

Spatio-Temporal Dynamics in Pipe Flow

by

David Christopher Moxey

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Mathematics Institute

June 2011



Contents

Li	st of I	Figures		iv
Li	st of]	[ables]		ix
Ao	cknow	ledgem	ents	x
De	eclara	tion		X
AJ	bstrac	t		xii
AI	brevi	ations		xiii
1	Intr	oductio	n	1
	1.1	Existin	g studies	3
	1.2	Spatio	-temporal intermittency	5
	1.3	Outlin	2	7
2	Con	putatio	onal Techniques	10
	2.1	Metho	d of Weighted Residuals	11
		2.1.1	Collocation Formulation	13
		2.1.2	Galerkin Formulation	15
	2.2	Spectra	al/hp Element Methods	17
		2.2.1	Diffusion Equation	18
		2.2.2	Constructing local modes	21
		2.2.3	Common choices for the expansion basis	26
		2.2.4	Elemental Operations	30

		2.2.5	Implementation	32
		2.2.6	Convergence Properties	35
	2.3	Tempo	oral Discretisations for Navier-Stokes	36
		2.3.1	Classical high-order schemes	37
		2.3.2	Splitting Schemes	38
		2.3.3	First-order splitting schemes for Navier-Stokes	39
		2.3.4	High-order schemes	41
		2.3.5	Pressure boundary condition	42
		2.3.6	Stiffly stable schemes	43
	2.4	Summ	ary	45
3	Nun	nerical	Methods for Pipe Flow	46
	3.1	Mathe	matical Framework	47
		3.1.1	Boundary conditions	49
		3.1.2	Hagen-Poiseuille Flow	50
	3.2	Forcin	g with Volumetric Flux	52
		3.2.1	Constant body forcing	54
		3.2.2	Forcing for transition experiments	55
		3.2.3	Direct flux condition	56
		3.2.4	Green's function flux forcing	58
		3.2.5	The Semtex Spectral Element Solver	61
		3.2.6	Implementation of volumetric flux forcing	62
	3.3	Multic	ore optimisations for parallel Fourier transforms	65
		3.3.1	Semtex implementation	66
		3.3.2	Multi-core systems	68
		3.3.3	Testing and results	70
	3.4	Summ	ary	72
4	Trai	nsitiona	l Dynamics of Pipe Flow	73
	4.1	Resolu	tion requirements	74
		4.1.1	Chebyshev-based meshes	75
		4.1.2	Spectral-element discretisation	76
		4.1.3	Resolution testing	79

Bi	bliogr	aphy		141
6	Sum	mary		139
	5.5	Conclu	usion and Discussion	137
	5.4	Critica	ll point	134
		5.3.2	Universality and connections to pipe flow	131
		5.3.1	Bond percolation	129
	5.3	Direct	ed percolation	128
		5.2.2	Results	124
		5.2.1	Methodology	120
	5.2	Puff sp	plitting	118
		5.1.2	Current state of lifetime statistics	117
		5.1.1	Lifetime statistics	114
	5.1	Decay	studies	111
5	Spre	ading t	urbulence in pipe flow	110
	4.6	Conclu	usions	108
		4.5.2	Results	106
		4.5.1	Methodology	104
	4.5	Patterr	n formation analysis	103
		4.4.4	Results	101
		4.4.3	Methodology	100
		4.4.2	Laminar lengths	98
		4.4.1	Intermittency factor	96
	4.4	Onset	of intermittency	96
		4.3.2	Results	91
		4.3.1	Methodology	91
	4.3	Transi	tion between localised and intermittent turbulence	91
		4.2.3	Results	89
		4.2.2	Identifying turbulence	87
		4.2.1	Simulation methodology	86
	4.2	Transi	tional Dynamics	83

List of Figures

Contours of the streamwise velocity through circular cross-sections of a pipe.	
The left-hand image depicts a laminar flow in which the fluid is steady and	
unchanging in time. This contrasts with the right-hand image, which is a	
snapshot of a highly fluctuating turbulent flow at $Re = 3,000.$	2
Visualisations of turbulent and intermittent flow in a cross-section of an	
$L = 25D$ pipe. Red indicates high values of the vorticity $\boldsymbol{\omega} = \nabla \times \boldsymbol{u}$ and	
blue low values.	5
Laminar-turbulent bands in experiments of Couette flow at $Re = 358$, taken	
from Prigent <i>et al.</i> (2003)	6
Global and local expansion bases Φ_k and ϕ_k for a three-element finite	
element expansion. χ^e is the linear projection mapping the standard element	
$\Omega_{\rm st}$ to Ω^e	22
Typical modal and nodal expansion bases defined inside the standard element	
Ω_{st}	27
Elemental mass matrix \mathbf{M}^{e} and Laplacian matrix \mathbf{L}^{e} for different choices of	
elemental modes ϕ_p at polynomial order $P = 20$. Dots indicate non-zero	
entries. (a) The most common choice for a modal expansion $\phi_p = \psi_p$. (b)	
The result of changing the interior modes seen in (a) to $J_{p-1}^{3,3}(\xi)$.	30
Solution to the diffusion equation using the parameters outlined in section	
2.2.5 with the initial condition $u(x) = x(1-x)$ for $(x, t) \in \Omega \times [0, 1]$. The	
distribution of Gauss-Lobatto-Legendre (GLL) points inside each element	
can be seen along the <i>x</i> -axis	33
Logarithmic plot of the error E_f measured using the energy norm for h - and	
<i>p</i> -type convergence (top and bottom lines respectively)	36
	Contours of the streamwise velocity through circular cross-sections of a pipe. The left-hand image depicts a laminar flow in which the fluid is steady and unchanging in time. This contrasts with the right-hand image, which is a snapshot of a highly fluctutating turbulent flow at $Re = 3,000$ Visualisations of turbulent and intermittent flow in a cross-section of an $L = 25D$ pipe. Red indicates high values of the vorticity $\boldsymbol{\omega} = \nabla \times \boldsymbol{u}$ and blue low values

3.1	Average velocity profiles for laminar (left) and turbulent (right) flows through	
	the cylindrical pipe, where averaging is performed over the axial and az-	
	imuthal directions. These profiles are normalised so that in each case the	
	flow rates $\int_C \boldsymbol{u} \cdot \mathrm{d}\boldsymbol{S}$ are equal	51
3.2	Schematic diagrams depicting flow being driven using (a) constant pressure	
	gradient and (b) constant volumetric flux.	53
3.3	Friction factor λ for fully developed pipe flow at a variety of Re. Dashed	
	line indicates the laminar Hagen-Poiseuille friction law $\lambda = 64/\text{Re}$ and the	
	solid line the Blasius law $\lambda = 0.3164 \text{Re}^{-1/4}$. Data points are taken from	
	the experimental results of McKeon et al. (2004) and show the transition	
	between the two laws	55
3.4	Geometry of the backwards facing step, for which the Stokes field possesses	
	a singular point at the step edge.	60
3.5	Exchange procedure for $Q = 2$ processors over $N_x = 6$ planes	67
3.6	Shared memory implementation of the algorithm of figure 3.5 for two nodes,	
	each possessing two processes. The first step involves each process perform-	
	ing the same local block transpose, where the blocks are written into a large	
	shared memory block and interleaved with the data from the other processor.	
	A single process on each machine then performs the all-to-all transpose	69
3.7	Execution time in milliseconds as a function of the number of processors P	
	for the original Semtex exchange (blue) and new multi-core exchange (red).	71
4.1	Examples of different types of mesh construction for the two-dimensional	
	circular cross-section of the pipe. Spectral-element meshes use polynomial	
	order $P = 12$. Elements which are filled in indicate the placement of GLL	
	nodal points	77
4.2	Trace of the axial velocity $u(0, t)$ for the instigation of turbulence in a pipe	
	of length $L = 20D$. The figure is split up into three regions: (a) initial	
	deformation of the turbulent profile; (b) onset of instability; (c) transition to	
	statistically steady turbulent flow	80
4.3	L^{∞} error of the near-wall axial velocity u_z^+ for various polynomial orders P	
	for four candidate meshes.	81

- 4.6 Dynamics of transitional flow from simulations of a L = 125D pipe. The central plot shows a space-time diagram with the streamwise direction horizontal and time increasing vertically upwards. The value of Re changes as indicated on the right. The magnitude of the transverse velocity q is sampled along the axis of the pipe and visualised in a frame moving at the average fluid velocity: $q(x - \overline{U}t, t)$. Colours are such that light corresponds to turbulent flow and black corresponds to laminar flow. Below and above are flow visualizations in vertical cross-sections through the pipe at the initial and final times, and over the 25D streamwise extents indicated by arrows.

4.7 Simulations highlighting the difference between localised turbulence at Re = 2,250 (left) and spatio-temporal intermittency at Re = 2,350 (right). 92

- 4.8 Streamwise velocity u and y-component of the transverse velocity v for the final state of figure 4.7 at Re = 2,250 (top) and Re = 2,350 (bottom). . . . 93

- 4.12 Representative space-time plots of the intermittency function $I(x, t; 8 \times 10^{-3})$ at a selection of three Reynolds numbers in a pipe of length L = 25D.103

4.13	Normalised single-point velocity pdfs for $Re = 2,350, 2,400, 2,450, 2,500$,	
	2,600, 2,800 and 3,000, as indicated by labels and alternating line types. For	
	$\text{Re} \geqslant 2{,}700,$ the distributions are nearly Gaussian, and for $\text{Re} \leqslant 2{,}600$ the	
	distributions are non-Gaussian	104
4.14	Above (a-c): Two-dimensional approximations of the complex pdf $\rho_1(r, \phi)$	
	of first-order Fourier modes. <i>Below</i> : Cross-sections of $\rho_1(r, \phi)$ at $\phi = 0$ for	
	Re = 2,200 (bottom curve), 2,300, 2,400, 2,500 and 2,800 (top curve), as	
	indicated by arrows and alternating line types	107
5.1	Spatio-temporal dynamics of puff splitting in a $L = 100D$ domain at	
	Re = 2,350. The turbulence intensity $q(x, t)$ as defined in equation (4.1) is	
	visualised in a frame of reference moving at speed $c = 0.914\overline{U}$. The white	
	dashed line indicates the splitting time detected by the criterion of section	
	5.2.1	119
5.2	Splitting dynamics in a pipe of length $L = 200D$ at Re = 2,350. Turbulence	
	intensity $q(x, t)$ as defined in equation (4.1) is visualised in a frame moving	
	at $c = 0.924\overline{U}$.	121
5.3	Snapshots of the streamwise vorticity ω_x through cross-sections of an $L =$	
	50D pipe at Re = 2,350 as a puff splits. Positive and negative vorticity are	
	denoted by blue and red respectively; black areas represent $\omega_x = 0$. At $t = 0$	
	(bottom image), Re is changed from 2,100 to 2,350, and further snapshots	
	are taken at $t = 50, 138, 198$ with the snapshot at $t = 438$ showing the final	
	two-puff configuration.	122
5.4	Schematic drawings depicting criterion used to determine splitting times.	
	Shaded rectangles indicate the region around the primary peak of the signal	
	to be ignored, of width $d_{\min} = 20D$	123
5.5	Initial condition number i against detected splitting time t_i for samples	
	obtained when Re increased from 2,150 to 2,350	125

5.6	Survival functions $S(t)$ for Re = 2,300 (top) and Re = 2,350 (bottom).	
	(n,r) = (90,51) and $(95,90)$ samples are obtained respectively using	
	Semtex (DNS1) and the code of Willis & Kerswell, 2009 (DNS2, courtesy of	
	Marc Avila). Dashed line represents the curve $S(t) = \exp[-(t - t_0)/\hat{\tau}(\text{Re})]$,	
	and shaded areas represent the 95% confidence interval associated with each	
	simulation.	126
5.7	Directed bond percolation on a (1+1)-dimensional square lattice utilising	
	periodic boundary conditions, where the percolation direction (time) is	
	vertically upwards. Active and inactive sites are denoted by black and white	
	dots respectively. Open bonds are represented by solid lines, and closed	
	bonds by dashed lines. Numbers along the bottom correspond to spatial	
	locations <i>i</i>	130
5.8	Examples of critical behaviour in directed bond percolation using $N = 500$	
	sites over $t = 2,000$ timesteps. Figures along the top row start with a	
	configuration possessing a single active site, and those below use a random	
	initial configuration	131
5.9	Space-time figures demonstrating a relationship between directed percolation	
	and pipe flow. In each case, the initial condition is an isolated puff, and Re	
	is instantaneously set to 2,700 (top) and 1,700 (bottom)	133
5.10	Characteristic lifetime τ as a function of Reynolds number for an isolated	
	puff. Coloured points represent experimental data (courtesy of Kerstin Avila,	
	Alberto de Lozar and Björn Hof), and solid black triangles those obtained	
	from Semtex (DNS1) and a hybrid spectral finite-difference code (DNS2).	
	The dashed line denotes the super-exponentially increasing fit for mean	
	decay times, and the solid line the decreasing fit for mean splitting times.	
	The crossover point at $\text{Re}_c = 2,040 \pm 10$, determines the transition between	
	transient and sustained turbulence in pipe flow.	136

List of Tables

4.1	Axial resolution requirements for the pipe expansion experiment of section	
	4.3.1 for various domain lengths L	93
5.1	Data obtained from DNS1 of the characteristic splitting time τ as a function	
	of Reynolds number, where each simulation observes r splittings from	
	<i>n</i> initial samples. t_1 denotes the first observed splitting time, and t_{max} the	
	maximum simulation time. Upper and lower confidence intervals are denoted	
	by $\pm \Delta \tau$	128

Acknowledgements

Firstly, I am incredibly grateful to my supervisor, Prof. Dwight Barkley, for his encouragement and support throughout the last four years. There is no question that without his invaluable advice, insight and knowledge this thesis would never have seen the light of day, and I feel enormously lucky and privileged to have had him as my supervisor.

My immense thanks go to Kerstin Avila, Alberto de Lozar, Marc Avila and Björn Hof of the University of Göttingen for their collaboration with us on the results of chapter 5. Many of the important conclusions drawn in this chapter would not have been possible without them. I also thank Prof. Hugh Blackburn for providing Semtex, the code used throughout this thesis to generate all of our results. I am also grateful for the comments of my examiners, Prof. Rich Kerswell and Prof. Sergey Nazarenko, which have greatly improved areas of this thesis.

This work would also not have been possible without the hundreds of thousands of computer hours provided by the University of Warwick Centre for Scientific Computing, HECTOR and GENCI-IDRIS, and I would like to thank the staff and administrators for their hard work maintaining these systems. I am also grateful to the Engineering and Physical Sciences Research Council for providing funding throughout my postgraduate studies.

Throughout both my undergraduate and postgraduate degrees, many people both inside and outside the Mathematics Institute at Warwick have made my time here unforgettable. Unfortunately I simply do not have the space to list you all! In particular however, I am incredibly grateful to have had the company of my fellow compatriates in B2.39 throughout the years: David White, Simon Cotter, Chris Cantwell, Damon McDougall, Tom Ranner and Andrew Duncan. All of you have been there to help too many times to mention, so thank you all! Many thanks go to Mark, Erin, Rod and Robby for providing many fun evenings and being around when I needed an ear.

Last, but absolutely not least, I dedicate this thesis to my partner Sarah and my mom, dad, brother, aunt and grandparents. This Ph.D. would never have been started – let alone completed – without the unwavering support and encouragement of every single one of you, and I thank you all immensely.

Declaration

I hereby declare that this thesis is a presentation of my own, original research, conducted under the supervision of Professor Dwight Barkley, except where stated otherwise in the text. It has not been submitted for award at any other institution or for any other qualification. Chapter 4 contains material from a paper published in the *Proceedings of the National Academy of Sciences* (Moxey & Barkley, 2010). Chapter 5 contains material from a paper to appear in *Science* (Avila *et al.*, 2011).

Abstract

When fluid flows through a channel, pipe or duct, there are two basic forms of motion: smooth laminar flow and disordered turbulent motion. The transition between these two states is a fundamental and open problem which has been studied for over 125 years. What has received far less attention are the intermittent dynamics which possess qualities of both turbulent and laminar regimes. The purpose of this thesis is therefore to investigate large-scale intermittent states through extensive numerical simulations in the hopes of further understanding the transition to turbulence in pipe flow.

We begin by reviewing the spectral-element code Semtex which is used to perform the simulations. We discuss modifications to this code to impose a constant flowrate to the flow through a pipe and to improve the computational efficiency on certain multicore architectures. We then move on to examine the reverse transition from turbulence to laminar flow in a long, 125 diameter periodic pipe, which unlike the forward transition does not depend on finite-amplitude perturbations to the flow and thus captures the natural dynamics contained within the transition. The Reynolds number Re is reduced from Re = 2,800 to Re = 2,250 over a long timescale, and by investigating the resultant spatio-temporal dynamics we discover that the transition can be characterised by three fundamentally different states separated by two Reynolds numbers. Below Re_c $\leq 2,300$, turbulence takes the form of equilibrium puffs which eventually decay. Above Re_i $\approx 2,600$, flow remains uniformly turbulent throughout the domain. Between these two values, the dynamics are an intermitent mixture of both turbulent and laminar regimes which take the form of unsteady alternating laminar-turbulent bands.

Finally, we concentrate on finding a more exact value for Re_c , which marks the onset of sustained turbulence in pipe flow. We examine the process through which isolated turbulent puffs split and find that, like decay, this process is stochastic and memoryless. By drawing comparisons with other simple stochastically driven systems – in particular, directed percolation – we compare the timescales for decay and splitting, and ascertain that $\text{Re}_c = 2,040 \pm 10$.

Abbreviations

CFL	Courant-Friedrichs-Lewy (condition)
CPU	Central Processing Unit
DFT	Discrete Fourier transform
DNS	Direct numerical simulation
DP	Directed percolation
iid	Independent identitically distributed
FFT	Fast Fourier transform
GLL	Gauss-Lobatto-Legendre (collocation points)
LE	Leading edge
LHS	Left hand side
MLE	Maximum likelihood estimator
MPI	Message Passing Interface
ODE	Ordinary differential equation
PDE	Partial differential equation
PDF	Probability density function
Re	Reynolds number
RHS	Right hand side
TE	Trailing edge
TW	Travelling wave

Chapter 1

Introduction

Fluid flows play an important role in a wide range of problems in science and engineering, be it in hugely complex atmospheric dynamics with vast numbers of degrees of freedom, the analysis of blood flowing through veins and capilliaries or using mixing properties to design efficient heat exchangers. The motion of a fluid is inherently complex, and in many cases appears random, which makes it both hard to predict and difficult to analyse. Understanding these dynamics has proven to be one of the most important and exhaustively studied problems in mathematics, physics and engineering, spanning more than two centuries of detailed investigation. Of all of the countless scenarios in which we can study fluid flow, in this thesis we consider one of the simplest, oldest and most puzzling – the flow of a fluid through a straight, cylindrical pipe.

In general, we may classify a fluid flow into three distinct states. In *laminar* flows, motion is smooth and regular, and often has symmetry in simple geometries. In *turbulent* flows, the fluid exhibits highly irregular and unpredictable motion which is usually characterised by the interaction of unsteady vortices over a wide variety of length scales. Figure 1.1 demonstrates the striking difference between these two states in pipe flow. One of the classical fluid dynamics problems is to understand precisely how and why turbulence arises from laminar behaviour. Key to this problem is the third type of flow, usually called *intermittent* or *transitional*, which describes fluids possessing characteristics of both laminar and turbulent regimes. Whilst fully laminar and turbulent flows have a rich research background, intermittent flows have, for the most part, received far less attention.



Figure 1.1: Contours of the streamwise velocity through circular cross-sections of a pipe. The left-hand image depicts a laminar flow in which the fluid is steady and unchanging in time. This contrasts with the right-hand image, which is a snapshot of a highly fluctuating turbulent flow at Re = 3,000.

The transitional dynamics of pipe flow were first investigated in the pioneering work of Reynolds (1883a),b which demonstrated the existence of a relationship between the onset of turbulence and a dimensionless parameter, known today as the ubiquitous *Reynolds number*

$$\operatorname{Re} = \frac{\overline{U}D}{\nu}.$$

Here *D* is the diameter of the pipe, v is the fluid's kinematic viscosity and \overline{U} is the mean (bulk) velocity of the flow. At low Re, the viscous effects of the fluid outweigh the inertial forces, causing disturbances to decay and the flow to revert to a laminar state. As Re increases, the balance of inertial and viscous forces changes so that a perturbation applied to the fluid grows in strength and eventually gives rise to a turbulent flow. This concept also translates to a number of other geometries, such as the flow between two solid planes or over a backwards-facing step, and as such the Reynolds number forms a key concept in fluid dynamics.

Curiously however, after being investigated for over 125 years, the question of a finding a *critical* Re below which turbulence cannot be sustained in a pipe is still an open problem. The main objective of this thesis then is to further investigate pipe flow in the hopes of clarifying this elusive critical Re.

1.1 Existing studies

Pipe flow, and fluid flow in general, has proven to be very resilient to mathematical analysis. Assuming the fluid is Newtonian and incompressible, the flow is governed by the Navier-Stokes equations

$$\frac{\partial \boldsymbol{u}}{\partial t} + (\boldsymbol{u} \cdot \nabla)\boldsymbol{u} = -\nabla p + \nu \nabla^2 \boldsymbol{u}, \qquad (1.1)$$

$$\nabla \cdot \boldsymbol{u} = \boldsymbol{0}, \tag{1.2}$$

where u is the fluid velocity and p is the pressure field. One of the simplest mathematical tasks is to determine whether, given a smooth initial condition, this partial differential equation (PDE) has a smooth solution for all times t > 0. The difficulty of this problem is particularly emphasised by the fact that it remains one of the six unsolved Clay Institute Millennium challenges, for which a solution commands a one-million dollar reward.

In pipe flow, the simplest analytic solution one may examine is that of laminar flow, seen in figure 1.1, which may be derived exactly from equation (1.1) (see section 3.1.2). A common mathematical technique for determining bifurcations in non-linear systems of PDEs which depend upon a parameter is to consider infinitesimal perturbations of a known solution and linearise the resulting equations of motion. However, even this relatively straight-forward approach has yielded few analytic results. Joseph & Carmi (1969) showed that any disturbance introduced to the fluid for Re < 81.49 decays to the laminar profile. Most evidence suggests that, in fact, laminar flow is stable to these perturbations at *any* Re. Whilst this has not been proven, Meseguer & Trefethen (2003) demonstrated this through direct numerical simulations (DNS) of the linearised Navier-Stokes equations and observed stability as high as Re = 10^7 . There are also a number of proofs in cases where restrictions on the type of perturbation are applied. For example, as stated by Crowder & Dalton (1971), the hydrodynamic stability of laminar flow subject to axisymmetric perturbations may be traced back to Sexl (1927), and some sixty years later stability under these conditions at any Re was proven by Herron (1991).

Despite the mathematical evidence suggesting laminar flow is quite resilient to perturbation, in the setting of a laboratory controlling the onset of turbulence has proven to be difficult. Reynolds was originally able to maintain laminar flow up to an impressive Re $\approx 12,000$ and later Pfenniger (1961) demonstrated laminar flow up to $Re = 10^5$. However, the pipe must be properly isolated from outside vibration and temperature fluctuations to achieve laminar flow at such high Reynolds numbers. The shape of the inlet must also be very carefully designed to avoid unnecessary perturbation of the flow. Given a pipe capable of sustaining laminar flow, the original idea behind Reynolds' work is the concept of a critical Reynolds number, Re_c , below which perturbations that are deliberately introduced to the flow eventually decay and turbulence cannot be sustained.

Most studies of turbulence expand on this concept by introducing finite-amplitude perturbations to the fluid, and observing how the flow evolves over time at different values of Re. Upon introduction of the perturbation, one may observe two outcomes: the flow may either transition to turbulence, or instantly revert to a laminar state. The regime that is observed depends on the amplitude of the perturbation, and the critical threshold needed to trigger transition obeys a power law as a function of Re (Darbyshire & Mullin, 1995, Hof et al., 2003, Peixinho & Mullin, 2007). However, the amplitude of the perturbation alone does not govern Re_c . Assuming that the perturbation is sufficiently strong, Reynolds' original experiments noted that around the transition point one typically observes isolated turbulent 'flashes' which grow and split to contaminate the downstream flow, and are separated by regions of laminar flow. Wygnanski & Champagne (1973) demonstrated that these flashes can be classified into two separate structures, which they deemed *puffs* and *slugs*. Puffs consist of a small pocket of turbulence which is embedded in the surrounding laminar flow, and are found close to the suspected transition point at Re $\approx 2,000$. Slugs are the high-Re equivalent found past Re \gtrsim 2,700 which rapidly grow to fill the whole pipe and cause turbulence to spread downstream of the perturbation. Understanding the transition to turbulence then hinges on observing the dynamics of puffs (and to a lesser degree slugs) and determining the mechanisms that govern their behaviour.

At lower Reynolds numbers puffs remain stable after the initial perturbation but eventually dissipate. The key observation of Faisst & Eckhardt (2004) is that the lifetimes of puffs are not dictated in a deterministic fashion. Instead, they are determined by exponentially distributed random variables, with their mean lifetime scaling as a function of Re. Therefore, if there exists some critical Reynolds number above which puffs have an infinite lifetime then this would represent the onset of sustained turbulence in pipe flow. Precisely how the mean lifetime scales has been a topic of great attention and has resulted in a wide variety



Figure 1.2: Visualisations of turbulent and intermittent flow in a cross-section of an L = 25D pipe. Red indicates high values of the vorticity $\boldsymbol{\omega} = \nabla \times \boldsymbol{u}$ and blue low values.

of studies (Faisst & Eckhardt, 2004, Peixinho & Mullin, 2006, Hof *et al.*, 2006, Willis & Kerswell, 2007, Hof *et al.*, 2008, de Lozar & Hof, 2009, Kuik *et al.*, 2010). The most recent statistical analysis of Avila *et al.* (2010) however, presents the strongest evidence yet that puff lifetimes remains finite at any Re, and thus it is impossible to determine the critical Reynolds number from considering the mean lifetime of puffs alone.

