioned in passing, and some maybe even forgotten by accident, as this would not do them justice. In a motivating lecture by Patrice Le Gal during the ICTW in Twente, he mentioned that people are like keys. They allow you to open certain doors, which you sometimes did not know they were there. I feel that during my PhD I have opened many doors in myself, and that is due to the people which I have had around me. Even if I like to spend time alone, I feel like I would not be able to be who I am without those people who form part of my everyday life. I strongly believe that the western idea of individuality, and that the human can be understood as an entity separated from his surroundings and his friends and social group is a lie. Instead, I believe that these things are an integral part of who we are. If you have gotten this far into reading the thesis, you probably are in this group of people, so thank you. I stand by the line of Bilbo Baggins on his 111th birthday “I don’t know half of you half as well as I should like; and I like less than half of you half as well as you deserve” [224].

On the scientific side, I do feel obligated to mention a few people by name. Detlef, for giving me the opportunity to work at the Physics of Fluids group, and giving me increasing levels of responsibility, which has allowed me to gain much more experience than I would ever imagine during a PhD. I hope you didn’t feel disappointed at any point, even if I complained a lot about many things, and you had to suffer my humour and cynicism, which sometimes can be hard to take. I know I learnt a lot about how to do and present science in your office and through the iPad-annotated PDFs of my drafts. I appreciate that these skills are very hard to teach.

Roberto, for all the support and discussions during the years. You have really taught me a lot, not only about numerics, but also about saviour-faire in the scientific world, and also in the world in general. When I thought solving a problem was impossible, somehow you always knew a very simple solution to it. I know that I will keep many, if not all of the lessons from these years for the rest of my life.

Chao, thank you for sitting in my committee, and also for the endless discussions about TC and RB. Working with you on the JFM paper was an intense learning experience, and helped me understand how to connect with experimentalists. Prof. Grossmann, for all the plentiful discussions full of physics during the PhD. I will always remember his motivating speech at the end of the Taylor-Couette workshop held in Enschede, and our posterior chat about the nature of the Navier-Stokes equa-

tions and how turbulence arises from them, which came at a time I was very frustrated with everything. And the other “sophos” from the group, Leen. I really enjoyed the discussions during the drinks and the defences about philosophy of science, history of fluid mechanics, the French, the British and Camus. Devaraj, speaking to you always gives me a different perspective on life, and for the moral and ethical advice throughout my PhD. Andrea, for some discussions and valuable insight into numerics and fluid mechanics in general. Jacco, for the advices on teaching and the nice discussions during the Copenhagen summer school.

I also have to thank the all my numerical colleagues during the years. Erwin, I really enjoy working with you, even if we think very differently on many issues. I think half, or more than half, of the work in this thesis would not have been possible without you. Thanks for proofreading all my work, as far as even doing so for the Dutch summary of this thesis. Yantao, it’s always a pleasure to work with you, your innate carefulness has saved me many problems. I hope that we will keep on collaborating in the future. Dennis Bakhuis, my first (and only) master student, I don’t think you appreciate how much I learnt while working with you, my graphs will never be the same. Janis, even if your stay here was short, I learnt quite a bit from you during your time here, I hope things work out better in the future. Vamsi and Xiaojue, the task to carry on the work is now yours, good luck with the science, and the non-scientific problems coming up! Richard, thanks for helping me getting started with everything my first year. My gratitude also goes to the TC experimentalists, Dennis van Gils, Sander H., Roeland and Ruben for all the discussions during the years. I also want to thank Sander W. for proofreading some articles, for discussions which made me question the fundamental assumptions behind my research, and getting me to reason about things which I took for granted. Finally, a big thank you to Bas and the tireless Joanita for the endless help and assistance provided.

I would also like to thank all the members of the committee for spending their time to come to Enschede and to sit there during my defense. I must also mention two people from outside the group: Sebastian Merbold and Hannes Brauckmann. Sebastian, for inviting us to Cottbus in November 2012, that was definitely a turning point in my PhD. And Hannes, thank you for all the discussions during the years. The situation could have made us antagonists, but I prefer to think after the Cottbus discussions it made us scientific colleagues. To end, I once again thank those countless people with which I’ve had discussions with throughout the years, be it in conferences, or in visits to their group (Chicago, Cottbus and Madrid, thanks for the hospitality). It is impossible to mention everyone, without omitting a single name, so I will stop here. You have gave me many ideas and insights which have influenced my work. Thanks also to ERC and PRACE for funding and computing resources, respectively. And to end I thank in a very sincere way all the people at SurfSARA for the infinite help during the years.