Another line of investigation which may aid in understanding the mechanism behind the onset of turbulence has been the discovery of unstable coherent structures which are alternative to that of the laminar profile. The work of Faisst & Eckhardt (2003) and Wedin & Kerswell (2004) confirmed that travelling waves (TWs) – symmetric structures which move through the pipe with constant wave speed – exist within the framework of pipe flow, and were later shown to exist experimentally by Hof *et al.* (2004). Further studies including those of Pringle & Kerswell (2007), Willis & Kerswell (2008), Willis & Kerswell (2009) and Pringle *et al.* (2009) have greatly expanded the number of known TW solutions. From a dynamical systems standpoint then, laminar flow is a globally stable solution below Re_c , and the perturbation of the fluid leads to a solution which wanders between the unstable solutions in phase space before eventually returning to laminarity. Above Re_c the flow undergoes a transition so that laminar flow is only *locally* stable and turbulent solutions are composed of a wide variety of TW solutions. However, since turbulent solutions only visit TWs around 10% of the time however (Kerswell & Tutty, 2007), alternative unstable solutions need to be found in order to reveal more of the phase space dynamics.

1.2 Spatio-temporal intermittency

The common theme throughout the majority of studies outlined in the previous section is the forward transition from laminar flow to turbulence, where observations are not only dependent on Re but on both the type of perturbation used to trigger transition and its amplitude. However, intermittent states, such as those seen in the experiments of Rotta (1956), have rich spatio-temporal behaviour which has for the most part not been investigated. The difference between fully turbulent flow and intermittent states is highlighted in figure 1.2; the top pipe shows a simulation at Re = 3,000, where the flow is uniformly turbulent and structure lengths are roughly uniform throughout the domain. At lower values of Re (here 2,300), areas of laminarity begin to appear, with structures growing considerably in size.

One line of research that is emerging in the investigation of critical Reynolds numbers is the consideration of alternating turbulent-laminar bands depicted in figure 1.3 which form within the flow. It has been established that in plane Couette flow (Prigent *et al.*, 2002, 2003, Barkley & Tuckerman, 2005, 2007), counter-rotating Taylor-Couette flow (Prigent *et al.*, 2002, 2003), and plane Poiseuille flow (Tsukahara *et al.*, 2005), near transition the system can exhibit a remarkable phenomenon in which turbulent and laminar flow form persistent alternating patterns on scales very long relative to both wall separation and the spacing between turbulent streaks. While the origin of these patterns remains a mystery, they are intimately connected with the lower limit of turbulence in shear flows. Key to the investigation of these states is the consideration of the *spatio-temporal* aspects of the flow. Investigating these states in pipe flow is a more challenging endeavour, mostly due to the strong advection of fluid down the length of the pipe.

The fundamentally different approach taken in this thesis then is to investigate the intermittent dynamics found in the *reverse* transition from turbulent to laminar flow, as opposed to most previous work which has focused on the forward transition problem. In doing this, we hope to observe the formation of turbulent-laminar bands, and examine how these bands evolve over time in order to determine their role in the transition process.

Whilst other experiments, for example Narasimha & Sreenivasan (1979), have investigated this relaminarisation process, the method used to reduce the Reynolds number usually



Figure 1.3: Laminar-turbulent bands in experiments of Couette flow at Re = 358, taken from Prigent *et al.* (2003).

involves a sharp expansion in the channel or pipe causing the flow to abruptly transition. This differs significantly from the approach we take here, in which Re is slowly and deliberately reduced over a long timescale in order to carefully examine the structures which form within the flow. Indeed, for the most part we are not interested in the full relaminarisation of the fluid, but in the states just before the transition. To replicate this process in an experimental setting demands pipes of many thousands of diameters in length, and also the ability to observe the velocity and pressure fields throughout the length of the pipe.

To overcome these limitations, we shall focus on performing a series of detailed computational experiments of pipe flow. Eggels et al. (1994) performed the fully-resolved turbulence computations in a pipe of length L = 5D. Later work by Shan *et al.* (1999) used a spectral element code to simulate puffs and slugs and expanded these domains to a much larger $L = 16\pi D$; however the computations here suffer from under-resolution in the axial direction, especially in the simulation of slugs at Re = 5,000. Only in the last decade has computing technology advanced to the point where, for example, puffs and slugs can be faithfully represented. Faisst & Eckhardt (2004) examined decay in a pipe of length L = 5Dwith better resolution, but too short to fully encompass a single puff. More recently however, Willis & Kerswell (2007), Pringle & Kerswell (2007) and Willis & Kerswell (2008) employed the use of an efficient hybrid finite-difference-spectral solver in domains of up to length L = 50D and are the first examples of fully-resolved simulations in domains large enough to contain puffs. Given that the main focus of this work is on examining the spatio-temporal properties of the flow, larger domains are required in order to observe the interactions of any smaller structures (such as puffs) within the pipe. We will therefore employ the use of much longer pipes of up to L = 125D, and in some shorter simulations up to L = 200D.

1.3 Outline

In chapter 2 we introduce the numerical techniques which will be used to perform the DNS of pipe flow. We begin by studying spectral-type methods, which describe a class of algorithms possessing exponential convergence of approximations to solutions of a PDE. The main goal of this chapter is to introduce the spectral element method, which combines the geometric flexibility of a finite element method with spectral-type accuracy. To motivate this, we show the implementation of the method for the one-dimensional diffusion equation. Finally, a number of high-order splitting schemes for the evolution of the Navier-Stokes equations are introduced, along with a discussion of the high-order pressure boundary condition which is needed to make these methods numerically stable.

Chapter 3 introduces the mathematical framework in which pipe flow is studied and the laminar equations of motion are derived. Following this, two numerical techniques useful for the study of pipe flow are outlined. The first is an algorithm which allows the flow to be driven by specifying a constant volumetric flux. This is particularly applicable in transitional flow problems as it prevents deviations in the mean flow \overline{U} and thus fixes Re. An efficient Green's function method is derived as a part of the mathematical formulation, meaning that unlike other methods the flowrate can be imposed exactly at each timestep. We then introduce Semtex, the code which is used in later chapters to perform pipe flow simulations, and show how this flux condition can be implemented. Finally, we outline the second numerical method which provides optimisation for the parallelisation of Semtex in the case where the underlying system architecture makes use of massively multi-core nodes.

These numerical techniques are then used extensively throughout chapter 4. The chapter begins with a resolution study, in which a number of simulations are performed in order to determine the size of small-scale structures found in turbulent flows, and thus ensure that the numerical domain is correctly resolved. We then continue with a series of simulations which focus on studying the reverse transition problem. Initially, a large-scale simulation in both space and time is performed, and the resulting spatio-temporal intermittency is examined. After observing the various states in this transition, we then proceed with three further studies, each of which is designed to highlight the various intermittent aspects of the flow and classify the various states found during transition.

Chapter 5 focuses on puffs and the role that they play in the transition process. Whilst the precise distribution of puff lifetimes has remained a contentious issue for a number of years, recently it was conclusively shown that the lifetime of a puff remains finite for all Re (Avila *et al.*, 2010), perhaps suggesting that all turbulence in a pipe is transient. Crucially however, the decay process neglects much of the spatial complexity found in chapter 4. In chapter 5, we focus on the opposite process known as puff splitting, in which a new puff is nucleated downstream of an existing puff. This phenomenon naturally incorporates the complex spatial interaction seen in the previous chapter and throughout much of the existing literature. We show that this process is also stochastic, and obeys a similar distribution to that of decay. Using this knowledge, we are able to draw upon conclusions found in a number of similar stochastically complex systems such as directed bond percolation, and from this obtain a critical Re denoting the phase transition from states dominated by decay to those where splitting becomes more probable.

Chapter 2

Computational Techniques

The investigation of many interesting physical systems usually leads to a mathematical representation involving a non-linear dynamical system, often in the form of a partial differential equation. In almost every case, it is impossible to explicitly derive non-trivial solutions to these systems, and so the development of techniques which allow us to approximate solutions is of great significance to both mathematicians and others in the physical sciences. Typically, these approximations involve the use of complex algorithms, and so we turn to computers in order to implement them through numerical means.

With fluid flows being amongst the most throughly examined non-linear systems, there are a vast array of numerical methods which can be used to approximate flow fields. Notably, the emphasis throughout this thesis is on obtaining solutions through direct numerical simulation; i.e. the direct integration of the partial differential equations governing the flow, at a scale which accurately reproduces all of the key features contained within it. The goal of this chapter is to outline a small number of these numerical methods. In the next chapter, we describe how these methods can be used in the setting of pipe flow.

We begin with a study of the method of weighted residuals, a generic framework which allows us to employ a wide variety of methods in order to computationally solve problems. In particular, we focus on *spectral* methods, which are generally known for their excellent convergence properties and accurate solutions. We study these methods in two different contexts: *collocation* methods rely on a fixed discretisation of the spatial domain, whereas *Galerkin* methods transform the PDE into a finite-dimensional functional analysis problem

without relying on explicit knowledge of the domain.

The bulk of this chapter concentrates on examining finite element methods, which are typically used to study problems defined in complex geometries. The inspiration behind this technique is to break the domain up into disjoint elements on which the equation can be more readily solved, and then recombine the contribution of each to produce the global solution according to certain constraints on elemental boundaries. In particular, we focus on the *spectral/hp* element method, which combines spectral-type accuracy with the geometric flexibility of finite elements, and is the method comprising the vast majority of the numerical results in later chapters. To demonstrate this scheme, we introduce the concepts necessary to formulate a one-dimensional implementation and use this to numerically solve the diffusion equation.

Finally, having made an overview of the schemes which are generally used in the spatial discretisation, we introduce a variety of timestepping schemes. Given the high-order spatial accuracy afforded by the discretisations mentioned above, we discuss the possible algorithms which can be employed to increase the accuracy of approximate flow fields as the equations of motion are integrated over time.

2.1 Method of Weighted Residuals

The basis of any spectral method relies heavily on the method of weighted residuals. It is well known that periodic functions on $[0, 2\pi]$ which belong to the set

$$L^{1}([0, 2\pi]) = \left\{ f : [0, 2\pi] \to \mathbb{R} \ \bigg| \ \int_{0}^{2\pi} |f(x)| \, \mathrm{d}x < \infty \right\}$$

can be written in terms of summations of infinitely many orthogonal functions – for example, in Fourier analysis, functions are represented in terms of trigonometric functions. In terms of numerical analysis, the basis of a spectral method is to use a truncated series containing *finitely* many functions in order to obtain an approximation to the function.

Consider a time-dependent linear differential equation over a domain Ω given by

$$\mathbf{L}(u) = 0, \tag{2.1}$$

and assume that the function can be accurately approximated by a partial expansion u^{δ} in terms of *modes* ϕ_k :

$$u(\boldsymbol{x},t) \approx u^{\delta}(\boldsymbol{x},t) = \sum_{k=0}^{N-1} \hat{u}_k(t)\phi_k(\boldsymbol{x}).$$
(2.2)

In general, an approximation of u by finitely many modes will not be exact, so we expect to find some difference between u and u^{δ} . Substituting (2.2) into (2.1), we therefore obtain a non-zero function R known as the *residual* which is dependent upon u^{δ} ,

$$\mathbf{L}(u^{\delta}) = R(u^{\delta}).$$

In the approximation (2.2), there is no distinct way of choosing the weights \hat{u}_k . However, by placing restrictions on the form of the residual, we can reduce the problem to solving a system of ordinary differential equations in terms of \hat{u}_k . The choice of restriction we use leads to a variety of different numerical schemes. First however, we need a definition in order to define the notion of orthogonality.

Definition 2.1.1 (L^2 inner product). The L^2 inner product $(\cdot, \cdot)_{L^2}$ for functions $f, g : \Omega \to \mathbb{C}$ is defined as

$$(f,g)_{L^2} = \int_{\Omega} f(\mathbf{x}) \overline{g(\mathbf{x})} \, \mathrm{d}\mathbf{x}$$

where \overline{z} denotes the complex conjugate of $z \in \mathbb{C}$.

One way to restrict the residual is through use of *test functions* $v_j(\mathbf{x})$ which are orthogonal to *R* under $(\cdot, \cdot)_{L^2}$;

$$(v_j, R)_{L^2} = 0, \quad j = 0, \dots, N.$$
 (2.3)

Two of the most popular choices for the test functions produce the *collocation* and *Galerkin* methods. Both have different convergence properties, but in general we expect that when coupled with a spectral-type modal expansion, each will exhibit exponential convergence of the error. The low errors found in these methods are particularly important in accurately reproducing the dynamics of systems which are sensitive to small perturbations.

2.1.1 Collocation Formulation

For the collocation method, we choose the test function to be a Dirac delta function $v_j(\mathbf{x}) = \delta(\mathbf{x} - \mathbf{x}_j)$, where the $\mathbf{x}_j \in \Omega$ are a set of points known as the *collocation points*. Substituting this into (2.3), we obtain

$$\int_{\Omega} R(\mathbf{x})\delta(\mathbf{x}-\mathbf{x}_j)\,\mathrm{d}\mathbf{x}=0$$

and so by definition, $R(x_j) = 0$. The differential equation is therefore satisfied exactly at each collocation point x_j .

Depending on the choice of modal functions, there are various subclasses of collocation methods. Typically, one chooses amongst common orthogonal bases depending upon the type of boundary conditions to be applied on the domain. In this thesis, we will concentrate on using collocation methods with periodic boundary conditions, on the interval $\Omega = [0, L]$.

Under these conditions, the predominant choice of modal function is formed from trigonometric basis functions, in which case we obtain the *Fourier collocation* method. Ω is discretised into a uniform grid of N + 1 points $x_j = jL/N$, $j \in \{0, ..., N\}$, and modes are taken to be $\phi_k(x) = e^{i\beta_k x}$ where $\beta_k = 2\pi k/N$. These functions form an orthogonal set, since

$$\sum_{j=0}^{N-1} e^{i\beta_k j} e^{i\beta_l j} = N\delta_{kl} = \begin{cases} 0, & k \neq l, \\ N, & k = l. \end{cases}$$

Assuming N is even, rewriting (2.2) in the form

$$U_j(t) := u^{\delta}(x_j, t) = \sum_{k=-N/2+1}^{N/2} \hat{u}_k(t) e^{i\beta_k j},$$

and using the orthogonality relation, we can determine the weights \hat{u}_k via the *discrete Fourier* transform (DFT)

$$\hat{u}_k = \frac{1}{N} \sum_{j=0}^{N-1} U_j(t) e^{i\beta_k j}$$

The DFT is the obvious analogue of the continuous transform,

$$\hat{f}(k) = \int_{\mathbb{R}} f(x)e^{-ikx} \,\mathrm{d}x$$

which itself is useful for solving PDEs analytically, mostly due to the transform of a differential operator:

$$\widehat{\frac{\mathrm{d}f}{\mathrm{d}x}}(k) = \int_{\mathbb{R}} f'(x)e^{ikx} \,\mathrm{d}x = ik \int_{\mathbb{R}} f(x)e^{-ikx} \,\mathrm{d}x = ik\,\widehat{f}(k)$$

In the discrete case, a similar identity holds:

$$\left. \frac{\partial u^{\delta}}{\partial x} \right|_{x=x_j} = \left(\frac{\partial}{\partial x} \sum_{k=-N/2+1}^{N/2} \hat{u}_k(t) e^{ikx} \right)_{x=x_j} = \sum_{k=-N/2+1}^{N/2} ik \hat{u}_k(t) e^{i\beta_k j}$$

so differentiation in spectral space is equivalent to multiplying each mode \hat{u}_k by ik. Notice that the Fourier modes are periodic, since

$$\hat{u}_{k+N} = \sum_{j=0}^{N-1} U_j e^{i\beta_{k+N}j} = \sum_{j=0}^{N-1} U_j e^{i\beta_k j} e^{2\pi ki} = \sum_{j=0}^{N-1} U_j e^{i\beta_k j} = \hat{u}_k$$

This property therefore requires the approximation u^{δ} to also be periodic. Also, in the case where the data points U_j are real, we may exploit additional Hermitian symmetry so that we need only calculate \hat{u}_k for $k \ge 0$:

$$\overline{\hat{u}_k} = \overline{\sum_{j=0}^{N-1} U_j e^{i\beta_k j}} = \sum_{j=0}^{N-1} \overline{U_j e^{i\beta_k j}} = \sum_{j=0}^{N-1} U_j e^{-i\beta_k j} = \sum_{j=0}^{N-1} U_j e^{i\beta_{-k} j} = \hat{u}_{-k}.$$

2.1.1.1 Non-linear Problems

Whilst the Fourier collocation method is extremely practical for linear problems, it becomes far less efficient when dealing with non-linear operators. For example, the spectral discretisation of the non-linear term

$$u\frac{\partial u}{\partial x}$$

results in an $N \times N$ matrix in general, and is therefore not practical to use when combined with an implicit timestepping scheme (see section 2.3) which requires inversion of the operator. An easy way around this problem is to maintain a physical grid *in addition to* an expansion in terms of modes. We then evaluate linear terms spectrally for optimal accuracy, and then transform the solution to physical space where the non-linear multiplication is performed.

Whilst in theory this procedure is more efficient, this process can only be considered if the transform technique is sufficiently quick to avoid the computational overheads. In addition, the multiplication in physical space leads to the appearence of aliasing errors in the solution. However, for a basis of Fourier functions in one dimension, both of these problems are easily overcome with use of the Fast Fourier Transform (FFT) and the classical 3/2 rule to dealias solutions (Canuto *et al.*, 1988).

This problem is investigated in further detail in section 3.3, where we consider a joint discretisation in which a three-dimensional domain is decomposed using the Fourier pseudo-spectral method in a single dimension and a spectral element scheme used in the remaining two. For large-scale problems requiring parallelisation, it is relatively easy in terms of implementation to decompose the domain over the Fourier modes. However, the transform between spectral and physical space commands large communication overheads, and thus we examine ways in which the transform can be optimised in the case where the underlying architecture comprises massively multi-core machines.

2.1.2 Galerkin Formulation

We motivate the study of Galerkin schemes by considering an example of the one-dimensional Poisson equation on the interval $\Omega = [0, 1]$, where $\mathbf{L} = \partial_{xx}^2 + g$, giving the equation

$$\frac{\partial^2 u}{\partial x^2} + g(x) = 0 \tag{2.4}$$

In order for this for the problem to be well-posed and thus have a unique solution, we must impose boundary conditions. We consider both Dirichlet and Neumann boundary conditions for this problem to show how both types are implemented in a Galerkin setting, so that

$$u(0) = g_D, \quad \frac{\partial u}{\partial x}(1) = g_N \tag{2.5}$$

where g_D and g_N are constants. With both equations (2.4) and (2.5), the problem is said to be in strong form. The first step in the Galerkin formulation is to multiply (2.4) by an

arbitrary test function v(x) and integrate, giving the integral equation

$$(v, \mathbf{L}(u))_{L^2} = \int_0^1 v \left(\frac{\partial^2 u}{\partial x^2} + g\right) \mathrm{d}x = 0.$$
(2.6)

Assuming u and v are sufficiently smooth, we may then integrate by parts to obtain

$$\underbrace{\int_{0}^{1} \frac{\partial v}{\partial x} \frac{\partial u}{\partial x} \, dx}_{a(u,v)} = \int_{0}^{1} v f \, dx + \left[v \frac{\partial u}{\partial x} \right]_{0}^{1}$$
$$= \underbrace{\int_{0}^{1} v f \, dx + v(1)g_{N}}_{f(v)} - v(0) \frac{\partial u}{\partial x}(0).$$

We make the further assumption that g(x) and thus f is well behaved in the sense that f(v) remains finite. Therefore we restrict the choice of solutions u to those lying in a trial space X, where

$$X = \{ u \mid u, u' \in L^2(\Omega), u(0) = g_D \}.$$

The space of possible test functions V is defined similarly, so that

$$V = \{ v \mid v, v' \in L^2(\Omega), v(0) = 0 \}.$$

With these choices, the problem is now said to be in *weak form*, and the weak solution to the original equation translates to the problem of finding $u \in X$ such that

$$a(u, v) = f(v), \quad \forall v \in V.$$

This problem is still infinite dimensional, and therefore not solvable through computational means. To overcome this, we take finite-dimensional subsets of X and V, denoted by X^{δ} and V^{δ} respectively. The problem then becomes to find $u^{\delta} \in X^{\delta}$ such that

$$a(u^{\delta}, v^{\delta}) = f(v^{\delta}), \quad \forall v^{\delta} \in V^{\delta}.$$

Whilst this scheme naturally incorporates the Neumann boundary conditions, it does not enforce any Dirichlet conditions. To this, we decompse the approximate solution u^{δ} into a homogeneous part $u^H \in V^{\delta}$ and a known component $u^D \in X^{\delta}$ so that

$$u^{\delta} = u^H + u^D$$

where $u^{H}(0) = 0$ and $u^{D}(0) = g_{D}$. This completes the classical Galerkin formulation, under which we choose the test space and trial space to be the same. Assuming certain properties of *a*, one can now apply the Lax-Milgram theorem (Lax & Milgram, 1954) to guarantee solutions to the problem.

2.2 Spectral/hp Element Methods

Either of the collocation or Galerkin methods defined above, when coupled with sensible choices of modes such as the Fourier scheme seen earlier, form extremely accurate numerical methods. They provide an ideal method for use in simple geometries in which uniform grids are easily applied. However, expanding this method to more complex domains generally requires the use of a non-uniform grid, greatly increasing the complexity of implementation and precluding the use of optimised algorithms such as the Fourier transform. Finite element methods comprise a class of algorithms which are designed to deal with awkward geometries by partitioning the domain into elements on which the solution can be more easily approximated. This partitioning scheme naturally allows us to resolve parts of the domain in greater detail by making their size smaller (*h*-refinement).

Finite element methods are in essence no different than the weighted residuals considered earlier. However, the choice of modal functions is now far more complex. Generally, when restricted to each element, these functions will be some polynomial of order P, which can be increased in order to improve accuracy (*p*-refinement). Typically a finite element method uses piecewise linear functions (i.e. P = 1) which greatly simplifies the mathematical formulation and is efficient to implement. The spectral/hp element method is a specialisation of the finite element method, and introduces a number of techniques leading to a more efficient implementation when high-order polynomials are employed. Their aim is to combine the accuracy found in spectral methods with the geometric flexibility of the general finite element method. In this section, we present a small selection of material from Karniadakis & Sherwin (2005), which provides a fully comprehensive guide to implementing

spectral/hp element methods in the setting of computational fluid mechanics.

2.2.1 Diffusion Equation

To motivate the study of spectral element methods, let us consider the one-dimensional diffusion equation with Neumann boundary conditions on the interval $\Omega = [a, b]$;

$$\frac{\partial u}{\partial t} = \alpha \frac{\partial^2 u}{\partial x^2} \qquad \text{in } \Omega$$
$$u_x(a,t) = g_a, u_x(b,t) = g_b \quad \text{on } \partial \Omega$$
$$u(x,0) = u_0(x) \qquad \text{initial condition}$$

The first step in the formulation of this problem is to construct the weak form of the problem, and as with the Helmholtz equation (2.4), the weak form is used to construct a Galerkin scheme. Multiplying by a test function v and integrating over the domain, we obtain

$$\int_{a}^{b} v \frac{\partial u}{\partial t} dx = -\alpha \int_{a}^{b} \frac{\partial u}{\partial x} \frac{\partial v}{\partial x} dx + \alpha \left[v \frac{\partial u}{\partial x} \right]_{x=a}^{x=b}$$
(2.7)

$$= -\alpha \int_{a}^{b} \frac{\partial u}{\partial x} \frac{\partial v}{\partial x} \, \mathrm{d}x + \alpha \left[v(b)g_{b} - v(a)g_{a} \right].$$
(2.8)

As before, we take approximations v^{δ} and u^{δ} to the test function and u respectively. Since these lie in the same function space, we take series expansions of both with respect to the same modes, so that

$$u^{\delta}(x,t) = \sum_{k=0}^{N-1} \hat{u}_k(t) \Phi_k(x), \qquad (2.9)$$

$$v^{\delta}(x) = \sum_{k=0}^{N-1} \hat{v}_k \Phi_k(x).$$
(2.10)

In this setting, the modes $\Phi_k(x)$ are referred to as the *global modes* and will be complicated. Indeed, for the most part they can remain unknown, and we shall see in the following section that we can instead define modes which are local to each element and can be used to implicitly construct the global modes as a part of the formulation. Substituting these approximations into (2.8) we obtain the expression

$$\int_{a}^{b} \left(\sum_{i=0}^{N-1} \hat{v}_{i} \Phi_{i} \right) \left(\sum_{j=0}^{N-1} \frac{\mathrm{d}\hat{u}_{j}}{\mathrm{d}t} \Phi_{j} \right) \mathrm{d}x = -\alpha \int_{a}^{b} \left(\sum_{i=0}^{N-1} \hat{v}_{i} \frac{\mathrm{d}\Phi_{i}}{\mathrm{d}x} \right) \left(\sum_{j=0}^{N-1} \hat{u}_{j} \frac{\mathrm{d}\Phi_{j}}{\mathrm{d}x} \right) \mathrm{d}x + \alpha \left[\hat{v}_{N-1} \Phi_{N-1}(b) g_{b} - \hat{v}_{0} \Phi_{0}(a) g_{a} \right].$$
(2.11)

Notice that in this derivation we have assumed that all global modes are zero at each endpoint of Ω , apart from Φ_0 which is non-zero at a, and Φ_{N-1} which is non-zero at b. This restriction allows us to simplify the problem, and is also useful in the implementation of Dirichlet boundary conditions, should they arise.

The left hand side of (2.11) then simplifies to

$$\int_{a}^{b} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} \hat{v}_{i} \frac{\mathrm{d}\hat{u}_{j}}{\mathrm{d}t} \Phi_{i}(x) \Phi_{j}(x) \,\mathrm{d}x = \sum_{i=0}^{N-1} \hat{v}_{i} \left[\sum_{j=0}^{N-1} \frac{\mathrm{d}\hat{u}_{j}}{\mathrm{d}t} \int_{a}^{b} \Phi_{i}(x) \Phi_{j}(x) \,\mathrm{d}x \right].$$
(2.12)

This can be re-written in matrix-vector form by defining an $N \times N$ matrix **M** and vectors $\hat{\boldsymbol{u}}_{g}, \hat{\boldsymbol{v}}_{g}$ by

$$\mathbf{M}_{ij} = \int_{a}^{b} \Phi_{i}(x) \Phi_{j}(x) \, \mathrm{d}x, \quad \hat{\boldsymbol{u}}_{g} = (\hat{u}_{0}(t), \dots, \hat{u}_{N-1}(t))^{\top}, \quad \hat{\boldsymbol{v}}_{g} = (\hat{v}_{0}, \dots, \hat{v}_{N-1})^{\top}.$$

The subscript g denotes that these coefficients relate to the global modes. By defining a matrix **L** and vector **B** by

$$\mathbf{L}_{ij} = \int_{a}^{b} \frac{\mathrm{d}\Phi_{i}}{\mathrm{d}x} \frac{\mathrm{d}\Phi_{j}}{\mathrm{d}x} \,\mathrm{d}x, \quad \boldsymbol{B} = (-g_{a}, 0, \dots, 0, g_{b})^{\top},$$

and substituting these definitions into equation (2.11) we obtain the system of ordinary differential equations (ODEs) given by

$$\hat{\boldsymbol{v}}_{g}^{\top} \left[\mathbf{M} \frac{\mathrm{d}\hat{\boldsymbol{u}}_{g}}{\mathrm{d}t} = -\alpha \mathbf{L}\hat{\boldsymbol{u}}_{g} + \alpha \boldsymbol{B} \right] \Rightarrow \mathbf{M} \frac{\mathrm{d}\hat{\boldsymbol{u}}_{g}}{\mathrm{d}t} = -\alpha \mathbf{L}\hat{\boldsymbol{u}}_{g} + \alpha \boldsymbol{B}.$$

A consequence of this is that we never need calculate the coefficients of the unknown test function. The matrices **M** and **L** are extremely common in finite and spectral element methods and as such are named.

Definition 2.2.1 (Mass and stiffness matrices). The matrix **M** is called the *mass matrix*, and the matrix **L** is called the *Laplacian* or *stiffness matrix*.

This equation gives us the problem in semi-discrete form, with a continuous derivative in time. To complete the discretisation, we require the use of a time-stepping scheme. To avoid constraints on the timestep Δt , we apply the implicit Euler scheme outlined in section 2.3. As a shorthand notation, let u^n denote the solution field at time $t_n = n\Delta t$. Then,

$$\mathbf{M}\frac{\hat{\boldsymbol{u}}_{g}^{n+1}-\hat{\boldsymbol{u}}_{g}^{n}}{\Delta t}=-\alpha\mathbf{L}\hat{\boldsymbol{u}}_{g}^{n+1}+\alpha\boldsymbol{B}\Rightarrow(\mathbf{M}+\alpha\Delta t\mathbf{L})\hat{\boldsymbol{u}}_{g}^{n+1}=\mathbf{M}\hat{\boldsymbol{u}}_{g}^{n}+\alpha\Delta t\boldsymbol{B}$$

In essence then, this is a difference equation; given the initial condition \hat{u}_g^n , inverting the matrix problem will then give us the solution for the global modes at time t_{n+1} , \hat{u}_g^{n+1} .

However to start this timestepping process, we need to interpolate the initial condition $u_0(x)$ in order to obtain coefficients $(\hat{u}_0)_k$ such that

$$u_0(x) \approx u_0^{\delta}(x) = \sum_{k=0}^{N-1} (\hat{u}_0)_k \Phi_k(x).$$
 (2.13)

This can be achieved by several distinct methods; in this setting however, the most intuitive approach is *Galerkin projection*, which is performed in a very similar fashion to the original Galerkin formulation. In an infinite dimensional setting, for each $v \in V$, we wish to find $u \in U$ such that

$$(u,v)_{L^2} = (u_0,v)_{L^2} \Longleftrightarrow \int_a^b u(x)v(x) \,\mathrm{d}x = \int_a^b u_0(x)v(x) \,\mathrm{d}x.$$

If we again follow a similar process by taking finite series approximations to u and v and substituting (2.10) and (2.13) into the above equation, we obtain

$$\mathbf{M}\hat{u}_{g}^{0}=\hat{f}$$

where $\hat{f}_i = \int_a^b u_0(x) \Phi_i(x) dx$. Inverting this matrix problem then gives us the transformed modes required to start the timestepping process. This is also one of the most computationally efficient ways to solve this problem.

Finally we require the inverse of this transformation, so that at the end of the simulation the

field can be obtained in physical space. Let us assume we desire the field at some final time $T = M\Delta t$. Then

$$u(x,T) \approx u^{\delta}(x,T) = \sum_{k=0}^{N-1} \hat{u}_k(T) \Phi_k(x).$$

The entire process then, is:

$$u^0 \xrightarrow{\text{Galerkin projection}} \hat{u}^0 \xrightarrow{\text{timestep}} \hat{u}^1 \xrightarrow{\text{timestep}} \cdots \xrightarrow{\text{timestep}} \hat{u}^M \xrightarrow{\text{reconstruct}} u^M.$$

2.2.2 Constructing local modes

Thus far, we have only considered a general Galerkin formulation of the diffusion equation in terms of global modes Φ_k . In an elemental method, it is far more natural to work *locally* on each element, and then perform a *global assembly*. This has the distinct advantage that the global modes never need to be explicitly constructed. This section describes the process of decomposing the domain into a mesh and choosing basis functions which are local to each element. Finally we show how integration can be performed locally, and then assembled to produce global quantities.

In the general *n*-dimensional formulation, the first step is to decompose the domain Ω into *elements*; that is, we find sets Ω^e with $1 \le e \le N_{el}$ and interior int Ω such that they partition Ω and only meet along their boundary;

$$\Omega = \bigcup_{e=1}^{N_{\rm el}} \Omega^e \quad \text{and} \quad \operatorname{int} \Omega^i \cap \operatorname{int} \Omega^j = \emptyset, \ i \neq j.$$

Definition 2.2.2 (Mesh). The collection of sets $\{\Omega^e\}$ is referred to as a *mesh* of Ω .

In one dimension where $\Omega = [a, b]$, the only choice for a mesh is the collection of intervals $\Omega^e = [x_e, x_{e+1}]$ where $a = x_0 < x_1 < \cdots < x_{N_{el}} = b$. Given this mesh, we rewrite the expansion (2.9) in the form

$$u^{\delta}(x,t) = \sum_{e=1}^{N_{el}} \sum_{p=0}^{P} \hat{u}_{p}^{e}(t) \phi_{p}^{e}(x)$$

where $\phi_p^e: \Omega^e \to \mathbb{R}$ is a *local mode*, and \hat{u}_p^e is its related local coefficient. Our goal then is



Figure 2.1: Global and local expansion bases Φ_k and ϕ_k for a three-element finite element expansion. χ^e is the linear projection mapping the standard element Ω_{st} to Ω^e .

to somehow relate the local coefficients \hat{u}_p^e to the global coefficients \hat{u}_k . Typically there will be more local coefficients than global, and hence we need additional restrictions in order to make the problem uniquely solvable.

Before considering this problem however, note that it is unnecessary to define a set of elemental basis functions ϕ_p^e on each element. Whilst technically possible, it introduces a number of additional computational overheads into the formulation. Instead, we exploit the simple elemental geometry to define a single set of local modes $\phi_p : \Omega_{st} \to \mathbb{R}$ on the *standard element* $\Omega_{st} = \{\xi \mid -1 \leq \xi \leq 1\}$, and define a bijective mapping $\chi^e : \Omega_{st} \to \Omega^e$ projecting Ω_{st} onto Ω^e . Local elemental modes can then be defined by the composite function $\phi_p^e = \phi_p \circ [\chi^e]^{-1}$. In the rest of this chapter, ξ denotes a location inside the standard element.

As a demonstration of this local-to-global mapping, let us consider a finite element problem
in one dimension where the polynomial order P = 1. Let $x_0 < x_1 < x_2 < x_3$ and $\Omega = [x_0, x_3] = \bigcup_{e=1}^3 \Omega^e$ with $\Omega^e = [x_{e-1}, x_e]$. The typical global basis functions for a finite element method with $N_{el} = 3$ are shown in figure 2.1, and are given by

$$\Phi_{0}(x) = \begin{cases} \frac{x - x_{1}}{x_{0} - x_{1}}, & x \in \Omega^{1} \\ 0, & x \notin \Omega^{1} \end{cases} \qquad \Phi_{1}(x) = \begin{cases} \frac{x - x_{0}}{x_{1} - x_{0}}, & x \in \Omega^{1} \\ \frac{x - x_{1}}{x_{1} - x_{2}}, & x \in \Omega^{2} \\ 0, & \text{otherwise} \end{cases}$$
$$\Phi_{2}(x) = \begin{cases} \frac{x - x_{1}}{x_{2} - x_{1}}, & x \in \Omega^{2} \\ \frac{x - x_{2}}{x_{2} - x_{3}}, & x \in \Omega^{3} \\ 0, & \text{otherwise} \end{cases} \qquad \Phi_{3}(x) = \begin{cases} \frac{x - x_{3}}{x_{2} - x_{3}}, & x \in \Omega^{3} \\ 0, & x \notin \Omega^{1}. \end{cases}$$

Each global mode is taken to be one at a single element boundary x_i and decays to zero linearly across neighbouring elements. In this case then, there are precisely four global degrees of freedom – that is, N = 4 in (2.9). Locally on each element however, by defining the local basis functions

$$\phi_0(\xi) = \frac{1-\xi}{2}, \quad \phi_1(\xi) = \frac{1+\xi}{2}$$

on Ω_{st} , with associated projection functions

$$\chi^{e}(\xi) = \frac{1-\xi}{2}x_{e-1} + \frac{1+\xi}{2}x_{e},$$

we see that each global mode can be written in terms of local modes. For example, for Φ_0 :

$$\Phi_0(x) = \begin{cases} \phi_0(\xi) = \phi_0([\chi^e]^{-1}(x)), & x \in \Omega^1, \\ 0, & \text{otherwise} \end{cases}$$

Locally to each element we have two coefficients, leading to a total of six elemental degrees of freedom. To uniquely determine the local coefficients in terms of the global ones (and vice-versa), we need to impose further conditions on the local modes. Ultimately this choice depends on the problem we are modelling, but in general we require that Φ_k remains continuous. The natural condition $\hat{u}_1^1 = \hat{u}_0^2$ and $\hat{u}_1^2 = \hat{u}_0^3$ ensures that this will always be the case, and guarantees that global and local coefficients match on a one-to-one basis.

For the more general one-dimensional problem, we define the two column vectors \hat{u}_g and \hat{u}_l denoting global and local coefficients respectively:

$$\hat{\boldsymbol{u}}_{g} = \begin{pmatrix} \hat{\boldsymbol{u}}_{0} \\ \hat{\boldsymbol{u}}_{1} \\ \vdots \\ \hat{\boldsymbol{u}}_{N-1} \end{pmatrix}, \quad \hat{\boldsymbol{u}}_{l} = \begin{pmatrix} \hat{\boldsymbol{u}}^{1} \\ \hat{\boldsymbol{u}}^{2} \\ \vdots \\ \hat{\boldsymbol{u}}^{N_{\text{el}}} \end{pmatrix}$$

where $u^e = (\hat{u}_0^e, \dots, \hat{u}_P^e)^{\top}$ is the list of all local coefficients for element Ω^e . Then \hat{u}_g and \hat{u}_l can be related by the *assembly matrix* **A** of size $(P + 1)N_{\text{el}} \times N$ so that

$$\hat{u}_l = A\hat{u}_g$$
.

In the finite element setting of figure 2.1, this system becomes

$$\begin{pmatrix} \hat{u}_{0}^{1} \\ \hat{u}_{1}^{1} \\ \hat{u}_{0}^{2} \\ \hat{u}_{1}^{2} \\ \hat{u}_{0}^{3} \\ \hat{u}_{1}^{3} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \hat{u}_{0} \\ \hat{u}_{1} \\ \hat{u}_{2} \\ \hat{u}_{3} \end{pmatrix}$$

Notice that the reverse operation of obtaining the global from local coefficients can also be performed by using the transpose of the assembly matrix, so that

$$\hat{u}_g = \mathbf{A}^{\top} \hat{u}_l.$$

The assembly matrix mostly provides a convenient way of performing operations locally on elements, and combining each elemental contribution in the correct fashion as to produce the correct result. A key operation in the Galerkin formulation is the use of integration of global modes in, for example, the construction of the inner product $(u^{\delta}, \Phi_k)_{L^2}$. The integral for each global mode can be stored in a vector I_g , so that

$$(I_g)_k = \int_{\Omega} u^{\delta}(x) \Phi_k(x) \,\mathrm{d}x.$$

Typically this integral will be evaluated numerically, and the process for this will be outlined

in the next section. Defining I_l to be the vector of local contributions

$$I_{l} = \begin{pmatrix} I^{1} \\ \vdots \\ I^{N_{\text{el}}} \end{pmatrix} \text{ where } I^{e} = \begin{pmatrix} \int_{\Omega_{\text{st}}} u^{\delta}(\chi^{e}(\xi))\phi_{0}(\xi)\frac{\mathrm{d}\chi^{e}}{\mathrm{d}\xi}\,\mathrm{d}\xi \\ \vdots \\ \int_{\Omega_{\text{st}}} u^{\delta}(\chi^{e}(\xi))\phi_{P}(\xi)\frac{\mathrm{d}\chi^{e}}{\mathrm{d}\xi}\,\mathrm{d}\xi \end{pmatrix}$$

we can therefore relate local and global integrals by the simple relation

$$I_g = \mathbf{A}^\top I_l.$$

To calculate the global integral over each mode, we may calculate the contributions locally of each element and combine their results according to the local-to-global mapping. By using χ^e to project each element onto Ω_{st} , we are able to use the elemental basis functions ϕ_p and hence never need to construct the complex global modes.

Similarly, the mass and Laplacian matrices can be constructed in an elemental fashion. First we define the local matrices \mathbf{M}^e and $\mathbf{L}^e \in \mathbb{R}^{(P+1) \times (P+1)}$ by

$$\mathbf{M}_{ij}^{e} = \int_{\Omega^{e}} \phi_{i}^{e}(x)\phi_{j}^{e}(x) \,\mathrm{d}x = \int_{\Omega_{\mathrm{st}}} \phi_{i}(\xi)\phi_{j}(\xi)\frac{\mathrm{d}\chi^{e}}{\mathrm{d}\xi}(\xi) \,\mathrm{d}\xi,$$
$$\mathbf{L}_{ij}^{e} = \int_{\Omega^{e}} \frac{\mathrm{d}\phi_{i}^{e}}{\mathrm{d}x}(x)\frac{\mathrm{d}\phi_{j}^{e}}{\mathrm{d}x}(x) \,\mathrm{d}x = \int_{\Omega_{\mathrm{st}}} \frac{\mathrm{d}\phi_{i}}{\mathrm{d}\xi}(\xi)\frac{\mathrm{d}\phi_{j}}{\mathrm{d}\xi}(\xi)\frac{\mathrm{d}\chi^{e}}{\mathrm{d}\xi}(\xi) \,\mathrm{d}\xi.$$

From these formulae it is also clear that it is only necessary to generate a single elemental mass matrix, since the derivative of each of the projection function χ^e is the constant

$$\frac{\mathrm{d}\chi^e}{\mathrm{d}\xi} = \frac{x_e - x_{e-1}}{2}$$

and can therefore be brought outside of the integral above. This enables us to take advantage of a significant computational optimisation when constructing the global counterparts **M** and **L**. We first assemble matrices \mathbf{M}_l and $\mathbf{L}_l \in \mathbb{R}^{(P+1)N_{\text{el}} \times (P+1)N_{\text{el}}}$, so that each \mathbf{M}^e or \mathbf{L}^e lies on the appropriate position along the diagonal. Then

$$\mathbf{M} = \mathbf{A}^{\top} \mathbf{M}_l \mathbf{A}.$$

Finally, note that A is generally a very sparse matrix whose only entries are 1, 0 and -1 in

a multi-dimensional formulation. For this reason it is usually undesirable to construct **A** explicitly, and instead a common sparse-matrix technique such as a mapping array is used to perform these scatter-gather operations.

2.2.3 Common choices for the expansion basis

The choice of local expansion basis functions ϕ_i is of great importance as it has a large effect on the potential efficiency and numerical accuracy of the code in two regards. Firstly, the computational cost of constructing the matrices initially is important, as they will require numerical integration, which is an expensive operation. Secondly, and more importantly in the context of this project given the large domains that need to be represented, is the resulting structure of matrices derived from the basis functions. Indeed, in large computational problems, the dimensions of the mass matrix may preclude it from being stored in computer memory, so the ability to exploit any structure of the matrices is essential.

Basis functions are classified into two general groups known as *modal* or *nodal* functions. To distinguish between these types, consider the following examples for $0 \le p \le P$:

$$\phi_p^A(x) = x^p, \tag{2.14}$$

$$\phi_p^B(x) = \ell_p(x) = \prod_{i,j=0, i \neq j}^P \frac{x - x_i}{x_j - x_i}$$
(2.15)

 ϕ_p^B is a basis using the Lagrange polynomials $\ell_p(x)$. It is an example of a nodal basis, since we first need to define a set of points $\{x_i \mid 0 \le i \le P\}$ from which $\ell_p(x)$ can be defined. ϕ_p^A is commonly called the *moment basis*, and is an example of a modal basis, where no such requirement is needed.

There are two computational costs we must consider. First, let us calculate the cost of constructing the elemental mass matrix \mathbf{M}^{e} . Each entry of the matrix may be obtained through

$$\mathbf{M}_{ij}^{e} = \int_{\Omega_{st}} x^{i} x^{j} \, \mathrm{d}x = \left[\frac{x^{i+j+1}}{i+j+1}\right]_{-1}^{1} = \begin{cases} \frac{2}{i+j+1}, & \text{if } p+q \text{ is even} \\ 0, & \text{if } p+q \text{ is odd.} \end{cases}$$

So in constructing this matrix we can ignore half of the entries. Similarly, if we assume



(a) Interior modes of the standard modal expansion basis $\psi_p(\xi)$, normalised so that $-1 \leq \psi_p(\xi) \leq 1$.



(b) Standard nodal basis using Lagrange polynomials $\ell_p(\xi)$, with nodal points located using the GLL quadrature points.

Figure 2.2: Typical modal and nodal expansion bases defined inside the standard element Ω_{st} .

that the modal points x_i are equally spaced in Ω_{st} , then the mass matrix produced by using $\phi_p^B(x)$ is full, with no non-zero entries in general. Thus, it is relatively expensive to construct the mass matrix using the functions defined thus far.

Typically these matrices will only need to be constructed once in a time-dependent problem such as the heat or Navier-Stokes equations, and as such the more important problem is that of inversion. The key to quick numerical inversion of matrices is to exploit any structure of the matrix. The only structure of note that either of these matrices are guaranteed to possess is that they are symmetric, since clearly

$$\mathbf{M}_{ii}^{e} = (\phi_{i}, \phi_{j})_{L^{2}} = (\phi_{j}, \phi_{i})_{L^{2}} = \mathbf{M}_{ii}^{e}.$$

Unfortunately, this structure is not generally useful for optimising inversion, although it allows for use of common methods such as LU decomposition and Cholesky factorisation. Another key factor which must be considered is the numerical error incurred in the inversion process. It is well-known that the condition number of a matrix A,

$$\kappa(\mathbf{A}) = \|\mathbf{A}\| \cdot \|\mathbf{A}^{-1}\|$$

provides a good indicator of the error that can be expected when a A is numerically inverted. $\|\cdot\|$ is a suitably chosen norm; for example, the L^2 matrix norm. The larger κ becomes, the larger the round-off error incurred in the process of inverting the matrix. For both choices of basis shown so far, κ becomes extremely large for even moderate values of P.

Finally, we must also consider the combination with the *h*-type refinement seen earlier; for example in figure 2.1. The condition on the local coefficients \hat{u}_p^e seen there would not necessarily guarantee that the global modes remained in C^0 . Instead, we would need a condition like

$$\sum_{p=0}^{P} \hat{u}_{p}^{e} \phi_{p}^{e}(1) = \sum_{p=0}^{P} \hat{u}_{p}^{e+1} \phi_{p}^{e+1}(-1),$$

which is tricky to implement, and breaks the symmetry of global matrices. This problem may be easily solved if we separate boundary modes from interior modes, which is enforced by the condition

$$\phi_p(-1) = \begin{cases} 1, & p = 0, \\ 0, & p \neq 0, \end{cases} \qquad \phi_p(1) = \begin{cases} 1, & p = 0, \\ 0, & p \neq 0, \end{cases}$$

then just as in figure 2.1, we can prescribe the condition

$$\hat{u}_{P}^{e} = \hat{u}_{0}^{e+1}$$

for adjacent elements. In summary, we require modes which:

- are orthogonal, or close to orthogonal, under the L^2 inner-product and have exploitable structure;
- produce matrices which have low condition number κ ;
- implement the interior-boundary separation.

One possible choice of mode which accomplishes most of these goals may be derived from a family of functions known as *Jacobi polynomials*.

Definition 2.2.3 (Jacobi polynomials). The solution to the one-dimensional ordinary differential equation:

$$\frac{\mathrm{d}}{\mathrm{d}x}\left[(1-x)^{1+\alpha}(1+x)^{1+\beta}\frac{\mathrm{d}}{\mathrm{d}x}u_p(x)\right] = \lambda_p(1-x)^{\alpha}(1+x)^{\beta}u_p(x)$$

where -1 < x < 1 and

$$\lambda_p = -p(\alpha + \beta + p + 1).$$

is defined to be the Jacobi polynomial $J_p^{\alpha,\beta}(x) := u_p(x)$.

Jacobi polynomials, under the correct weighting function, are orthogonal, since

$$\int_{\Omega_{\rm st}} (1-x)^{\alpha} (1+x)^{\beta} P_p^{\alpha,\beta}(x) P_q^{\alpha,\beta}(x) \,\mathrm{d}x = C_p^{\alpha,\beta} \delta_{pq} \tag{2.16}$$

where δ_{pq} is the standard Kronecker delta, and

$$C_p^{\alpha,\beta} = \frac{2^{\alpha+\beta+1}}{2p+\alpha+\beta+1} \cdot \frac{\Gamma(p+\alpha+1)\Gamma(p+\beta+1)}{p!\,\Gamma(p+\alpha+\beta+1)}.$$

Combining this result with the properties desired above, we construct the modes $\psi_p(\xi)$ visualised in figure 2.2(a):

$$\phi_p(\xi) = \psi_p(\xi) = \begin{cases} \frac{1-\xi}{2}, & p = 0, \\ \frac{(1-\xi)(1+\xi)}{4} J_{p-1}^{1,1}(\xi), & 0$$

When p = 0 and p = P, the modes are linear and identical to those seen in figure 2.1 and represent the boundary modes. The weighting functions shown in the orthogonality relationship are incorporated into the choice of the interior modes to decouple them as much as possible.

The results of applying this basis to \mathbf{M}^{e} is shown in figure 2.3(a). Although \mathbf{M}^{e} is still not completely banded as the boundary modes are not orthogonal to the interior modes, the



Figure 2.3: Elemental mass matrix \mathbf{M}^e and Laplacian matrix \mathbf{L}^e for different choices of elemental modes ϕ_p at polynomial order P = 20. Dots indicate non-zero entries. (a) The most common choice for a modal expansion $\phi_p = \psi_p$. (b) The result of changing the interior modes seen in (a) to $J_{p-1}^{3,3}(\xi)$.

degree of coupling is quite weak, and the resultant matrix has structure which is easy to exploit. This is emphasised even more in the Laplacian matrix \mathbf{L}^{e} which is almost diagonal. These matrices are much sparser and can be stored in a very condensed format, avoiding one of the key issues of size.

Figure 2.3(b) shows a choice of basis similar to $\psi_p(\xi)$ in that it uses the same boundary modes, but with Jacobi function $J_{p-1}^{3,3}(\xi)$. The resultant mass matrix is mostly full, and emphasises that ψ_p represents a good choice of modal basis for this application.

2.2.4 Elemental Operations

The final challenge in implementing the one-dimensional spectral-element method is integration and differention of the modal functions, which forms a key part in the Galerkin formulation of the problem (e.g. in constructing \mathbf{M}^{e} and \mathbf{L}^{e}). In this simple example of the diffusion equation it is technically possible to create analytic formulae to calculate integrals such as

$$\int_{\Omega_{\rm st}}\phi_i(\xi)\phi_j(\xi)\,\mathrm{d}\xi.$$

However, the process is both tedious and prone to error. In addition, an increase of the polynomial order inside an element requires us to have pre-computed the formulae by hand. Instead, we focus on evaluating these quantities numerically so that the process can be fully automated. The concept of numerical integration of functions is relatively simple. One picks a set of points $\{\xi_i \in \Omega_{st} \mid 1 \leq i \leq Q\}$ and a set of weights $\{w_i \in \Omega_{st} \mid 1 \leq i \leq Q\}$ and approximates

$$\int_{\Omega_{\rm st}} f(\xi) \,\mathrm{d}\xi \approx \sum_{i=1}^{Q} w_i f(\xi_i). \tag{2.17}$$

Typically the nodes are placed in an equidistant fashion, which leads to classical expansions such as the trapezoidal rule. In general however, given that we have a choice of Q nodes and weights, giving 2Q unknowns, it should be possible to obtain exact values for (2.17) if f is a polynomial of order 2Q - 1. This remarkable and somewhat counter-intuitive discovery was made by Gauss and is very well documented as an extremely accurate method for integration.

To derive these weights and nodal positions exactly, we may solve the system of equations

$$\sum_{i=1}^{Q} w_i \xi_i^j = \int_{\Omega_{\mathrm{st}}} \xi^j \, \mathrm{d}\xi, \quad 0 \leq j \leq 2Q - 1.$$

Unfortunately, this is a somewhat difficult system to solve and a much better method relies on sets of orthogonal polynomials – in this case, a subfamily of the family of Jacobi polynomials.

Definition 2.2.4 (Legendre polynomials). The polynomials $L_p(\xi) = J_p^{0,0}(\xi)$ are called *Legendre polynomials*.