The life of Rodolfo

*Soy América y Europa, mi raíz es bereber.
Soy Atlántico y simiente, soy alisio, sal y miel.
Soy de aquí, soy de una tierra, donde el mar cautiva el Sol...
Con el Mundo por frontera, ser canario es mi razón.*

—BENITO CABRERA – SOY DE AQUÍ

*I might not have gone where I intended to go, but I think I have ended up where I
needed to be.*

—DOUGLAS ADAMS

After a series of convoluted journeys around the planet by both clans, a male member of the Ostilla family called Roberto and a female member of the Mónico family called Fani met each other in Venezuela. They moved to the Canary Islands, Spain, and had two children. The elderly child, a boy, was born on August 1st 1988. According to legend, official government records have his name written down as Rodolfo. However, the real spelling of his name remains concealed, and over 300 variations have been employed during his life. He believes deep inside that if he ever becomes famous enough, there will probably be a wikipedia page on all the diverse spellings. The younger child, born in 1995, was female. Her parents decided to spare her a fate like her brother, and called her Daniela.

For the next years, Rodolfo lived 20 metres from the beach, and 200 metres from the hospital where he was born, in Puerto de la Cruz. The influence of water and the sea was very important in his every day life. He really enjoyed watching and listening to the waves, which probably explains why he eventually became a fluid scientist, even if his simulated waves do not provide the same sensations. He went to the Trinity School from three to seventeen, where he developed a passion for correcting and annoying teachers. He has some very fond memories of this time, especially of some of the teachers which went way out of their to feed his intellectual curiosity.

When choosing a university, he decided not to go too far away from home, but to still leave the fortunate islands. He chose Sevilla in the south of Spain to study aerospace engineering, with the hope of designing a space ship some day. This was because one of his childhood dreams was setting up an interplanetary travel agency. A traumatic childhood experience involving the November 92 coup d'état in Caracas also probably played a role, where the palace beside his house was bombed by aeroplanes, and he spent the next months drawing buildings engulfed in flames. Rodolfo

really enjoyed his time there. In Sevilla, he took up juggling, he made really good friends, and played bass guitar in a couple of concerts. He was very puzzled as he found out that he was African, and had a weird Spanish accent. He thought this remarkable, as it seemed to him that the people around him were the ones speaking weird, after all, everybody in his life had spoken like he did. The exams were hard and many of the teachers were very strict. He of course, continued with his regular art of annoying questions, improving these skills every day. Skills he presently uses in every Physics of Fluids seminar and Masters defence ceremonies. He participated in a rocket building society, which was the first time he felt he could use his knowledge to build something. However, he found out that building a space ship involves too much paper work, and got bored with a CubeSat project. Instead, he was attracted by the myth surrounding the fluid dynamics course, and turbulence in particular, which eventually would seal this book's fate.

He decided to leave Sevilla after four years to pursue an MSc in Aerodynamics in Cranfield University, UK. Cranfield was not the best of times for Rodolfo, as he did not enjoy the atmosphere of the university, and their attitude to tackling problems. He also seemed incapable of facing the British weather even after months of intense Monty Python and his natural cynicism. He considers himself very lucky to have had a caring and motivated supervisor for his MSc thesis, and to have made very good friends, which alleviated the stay there.

Moving back to Madrid for a six month internship at the European Space Astronomy Center, part of ESA, Rodolfo finally fulfilled the dream of looking at the stars and being paid for it. After these six months, and seeing his disappointment with his MSc, a soul under the name of José Gordillo put him in contact with Detlef Lohse, something for which he will always be thankful, and thus began his move to Enschede. He regards his PhD student period as a great period in his life, where he grew a lot, both scientifically and as a person. He regularly attended the juggling society, who made it easier for him to feel home in Enschede and helped him learn the local language. Paradoxically, he is a Spanish guy who learnt how to scuba dive and windsurf while living in the Netherlands. He regards his inability to ride the unicycle properly as the greatest failure of his PhD. It is unclear what the future has in mind for our absurd hero, but he will always remember to look at the bright side of life.

*Caminante, son tus huellas el camino y nada más;
Caminante, no hay camino, se hace camino al andar.
Al andar se hace el camino, y al volver la vista atrás
se ve la senda que nunca se ha de volver a pisar.
Caminante no hay camino, sino estelas en la mar.*

—ANTONIO MACHADO