Legendre polynomials satisfy many interesting properties. They are clearly orthogonal in [-1, 1] by equation (2.16). More importantly, the set $\{L_p(\xi) \mid p \in \mathbb{N} \cup \{0\}\}$ forms a complete set of functions on [-1, 1]. Using this property, we can prove the following theorem:

Theorem 2.2.5 (Gaussian Quadrature). Let $\xi_{i,Q}^{\alpha,\beta}$ be the zeros of the Jacobi polynomial $J_Q^{\alpha,\beta}(\xi)$ of order Q so that

$$J_Q^{\alpha,\beta}(\xi_{i,Q}^{\alpha,\beta}) = 0, \quad 1 \leq i < Q.$$

Then (2.17) is an exact relation if f is a polynomial with deg f < 2Q if

$$\xi_i = \xi_{i,Q}^{0,0}, \qquad w_i = \frac{2}{1-\xi_i^2} \left[\left. \frac{\mathrm{d}}{\mathrm{d}\xi} (L_Q(\xi)) \right|_{\xi=\xi_i} \right]^{-2}.$$

In a spectral element formulation, we do not use these classical Gaussian quadrature points, but instead a slightly modified set known as *Gauss-Lobatto-Legendre* (GLL) points which incorporate nodal points at the boundary. These are defined by

$$\xi_i = \begin{cases} -1, & i = 1, \\ \xi_{i,Q-2}^{1,1}, & 1 < i < Q, \\ 1, & i = Q, \end{cases} \quad w_i = \frac{2}{Q(Q-1)[P_Q^{0,0}(\xi_i)]^2}.$$

In this case, (2.17) is only exact if deg $f \leq 2Q-3$. By using this choice of nodal points with a nodal basis such as the Lagrange polynomials $\ell_p(\xi)$ as depicted in figure 2.2(b), the elemental mass matrix becomes diagonal, improving the efficiency of the code. (Unfortunately this is not the case for the Laplacian matrix). The Laplacian matrix requires knowledge of the derivatives of the modal functions ϕ_p at the collocation points. This can either be accomplished analytically, or through numerical means using a similar scheme and is detailed fully in Karniadakis & Sherwin (2005).

Finally then, we can draw some conclusions on the number of quadrature points needed in the setting of the diffusion equation. For the construction of the mass matrix, which is simply the product of two polynomials of order P, the minimum number of points required to integrate this exactly (up to machine precision) is $Q_{\min} \ge P + \frac{3}{2}$; hence P + 2 points is enough. Since the diffusion equation is linear, we do not expect that any of the terms involved will require more points than this. However in the case of the Navier-Stokes equations, the convective term $(\mathbf{u} \cdot \nabla)\mathbf{u}$ will in general be represented by a polynomial of order 3P, and thus will require $Q_{\min} \ge \frac{3}{2}(P + 1) - i.e.$ a minimum of 3P/2 + 2 terms.

2.2.5 Implementation

The culmination of the previous sections then is the implementation of the one-dimensional spectral-element method for the diffusion equation. To test the validity of the method, we



Figure 2.4: Solution to the diffusion equation using the parameters outlined in section 2.2.5 with the initial condition u(x) = x(1-x) for $(x,t) \in \Omega \times [0,1]$. The distribution of Gauss-Lobatto-Legendre (GLL) points inside each element can be seen along the *x*-axis.

pick the following parameters:

- $\Omega = [0, 1];$
- Homogeneous Neumann boundary conditions: $g_0 = g_1 = 0$;
- Standard modal basis functions $\psi(\xi)$ inside each element;
- Initial condition $u_0(x) = x(1-x)$;
- $\Delta t = 10^{-3}$ and diffusion coefficient $\alpha = 1$, with the final time T = 1.

The program runs in the following order:

- 1. Initial set-up: allocate memory for arrays; define element boundaries; create mapping array used to represent the assembly matrix **A**; generate quadrature points and weights.
- 2. Generate elemental mass and Laplacian matrices \mathbf{M}^{e} and \mathbf{L}^{e} .
- 3. Assemble the global matrices **M** and **L**, and use these to construct the LHS matrix $(\mathbf{M} + \alpha \Delta t \mathbf{L})$. Calculate the LU decomposition of this matrix.
- 4. Inside each element, construct the initial condition u_e^0 and apply Galerkin projection to obtain the initial global coefficients \hat{u}_g^0 .

- 5. Enter the timestepping loop:
 - Construct the global RHS: $-\mathbf{M}\hat{u}_{g}^{n} + \alpha \Delta t \mathbf{B}$.
 - Invert the global problem to obtain \hat{u}_g^{n+1} .
- 6. Once the desired number of iterations has been reached, use A to reconstruct the solution inside each element.

The resultant solution field can be seen in figure 2.4. We first check to see that the approximate solution converges to the correct constant function. Recall that the exact solution to the diffusion equation with homogeneous Neumann boundary conditions is well known:

$$u(x,t) = a_0 + \sum_{n=1}^{\infty} a_n e^{-n^2 \pi^2 t} \cos(n\pi x)$$

where

$$a_0 = \int_0^1 u_0(x) \, dx$$
 and $a_n = 2 \int_0^1 u_0(x) \cos(n\pi x) \, dx$.

Hence for $u_0(x) = x(1-x)$ we calculate

$$a_0 = \frac{1}{6}$$
 and $a_n = \begin{cases} -\frac{4}{n^2 \pi^2}, & n \text{ even,} \\ 0, & n \text{ odd} \end{cases}$

which gives the solution

$$u(x,t) = \frac{1}{6} - \sum_{n=1}^{\infty} \frac{1}{n^2 \pi^2} e^{-4n^2 \pi^2 t} \cos(2n\pi x).$$

We also have that $\lim_{t\to\infty} u(x,t) = \frac{1}{6} \approx 0.168$ for each $x \in \Omega$ since

$$\left|\sum_{n=1}^{\infty} \frac{1}{n^2 \pi^2} e^{-4n^2 \pi^2 t} \cos(2n\pi x)\right| \leq \frac{e^{-4\pi^2 t}}{\pi^2} \sum_{n=1}^{\infty} \frac{1}{n^2} = \frac{1}{6} e^{-4\pi^2 t} \xrightarrow{t \to \infty} 0.$$

This clearly corresponds to the results seen in figure 2.4.

2.2.6 Convergence Properties

In order to test the convergence properties of the spectral element method, we use the same parameters of the previous section, but choose a different initial condition $u_0(x) = \cos(2\pi x)$, since the function x(1-x) does not satisfy the Neumann boundary conditions. This has the additional bonus of greatly simplifying the exact solution, since now

$$a_n = \begin{cases} \frac{1}{2}, & n = 2, \\ 0, & n \neq 2 \end{cases}$$

by orthogonality of the cosine function. Hence the new exact solution is

$$u(x,t) = e^{-4\pi^2 t} \cos(2\pi x).$$

To minimise the $O(\Delta t)$ timestepping error inherent to the implicit Euler timestepping scheme, we also set $\Delta t = 10^{-8}$ and change the final time to $T = 10^{-6}$.

One of the simplest methods for obtaining convergence statistics is through the error term $E_f = ||u^{\delta}(\cdot, T) - u(\cdot, T)||$, where the norm $|| \cdot ||$ is problem-specific. Typically choices include the discrete L^{∞} or L^2 norms. However, since this problem utilises a Galerkin scheme, a more appropriate norm can be defined which provides an estimate for the average error.

Definition 2.2.6 (Energy Norm). Let a(u, v) be a symmetric bilinear form. Then we define the *energy norm* $\|\cdot\|_E$ by

$$\|u\|_E = \sqrt{a(u,u)}.$$

For the Galerkin scheme considered for the diffusion equation,

$$a(u,v) = \int_{\Omega} \frac{\partial u}{\partial x} \frac{\partial v}{\partial x} \, \mathrm{d}x \Rightarrow \|u\|_E = \int_{\Omega} \left(\frac{\partial u}{\partial x}\right)^2 \mathrm{d}x.$$

To test each refinement method, two sets of simulations were run:

- *p*-type: Set $N_{\rm el} = 3$ and vary $1 \le P \le 200$.
- *h*-type: Set P = 1 and vary $1 \le N_{\text{el}} \le 200$.



Figure 2.5: Logarithmic plot of the error E_f measured using the energy norm for *h*- and *p*-type convergence (top and bottom lines respectively).

Figure 2.5 collates the data from the above simulations on a logarithmic scale, and shows some of the excellent convergence properties of the spectral-element method. *h*-type convergence is linear; that is, the error scales as $O(N_{el})$. This is to be expected, since each additional element only increases the number of global modes by P = 1. *p*-type convergence on the other hand is exponential – as the polynomial order is increased linearly, the error decreases exponentially. This exponential convergence greatly improves the quality of approximate solutions, and shows why the spectral-element method is well-suited to problems requiring extremely fine spatial resolution.

2.3 Temporal Discretisations for Navier-Stokes

So far, we have introduced a highly accurate spectral element scheme. However, the use of such schemes is generally preserved for spatial co-ordinates in which a great deal of accuracy is required. The discretisation of many non-linear systems often leads to a Courant-Friedrichs-Lewy (CFL) condition dictating the stability of solutions, which enforces restrictions on the timestep Δt and thus rendering it very small. We do not therefore require exponential convergence of solutions temporally; for even the most detailed turbulence simulations,

 $O(\Delta t^3)$ convergence is sufficient.

These high-order schemes are common throughout the numerical analysis of PDEs. However, the unique pressure coupling inherent in Navier-Stokes makes the development of high-order timestepping techniques difficult. In this section, we introduce some basic high-order timestepping routines and show how they can be used in the discretisation of Navier-Stokes.

2.3.1 Classical high-order schemes

We begin by considering the problem of discretising the temporal domain for the PDE

$$\frac{\partial \boldsymbol{u}}{\partial t} = F(\boldsymbol{u}(\boldsymbol{x}), \boldsymbol{x}, t)$$

for $t \in [0, T]$. Virtually all of the time-stepping schemes proceed with the assumption that the discretisation of the domain is equally spaced, so that we approximate the solution at times $t_n = n\Delta t$, where Δt is commonly referred to as the *time-step*.

There are a vast array of time-stepping schemes which take the solution field at time t_n , $u^n := u(x, t_n)$ to u^{n+1} . They can be generally classified into two groups. *Explicit* schemes only rely on knowledge of u^n and are cheap to implement. On the other hand, *implicit* schemes require solving an equation which is dependent upon both u^n and u^{n+1} . One of the most popular schemes is the θ -method, which uses a mixture of both implicit and explicit terms.

Definition 2.3.1 (θ -method). Let $F^n = F(u^n, x, t_n)$. Then the θ -method time-stepping scheme for obtaining u^{n+1} solves the (implicit) equation

$$\frac{\boldsymbol{u}^{n+1}-\boldsymbol{u}^n}{\Delta t}=(1-\theta)\boldsymbol{F}^n+\theta\boldsymbol{F}^{n+1}.$$

In particular, there are three special choices for θ :

$$\theta = \begin{cases} 0, & \text{explicit Euler,} \\ \frac{1}{2}, & \text{Crank-Nicolson,} \\ 1, & \text{implicit Euler.} \end{cases}$$

It can be shown that both Euler methods incur errors of $O(\Delta t)$ in the approximation. The

benefit of implicit schemes is their improved stability, with fewer restrictions on the timestep Δt and hence they are often preferred over explicit schemes. For example, when applied to the diffusion equation, the explicit Euler scheme only remains stable if the CFL condition

$$\frac{\Delta t}{\Delta x^2} \le \frac{1}{2}$$

is satisfied. Of particular significance is Crank-Nicholson time-stepping, which only incurs a high-order $O(\Delta t^2)$ temporal error. Due to its relative simplicity and high-order accuracy, it is highly popular.

Sometimes however, we require even higher-order methods. In this case, we generally employ the use of *multi-step* methods, which approximate the time derivative through use of previous time iterations. Two of the most widely used multi-step methods involve approximations of the form

$$\frac{\boldsymbol{u}^{n+1}-\boldsymbol{u}^n}{\Delta t}=\sum_{q=0}^{J-1}\beta_q \boldsymbol{F}^{n+1-q}.$$

The choice of β_q determines both whether the scheme is explicit or implicit. If $\beta_0 = 0$ then the scheme is explicit and is known as the *Adams-Bashforth* method; otherwise the scheme becomes implicit and is known as the *Adams-Moulton* method. The scheme can in theory be extended to arbitrary order by increasing the integration order J at the expense of storing past timesteps in memory.

2.3.2 Splitting Schemes

None of the schemes above are directly applicable to the Navier-Stokes equations

$$\frac{\partial \boldsymbol{u}}{\partial t} + (\boldsymbol{u} \cdot \nabla)\boldsymbol{u} = -\nabla p + \nu \nabla^2 \boldsymbol{u}, \qquad (2.18)$$

$$\nabla \cdot \boldsymbol{u} = \boldsymbol{0}, \tag{2.19}$$

due to the coupling of the velocity field u with the pressure field p in equation (2.18). Additionally, the problem of enforcing the incompressibility condition (2.19) is not addressed. A key aspect of any numerical scheme for Navier-Stokes is the ability to calculate u without knowing the pressure field p. One of the approaches for creating algorithms for time-stepping Navier-Stokes is the *splitting* scheme. We introduce the concept by examining the simple one dimensional problem

$$\frac{\partial u}{\partial t} = -(\mathcal{L}_1 + \mathcal{L}_2)u$$
$$u(x, 0) = u_0(x),$$

where each \mathcal{L}_i is a (potentially non-linear) differential operator. We may obtain an approximate solution by solving the two equations:

$$\frac{\partial \hat{u}}{\partial t} + \mathcal{L}_1 \hat{u} = 0 \text{ in } (t_n, t_{n+1})$$
$$\frac{\partial \hat{\hat{u}}}{\partial t} + \mathcal{L}_2 \hat{\hat{u}} = 0 \text{ in } (t_n, t_{n+1})$$

where we set $\hat{u}(t_n) = u^n$ and $\hat{u}(t_n) = \hat{u}^{n+1}$. The estimated solution at time t_{n+1} is $u^{n+1} = \hat{u}(t_{n+1})$. This simple method, introduced by Yanenko (1971), incurs a *splitting* error of $O(\Delta t)$.

To each of these equations we apply a timestepping scheme in order to solve the overall system. Since the error incurred through splitting is $O(\Delta t)$, we use implicit Euler which maintains the same level of accuracy, giving

$$\frac{\hat{u} - u^n}{\Delta t} + \mathcal{L}_1 \hat{u} = 0,$$
$$\frac{u^{n+1} - \hat{u}}{\Delta t} + \mathcal{L}_2 u^{n+1} = 0.$$

To ensure the problem is well-posed, boundary conditions must be properly applied to each sub-step of the problem. The nature of these conditions is problem-specific, and can often be regarded as artificial in the sense that whilst they are often derived from the original equations, they are not mathematically required to solve the overall system.

2.3.3 First-order splitting schemes for Navier-Stokes

The splitting scheme outlined above provides a method for solving the Navier-Stokes equations by omitting the pressure term ∇p in the first substep, an idea first introduced by Temam (1969). Initially, we ignore the incompressibility constraint and solve for an intermediate velocity field \hat{u} given the initial condition u^n . The second sub-step then imposes a pressure correction which projects \hat{u} onto a divergence-free space. The scheme then is

$$\frac{\hat{\boldsymbol{u}} - \boldsymbol{u}^n}{\Delta t} + (\hat{\boldsymbol{u}} \cdot \nabla)\hat{\boldsymbol{u}} = \nu \nabla^2 \hat{\boldsymbol{u}}$$
(2.20a)

$$\frac{\boldsymbol{u}^{n+1} - \hat{\boldsymbol{u}}}{\Delta t} = -\nabla \overline{p}^{n+1} \tag{2.20b}$$

where (2.20a) is subject to the no-slip boundary condition

$$\hat{\boldsymbol{u}}(\partial\Omega) = \boldsymbol{0}$$

and (2.20b) satisfies

$$\nabla \cdot \boldsymbol{u}^{n+1} = \boldsymbol{0} \tag{2.21a}$$

$$\boldsymbol{u}^{n+1} \cdot \boldsymbol{n} = \boldsymbol{0}. \tag{2.21b}$$

where *n* is the normal unit vector to $\partial \Omega$. It is important to note that the field \overline{p} is *not* the exact pressure field *p*. In order to obtain an explicit equation for the pressure, we take the divergence of (2.20b) to give the Poisson equation

$$\nabla^2 \overline{p}^{n+1} = \nabla \cdot \left(\frac{\hat{\boldsymbol{u}}}{\Delta t}\right). \tag{2.22}$$

The choice of pressure boundary condition here is critical. For this first order scheme, we consider the projection of (2.20b) in the normal direction to obtain the equation

$$\frac{\partial \overline{p}^{n+1}}{\partial n} = 0 \text{ on } \partial \Omega.$$
(2.23)

However, by applying the divergence operator to (2.18), we see that the exact pressure p will in fact satisfy the equations

$$\nabla^2 p = -\nabla \cdot [(\boldsymbol{u} \cdot \nabla)\boldsymbol{u}] \quad \text{in } \Omega,$$
$$\frac{\partial p}{\partial n} = \boldsymbol{n} \cdot \nu \nabla^2 \boldsymbol{u} \qquad \text{on } \partial \Omega.$$

It is clear that these equations differ from one another, and hence \overline{p}^{n+1} can only be an approximation to p. However, Temam (1979) showed that this does not affect the convergence

of the scheme. Error analysis shows that the velocity field is subject to an error of $O(\Delta t)$, whereas the pressure has error $O(\sqrt{\Delta t})$, although with modifications it can be corrected to give $O(\Delta t)$ accuracy.

The choice of boundary condition on the pressure-correction substep is critical in obtaining a scheme which is numerically accurate. In the next section we will extend this splitting method to implement a higher-order timestepping scheme.

2.3.4 High-order schemes

There have been many attempts to introduce high-order splitting schemes to Navier-Stokes. Classically, we re-write (3.1a) as

$$\frac{\partial \boldsymbol{u}}{\partial t} = -\nabla p + \nu \mathbf{L}(\boldsymbol{u}) + \mathbf{N}(\boldsymbol{u})$$
(2.24)

where L and N represent the viscous dissipation and non-linear advection operators of Navier-Stokes respectively. Upon integrating (2.24) between t_n and t_{n+1} we obtain the equation

$$\boldsymbol{u}^{n+1} - \boldsymbol{u}^n = -\int_{t_n}^{t_{n+1}} \nabla p \, \mathrm{d}t + \nu \int_{t_n}^{t_{n+1}} \mathbf{L}(\boldsymbol{u}) \, \mathrm{d}t + \int_{t_n}^{t_{n+1}} \mathbf{N}(\boldsymbol{u}) \, \mathrm{d}t.$$

We then re-write these three integrals as

$$\int_{t_n}^{t_{n+1}} \nabla p \, \mathrm{d}t = \Delta t \, \nabla \overline{p}^{n+1}, \qquad (2.25a)$$

$$\int_{t_n}^{t_{n+1}} \mathbf{N}(\boldsymbol{u}) \,\mathrm{d}t = \Delta t \sum_{q=0}^{J_e-1} \beta_q \mathbf{N}(\boldsymbol{u}^{n-q}), \qquad (2.25b)$$

$$\int_{t_n}^{t_{n+1}} \mathbf{L}(\boldsymbol{u}) \, \mathrm{d}t = \Delta t \sum_{q=0}^{J_i-1} \gamma_q \mathbf{L}(\boldsymbol{u}^{n+1-q}).$$
(2.25c)

In these equations, \overline{p}^{n+1} is the pressure field ensuring that the final velocity is incompressible at time t_{n+1} , and β_q and γ_q are the coefficients giving Adams-Bashforth/Adams-Moulton schemes of order J_e and J_i respectively. From these equations, we obtain the three-step splitting scheme

$$\frac{\hat{\boldsymbol{u}} - \boldsymbol{u}^n}{\Delta t} = \sum_{q=0}^{J_e - 1} \beta_q \mathbf{N}(\boldsymbol{u}^{n-q}), \qquad (2.26a)$$

$$\frac{\hat{\hat{\boldsymbol{u}}} - \hat{\boldsymbol{u}}}{\Delta t} = -\nabla \overline{p}^{n+1}, \qquad (2.26b)$$

$$\frac{u^{n+1} - \hat{\hat{u}}}{\Delta t} = \nu \sum_{q=0}^{J_i - 1} \gamma_q \mathbf{L}(u^{n+1-q}).$$
(2.26c)

Following the classical high-order approach, a timestep is performed as follows. First, we calculate \hat{u} from (2.26a). Then, assuming that $\nabla \cdot \hat{\hat{u}} = 0$, we take the divergence of (2.26b) to obtain the Poisson equation

$$\nabla \cdot \left(\frac{\hat{\boldsymbol{u}}}{\Delta t}\right) = \nabla^2 \overline{p}^{n+1}.$$
(2.27)

Using this pressure field to calculate $\hat{\hat{u}}$, we then solve (2.26c) for the solution u^{n+1} .

2.3.5 Pressure boundary condition

In practice, to solve these equations we require boundary conditions for equations (2.27) and (2.26c). The latter may be easily derived by imposing the desired Dirchlet boundary condition $u^{n+1} = u_0$ on $\partial \Omega$. However, the more artificial pressure boundary condition must be derived. In the classical scheme, we assume that $\hat{\hat{u}} \cdot n = u_0 \cdot n$, leading to the boundary condition

$$\frac{\partial \overline{p}^{n+1}}{\partial n} = -\frac{u_0 \cdot n - \hat{u} \cdot n}{\Delta t}$$

on $\partial\Omega$. However, whilst this choice of boundary condition is very efficient, it leads to large splitting errors, as shown by Marcus (1984) in a study of Taylor-Couette flow using this scheme. The exact pressure boundary condition may be derived from the semi-discrete form of (2.24);

$$\frac{\partial \overline{p}^{n+1}}{\partial n} = n \cdot \left[\sum_{q=0}^{J_e-1} \beta_q \mathbf{N}(\boldsymbol{u}^{n-q}) + \nu \sum_{q=0}^{J_i-1} \gamma_q \mathbf{L}(\boldsymbol{u}^{n+1-q}) \right].$$
(2.28)

Unfortunately, this equation is coupled due to the implicit linear term at the (n + 1)-th time level. Approximating this term by an explicit expansion also leads to numerical inaccuracies.

One of the most critical observations in the high-order scheme is the identification of a pressure boundary condition which gives low splitting errors. Orszag *et al.* (1986) finally derived such a condition by rewriting the viscous operator L in (2.28) as the summation of a rotational and irrotational part ∇Q estimated by explicit and implicit schemes respectively. This then gives

$$\frac{\partial \overline{p}^{n+1}}{\partial n} = n \cdot \left[\sum_{q=0}^{J_e-1} \beta_q \mathbf{N}(\boldsymbol{u}^{n-q}) + \underbrace{\nu \sum_{q=0}^{J_i-1} \gamma_q \nabla Q^{n+1-q}}_{\text{irrotational}} - \underbrace{\nu \sum_{q=0}^{J_e-1} \beta_q [\nabla \times (\nabla \times \boldsymbol{u})]^{n-q}}_{\text{rotational}} \right]$$

A straightforward analysis shows that the numerical errors which occur in the previous choices are not present with this condition. Since we require that $Q^{n+1} = 0$ to satisfy the incompressibility condition, we can therefore take $\gamma_0 = 0$, and this then removes the implicit coupling of the boundary condition.

2.3.6 Stiffly stable schemes

Whilst the classical high-order splitting scheme coupled with the high-order pressure boundary condition produce a time-stepping scheme capable of producing very accurate results, the use of Adams-Bashforth and Adams-Moulton leads to large restrictions in Δt in order to avoid numerical instability, especially at higher orders. The stability diagram for this method can be found in Karniadakis & Sherwin (2005). To address this, Karniadakis *et al.* (1991) adapted the classical scheme to incorporate a *stiffly stable* multi-step method.

Stiffy stable schemes, detailed thoroughly in Gear (1971), are used in the study of equations which contain solutions which grow at widely differing rates. For example, consider the system of ODEs

$$y'(t) = 998y(t) + 1998z(t),$$

$$z'(t) = -1000y(t) - 1999z(t),$$

which, under the initial conditions y(0) = 1 and z(0) = 0, have solutions

$$y(t) = 2e^{-t} - e^{-1000t},$$

$$z(t) = -e^{-t} + e^{-1000t}.$$

Standard explicit methods will impose severe restrictions on the choice of timestep due to the very quickly decaying e^{-1000t} term. In the large-time limit however, this term is extremely close to 0 and can be safely ignored. Whilst classical implicit methods give an obvious way to circumvent this problem in this simple system, in more complex non-linear systems they cannot easily be applied without encountering stability problems. Stiffly stable schemes are designed to provide high-order methods capable of handling these systems.

In Karniadakis *et al.* (1991), the stiffly stable scheme is applied to Navier-Stokes by integrating equation (2.24) from time level $n - J_i$ to n + 1, resulting in the discretisation

$$\frac{\gamma_0 \boldsymbol{u}^{n+1} - \sum_{q=0}^{J_i - 1} \alpha_q \boldsymbol{u}^{n-q}}{\Delta t} = -\nabla \overline{p}^{n+1} + \sum_{q=0}^{J_e - 1} \beta_q \mathbf{N}(\boldsymbol{u}^{n-q}) + \nu \mathbf{L}(\boldsymbol{u}^{n+1})$$

 α_q and $\gamma_0 = \sum_{i=0}^{J_i-1} \alpha_i$ denote the standard stiffly-stable coefficients of Gear (1971); however β_q are not the same explicit coefficients as used in earlier splitting schemes. All of these are readily available in Table IV of Karniadakis *et al.* (1991) and Karniadakis & Sherwin (2005) up to third order.

As before, we may then use operator splitting to obtain the three-step scheme

$$\frac{\hat{\boldsymbol{u}} - \sum_{q=0}^{J_i - 1} \alpha_q \boldsymbol{u}^{n-q}}{\Delta t} = -\sum_{q=0}^{J_e - 1} \beta_q \mathbf{N}(\boldsymbol{u}^{n-q}),$$
$$\frac{\hat{\boldsymbol{u}} - \hat{\boldsymbol{u}}}{\Delta t} = -\nabla \overline{p}^{n+1},$$
$$\frac{\gamma_0 \boldsymbol{u}^{n+1} - \hat{\boldsymbol{u}}}{\Delta t} = \nu \mathbf{L}(\boldsymbol{u}^{n+1}).$$

Coupled with the high order pressure boundary condition, this scheme gives a highly accurate and stable scheme for timestepping Navier-Stokes.

2.4 Summary

In this chapter, we have introduced the fundamental tools required to construct spectral schemes and the spectral/hp element scheme, along with a high-order timestepping scheme for the Navier-Stokes equations. In the next chapter, we study these schemes in the context of pipe flow, and introduce Semtex, the joint spectral element–Fourier pseudospectral solver which will be used to perform pipe flow simulations throughout the rest of this thesis. We study how a volumetric flux condition can be used to drive flow through a pipe, and introduce a computational technique for reducing communication time required for parallelising the discrete Fourier transform.

Chapter 3

Numerical Methods for Pipe Flow

In the previous chapter, the spectral/hp element method was introduced along with a highorder timestepping scheme for Navier-Stokes. The focus in this chapter is to demonstrate how these schemes are implemented in the Semtex spectral element code, which will be used throughout the rest of this thesis to perform a wide variety of pipe flow simulations. Additionally, a number of modifications to the underlying code are introduced in order to specialise the code for studying pipe flow problems.

Firstly, we define the mathematical framework under which we will study fluid flows, and examine the Navier-Stokes equations which determine the evolution of a fluid. In order to make the problem well-posed, we also consider the boundary conditions which are required to drive the flow. Using these conditions, and some elementary simplifications to the underlying equations, we derive the exact equations of motion for a laminar flow.

Many numerical implementations, including Semtex, assume that periodic boundary conditions are used in the axial direction, necessitating the use of a constant body force in order to drive the flow. However, as will be explained in the following sections, this is often not preferred for transitional experiments. The second section of this chapter therefore introduces an alternative method which requires that the flow had constant volumetric flux, and explains how it can be efficiently integrated into a timestepping scheme, either through an exact condition or through a Green's function correction technique. We then introduce Semtex and demonstrate how this method is integrated into the underlying code. In the final section of this chapter, we consider the parallelisation of the discrete Fourier transform used by Semtex to decompose a domain over multiple nodes. This is an expensive process due to the all-to-all communication required in constructing each Fourier mode. Whilst this heavy communication cannot be avoided, any features of the underlying network topology should be exploited to enable maximum throughput of data. In particular, we introduce a method which provides significant scaling improvements for clusters utilising a network of massively multi-core nodes.

3.1 Mathematical Framework

Whilst there are many variants of the pipe flow problem, the classical one which is both the simplest and most studied assumes that the pipe is smooth and the flow through it incompressible and viscous, so that the velocity field is a solution to the incompressible Navier-Stokes equations. In this section, we outline these non-linear equations along with their linear counterpart, the Stokes equations, in both the cylindrical and Cartesian co-ordinate systems. The motivation for these choices of co-ordinates (in particular the Cartesian system) will be discussed at greater length in chapter 4. We then move on to derive the exact equations of motion for a laminar flow.

Definition 3.1.1 (Navier-Stokes equations). The *incompressible Navier-Stokes equations* are defined on a domain $\Omega \subseteq \mathbb{R}^n$ by the non-linear system of partial differential equations

$$\frac{\partial \boldsymbol{u}}{\partial t} + (\boldsymbol{u} \cdot \nabla)\boldsymbol{u} = -\frac{1}{\rho}\nabla p + \nu\nabla^2 \boldsymbol{u} + \boldsymbol{f}, \qquad (3.1a)$$

$$\nabla \cdot \boldsymbol{u} = \boldsymbol{0}, \tag{3.1b}$$

where u is the velocity field, p is the scalar pressure field, f is a (potentially time-dependent) forcing term, ρ is the fluid density and v is the kinematic viscosity of the fluid. Equation (3.1a) is known as the *momentum equation* and (3.1b) as the *incompressibility condition* or *continuity equation*.

Definition 3.1.2 (Stokes equations). Neglecting the non-linear advective term $(u \cdot \nabla)u$ from

equation (3.1a) yields the linear system

$$\frac{\partial \boldsymbol{u}}{\partial t} = -\frac{1}{\rho} \nabla p + \nu \nabla^2 \boldsymbol{u} + \boldsymbol{f}, \qquad (3.2a)$$

$$\nabla \cdot \boldsymbol{u} = \boldsymbol{0}, \tag{3.2b}$$

known as the Stokes equations.

We begin with an elemental simplification of (3.1a). The incompressibility condition (3.1b) implies that the fluid density ρ remains constant, and since we are not concerned with its value in this work, we define $\rho = 1$ without loss of generality.

Throughout this work, we will consider a cylindrical pipe of length L and radius R and diameter D = 2R. The domain $\Omega \subset \mathbb{R}^3$ is therefore

$$\Omega = \{ (x, y, z) \mid 0 \le x \le L, y^2 + z^2 \le R^2 \} \subseteq \mathbb{R}^3.$$

In this Cartesian framework, y and z describe the transverse directions (i.e. those lying in circular cross-sections) and x aligns with the streamwise or axial direction. The threedimensional velocity field u is written as

$$\boldsymbol{u}(\boldsymbol{x},\boldsymbol{y},\boldsymbol{z},t) = \boldsymbol{u}(\boldsymbol{x},t) = (\boldsymbol{u}(\boldsymbol{x},t),\boldsymbol{v}(\boldsymbol{x},t),\boldsymbol{w}(\boldsymbol{x},t))$$

where again u represents the axial component of velocity and v, w align with y, z respectively.

For analytical purposes it is useful to consider the equations in the far more natural setting of the cylindrical polar co-ordinate system. This is defined in the usual way by the transformation

$$x = x, \quad y = r\cos\theta, \quad z = r\sin\theta,$$
 (3.3)

for $r \ge 0$ and $\theta \in [0, 2\pi)$. The domain thus becomes

$$\Omega = \{ (x, r, \theta) \mid x \in [0, L], r \in [0, R], \theta \in [0, 2\pi) \}.$$

Under this co-ordinate system, we refer to the axial component of velocity by u_x , and the transverse components by u_r and u_θ respectively. Applying (3.3) to (3.1) gives the transformed system of equations:

$$\begin{aligned} \frac{\partial u_x}{\partial t} + (\boldsymbol{u} \cdot \nabla)u_x &= -\frac{\partial p}{\partial x} + \nu \nabla^2 u_x, \\ \frac{\partial u_r}{\partial t} + (\boldsymbol{u} \cdot \nabla)u_r &= -\frac{\partial p}{\partial r} + \nu \left(\nabla^2 u_r - \frac{u_r}{r} - \frac{2}{r^2} \frac{\partial u_\theta}{\partial \theta} \right), \\ \frac{\partial u_\theta}{\partial t} + (\boldsymbol{u} \cdot \nabla)u_\theta + \frac{u_r u_\theta}{r} &= -\frac{\partial p}{\partial \theta} + \nu \left(\nabla^2 u_\theta - \frac{2}{r^2} \frac{\partial u_r}{\partial \theta} - \frac{u_\theta}{r^2} \right), \\ \frac{\partial u_x}{\partial x} + \frac{1}{r} \frac{\partial}{\partial r} (r u_r) + \frac{1}{r} \frac{\partial u_\theta}{\partial \theta} = 0, \end{aligned}$$

where, in this co-ordinate system,

$$\boldsymbol{u} \cdot \nabla = u_r \frac{\partial}{\partial r} + \frac{u_\theta}{r} \frac{\partial}{\partial \theta} + u_x \frac{\partial}{\partial x},$$
$$\nabla^2 = \frac{1}{r} \frac{\partial}{\partial r} \left(r \frac{\partial}{\partial r} \right) + \frac{1}{r^2} \frac{\partial^2}{\partial \theta^2} + \frac{\partial^2}{\partial x^2}$$

3.1.1 Boundary conditions

In addition to the incompressibility condition, it is necessary to prescribe boundary conditions so that the problem is well-defined. In this section we outline the basic conditions that are typically used for pipe flow.

Definition 3.1.3 (No-slip condition). A fluid satisfies the *no-slip condition* at a wall $W \subset \Omega$ if for every $t \ge 0$,

$$u(W) = \{0\}. \tag{3.4}$$

The no-slip condition essentially stipulates that, where the fluid comes into a contact with a wall, friction dictates that it remain motionless. In the setting of pipe flow, this condition is applied on the surface

$$W = \{ (x, y, z) \mid y^2 + z^2 = R^2 \} \subset \Omega.$$

If the pipe is of finite length $(L \neq \infty)$, we also need two more conditions on each end of the pipe; the *inflow* at x = 0 condition specifies the velocity profile of fluid entering the pipe, and the *outflow* condition specifies how fluid leaves the pipe at x = L.

If the mathematical model is strictly intended to duplicate physical experiments, then one typically chooses a parabolic inflow condition, chosen to represent laminar flow (see section 3.1.2) and a homogeneous Neumann boundary condition

$$\left. \frac{\partial \boldsymbol{u}}{\partial \boldsymbol{x}} \right|_{\boldsymbol{x}=\boldsymbol{L}} = 0, \tag{3.5}$$

for the outflow. Another possibility for these boundary conditions is a *periodic* condition;

$$\boldsymbol{u}(0, \boldsymbol{y}, \boldsymbol{z}, t) = \boldsymbol{u}(L, \boldsymbol{y}, \boldsymbol{z}, t) \quad \forall t \ge 0.$$

Under this regime, any fluid particles leaving the pipe at x = L re-enter the pipe at x = 0.

When using an inflow-outflow condition for turbulence experiments, we must choose a domain which is large enough to allow the flow to transition from laminar to turbulent flow upon a perturbation at the pipe wall. Therefore in a numerical experiment where increasing L necessitates more computational overheads, this condition is not preferred. The periodic boundary condition gives us an alternative method to simulate turbulent flow in a much more computationally-friendly manner, and using the methods of the previous section, allows us to form an efficient and highly accurate scheme for integrating Navier-Stokes.

One topic yet to be discussed is the force required to drive the flow through the pipe. The Neumann outflow condition above, when coupled with an appropriate inflow condition, causes a pressure gradient Δp to form between the start and end of the pipe, and thus we may neglect the forcing term f from the equations. For a periodic flow however, we must supply an external force f, since otherwise the flow will either remain at rest indefinitely or monotonically decay to zero. Since the value of Re depends on the mean flow, and thus f, we must derive a relationship between f and Re in order to perform a simulation at any desired Reynolds number before the start of a simulation. This topic will be further addressed in section 3.2.

3.1.2 Hagen-Poiseuille Flow

Almost all of the complexity of the Navier-Stokes equations can be traced back to the nonlinear term $(\boldsymbol{u} \cdot \nabla)\boldsymbol{u}$, which not only strongly couples the systems of equations together, but



Figure 3.1: Average velocity profiles for laminar (left) and turbulent (right) flows through the cylindrical pipe, where averaging is performed over the axial and azimuthal directions. These profiles are normalised so that in each case the flowrates $\int_C \mathbf{u} \cdot d\mathbf{S}$ are equal.

also adds quadratic-order non-linearity to each component of the vector field. However, by making some simple experimental observations, we can make assumptions about the types of flow that we expect to see, and therefore simplify the system in some way in order to find analytic solutions.

The independent investigations of Hagen (1839) and Poiseuille (1840) represent the first attempt at finding exact flow fields in pipes. In these works, the primary motivation was the study of blood flow through capillaries. On these scales, average velocities are typically slow and thus turbulence is not usually observed. In this laminar setting, the problem becomes far easier to model. In this section, we briefly provide a modern derivation of the equations of motion they discovered. We begin with two simplifications. Firstly, we assume the flow is steady;

Definition 3.1.4 (Steady flow). A flow *u* satisfying (3.1) is said to be *steady* if $\partial_t u = 0$ at all times $t \ge 0$.

Secondly, we assume that the flow is confined to the axial direction, and is axisymmetric so that

$$u(x, r, \theta, t) = (u(r), 0, 0).$$

This is an example of a *shear flow*; a flow which is driven using some forcing term in a single direction and otherwise remains steady. Substituting these conditions into (3.1a), we obtain

the simplified system

$$\frac{\partial p}{\partial r} = \frac{\partial p}{\partial \theta} = 0$$
$$\frac{1}{r}\frac{\partial}{\partial r}\left(r\frac{\partial u}{\partial r}\right) = -\frac{\Delta p}{\nu} + u\frac{\partial u}{\partial x}$$

where we assume that $\Delta p = -\partial_z p$ is the constant pressure gradient driving the flow. Additionally, from the continuity equation (3.1b) we have that $\partial_x u = 0$ and hence we must solve

$$\frac{1}{r}\frac{\partial}{\partial r}\left(r\frac{\partial u}{\partial r}\right) = -\frac{\Delta p}{\nu}.$$

Integrating twice, we obtain the general solution in terms of two constants of integration, C_1 and C_2

$$u(r) = -\frac{\Delta p r^2}{4\nu} + C_1 \log r + C_2$$

Since we insist that our solution is well-defined at r = 0, it is clear that $C_1 = 0$, and thus by applying the no-slip boundary condition (3.4) at r = R we obtain the solution

$$u(r) = \frac{\Delta p}{4\nu} (R^2 - r^2).$$

This simple parabolic velocity profile is shown on the left in figure 3.1. Hagen-Poiseuille flow is known as *laminar* – that is, the flow is separated into parallel layers which do not interact with one another. The right hand side of figure 3.1 demonstrates how the profile alters for a turbulent flow, which is distinctly more plug-shaped.

3.2 Forcing with Volumetric Flux

In physical systems there are two commonly adopted methods used to drive flow through a pipe. The most obvious and intuitive approach is to place a header tank at the inflow of the pipe in which the level of water is kept constant, and leave the outflow of the pipe unplugged as depicted in figure 3.2(a). This forces a pressure gradient to be formed between the start and the end of the pipe, and by keeping a constant water level in the tank, this gradient remains unchanged. Figure 3.2(b) shows an alternative arrangement whereby a piston is placed at one end of the pipe. When the piston is moved to the right, fluid is drawn through



(a) Flow driven using a constant pressure gradient. The water level *h* is kept constant so that the pressure gradient Δp remains constant.



(b) Flow driven using a constant volumetric flux. The piston is moved to the right at a constant speed \overline{U} , resulting in constant flow rate through cross-sections of the pipe.

Figure 3.2: Schematic diagrams depicting flow being driven using (a) constant pressure gradient and (b) constant volumetric flux.

the pipe in a similar fashion to a syringe. If the piston is moved at a constant speed then, given that the fluid is incompressible, we implicitly impose a constant flux condition through circular cross-sections of the pipe.

In the following sections, we demonstrate how both of these methods can be implemented numerically in a pipe of radius R and length L, so that the fluid velocity is governed by the incompressible Navier-Stokes equations (3.1) under a cylindrical polar co-ordinate system. In this form, the fluid velocity u is represented as (u_x, u_r, u_θ) , and requires periodic boundary conditions in both axial and azimuthal directions. Additionally, we require numerical boundary conditions at the pipe axis. For the analysis which follows, the only condition of significance is that of the radial velocity, for which we insist that $u_r(r = 0) = 0$. A discussion of the high-order boundary conditions required in this formulation is given in Blackburn & Sherwin (2004).

3.2.1 Constant body forcing

We begin by examining the pressure gradient method, which is perhaps the easiest method to implement in both a physical and numerical setting. Computationally, we make use of a constant body force $f = (f_x, 0, 0)$ which is appended to the momentum equation. However, determining the correct value of f for any desired Re requires careful examination.

It is well-known that the pressure drop Δp in fluid of density ρ flowing through a pipe of length L and diameter D at mean speed \overline{U} may be expressed in terms of the Darcy-Weisback relation:

$$\Delta p = \lambda \frac{\rho L \overline{U}^2}{2D}$$

where λ , the *friction factor*, is the dimensionless constant relating the friction at the pipe wall to the loss in pressure. The axial body force, or alternatively pressure drop per unit length, is therefore $f_x = \Delta p/L$.

Determining the friction factor is not a straightforward task, as it depends on the regime of the flow (laminar, fully developed turbulence or transitional) as well as the Reynolds number. By examining the solution to Navier-Stokes in the case of a laminar flow, λ may be exactly derived as 64/Re. However, the value of λ for a turbulent flow differs substantially and must be derived using knowledge of the various layers typically found in a turbulent flow. Assuming the mean turbulent flow resembles a one-seventh power law, Blasius (1913) derived the simple relationship $\lambda = 0.3164 \text{Re}^{-1/4}$, which is widely regarded to be the most accurate for $2 \times 10^3 \leq \text{Re} \leq 10^5$ but is inaccurate beyond this. For higher Re, several alternative formulae are available. The most widely used may be attributed to Prandtl. More recently McKeon *et al.* (2005) derived the implicit formula

$$\frac{1}{\sqrt{\lambda}} = 1.920 \log(\text{Re}\sqrt{\lambda}) - 0.475 - \frac{7.04}{\left(\text{Re}\sqrt{\lambda}\right)^{0.55}}$$

Using the Blasius rule for simplicity and assuming $\overline{U} = D = \rho = 1$, we obtain the relation

$$f_x = \frac{\lambda}{2} = \frac{0.1582}{\mathrm{Re}^{1/4}}$$

which can be used to drive turbulent flow at a prescribed Reynolds number.



Figure 3.3: Friction factor λ for fully developed pipe flow at a variety of Re. Dashed line indicates the laminar Hagen-Poiseuille friction law $\lambda = 64/\text{Re}$ and the solid line the Blasius law $\lambda = 0.3164\text{Re}^{-1/4}$. Data points are taken from the experimental results of McKeon *et al.* (2004) and show the transition between the two laws.

3.2.2 Forcing for transition experiments

For the study of transitional flows, the use of a constant pressure gradient is not preferred. Recall that Re is defined using the bulk velocity

$$\overline{U} = \frac{1}{\mu(C)} \int_C \langle \boldsymbol{u} \rangle_x \cdot \mathrm{d}\boldsymbol{S}$$

where C is the disc $\{(r, \theta) \mid r \leq R\}, \mu(C) = \pi R^2$ denotes its area and $\langle \cdot \rangle_x$ denotes

$$\langle g \rangle_x = \frac{1}{L} \int_0^L g(x, r, \theta) \,\mathrm{d}x.$$

If f remains fixed, simulations involving turbulence will lead to fluctuations in \overline{U} , allowing the Reynolds number to vary with time. This variation is usually unimportant for high-Re simulations, but at transition it is necessary to maintain a fixed Re.

Another somewhat more serious issue however is the lack of a relationship between λ and Re in the transitional regime, making it difficult to prescribe f prior to starting a simulation.

Figure 3.3 shows the laminar Hagen-Poiseuille friction law and the Blasius friction law over a wide range of Re, along with experimental data which bridges the gap between the two. At Re $\approx 2,000$, λ abruptly changes as the flow becomes turbulent. Whilst it is possible to estimate λ from this graph or others like it, Re would need to be manually determined at the end of any simulation.

One possible solution to this is given by the piston driving approach described earlier. Since the fluid is incompressible, the volumetric flux Q(u) defined by the functional

$$Q(\boldsymbol{u}) = \frac{1}{\pi R^2} \int_C \langle \boldsymbol{u} \rangle_x \cdot \mathrm{d}\boldsymbol{S} = \frac{1}{\pi R^2} \iint_C \boldsymbol{n} \cdot \langle \boldsymbol{u} \rangle_x \, \mathrm{d}\boldsymbol{S} = \frac{1}{\pi R^2} \iint_C r \langle \boldsymbol{u}_x \rangle_x \, \mathrm{d}r \, \mathrm{d}\theta,$$

is required to be fixed at a prescribed value \overline{Q} . This ensures that \overline{U} remains constant, and hence the Reynolds number does not change.

Enforcing this computationally is most obviously achieved by making f time-dependent, measuring the flux and increasing or decreasing it by small amounts at each timestep. However, determining the amount to increase it by is non-trivial. In the following two sections, we derive an exact condition analytically, and also introduce a Green's function style technique which allows this method to be efficiently integrated into a numerical code.

3.2.3 Direct flux condition

It is possible to derive an exact value for the correct value of f_x which enforces the constant flux condition. Orlandi & Fatica (1997) state that this value can be obtained by applying the transformation $v_r = ru_r$, $v_\theta = ru_\theta$, $v_x = u_x$, thus re-writing the axial component of the Navier-Stokes equations as

$$\underbrace{\overbrace{\partial v_x}^{(i)}}_{\partial t} + \underbrace{\overbrace{r \frac{\partial}{\partial r}(v_r v_x)}^{(ii)} + \overbrace{r^2 \frac{\partial}{\partial \theta}(v_\theta v_x)}^{(iii)} + \underbrace{\overbrace{\partial v_x^2}^{\partial v_x}}_{(v \cdot \nabla) v_x}}_{(v \cdot \nabla) v_x} = \underbrace{\overbrace{(v \cdot \nabla) v_x}^{(v)} + \underbrace{\overbrace{\partial v_x}^{(v)} + \underbrace{\overbrace{\partial v_x}^{(v)}}_{(v \cdot \nabla) v_x}}_{\nabla^2 v_x} + \underbrace{\overbrace{\partial v_x}^{(vii)} + \underbrace{\overbrace{\partial v_x}^{(vii)} + \underbrace{\overbrace{\partial v_x}^{(vii)}}_{\nabla^2 v_x}}_{\nabla^2 v_x} + \underbrace{\overbrace{\partial v_x}^{(vii)} + \underbrace{\overbrace{\partial v_x}^{(vii)} + \underbrace{\overbrace{\partial v_x}^{(vii)} + \underbrace{\overbrace{\partial v_x}^{(vii)} + \underbrace{\overbrace{\partial v_x}^{(viii)} + \underbrace{\overbrace{\partial v_x}^{(vii)} + \underbrace{\overbrace{\partial v_x}^{(vii)} + \underbrace{\overbrace{\partial v_x}^{(vii)} + \underbrace{\overbrace{\partial v_x}^{(vii)} + \underbrace{i_x}^{(vii)} + \underbrace{i_x}^{(vi$$

We then follow the approach of Kim (2002), whereby we divide this equation by the volume of the domain $(\pi R^2 L)$ and integrate both sides. First, let $\langle \cdot \rangle_{x,\theta}$ denote the averaging operation

$$\langle g \rangle_{x,\theta} = \frac{1}{2\pi L} \int_0^{2\pi} \int_0^L g(x,r,\theta) \,\mathrm{d}x \,\mathrm{d}\theta.$$

Then, we evaluate each term:

(i)
$$\frac{1}{\pi R^2 L} \int_0^{2\pi} \int_0^R \int_0^L \frac{\partial v_x}{\partial t} r \, dx \, dr \, d\theta = \frac{1}{\pi R^2} \frac{\partial}{\partial t} \left(\int_0^{2\pi} \int_0^R \langle v_x \rangle_x r \, dr \, d\theta \right) = \frac{\partial Q}{\partial t}.$$

(ii) $\int_{0}^{R} \frac{1}{r} \frac{\partial}{\partial r} (v_r v_x) r \, dr = \left[v_r v_\theta \right]_{0}^{R} = 0$ due to the no-slip and axis boundary conditions on u_r .

(iii) This term immediately evaluates to $\left[v_{\theta}v_{x}\right]_{0}^{2\pi} = 0$ since $v_{\theta}(\theta = 0) = v_{\theta}(\theta = 2\pi)$.

- (iv) By a similar argument, this term is also zero since v_x is periodic.
- (v) Since we assume that p(x = 0) = p(x = L) by the use of a forcing term, this term trivially integrates to zero.
- (vi) We re-write this term as:

$$\frac{1}{\operatorname{Re}} \cdot \frac{1}{\pi R^2 L} \int_0^R \int_0^{2\pi} \int_0^L \frac{1}{r} \frac{\partial}{\partial r} \left(r \frac{\partial v_x}{\partial r} \right) r \, \mathrm{d}x \, \mathrm{d}\theta \, \mathrm{d}r = \frac{2}{R^2 \cdot \operatorname{Re}} \int_0^R \frac{\partial}{\partial r} \left(r \frac{\partial \langle v \rangle_{x,\theta}}{\partial r} \right) \, \mathrm{d}r$$
$$= \frac{2}{R \cdot \operatorname{Re}} \left. \frac{\partial \langle v_x \rangle_{x,\theta}}{\partial r} \right|_{r=R}$$

- (vii) By applying periodic boundary conditions, we again find both these terms evaluate to zero.
- (viii) This term remains unchanged under integration since f_x remains constant in space.

If we require the flux remain unchanged then $\partial_t Q = 0$, and so the condition

$$f_x = -\frac{2}{R \cdot \text{Re}} \left. \frac{\partial \langle v_x \rangle_{x,\theta}}{\partial r} \right|_{r=R}$$

as stated (but not derived) by Veenman (2004), imposes this condition. This method is readily implemented in the case of pipe flow. The last term in this equation is more commonly known as the *wall-shear stress*:

Definition 3.2.1 (Wall-shear stress). The *wall shear stress* τ_w is defined by the relation

$$\tau_w = \nu \left. \frac{\partial \langle u \rangle_{x,\theta}}{\partial r} \right|_{r=R}$$

where u denotes the velocity of the fluid in the direction of flow.

However, in more complex geometries, the calculation of the wall-shear stress can be expensive, and moreover the derivation above will need to be changed on a case-by-case basis depending on the boundary conditions.

3.2.4 Green's function flux forcing

A more generalised approach to enforcing volumetric flux is very briefly outlined in Chu *et al.* (1992), which makes use of an efficient Green's function method in order to impose the flux condition. Whilst this method will be expanded upon here in the context of pipe flow in a domain of length L, it is trivially extended to other fluid problems. For example, in the next section we demonstrate how it can be implemented in the context of a spectral element code.

We begin by taking the incompressible Navier-Stokes equations written in the form of equation (2.24), so that u is a solution of

$$\frac{\partial \boldsymbol{u}}{\partial t} = -\nabla p + \nu \mathbf{L}(\boldsymbol{u}) + \mathbf{N}(\boldsymbol{u}), \qquad (3.6)$$

$$\boldsymbol{u}(\boldsymbol{x},0) = \boldsymbol{u}^n(\boldsymbol{x}). \tag{3.7}$$

One timestep of this system therefore gives the solution obtained at time $t_{n+1} = (n+1)\Delta t$. In addition, let us assume periodic boundary conditions

$$\boldsymbol{u}(x=0,r,\theta,t) = \boldsymbol{u}(x=L,r,\theta,t), \quad \forall r \in [0,R], \theta \in [0,2\pi), t \ge 0$$

and impose a no-slip boundary condition on the pipe wall. Under this formulation, in order to uniquely determine p we necessarily assume that the initial condition u^0 monotonically decays to the trivial zero solution since no external forcing term is added to the system.

An alternative approach to imposing the volumetric flux is through the use of a correction
$\tilde{u}(r, \theta, x, t)$, so that $\overline{u} := u + \tilde{u}$ with associated pressure field \overline{p} is a solution of (3.6) and has flux $Q(\overline{u})$ fixed at some desired value \overline{Q} . In many cases this approach is hindered by the nonlinear operator N(u). However the use of a splitting scheme, such as the three-step high-order splitting scheme in section 2.3.6, offers a solution to this problem.

Let us consider a general two-step splitting scheme in a continuous setting, so that the equations to solve are:

$$\begin{cases} \frac{\partial u}{\partial t} = \mathbf{N}(u) & \frac{\hat{u}(x, \Delta t)}{u(x, 0) = u^n(x)} & \begin{cases} \frac{\partial u}{\partial t} = -\nabla p + \nu \mathbf{L}(u) \\ u(x, 0) = \hat{u}(x, \Delta t) \end{cases}$$

where the solution of the first equation is the intermediate solution \hat{u} and of the second is the desired field u^{n+1} . In this setting, the first set of equations performs the non-linear advection, and the second provides the pressure correction and linear dissipation. This second set of equations is in fact the definition of the linear Stokes problem, and so we may append any other solution of these equations without penalty (up to roundoff error), therefore bypassing the nonlinearity of the original equations. Consider then a correction field \tilde{u} which satisfies the forced Stokes equations

$$\frac{\partial \tilde{\boldsymbol{u}}}{\partial t} = -\nabla \tilde{p} + \nu \mathbf{L}(\tilde{\boldsymbol{u}}) + \alpha^n \hat{\boldsymbol{x}}$$

and is subject to the usual incompressibility condition and identical boundary conditions to the original equation. α^n represents a (yet undetermined) body forcing term which is calculated at each time level *n*. In order to satisfy (3.7), \tilde{u} necessarily has initial condition $\tilde{u}(x, 0) = 0$.

The dependence of \tilde{u} on α^n must be removed for the scheme to be computationally useful, since otherwise the Stokes solution must be computed at each timestep. Assuming that the flow is always driven (and hence $\alpha^n \neq 0$), define $v = \tilde{u}/\alpha^n$ so that v solves the equations

$$\frac{\partial \boldsymbol{v}}{\partial t} = -\nabla \hat{p} + \nu \mathbf{L}(\boldsymbol{v}) + \hat{\boldsymbol{x}}, \qquad (3.8)$$

$$\boldsymbol{v}(\boldsymbol{x},0) = \boldsymbol{0}; \tag{3.9}$$

i.e. the Stokes equations with unit forcing. To correct the velocity field and impose the flux



Figure 3.4: Geometry of the backwards facing step, for which the Stokes field possesses a singular point at the step edge.

condition, we then calculate $\overline{u}^n = u^n + \tilde{u}^n = u^n + \alpha^n v^1$, and so we may compute v^1 before the start of the simulation. Finally, in order to determine α^n , taking the flux of both sides of the previous equation yields

$$Q(\overline{\boldsymbol{u}}^n) = Q(\boldsymbol{u}^n) + \alpha^n Q(\boldsymbol{v}^1) \Rightarrow \alpha^n = \frac{\overline{Q} - Q(\boldsymbol{u}^n)}{Q(\boldsymbol{v}^1)}.$$
(3.10)

The procedure then is to initially solve (3.8) to obtain $v^1(x) = v(x, \Delta t)$ and its flux $Q(v^1)$ before the main simulation occurs. Given a solution u^{n+1} which solves (3.6) and (3.7) and is obtained through a splitting scheme, we calculate α^{n+1} by

$$\alpha^{n+1} = \frac{\overline{Q} - Q(\boldsymbol{u}^{n+1})}{Q(\boldsymbol{v}^1)}$$

The corrected velocity field $\overline{u}^{n+1} = u^{n+1} + \alpha^{n+1}v^1$ then has has flowrate exactly equal to \overline{Q} . We then use \overline{u}^{n+2} instead of u^{n+1} as the initial condition in (3.7) at timestep n + 1 to obtain a solution with desired flowrate at timestep n + 2, and so on. This method is extremely efficient; both v^1 and $Q(v^1)$ can be pre-computed, and so all that is needed at each timestep is to calculate $Q(u^{n+1})$ and perform a simple vector addition operation.

In certain domains which possess sharp corners, a solution to the Stokes equation cannot be guaranteed to exist even if the Navier-Stokes equations have a solution, and thus this flowrate correction technique cannot be used. The simplest example of this is the backwards facing step, depicted in figure 3.4, for which the Stokes solution has a singularity at the step edge. One possible solution, proposed by Lund *et al.* (1998) and demonstrated in Kopera (2011) is the use of a so-called copy boundary condition. The domain is partitioned into two pieces; the regular backwards facing step geometry Ω_{bfs} and an extended inlet Ω_p . Ω_p is a self-contained periodic channel, with fluid exiting at E_2 re-entering at E_1 . The flow in Ω_p is driven using the constant flowrate condition. The velocity field at E_2 is then 'copied' to form the inlet condition of Ω_{bfs} . This is particularly useful for simulations which require a turbulent inflow condition, since a turbulent initial field in the inlet channel is easily obtained. Care must be taken, however, to ensure that the channel is sufficiently long so that any periodic features of the flow dissipate before entering the main domain. The implementation of this boundary condition in Semtex is discussed in Cantwell (2009).

3.2.5 The Semtex Spectral Element Solver

Spectral-element solvers are rarely used in three dimensional problems with 'simple' geometries such as cylindrical pipes. There are several major reasons for this; most prevelant is possibly the complexity of implementation and length of time required to develop a suitable code. Large-scale and high dimensional computational problems (such as a sufficiently resolved integration of Navier-Stokes) also prove difficult to parallelise under a spectral element regime due to the underlying formulation.

The solution to both of these problems is to use a joint discretisation. Such schemes typically use one of the spectral-type methods outlined earlier which exploit any homogeneous properties of the domain, and a spectral element discretistation in dimensions where the geometry is more complex. Typically this leads to both a simplification in implementation and additionally speeds up execution time.

Semtex is a spectral/hp element solver code using the classical nodal basis functions seen in figure 2.2(b) which is designed under this paradigm. It is written in C++ and, whilst the focus throughout the code is the direct integration of Navier-Stokes, the framework is general enough to implement solvers for other PDEs. To this end, it comes in several distinct parts:

- Veclib: a general linear algebra library responsible for performing a variety of common operations on matrices and vectors.
- Femlib: another library which performs many standard operations required for the spectral-element method; generating quadrature points and weights, constructing projection functions χ^e , managing message handling for parallelisation, etc.

- A number of C++ classes which represent the key parts of the elemental scheme (i.e. the mesh, elements and fields for storage) which rely on both of these libraries.
- Drivers for solving various PDEs; in particular a DNS application which implements the stiffly-stable timestepping scheme of section 2.3 to enable high-order timestepping of up to $O(\Delta t^3)$.
- A variety of utilities which manipulate the data output from the driver applications, for example, to compute vorticity fields and extract data from particular points of the domain.

The code is only suited to problems which are two- and three-dimensional. In two dimensions, Semtex uses a spectral element mesh which is composed entirely of quadrilateral elements. In the third dimension, the Fourier pseudo-spectral method outlined in section 2.1.1 is used. In addition, either a Cartesian or cylindrical polar co-ordinate system can be used, with periodic boundary conditions in the homogenenous direction. In the Cartesian formulation the Fourier discretisation in a domain of length L_x is therefore given by

$$\boldsymbol{u}(\boldsymbol{x},t) = \sum_{k=-N/2+1}^{N/2} \hat{\boldsymbol{u}}_k(y,z,t) e^{2\pi i k x/L_x},$$
(3.11)

and equivalently in polar form,

$$\boldsymbol{u}(x,r,\theta,t) = \sum_{k=-N/2+1}^{N/2} \hat{\boldsymbol{u}}_k(x,r,t) e^{ik\theta}.$$
(3.12)

3.2.6 Implementation of volumetric flux forcing

The implementation of the volumetric flux forcing outlined in section 3.2 is relatively simple. Before the start of the main timestepping loop, the unit-forcing Stokes solution v^1 is calculated. In Semtex, this is done automatically by using the same high-order scheme that solves the Navier-Stokes, and in principle other codes may adopt this approach. Alternatively, a user-specified field may be supplied in place of v^1 . To complete the implementation, we must calculate the flux $Q(v^1)$ and, more generally, at each timestep $Q(u^n)$. There are several different cases depending upon the dimension of the domain, direction of flow and the choice of co-ordinate system. In a three dimensional setting, the data are mainly stored in their Fourier transformed state throughout the simulation for efficiency, since most spectral-element operations can be done using this data, and transforms are only required in the pseudo-spectral construction of N(u). Since the transform requires the use of expensive inter-processor communication in order to calculate derivatives in the direction of the Fourier co-efficients, it is highly desirable to find a method for calculating the flux using the Fourier transformed data.

For a Cartesian co-ordinate system, the flow may traverse in the direction of the Fourier modes, which is the applicable case for pipe flow simulations. The zero-th Fourier mode contains the two-dimensional mean velocity field. Since the flow is incompressible, the flux of this Fourier mode can therefore be used to calculate the flux of the velocity field through any cross-section C. Hence, to determine the flux of any given field, we calculate the area $\mu(C)$ of the cross-section C and integrate over the zero-th mode. Of course, integration on the GLL mesh is an essential operation for the implementation of any spectral element method, and so we use the in-built Semtex routines. If the simulation is being run in parallel across multiple machines, then due to the way the data are stored in memory (detailed further in section 3.3), the root process holds the zero-th Fourier mode and will therefore calculate the forcing constant α^n . It then communicates this to the remaining processes, which then perform the necessary velocity field correction.

For the purposes of completeness, we also present a formulation for flows which propogate in the non-Fourier direction, since many other interesting geometries (such as the backwardsfacing step of figure 3.4 or a pipe with a stenosis) can be simulated using this method. For the two dimensional case, we accept as input a line defined by $E = \bigcup_{e=1}^{n} E_i^e \subset \Omega$ where E_i^e denotes the *i*-th edge of element Ω^e . We proceed under the assumption that the choice of *E* is 'sensible'; that is, it spans the domain, is perpendicular to the flow direction and is continuous. The flux may then be calculated using the in-built Semtex functions and a small amount of mathematical manipulation depending on the co-ordinate system in use.

In a three-dimensional setting, the line *E* is combined with the homogeneous Fourier direction so that we measure the flux through a cross-section $C = [0, L_x] \times E$. This implies that the unit normal to *E* is of the form $\mathbf{n} = (n_1, n_2, 0)$. Given the Fourier transformed fields defined in (3.11) we can still calculate the flux of the physical field \mathbf{u} using the zero-th

Fourier mode:

$$\mu(C)Q(u) = \int_{C} u \cdot dS = \sum_{k=-N/2+1}^{N/2} \int_{C} \hat{u}_{k}(y, z, t)e^{2\pi i k x/L_{x}} \cdot \mathbf{n} \, dS$$

$$= \sum_{k=-N/2+1}^{N/2} \int_{0}^{L_{x}} \int_{E} (\hat{u}_{k}(y, z, t)n_{1} + \hat{v}_{k}(y, z, t)n_{2})e^{2\pi i k x/L_{x}} \, dy \, dz \, dx$$

$$= \sum_{k=-N/2+1}^{N/2} \left(\int_{0}^{L_{x}} e^{2\pi i k x/L_{x}} \, dx \right) \left(\int_{E} \hat{u}_{k}(y, z, t)n_{1} + \hat{v}_{k}(y, z, t)n_{2} \, dy \, dz \right)$$

$$= L_{x} \int_{E} \hat{u}_{0}(y, z, t)n_{1} + \hat{v}_{0}(y, z, t)n_{2} \, dy \, dz$$

since

$$\int_0^{L_x} e^{2\pi i k x/L_x} \, \mathrm{d}x = \begin{cases} L_x, & k = 0, \\ 0, & \text{otherwise} \end{cases}$$

The last integral over E is readily calculated using standard Semtex routines. Similarly, in a cylindrical geometry, $C = \{(x, r, \theta) \mid (x, r) \in E, \theta \in [0, 2\pi]\}$ and we find that

$$\Phi(\mathbf{u}) = \frac{2\pi}{\mu(C)} \int_E \left(\hat{u}_0(x, r, t) n_1 + \hat{v}_0(x, r, t) n_2 \right) r \, \mathrm{d}x \, \mathrm{d}x$$

where $u = u_r$ and $w = u_x$.

In summary then, the implementation runs as follows:

- Before running the simulation, solve the Stokes problem with unit forcing at a given viscosity v for a single timestep giving a field v¹.
- The root processor calculates Φ(v¹) and communicates this to all other processors.
 Each processor also calculates μ(C).
- At the end of each timestep, the root processor calculates α(t_{n+1}) and scatters this to the other processors. Each processor then performs the correction u
 ⁿ⁺¹ = uⁿ⁺¹ + αⁿ⁺¹v¹ on each processor.

3.3 Multicore optimisations for parallel Fourier transforms

Fourier transforms form a key part of many spectral schemes. They are particularly important in the pseudo-spectral scheme, since non-linear terms need to be multiplied in physical space, involving the use of two (or more) transforms per field per timestep. An efficient discrete Fourier transform (DFT) algorithm is therefore essential for most spectral codes. The Fast Fourier Transform (FFT), an adaptation of the DFT for data sets of $n = 2^m$ points, typically requires $O(n \log_2 n)$ operations. This is significantly faster than a naive DFT of complexity $O(n^2)$, and is essentially the key algorithm which makes quick spectrally-accurate codes cheap enough to implement.

However, in large domains (especially those which are three-dimensional) with high resolution requirements, the spatial discretisation often proves too large to fit in memory. More importantly, the number of operations required to perform the Fourier transforms means the use of more than a single processor is essential. To this end, the domain must be split over several processors, each storing a smaller piece of the domain.

As an example of how this impacts computational efficiency, consider a one-dimensional problem on the domain [0, 1], so that the discrete Fourier transform of a function $f : [0, 1] \rightarrow \mathbb{R}$ is given by

$$\hat{f}_k = \sum_{n=0}^{N-1} f(x_n) e^{2\pi i k x}, \quad -N/2 + 1 \le k < N/2,$$

where $x_n = n/N$. Now suppose that *N* is very large, and the domain [0, 1] is split evenly over $1 \le p \le P$ processors of equal capability. Furthermore, assume that *P* divides *N*, so that the *p*-th processor holds all spatial points {(p-1)N/P, ..., pN/P-1}. In this topology, each processor holds N/2P complex Fourier modes so that the memory requirements are the same on each.

Unfortuantely, in order to calculate the Fourier modes local to any processor p, the value of f is required at *all* of the spatial points in [0, 1]. In other words, every processor must communicate with every other processor. Obviously this scenario is the worst possible in terms of communication.

In the higher-dimensional joint discretisation setting of Semtex, we incur this overhead because the direction of Fourier data is the easiest to parallelise. The purpose of this section

is to investigate ways in which the transformation process can be made more efficient. In the first part, the general procedure used in Semtex is outlined. Next, an improvement of the algorithm in the case of multi-core nodes is introduced and its implementation explained in the context of Semtex. Finally, some benchmarks are provided to show the possible scalings of this algorithm for problems which are heavily Fourier bound.

3.3.1 Semtex implementation

The use of the hybrid Fourier-spectral element method used in Semtex makes it readily parallelisable in a similar fashion to the one-dimensional case described above. As is standard with scientific codes, Semtex uses the *Message Passing Interface* (MPI) library to parallelise the code, which is both efficient and abstracts many of the technical aspects of implementing codes which involve multi-processor programming.

Firstly, we describe the global ordering of the data. Given a polynomial order Q and $N_{\rm el}$ elements, each two-dimensional spectral-element plane of the discretised system can be stored using $QN_{\rm el}^2$ double-precision numbers. Given N_x points in the homogeneous direction, we can then order these planes so that, for points $x_j = j\Delta x$ with $0 \le j < N_x$, the *j*-th plane is stored at memory location $j \times QN_{\rm el}^2$.

The parallelisation procedure is precisely the same as above. Given *P* processors, processor *p* holds planes $\{(p-1)N_x/P, \ldots, pN_x/P-1\}$, where we assume that *P* divides N_x . Given this topology, one reasonably efficient method for performing the Fourier transform is to perform a parallel block-transpose of the data, which reduces the number of messages that need to be passed between processes, although the total message size remains unchanged. We now describe how this process works.

Let $N_p = QN_{el}^2$ denote the number of points per plane, and suppose further that P divides N_p . (If it does not, we may zero-pad the data so that it does.) Then, we divide each plane into P blocks containing $N_b = N_p/P$ data points. The aim of the transpose is to re-order the data across processors so that each processor holds blocks $\{B_{0,p}, \ldots, B_{N_x-1,p}\}$ where $B_{k,b}$ denotes the b-th block on the k-plane.

Figure 3.5 shows this process for a simple case where Q = 2 and each processor has three planes (so that $N_x = 6$). Overall, the process is reasonably straightforward, although quite



Figure 3.5: Exchange procedure for Q = 2 processors over $N_x = 6$ planes.

memory intensive. At the start, each processor contains a $N_x/P \times P$ block matrix which is stored in row-major order, so that memory is traversed fastest by increasing the plane index.

The first step is for each processor to perform a block transpose of the local plane data. With the data in the correct ordering, an all-to-all exchanged is performed. In this way, row pof the transposed data is sent to processor q for all $1 \le q$, $p \le P$ with $q \ne p$. After this process is performed, each processor contains all of the data for block p in ascending plane order. Finally, the $P \times N_x/P$ matrix can be represented as a block column vector since the blocks are stored contiguously in memory.

After the exchange is performed, multiple one-dimensional Fourier transforms are performed on the data. This is most efficiently done using a Fortran routine, since the resultant matrix in column-major form is traversed slowest by plane number. Whilst Semtex ships with many FFT routines capable of performing this task, experimentation has shown that the algorithms presented by Temperton (1992) are both the fastest and conveniently allow for domains utilising $N_x = 2^p 3^q 5^r$ points. Once the Fourier transformed data has been (presumably) modified and the inverse transform taken, the final step is to perform an inverse exchange so that the data are ordered correctly. This procedure is precisely the same but reversed, so that the inter-processor exchange is done first, followed by the intra-processor exchange.

3.3.2 Multi-core systems

Modern architectures are rapidly progressing towards the multi-core paradigm, where a typical computer will possess one or more CPUs, each of which is composed of two or more cores. All of the CPUs also have access to the same shared memory space. A collection of these computers (from here-on referred to as *nodes*), combined with a fast interconnect, provide an affordable and extremely powerful super-computer. This trend is set to continue, at least for the foreseeable future.

Unfortunately, this paradigm has serious repercussions for the algorithm of the previous section. Firstly, in MPI codes each process runs independently. Since all processes local to a common node access the same memory, the internal transpose step proves extremely bus-intensive. In addition, whilst MPI has optimisations for communication between processes that reside on the same node, message passing will still be inefficient and somewhat unnecessary. Finally, since each process indepedently sends information to other nodes over the interconnect, this increases delays due to latency.

This problem is especially apparent on HECToR, the UK national supercomputer. Their latest machine is composed of 1,856 24-core nodes. Each node possesses two 12-core AMD Opterons CPUs, and each of these is actually two 6-core CPUs cast onto a single die. Currently, only *one* of these 6-core processors is connected to the interconnect¹, leading to disastrous scalings for codes involving all-to-all process communication.

One technique that has proved reasonably successful in countering this problem is the use of *System V shared memory*, whereby a large chunk of memory is allocated to accommodate the data for all processes which share a common node. One process per node is then responsible for handling communications for all local processes. Hence, whilst the size of the message is the same, the number of messages that are sent are reduced substantially for large multi-core systems. This improves communication time, since far less latency is incurred.

Figure 3.6 shows schematically how the algorithm of the previous section can be adapted to this shared-memory technique for a problem with two nodes and two processes local to each node. Conceptually the idea is the same as the previous section; however care must be taken to place the data in the correct ordering so that each process has the correct information when

¹As of 2011, after this work was performed, this is no longer the case as the machine has been upgraded with the Gemini interconnect.





the exchange is complete.

We make the assumption that each node has the same number of cores. Let us assume there are N nodes each containing M cores (so that P = MN), and let p_j^i denote the *i*-th core on the *j*-th node. As before, each plane is split into blocks $B_{k,b}$. Process p_j^0 is (without loss of generality) assigned to be node *j*'s representative and handles all inter-node communication.

Initially, a chunk of shared memory is allocated of size $N_x M/P \times P$, which is large enough to contain the planar data on each process local to that node. Each local process then blocktransposes its data; however, this time, the transpose is done with chunks consisting of Mblocks. In figure 3.6 for example, process p_0^0 performs

$B_{0,0}$ $B_{0,1}$	$B_{0,2}$ $B_{0,3}$	$B_{0,0}$	$B_{0,1}$	$B_{1,0}$	$B_{1,1}$
$B_{1,0}$ $B_{1,1}$	$B_{1,2}$ $B_{1,3}$	$B_{0,2}$	$B_{0,3}$	$B_{1,2}$	<i>B</i> _{1,3}

Process p_i^j writes row k of this matrix into row $kN_x/P + i$, so that the rows of the local processes' data are interleaved with one another. Note that it is entirely unnecessary to transpose the matrix in-place and then copy it into shared memory, as this can be accomplished in a single step.

The data is now in the correct order so that when the representatives of each node perform an all-to-all inter-node transpose, the resulting block matrix has columns containing the appropriate data for each local process. Each process then copies the appropriate column from the matrix ready to Fourier transform as before. The inverse exchange works in the same way, but in reverse.

3.3.3 Testing and results

The first step in implementation was to verify the validity of the new exchange algorithm. Testing was performed on phase2b of the HECToR facility, which uses the 24-core nodes described in the previous section. A small test program was written to perform exchanges using both the existing and multi-core algorithm for a variety of different N_x , P and N_b . In every case, arrays obtained through the new exchange were verified to be identical to the original exchange algorithm.

In order to measure timings and examine the scaling properties of the new algorithm as the



Figure 3.7: Execution time in milliseconds as a function of the number of processors P for the original Semtex exchange (blue) and new multi-core exchange (red).

number of processors varies, a further series of simulations were performed using this test code. In particular, we focused on investigating how the algorithm scales for Fourier-bound problems which require large communications overheads. Three simulations were performed using $N_x = 1,344, 1,920$ and 2,688 planes, and the size of data blocks were chosen so that the amount of data transferred over the interconnect was roughly equivalent to the fully-resolved pipe flow simulations of the following chapter. For each value of N_x , 200 exchanges were performed and from this we obtain the mean execution time for a single exchange. This was repeated for a large number of processors so that the number of nodes used varied from one (P = 24) to eight (P = 192).

The results for $N_x = 1,344$ can be seen in figure 3.7. Initially, execution time of the multicore exchange is worse than the normal MPI exchange. This can be attributed to the final stage of the new exchange algorithm, since to fit within the existing framework of Semtex, we must copy data back from the shared memory buffer to each processes' individual memory region. Initially the cost of this expensive operation outweighs the reduced number of messages. However, as P is increased, the difference between the old and new exchanges becomes apparent, so that for 8 nodes the exchange is completed nearly 8 times quicker than the existing exchange. Similar scalings are observed for $N_x = 1,920$ and 2,688. However, the most important aspect of this algorithm is that it makes simulations possible on this architecture that require large numbers of Fourier collocation points, such as those in long domains or at very high Re. With the existing exchange algorithm, communication time can be kept low by choosing a small number of processors. However, the number of planes per processor is then large and this then becomes the bottleneck. More processors solves this problem, but then the exchange algorithm dominates simulation times. Whilst the new multi-core algorithm still requires substantial amounts of time, it does at least offer a much better improvement in scaling.

There is however substantial room for improvement in optimising this algorithm. The most obvious modification is to use shared memory regions throughout the program instead of solely in the exchange algorithm, which then substantially reduces the amount of data which must be transferred over the memory bus of each node. However, such a task is a difficult undertaking in the context of Semtex, and would require substantial code modification to the core routines. Ultimately though, the proof-of-concept code produced here demonstrates how aspects of the multi-core paradigm can be exploited to improve program performance for spectral codes, and is a technique that can be readily employed in many other applications.

3.4 Summary

In this chapter, we have introduced the mathematical framework for performing numerical simulations of pipe flow and the Semtex code which will be used to evaluate them. In particular, we have summarised the common methods used to drive the flow by imposing a volumetric flux condition. To impose an exact flow rate, we have introduced a highly efficient Green's function approach which allows the condition to be implemented as a natural part of the formulation. Finally, we have outlined a method for optimising the parallelisation of the Fast Fourier Transform in the case where the underlying network topology consists of massively multi-core nodes. The focus of the rest of this thesis is to apply these numerical methods to simulations of pipe flow. In the next chapter, we aim to determine a mesh which is capable of sufficiently resolving the small-scale structures in turbulent flow. We then move to performing simulations of the reverse transition from turbulence to laminarity in order to observe and classify the resultant states.

Chapter 4

Transitional Dynamics of Pipe Flow

The previous two chapters have introduced the spectral/hp element method and a high-order timestepping scheme for the Navier-Stokes equations which, when combined, provide a highly accurate and efficient scheme for the numerical investigation of a number of fluid problems. In this chapter, we will focus on applying these computational methods to studying various aspects of the transition to turbulence in pipe flow.

The first section of this chapter investigates a number of potential meshes for use in pipe flow simulations. Ensuring that the spatial domain is correctly resolved is a key aspect of any numerical study, but is particularly important in turbulent pipe flow simulations (amongst other fluid problems) because even minor under-resolution can lead to significant errors in the resulting velocity and pressure fields. We therefore study a classical Chebyshev-Fourier collocation mesh and demonstrate how the structure can be duplicated in a spectral element mesh. We then perform a detailed numerical study in order to determine the minimum resolution requirements for later numerical studies in order to reduce unnecessary computational overheads.

The rest of this chapter is dedicated to computational investigations of pipe flow. Initially we present the results of a large-scale, long-time simulation in a L = 125D pipe which unlike the vast majority of existing studies, aims to examine the *reverse* transition as flow relaminarises from a turbulent initial condition. Unlike the forward transition problem, which depends on a finite-amplitude perturbation of a laminar flow, the reverse transition reveals states which form naturally within the transition. We then move on to investigate and classify the intermittent states found within this transition.

4.1 **Resolution requirements**

One of the key requirements for an accurate turbulence simulation is adequate resolution of the domain under consideration. The precise conditions are difficult to derive explicitly and are not reproduced here; however, they are derived in detail in Grötzbach (1983) for the Raleigh-Bernard convection problem and are further adapted for pipe flow in Eggels *et al.* (1994). Summarising these works, to perform a fully-resolved simulation we must:

- (i) have a large enough domain to contain large-scale vortices and structures;
- (ii) correctly reproduce the steep gradient of the solution field found at no-slip boundaries;
- (iii) ensure a small enough mean grid width to reproduce small-scale structures.

The first of these criteria is easily satisfied by using existing experimental knowledge of large-scale structures found in pipe flow. However, it is important to note that the presence of these structures will severely impact on the amount of computer time needed to perform the simulation.

Determining precise conditions to satisfy the second criterion involves examining in finer detail the velocity field near the pipe wall. To do this, one defines an appropriate dimensionless scaling which represents the size of structures that can be found around the boundary of the pipe.

Definition 4.1.1 (Wall unit). Let y denote the distance from the wall. Then a *wall unit* y^+ is defined as

$$y^+ = \frac{u_\tau y}{v}$$

Here u_{τ} is the *friction velocity*,

$$u_{\tau} = \sqrt{\frac{\tau_w}{\rho}}$$

and τ_w is the wall-shear stress as defined in the previous chapter.

In a similar fashion, we can use the friction velocity to transform the axial flow field near the wall by writing

$$u^+ = \frac{u}{u_\tau}.$$

In these units, fully developed turbulence can be described in a universal fashion. The boundary layer itself can be decomposed into various sublayers which scale according to a variety of laws.

Of particular importance is the viscous sublayer, which has a width of around five wall units. Eggels *et al.* (1994) states that a mesh will be sufficiently refined if at least three nodal points are placed inside this sublayer. In addition, if the mean grid spacing Δ satisfies $\Delta \leq \pi \eta$, where η is the Kolmogorov length scale (the scale of the smallest structures that must be resolved), then the third criterion will be satisfied. An alternative computational investigation of pipe flow performed by McIver *et al.* (2000) using a spectral element formulation additionally states that the first non-boundary collocation point should be placed such that $y^+ < 1$.

There is little or no documentation about the resolution that is required for transitional or intermittent dynamics. Mostly it is assumed that if the domain is resolved correctly for isotropic turbulence then this is sufficient for transitional simulations. To that end, this is the approach adopted here, although it is not apparent whether this is the correct approach to take.

In the following sections, we will attempt to identify a suitably resolved spectral element mesh. We begin by introducing the Chebyshev collocation scheme which is often used in spectral discretisations of pipe flow in the radial direction. We then move on to discuss two different mesh constructions which emulate the placement of the Chebyshev nodal positions. With the use of fully-developed turbulence simulations, we attempt to identify a suitable candidate mesh for further use.

4.1.1 Chebyshev-based meshes

Due to the simple geometry of a cylinder, spectral collocation schemes are especially useful for simulating pipe flows due to their excellent convergence properties. However, we must choose basis functions which correctly enforce the various boundary conditions. In chapter 2

it was shown that Fourier trigonometric functions make an efficient choice of modal functions in problems requiring the use of periodic boundary conditions. In the context of pipe flow then, we may naturally apply these modes in both the azimuthal and axial directions. At the pipe wall however, we must impose a no-slip condition, which requires the use of a Dirichlet boundary condition, and so alternative modes must be used.

Definition 4.1.2 (Chebyshev polynomials). The Chebyshev polynomials $T_k : [-1, 1] \rightarrow \mathbb{R}$ are defined by

$$T_k(x) = \cos(k\cos^{-1}x), \quad k \ge 0,$$

and form an orthogonal set of functions under the L^2 inner product with the weighting $(1-x^2)^{-1/2}$.

Chebyshev polynomials are widely used in spectral collocation schemes. Whilst differentiation is not as readily performed as in the Fourier case, they do allow for Dirichlet boundary conditions to be easily implemented. Additionally, by selecting the nodal points

$$x_j = \cos\left(\frac{j\pi}{N_r}\right),\,$$

we obtain the computationally useful equality

$$T_k(x_j) = \cos\left(\frac{\pi j k}{N_r}\right)$$

and so the fast Fourier transform can be used to efficiently convert to and from spectral space, making tasks such as differentiation both easy and highly accurate.

In the context of pipe flow, a typical choice $N_r = 40$ radial points and $N_{\theta} = 48$ azimuthal points can be seen in figure 4.1(a), which allows computations to be well-resolved up to Re \approx 3000. As well as being computationally efficient, the density of nodal points increases towards the pipe wall which assists in the extra resolution required in the boundary layer.

4.1.2 Spectral-element discretisation

Whilst the use of an elemental method seems somewhat of an odd choice to use in the setting of pipe flow, their non-uniform properties allow us to increase resolution in specific parts



(a) Chebyshev-Fourier mesh with (b) Spectral-element mesh using $N_{\rm el} = N_r = 40$ and $N_{\theta} = 48$. 64 and two-fold reflective symmetry.



(c) Spectral-element mesh using $N_{\rm el} = 44$ and rotational construction.

Figure 4.1: Examples of different types of mesh construction for the two-dimensional circular cross-section of the pipe. Spectral-element meshes use polynomial order P = 12. Elements which are filled in indicate the placement of GLL nodal points.

of the domain. By constructing a mesh for the two-dimensional circular cross-section of the pipe, we may place smaller elements around the boundary layer, where finer resolution is required, and larger elements towards the centre of the pipe where typically the scales of structures are much larger. To a certain extent, this mitigates the problem highlighted in figure 4.1(a), where the azimuthal discretisation leads to a build-up of nodal points at the centre of the pipe.

Another significant advantage is the ability to use a Cartesian co-ordinate system. Under a cylindrical polar discretisation, we must employ the use of purely computational boundary conditions along the pipe axis, such as those derived in Blackburn & Sherwin (2004). Unless carefully implemented, this boundary condition may lead to additional computational errors

not found under a Cartesian formulation. This is particularly highlighted in McIver *et al.* (2000) which investigates the errors found in Cartesian and cylindrical formulations in domains of comparable resolution.

Elements can generally be any convex polygon. Here however, we will consider the quadrilateral elements used in Semtex. In particular, meshes used in Semtex and other elemental codes are often required to be *conformal*; that is, any two intersecting elements must possess a common edge. A strategy must therefore be employed to reduce the number of elements found towards the centre of the pipe. There are three problems we must overcome:

- accurately representing the outer curve of the disk;
- partitioning the disk using quadrilateral elements to form a conformal mesh;
- using as few nodal points as possible to resolve the domain accurately.

An easy solution to the first point is to use many elements to approximate the curve of the circle. However, this is computationally inefficient and is not easily adapted to be conformal. Instead, it is desirable to curve the edges of elements, which provides a better fit to the pipe wall. This may be accomplished by altering the bijection χ^e which maps the standard region Ω_{st} to the *e*-th element Ω^e . This process requires careful thought so that the structure of the resulting mass and Laplacian matrices remains relatively undisturbed, and therefore it is not reproduced here. Karniadakis & Sherwin (2005) investigate this problem in far greater detail.

Figures 4.1(b) and 4.1(c) show two different styles of construction that can be employed, both of which have been used to perform the simulations throughout this thesis. Whilst both have forms of reflective and rotational symmetry, we define figure 4.1(b) to be an example of a 'reflective' mesh since the circle is split into four identical quarters. On the other hand, the mesh of figure 4.1(c) is designed to emulate the placement of the Chebyshev collocation points but remove the large build-up of points around the origin. In addition it curves internal elements to reduce the mean grid spacing and hence is deemed a 'rotational' mesh.

In the following section, we will demonstate that the rotational mesh style has many advantages over the reflective mesh. The most obvious of these is that the number of elements $N_{\rm el}$ is nearly half that of the reflective mesh. Additionally, the rotational mesh has far fewer elements containing obtuse angles which cause numerical errors to become larger under the spectral element formulation. Finally, it is easier to generalise the mesh depending on the level of resolution required. We define figure 4.1(c) to be a $K_{8,5}$ mesh, where in general, $K_{\theta,r}$ denotes a mesh with θ elements placed azimuthally and r rings (excluding the four central elements) placed radially. Obviously this notation is not well-defined in the sense that the placement of elements is not unique, but provides a useful shorthand to refer to each mesh.

4.1.3 **Resolution testing**

In order to satisfy the third requirement of the previous section, a number of short experiments have been performed to determine suitable choices for the polynomial order P and element mesh $K_{\theta,r}$. Suitable in this sense means that the parameters are chosen which are small enough to minimise computational effort, whilst providing a sufficiently resolved domain up to Re \approx 3,000 and hence provide the ability to simulate fully-turbulent flows.

In these simulations, the domain length and axial resolution are fixed at L = 10D and $N_x = 386$ respectively, providing a domain capable of generating data over long timescales in short periods of execution time, and also allowing Δ to remain sufficiently small. Periodic boundary conditions are taken in the axial direction and the flow is driven using the volumetric flux technique of section 3.2. A fully turbulent flow is simulated for a short period of time using a very high-resolution mesh. The resulting flow field provides an initial condition to test any given candidate mesh $K_{r,\theta}$.

A spectral interpolation technique is used to project the high-order flow field onto each candidate mesh, and as a result it is likely that some small numerical errors may appear. For instance, some meshes may contain nodal positions near to or on the boundary which do not appear on the high-resolution mesh. In order to alleviate these problems, the field is allowed to 'burn in' for 100-200 time units, providing sufficient time for transient effects resulting from the interpolation to decay.

After the burn-in period, a large number of quantities are recorded; of particular interest are:

- regular outputs of field data;
- time-averaged field data;



Figure 4.2: Trace of the axial velocity u(0, t) for the instigation of turbulence in a pipe of length L = 20D. The figure is split up into three regions: (a) initial deformation of the turbulent profile; (b) onset of instability; (c) transition to statistically steady turbulent flow.

- the forcing term $\alpha(t)$ used to impose the constant volumetric flux condition;
- shear stress data around the boundary to determine τ_w .

By comparing these quantities to existing literature, or the results from the high-resolution mesh, we can draw conclusions pertaining to the suitability of any given candidate mesh.

The choice of high-resolution mesh is critical, since it is the benchmark to which all other meshes are compared. Initial testing revealed that the reflective mesh of figure 4.1(b) using P = 12 provides a sufficient mesh to provide a good comparison. Indeed, comparison against literature (in particular McIver *et al.*, 2000, which also uses this mesh) shows it remains well-resolved up to Re \approx 5,000.

To generate the initial turbulent field, we use a standard technique seen in much of the existing literature. An initial laminar condition u_0 is generated with the mean profile of a turbulent field. To this, we add normally-distributed white noise with a small amplitude. A good approximation to the mean profile can be obtained by using a one-seventh power law, so that the initial condition becomes

$$u_0(x, r, \theta) = (C(R - r)^{\frac{1}{7}}, 0, 0) + 10^{-4} n(x, r, \theta)$$



Figure 4.3: L^{∞} error of the near-wall axial velocity u_z^+ for various polynomial orders *P* for four candidate meshes.

where C is a constant ensuring that the flux $Q(u_0)$ is approximately equal to the bulk velocity and $\mathbf{n}(x, r, \theta)$ is normally distributed noise with zero mean. Figure 4.2 shows the resulting transition to turbulence for this initial condition in a pipe of length L = 20D and diameter D = 1. Part (a) of the figure shows a deformation of the initial mean-turbulent profile into one approaching laminarity. However, after around 75 time units, part (b) shows the sudden growth of the small perturbations, leading to the initial onset of turbulent flow. These perturbations grow most rapidly in the axial component of velocity, and as such the flowrate restriction forces the mean velocity to remain near its maximum of $u_{\text{max}} = 2$. After a short period of apparent stability, the disturbance grows large enough in the transverse velocity that a sharp drop in mean velocity to a statistically steady turbulent state is seen, as shown in part (c) of the figure.

One remaining question is that of the length of time required to generate suitable averages for the field data and wall-shear stress. Here, each simulation of a candidate mesh gathers precisely 5×10^2 time units of data after burn-in, a quantity which is regarded to be far more than is statistically necessary. For the high-resolution mesh, we use a field obtained by averaging over 2×10^3 time units to obtain a more accurate temporal average.

A classical test for measuring the accuracy of numerical turbulence results is a semi-

logarithmic plot of the non-dimensionalised mean axial wall velocity u^+ as a function of wall units y^+ . This curve has a distinctive shape and is very sensitive to numerical error. To obtain a general overview as to the suitability of each mesh, a velocity field $\tilde{u}^+(y^+)$ is generated and the L^{∞} error $\|\tilde{u}^+ - u^+\|_{\infty}$ is calculated, where u^+ is the field obtained from a simulation on the high-resolution mesh. Figure 4.3 shows these errors for polynomial orders $7 \le P \le 12$ on the reflective mesh of figure 4.1(b) and a variety of rotational meshes: namely, $K_{6,5}$ and $K_{8,r}$ for r = 4, 5.

There are a number of immediate conclusions that can be drawn from this plot. Firstly, increasing *P* broadly decreases the error for all meshes, demonstrating that this is a reasonably suitable metric by which to compare them. In addition, the reflective mesh (for which P = 12 is the high-resolution comparison mesh) shows consistently similar results for $P \ge 10$, indicating its suitability for comparison at P = 12.

The worst performing mesh is clearly $K_{6,5}$, which shows persistent under-resolution even at high polynomial order. The radial spacings here are precisely those used in the $K_{8,5}$ mesh which performs significantly better at all polynomial orders. This indicates that 6 azimuthal elements are insufficient to correctly resolve the domain, and we therefore require $\theta \ge 8$ azimuthally placed elements. To determine the suitability of the radial resolution, we consider another mesh $K_{8,4}$ which has fewer radial elements. At P = 7 and 8, the mesh performs similarly to the $K_{6,5}$ mesh, and indicates the domain is significantly under-resolved in the radial direction. However, for $P \ge 10$, the error is more similar to that seen in the $K_{8,5}$ mesh indicating that the errors seen are more likely to arise from azimuthal under-resolution.

These results demonstrate that the $K_{8,4}$ mesh at P = 12 provides an ideal compromise of azimuthal and radial resolution, providing a very close approximation to the comparison mesh. Figure 4.4 shows a standard semi-logarithmic plot of u^+ against wall units y^+ for the comparison mesh (solid blue line) and the $K_{8,4}$ (dots) mesh which demonstrate the accuracy of the fit across the whole of the curve.

The significant advantage of this mesh is the improvement in computational speed. For any mesh, the total number of degrees of freedom in the system is given by $N_{dof} = 4N_x N_{el} P^2$. (The multiplication by four arrives from the three components of the velocity field and the pressure field.) For the high resolution mesh, $N_{el} = 64$, whereas the $K_{8,4}$ mesh only has $N_{el} = 36$ elements, thus reducing the overall number of degrees of freedom by a factor of



Figure 4.4: Axial component of the near-wall velocity u_z^+ in a full-turbulence simulation at Re = 3000 using two different types of meshes.

around two at any given P.

4.2 Transitional Dynamics

The investigation of pipe flow has its origins with Reynolds (1883a),b, whose famous experiments demonstrated the classical relationship between the onset of turbulence and the dimensionless parameter known today as the *Reynolds number*

$$\operatorname{Re} = \frac{\overline{U}D}{\nu}.$$

Reynolds discovered that above some critical threshold, the underlying Hagen-Poiseuille equations will, if sufficiently disturbed, give rise to turbulence. Similarly, below this threshold the parabolic Hagen-Poiseuille velocity profile is recovered. His experiments demonstrated that this transition occured around Re $\approx 2,000$, although he was not able to determine a precise value for the transition point. Finding the exact value for this point has remained a topic of fundamental importance and remains an open problem. In chapter 5 we will investigate this problem in far greater detail.



Figure 4.5: Schematic traces of the axial velocity u for a puff (left) and slug (right) along the pipe axis x advected from left to right. Leading edges (LE) and trailing edges (TE) are noted on each figure. u_{max} signifies the maximum attainable laminar velocity (equal to 2 if $\overline{U} = 1$).

In practice, pipe flow is extremely susceptible to perturbations. Reynolds was originally able to maintain laminar flow up to an impressive Re $\approx 12,000$, but this value has been significantly extended; Pfenniger (1961) demonstrated laminar flow up to 10^5 . However, extreme measures need to be taken to isolate the system from external noise and to reduce disturbances at the in-flow. The sensitivity of pipe flow is aptly described by Wygnanski & Champagne (1973), who noted that placing a standard incandescent light bulb 10 feet from their experimental apparatus led to a significant change in their observed data.

Mathematically however, the laminar equations of motion remain extremely resilient to the application of infinitesimal perturbations. Unlike many dynamical systems, in which the underlying stable profile undergoes a bifurcation to instability upon the change of a parameter, it is believed that the base flow profile remains linearly stable (Meseguer & Trefethen, 2003), with proofs existing in the case where the infinitesimal perturbation is axisymmetric (Herron, 1991). The conflicting nature of these theoretical and experimental observations has lead pipe flow to become a topic of great interest for mathematicians, physicists and engineers alike, motivating a large number of studies over the last 125 years.

Most of these studies have focused on the investigation of coherent structures found in the flow as the flow transitions from laminarity to turbulence. Reynolds' original experiments noted that for Re $\approx 2,300$, applying a perturbation to the fluid leads to isolated 'flashes' in the fluid which grow and split to contaminate the downstream flow. The work of Wygnanski & Champagne (1973) greatly expanded the understanding of these flashes and classified them into two distinct states. *Puffs* are small regions of intense turbulence found around

transitional Re, which maintain a length of approximately 25D and are categorised by their distinctive trailing edge (TE), schematically defined in the left side of figure 4.5. *Slugs* are the high-Re equivalent of puffs, with a far sharper leading edge (LE) and are found for Re which is typically associated with fully-developed turbulence.

Both of these structures can be generated by sufficiently disturbing the laminar flow profile with a finite-amplitude perturbation at various Re. Puffs are found for values as low as Re \geq 1,700, and at these lower Re will often decay quickly. For Re \geq 2,300, splitting events occur with new puffs being nucleated downstream from the leading edge of existing puffs (Nishi *et al.*, 2008). Between these two values, puffs remain isolated and of a constant length, and are thus deemed 'equilibrium' puffs; a topic investigated at length by Wygnanski *et al.* (1975) and further by Bandyopadhyay (1986). Slugs evolve in a radically different fashion, with the TE moving far slower than the mean flow and the LE advected much faster, so that any disturbance rapidly contaminates the downstream flow (Duguet *et al.*, 2010).

Until recently, the simulation of puffs and slugs has proven beyond the reach of computation due to the large domains required to adequately resolve them. The first fully-resolved turbulence computations at Re = 3,500 were performed by Eggels *et al.* (1993) in a pipe of length L = 5D and in Eggels *et al.* (1994) these were compared to existing experimental data. Later work by Shan *et al.* (1999) used a 1D/2D joint spectral-element–Fourier code to simulate puffs and slugs and expanded these domains to a much larger $L = 16\pi D$; however the computations here suffer from an under-resolution in the axial direction, especially in the simulation of slugs at Re = 5,000. Faisst & Eckhardt (2004) examined decay rates of puffs in a fully-resolved domain of length L = 5D; however too short to fully encompass a single puff. Most recently, work by Willis & Kerswell (2007) employed the use of an efficient hybrid finite-difference–spectral solver in domains of length L = 50D and are the first examples of fully-resolved simulations in domains large enough to faithfully represent puffs. Further work has also made significant use of computational resources in order to simulate puffs and slugs (Willis & Kerswell, 2008, Hof *et al.*, 2010, Avila *et al.*, 2010, Duguet *et al.*, 2010).

However, the common theme amongst most of these works is that they concentrate on the transition from a laminar to turbulent flow. This requires the use of a perturbation and thus the transition may depend upon both the type and amplitude of the disturbance applied. In

addition, most studies have concentrated on investigating the temporal aspects of the flow; for instance, measuring the average time taken for turbulence to decay. Complex spatio-temporal interactions, such as the interaction of two puffs, have been for the most part overlooked.

Another line of research has emerged concerning alternating turbulent-laminar flow states which form in subcritical shear flows over large length scales (Prigent *et al.*, 2002, 2003, Barkley & Tuckerman, 2005, 2007, Tsukahara *et al.*, 2005, Manneville, 2009). This work has established that in plane Couette flow (Prigent *et al.*, 2002, 2003, Barkley & Tuckerman, 2005, 2007), counter-rotating Taylor-Couette flow (Prigent *et al.*, 2002, 2003), and plane Poiseuille flow (Tsukahara *et al.*, 2005), near transition the system can exhibit a remarkable phenomenon, in which turbulent and laminar flow form persistent alternating patterns on scales very long relative to wall separation and the spacing between turbulent streaks. While the origin of these patterns remains a mystery, they are intimately connected with the lower limit of turbulence in shear flows.

In this section then, we will attempt to quantify whether pipe flow is also capable of supporting these alternating bands, and, if they exist, investigate the form that they take.

4.2.1 Simulation methodology

The general methodology for these series of simulations is to examine a pipe of length L which is sufficiently long to encompass large-scale structures. Periodic boundary conditions are used at each end of the pipe, removing the need for inflow and outflow boundary conditions and thus the experimental constraint of a domain of the order of thousands of diameters. Whilst it is possible to argue that the choice of periodic boundary conditions allows us to use an arbitrarily small domain length, it must be large enough to encompass the structures discussed thus far. At a minimum then, we require $L \ge 25D$ so that the domain may resolve a single puff. However, if larger-scale structures exist then they will not be represented. Here then, we opt to use a pipe of length $L = 40\pi D \approx 125D$. Whilst extremely computationally expensive, this very large choice of length allows us to study very large-scale structures, should they exist, in the transitional regime.

The simulation is started with a uniformly turbulent state at Re = 5,000 using the reflective mesh of section 4.1.3 at polynomial order P = 12. Whilst extremely expensive, this

increased resolution allows us to use the perturbation technique shown in figure 4.2 to generate the turbulent initial condition. The flow is driven using a constant volumetric flux so that the bulk velocity $\overline{U} = 1$, ensuring that the simulations are performed at a prescribed Reynolds number.

Since the turbulent flow is more resilient to large alterations in parameters (e.g. mean flow speed) at higher Re and thus less prone to abrupt relaminarisation, we rapidly reduce Re to 4,000 and 3,000 in quick succession, simulating each for around 300 time units (D/\overline{U}) . Nearer the transitional regime, we reduce the polynomial order to P = 11 and decrease Re in finer increments; initially, in steps of 200 and then in smaller steps of 50. We also increase the simulation length to examine lower Re cases in greater detail, with each case being run for 2,000 time units. This simulation represents the longest continuous fully-resolved numerical simulation of transitional pipe flow to date.

4.2.2 Identifying turbulence

Given the large number of degrees of freedom, the field data for this simulation is too large to record frequently. Instead, at any given time t_n , we sample this data at a subset of points $H = \{(x_k, 0, 0, t_n) \in \Omega \times [0, \infty) \mid 0 \le k \le N_x\}$ called *history points*, where $x_k = k\Delta x$ denotes the position of each two-dimensional circular Fourier plane along the axis of the pipe. Since the mean flow obtains its maximum at the centre of the pipe, this allows us to obtain a general overview of the state of the system without recording all field data.

Given the existence of possible intermittent states, it is also necessary to construct a quantity to identify the intensity of turbulent areas throughout the pipe length. A classical choice as used in Pope (2000) is the vorticity $\boldsymbol{\omega} = \nabla \times \boldsymbol{u}$. However, this cannot be constructed from the data contained in H. Instead, we consider a scalar based upon (v, w), the transverse components of velocity. Whenever the flow is laminar, we must have that v = w = 0; therefore in areas of low turbulence, v and w must be close to zero.

We thus construct the scalar field $q: H \to \mathbb{R}$

$$q(x_k, t_n) = \sqrt{v(x_k, t_n)^2 + w(x_k, t_n)^2},$$
(4.1)

which is related to the transverse kinetic energy $E_{\text{transverse}} = \frac{1}{2}(v^2 + w^2)$, an important

quantity for studying the transport and dissipation of turbulent flows. This satisfies the condition that q = 0 whenever the fluid is laminar. Composing this with the square root intensifies areas of turbulence and allows us to see areas of very-near laminar behaviour.

Since $\overline{U} = 1$, a fluid particle along the axis will travel through the pipe once every 125 time units on average, and so it is desirable to view the data in a moving frame of reference. Given a speed *c*, this can be achieved through use of the Galilean transformation

$$q(x,t) \mapsto q(x-ct,t).$$

This transformation would be somewhat difficult to implement in a standard inflow-outflow problem, since it requires changing the no-slip boundary condition so that $u(\partial \Omega) = (0, 0, c)$, and hence assumes a prior knowledge of the speed c.

Another advantage of using periodic boundary conditions then is that it allows us to apply the transformation for any given c when post-processing the data. Expanding q as a Fourier series,

$$q(x+ct,t) = \sum_{k \in \mathbb{Z}} \hat{q}_k(t) e^{ik(x-ct)} = \sum_{k \in \mathbb{Z}} [\hat{q}_k(t)e^{-ikct}] e^{ikx},$$

and so the transformation can be obtained through a simple multiplication of the Fourier modes of q.

Although c can be any arbitrary constant, there are two choices which would be suitable for this simulation. We may pick the bulk velocity \overline{U} by which all velocites are nondimensionalised; otherwise, we could pick the mean puff velocity U_{puff} . The latter option is not suitable for these simulations, since slugs are the dominant structures found at the initial condition where Re = 3,000. In addition, even where puffs exist at lower Re, their position (and therefore speed) is somewhat difficult to determine, as well as being Reynolds number dependent (de Lozar & Hof, 2009). A better choice is therefore $c = \overline{U}$, which provides a consistent frame of reference to compare the states over a range of Re.

4.2.3 Results

The results of the above data extraction process can be seen in figure 4.6. The main figure shows the space-time contour of q moving in a frame of reference moving at a speed \overline{U} .



Figure 4.6: Dynamics of transitional flow from simulations of a L = 125D pipe. The central plot shows a space-time diagram with the streamwise direction horizontal and time increasing vertically upwards. The value of Re changes as indicated on the right. The magnitude of the transverse velocity q is sampled along the axis of the pipe and visualised in a frame moving at the average fluid velocity: q(x - Ut, t). Colours are such that light corresponds to turbulent flow and black corresponds to laminar flow. Below and above are flow visualizations in vertical cross-sections through the pipe at the initial and final times, and over the 25D streamwise extents indicated by arrows.

From the bottom of the figure moving upwards (i.e. as time increases), Re is decreased in discrete steps at the times indicated by ticks on the right of the figure. Three uniquely identifiable regimes can be visually identified.

At Re $\gtrsim 2,600$, the flow remains uniformly turbulent, with no large-scale variation evident in the turbulent structure throughout the pipe as visualised in the bottom portion of the figure. This is easily contrasted with the behaviour seen as Re is reduced below 2,600. Regions of nearly laminar flow (identified by dark colours in the figure) are seen to spontaneously form and disappear within large areas of turbulence. Slightly downstream of these regions one finds a sharp increase in the turbulent activity, similar to that seen diagramatically in figure 4.5. Furthermore, this change in behaviour occurs immediately upon the change of Reynolds number from 2,600 to 2,500. At this point, the fluid is in an *intermittent* state, possessing qualities of both turbulence and laminarity.

As Re is reduced further, the small laminar patches that appear at Re = 2,500 grow steadily in size, giving rise to a more regular alternation of turbulent and laminar flow. The turbulentlaminar states are not, however, steady; there are splitting and extinction events which occur in a sporadic fashion, sometimes in quick succession. (For an example, see $t \approx 8,600, x \approx$ 100*D* for an extinction and $t \approx 8,000, x \approx 80D$ for a splitting event.)

These extinctions and splittings slow as Re is reduced so that the turbulent patterns become almost steady, and at the end of the simulation there are four clear turbulent structures separated by relatively laminar flow. The upper flow visualisation reveals that these turbulent structures are, in fact, equilibrium puffs, identical to those observed by Wygnanski & Champagne (1973) and commonly seen in many other experiments.

The 'slant' of the puffs is due to the fact that the frame of reference here is taken using the bulk velocity \overline{U} . Puffs will typically travel at speeds slower than the bulk velocity as computed in Shan *et al.* (1999), since they require energy from the mean flow to pass through them in order to sustain turbulence. Typical speeds for Re = 2,250 are around $U_{\text{puff}} \approx 0.95\overline{U}$, demonstrated experimentally by de Lozar & Hof (2009), and have been validated to be approximately the same here, giving another indication that the domain is correctly resolved.

Figure 4.6 shows that pipe flow exhibits a variety of previously unseen behaviour. Firstly, it conclusively demonstrates the spontaneous emergence of (equilibrium) puffs from a

uniformly turbulent state. The dynamics of the flow at transition then naturally incorporate puffs, and they are not a result of random perturbations to the flow. Moreover, states involving trains of two or more puffs are shown to exist that do not end in the annihilation of all but one puff, conclusively demonstrating that the turbulent-laminar alternation seen in other flows also exists in pipe flow. Spatio-temporal intermittency therefore plays a key role in the transition process.

The rest of this chapter will be dedicated to investigating the states and transitions contained in figure 4.6. Whilst the onset of intermittency is intuitively and visually clear, it still remains for it to be quantified. In addition, a second transition, which is not clearly identified in the figure, can be identified which distinguishes the unsteady intermittent flashes to the more ordered behaviour seen towards the end of the simulation. We shall investigate the latter transition first.

4.3 Transition between localised and intermittent turbulence

In figure 4.6, there is a striking difference between the unsteady, constantly changing turbulent flashes seen at Re = 2,350 and the more ordered puff-like behaviour seen at Re = 2,250. The splitting and extinction events which occur at higher Re seemingly happen at random, and therefore the major question that may be asked is: "is there a critical Re below which puffs remain isolated?" The work presented in this section represents an initial attempt at answering this question, based on the dynamics seen in figure 4.6.

4.3.1 Methodology

The methodology for this series of simulations will closely follow the domain expansion approach adopted for Couette flow by Barkley & Tuckerman (2005), in which a similar pattern of alternating laminar-turbulent bands is examined. The initial condition is taken to be a single equilibrium puff in a domain of length L = 25D with periodic boundary conditions, where the domain is resolved with the same density of collocation points as in the previous section and a constant volumetric flux $\overline{U} = 1$ is enforced. Initially, the puff is evolved for 1,000 time units to ensure that it remains stable and does not abruptly relaminarise.

L (inclusive)	N_{x}
25 <i>D</i>	384
$30D \rightarrow 35D$	512
$40D \rightarrow 45D$	768
$50D \rightarrow 70D$	1,024
$75D \rightarrow 90D$	1,536

Table 4.1: Axial resolution requirements for the pipe expansion experiment of section 4.3.1 for various domain lengths L.

After this initial period, L is increased by five diameters, and the solution field is stretched to fit the new domain length. The field is then timestepped for 500 time units, allowing enough time for the solution to stabilise after the stretching procedure and any transient effects to disappear. This lengthening procedure is applied repeatedly at 500 time unit intervals until L = 90D, at each stage ensuring that the domain remains correctly resolved by increasing the axial resolution by use of the values stated in table 4.1.

4.3.2 Results

The results of this simulation can be seen in the space-time diagrams of figure 4.7. The difference between the two states is immediately apparent. At Re = 2,250, a single localised turbulent puff persists independently as L is increased. On the other hand, at Re = 2,350 any change in L results in splittings so that puff structures fill the entire length of the pipe and the ratio of turbulent to laminar flow is kept constant.

We begin by looking in more detail at Re = 2,250. Once *L* is sufficiently large, the predominant state is that of a single localised puff, and this state remains unchanged throughout the entire simulation. Surrounding the turbulent puff is not only laminar flow; far downstream, fully developed parabolic Hagen-Poiseuille flow is recovered. This is demonstrated by the trace of the axial velocity in the upper part of figure 4.8 which shows that the centre-line velocity asymptotes to $u_{\text{max}} = 2$.

Note that, as is clearly shown in figure 4.6, states exist at Re = 2,250 which contain multiple puffs, and of course since these simulations are run in a periodic domain, the solution here actually corresponds to multiple puffs in an infinitely long pipe. However, given the long domain length and the re-laminarisation downstream of the puff, the simulation here



Figure 4.7: Simulations highlighting the difference between localised turbulence at Re = 2,250 (left) and spatio-temporal intermittency at $\vec{Re} = 2,350 \text{ (right)}.$



Figure 4.8: Streamwise velocity u and y-component of the transverse velocity v for the final state of figure 4.7 at Re = 2,250 (top) and Re = 2,350 (bottom).

suggests that the long-time, large-domain state at Re = 2,250 is persistent localised patches of turbulence travelling at a fixed speed and maintaining a constant size, residing in a background of laminar flow.

By contrast, the behaviour seen on the right of figure 4.7 at Re = 2,350 shows far more complex dynamics than those at Re = 2,250. The single puff at L = 25D undergoes splittings at L = 50D and L = 75D so as to maintain the spatial proportion of turbulent to laminar flow. Whilst this splitting is unsteady and unordered (especially during the last 500 time units), the lower portion of figure 4.8 shows the distinctive velocity traces of three puffs travelling closely in the pipe. Note that the maximum parabolic velocity $u_{max} = 2$ is never attained; instead, the puffs appear to maintain a (relatively) fixed distance from one another, never allowing the flow to fully re-laminarise. Further work by Samanta *et al.* (2010) on this topic has confirmed that puffs maintain a fixed interaction distance $d_{min} \approx 20D$ at Re = 2,350 so that in a pipe of length 90D, there should be a maximum of three puffs,
corresponding to the number seen in these studies.

The striking difference in the two simulations performed here lends credibility to the idea of a transition point between *intensive* and *extensive* dynamics; two terms which are often referred to in the context of thermodynamical systems. An extensive quantity is dictated by the size of a domain, whereas intensive quantities remain scale invariant. At Re = 2,350, turbulence can be viewed as an extensive quantity, since the amount of turbulent flow is dictated by L and not the initial conditions. However at Re = 2,250, the turbulence becomes an intensive quantity since changing L has no effect on the proportion of turbulent flow in the pipe. Somewhere between these two values then, there should exist a critical point which may be approximated as 2,300 from this data. Beyond this value, the dynamics become extensive leading to the splitting of puff states, causing turbulence to spread throughout the domain and thus leading to a sustained turbulence state. Therefore the intensive-extensive transition point can be classified as the critical Reynolds number and thus $Re_c \approx 2,300$.

Whilst this is a valid conclusion from the available data, it is not however entirely correct, since as will be explained in chapter 5 we must not only consider a single simulation, but in fact multiple simulations each of which uses a different puff state as initial condition. The reason for this stems from the investigation of the lifetimes of puffs, in which the question asked is, "does there exist a critical Reynolds number above which puffs do not decay?" The lifetime studies of Faisst & Eckhardt (2004) conclusively demonstrated that the decay of puffs is governed by a memoryless process which appears stochastic and is characterised by a mean lifetime. The precise scaling of this lifetime as a function of Re has received much attention since the original publication; however, further work by Hof *et al.* (2006), Hof *et al.* (2008) and Avila *et al.* (2010) amongst many others shows that given sufficient time there is always a non-zero probability of decay.

The notion of a stochastic process which is embedded within a deterministic system is somewhat strange since, from the standpoint of dynamical systems, any given initial condition has a prescribed and unchanging evolution through the phase space of possible solutions and no random effects are introduced to the solution in the manner of a stochastic PDE. However, the complexity of this space means that the evolution of two very similar, but distinct states can be radically different. In this case, the behaviour of the decay process presents itself as observing certain properties which can be deemed to be stochastic, thus the use of the terminology.

Since the nature of the decay process is stochastic, it therefore seems that to correctly interpret these results a large number of additional simulations must be performed, and this is done in chapter 5. However, it is important to note that even with this oversight, the results presented here still lend credibility to the idea of an intensive-extensive transition. By considering an ensemble of initial conditions, we will determine in chapter 5 what the nature of this point is, and whether it is indeed the case that all puffs split above Re_c , or whether some other stochastic effects must be taken into account.

The one clear observation from these simulations is that whilst lifetime studies seem to suggest that all localised puffs will eventually relaminarise regardless of Re, the data at Re = 2,350 in figure 4.7 shows that this is not likely to be the case. Even at Re = 2,250, the characteristic lifetime for the decay of a puff is above $\tau = 10^{10}$ (de Lozar & Hof, 2009), nine orders of magnitude greater than the time taken to observe a split in the simulation at Re = 2,350. From these results, we may infer that the upper bound on this transition point is Re_c \leq 2,350, but further simulations which take into account any stochastic properties of the splitting process must be performed in order to determine a precise value.

4.4 **Onset of intermittency**

The most visually apparent transition point of figure 4.6 is the appearance of laminar spots from the fully turbulent initial condition. As Re increases upward from 2,300 (moving downwards in figure 4.6), intermittent puff splittings and extinctions give rise to an increased proportion of turbulent flow in which puffs are not clearly identifiable. One continues to observe more intense turbulence on the downstream sides of laminar regions much like the trailing edge of a puff. Intermittent events are visually apparent until Re = 2,500, and minor fluctuations can be detected at Re = 2,600, but not at higher Re.

A natural question to ask then is "is there a critical Reynolds number below which the flow becomes intermittent?" Unfortunately this question, whilst simple to state, is difficult to answer and somewhat ill-posed. Firstly, whilst the onset of intermittency is easy to identify by eye in figure 4.6, we must address what is meant by 'intermittency' and thus identify a quantitative method for identifying flows which are intermittent. In addition, this method

should hopefully be independent of any artificial parameters so that the transition as a function of Re can be obtained.

In the sections below, we describe two different techniques for identifying intermittent flows, and outline the methodology for an experiment to identify a transition point. We then conclude with the results of this experiment along with an analysis of the data in hopes of identifying the transition point.

4.4.1 Intermittency factor

The intermittency factor is a classic metric representing the balance of turbulent to laminar flow which is present in a fluid at any given time. It has not typically been used in pipe flow studies, with the notable exception of Rotta (1956) who examined various unsteady states around transition using the intermittency factor. In order to define the intermittency factor, we must select a function which is zero around areas of laminar (or nearly laminar) flow. As in the previous sections, we adopt the use of the function q which only relies on the history points H.

Definition 4.4.1 (Intermittency factor). Let $q : H \to \mathbb{R}$ be defined as in section 4.2.2. Then, given a parameter $q_* > 0$, the *intermittency factor* γ_{q_*} is defined as

$$\gamma_{q_*} = \langle I(\cdot, \cdot; q_*) \rangle$$

where the mean is taken over all points in H and $I : H \to \{0, 1\}$ is the *intermittency function*

$$I(x_k, t_n; q_*) = \begin{cases} 1, & q(x_k, t_n) > q_*, \\ 0, & \text{otherwise.} \end{cases}$$

Under this definition, intermittent flow is characterised by states where $0 < \gamma q_* < 1$; i.e. the flow is neither fully turbulent nor fully laminar. Whilst γ is dependent on the choice of parameter q_* , by studying a variety of suitably small parameters we can (to a certain degree) make our analysis independent of q_* .



Figure 4.9: Intermittency function $I(x, t; 1 \times 10^{-2})$ for velocity traces of a single puff (top) and multiple puffs (bottom).

Figure 4.9 shows the intermittency factor I for two velocity traces along the axis of the pipe corresponding to the final conditions of figure 4.7. The intermittency function picks out areas of laminarity reasonably well. However, due to the fact that the transverse velocities cross zero and also the leading edge of the puff is not well-defined, turbulence can sometimes be mis-identified as laminar flow. There are various techniques which may be adopted to reduce this effect. As with other noisy signals, the data may be smoothed using a simple window-average technique or by more sophisticated convolution methods, hence reducing the number of crossing points. Alternatively we may use a binning procedure to make the spatial or temporal spacing coarser, removing small pockets of laminar flow. However, both of these techniques add additional parameters such as the width of the smoothing window or bin size, thus even when the field is uniformly turbulent, γ will not saturate to precisely one.

4.4.2 Laminar lengths

Whilst the intermittency factor provides a general overview of the proportion of turbulent to non-turbulent flow, it gives no indication as to the size of any laminar areas that may form within the fluid. Intuitively, as puff-like structures start to dominate the flow field, the size of laminar regions should increase and give another indication as to the onset of intermittency. This technique was adopted by Manneville (2009) to identify the appearance of laminar areas in two-dimensional cross-sections of Couette flow. It is also relatively simple to implement in pipe flow, since we need only consider the flow along the axis of the pipe. Here we define an implementation for pipe flow in such a way that it relies only on history points H.

Definition 4.4.2 (Laminar length). Let $I(x_k, t_n; q_*)$ be the intermittency function given in definition 4.4.1. At any given time t_n , suppose that there exist natural numbers i and ℓ such that $0 \le i < N_x$, $\ell \ge 1$, and

$$I(x_i, t_n; q_*) = I(x_{i+1}, t_n; q_*) = \dots = I(x_{i+\ell}, t_n; q_*) = 0.$$

Then the fluid is said to possess a *laminar length* of size ℓ .

Remark 4.4.3. Note that laminar lengths are still well-defined using the definition above even if $i + \ell > N_x$, since periodic boundary conditions are used in the axial direction and so $x_{i+N_x} = x_i$. Indeed, care must be taken to include periodicity when calculating laminar lengths of the step function *I*.

From these laminar lengths, it is possible to form a discrete probability distribution. Let $c_{\ell,n}$ denote the number of laminar lengths of size ℓ in the flow at time t_n . Given a simulation with N time samples, these weights then give a probability mass function $L : \{1, \ldots, N_x - 1\} \rightarrow [0, 1]$ where

$$L(\ell) = \sum_{n=1}^{N} \frac{c_{\ell,n}}{Z}$$
 and $Z = \sum_{n=1}^{N} \sum_{\ell=1}^{N_x - 1} c_{\ell,n}$

Given any particular simulation at a given Re, we can calculate L and its associated random variable $X \sim L$. From this we may calculate the mean and variance of X,

$$\mathbb{E}(X) = \sum_{\ell=1}^{N_x - 1} \ell \mathbb{P}(X = \ell) = \sum_{\ell=1}^{N_x - 1} \ell L(\ell), \quad \sigma^2(X) = \mathbb{E}(X^2) - [\mathbb{E}(X)]^2.$$



Figure 4.10: Visualisation of the laminar length distribution $L(\ell)$ for in a pipe of length L = 25D at Re = 2,200 (left) and Re = 2,600 (right).

For intermittent and fully turbulent flows, the laminar length distribution will be hugely peaked around $\ell = 1$, as demonstrated in figure 4.10. Even in the case of a single puff, the leading edge is unsteady and may, at points, 'break up' a large, mostly laminar length into many parts. However, as the flow becomes more intermittent, the distribution should become less sharply peaked as the number of larger areas increases. Of particular interest then is the variance, which measures the spread of the distribution and should therefore be sensitive to a transition from turbulence to intermittency.

4.4.3 Methodology

The methodology of this experiment is somewhat different than the two experiments described earlier in that we do not perform a single simulation, changing the Reynolds number at prescribed times. Instead, to begin, a uniformly turbulent flow is simulated at Re = 3,000and is evolved for 10,000 time units. The history data are sampled at intervals of 50 timesteps (10 times per time unit), and these form the samples from which the intermittency function and laminar length distributions are calculated.

We then start another experiment at a lower Re using a randomly chosen snapshot from this simulation, which is burned in for O(200) time units so that any transient effects of

the change of Re disappear and the simulation forgets its initial condition. The burn-in time is chosen both heuristically from figure 4.6 and upon the evidence of Eggels *et al.* (1994). Additional work by Avila *et al.* (2010) has shown that samples can be taken from a continuous simulation at as little as 10 time units apart and, upon a change of Re, forget their initial conditions. This choice of burn-in time is therefore highly likely to provide sufficient time for transient effects to dissipate. We continue this process for the whole series of Reynolds numbers we wish to simulate. Finally then, at each Re, we obtain 10^5 samples, large enough to calculate reasonable averages for both the intermittency factor and laminar length distribution.

Throughout all of the simulations, the domain length L is fixed at 25D for two reasons. Firstly, the computational costs of running a fully-resolved 10,000 time unit simulation at any given domain length are extremely large. This choice of L allows us to use fewer processors and hence run multiple simulations concurrently. Of equal significance however, the domain has deliberately been chosen to coincide with the length scale of a single puff. Whilst this choice is not particularly important for these simulations, the next section will re-use this simulation data and hence the choice will be further explained there.

4.4.4 Results

The results of the above simulations can be seen in figure 4.11 which depicts the intermittency factor and laminar length scales for $2,350 \le \text{Re} \le 3,000$. We first consider the intermittency factor. It is clear that, regardless of the choice of parameter q_* , γ increases continuously until Re $\approx 2,600$ where it saturates to $\gamma \approx 1$. Based on the intermittency factor alone then, we draw the conclusion that there exists a critical point at Re_i = 2,600 signifying to the transition to intermittency as Re decreases.

Figure 4.12 shows the gradual increase in size of laminar areas (coloured as black) as Re is reduced. In (c) at the transition point, laminar flow is barely distinguishable from background turbulence. Reducing Re to 2,450, we see a large increase in the size of laminar areas. Not indicated in the plot is the increased turbulent activity as puff-like structures begin to dominate the flow. As Re is reduced through to 2,300 in plot (a), these areas increase in size to that of regular puff structures.



Figure 4.11: Intermittency factor γ (top) and the variance of the laminar length distribution σ (bottom) for indicated Reynolds numbers. Each curve represents a different threshold, from $q_* = 4 \times 10^{-3}$ (top curve) to $q_* = 1 \times 10^{-2}$ (bottom curve) in increments of 2×10^{-3} .

Further evidence for the transition is given by studying the standard deviation σ of the laminar length distribution, depicted in bottom half of figure 4.11. Again at Re_i \approx 2,600, the flow undergoes a transition and thus the underlying distribution spreads as larger laminar areas appear within the flow. As Re reduces further, these areas dominate the flow and lead to a significant increase in σ . Figure 4.10 demonstrates this spreading process at two Reynolds numbers above and below Re_i.

Another well-known metric for characterising turbulent flows is through the use of onedimensional probability density functions of a single velocity component. Figure 4.13 shows f(v), the normalised pdf of the transverse velocity v, sampled at a single point along the pipe axis. It is well-established that f is Gaussian for sufficiently turbulent flow (in particular, see figure 21 of Tavoularis & Corrsin, 1981), although the mechanism by which this occurs remains unclear.



Figure 4.12: Representative space-time plots of the intermittency function $I(x, t; 8 \times 10^{-3})$ at a selection of three Reynolds numbers in a pipe of length L = 25D.

The probability density functions of figure 4.13 show similar behaviour whenever Re \geq 2,700, with Gaussian distributions being represented by quadratic curves. However, upon reducing Re below 2,600, the peaks of the pdf sharpen and the tails widen, so that the distributions become more indicative of a power-law type distribution. There is a very small deviation from the normal distribution at Re_i = 2,600, although this is hard to see in the figure.

4.5 Pattern formation analysis

Whilst the previous section succinctly captures the dynamics of the onset of intermittency, neither the intermittency factor nor the one-dimensional density functions encapsulate other



Figure 4.13: Normalised single-point velocity pdfs for Re = 2,350, 2,400, 2,450, 2,500, 2,600, 2,800 and 3,000, as indicated by labels and alternating line types. For Re \ge 2,700, the distributions are nearly Gaussian, and for Re \le 2,600 the distributions are non-Gaussian.

transition points, such as the intensive-extensive transition of section 4.3. Following the approach of Tuckerman *et al.* (2008), we investigate an alternative analysis of the data set of the previous section which allows us to incorporate this bifurcation point, and provides a compelling view of the essence of all three flow regimes and their respective transitions.

4.5.1 Methodology

As mentioned in the previous section, this experiment re-uses the L = 25D data set of the previous section. However, we extend the range of Reynolds numbers to include simulations from $2,250 \le \text{Re} \le 3,000$, incorporating Reynolds numbers at which localised behaviour was observed.

We begin by taking the discrete Fourier transform of q along the axis of the pipe;

$$\hat{q}_k(t_n) = \sum_{j=0}^{N_x-1} e^{ikx_j} q(x_j, t_n), \quad k = -N_x/2 + 1, \dots, N_x/2.$$

From a pattern formation perspective, the Fourier modes encode information about the

structure of the solution. In particular, if periodic features of wavelength $n\pi/L$ exist in the function this will be highlighted by an increase in the magnitude of the *n*-th Fourier coefficient.

In a pipe flow context, by choosing L = 25D so that the length of a pipe directly relates to the scale of a single puff, any disturbance to the fluid which can be accurately represented by q and has a wavelength of the length of the pipe should correspond to an increase in $|\hat{q}_1|$.

Similarly, in larger pipes, trains of puffs will be identifiable by looking at the Fourier mode corresponding to the number of puffs seen. However, as aptly demonstrated in figure 4.7, at certain Reynolds numbers puffs will remain in a localised state and although $|\hat{q}_1|$ will show an increase, the parameter will become far less sensitive as the structure is split over many more modes. Choosing L = 25D allows us to examine puff-scale structures without this additional complication.

Let us assume that \hat{q}_k is a complex random variable distributed according to the twodimensional probability density function $\rho_k(r, \theta)$ where (r, θ) describes a point $z = re^{i\theta} \in \mathbb{C}$. In general, approximating this PDF requires many samples from the probability space, which is expensive to generate with DNS.

One clear optimisation can be obtained by noting that q is measured at points along the centre-line of the pipe, indicating that θ simply encodes the axial position of the puff. Since the simulations use periodic boundary conditions, without loss of generality we may translate the puff to any streamwise location, and this implies that ρ_k is axi-symmetric; that is, $\rho_k|_{\theta \in [0,2\pi)}$ is constant.

To construct an approximation to ρ_k then, we take $1 \leq j \leq N_{\text{samp}}$ samples \hat{q}_k^j from the distribution and construct a step function $\overline{\rho}_k$ which is constant on annuli

$$A_i = \{ (r, \theta) \mid r_i \leq r < r_{i+1} \}.$$

Here we take $r_i = i \Delta r$ for $i \in \{0, 1, 2, ...\}$; i.e. a uniform radial grid with spacing Δr . Necessarily $\overline{\rho}_k|_{A_i} = N_i/Z_i N$, where N_i denotes the number of samples lying in A_i and Z_i is the appropriate weighting factor for normalisation. There are many ways to calculate Z_i corresponding to the measure we wish to use for the probability space. However, the most intuitive approach is to insist that ρ_k is approximately equal to $\overline{\rho}_k$ on each annulus A_i . In this case, integrating $\overline{\rho}_k$ over A_i yields

$$\int_{A_i} \overline{\rho}_k \, \mathrm{d}A = \frac{2\pi N_i}{NZ_i} \int_{r_i}^{r_{i+1}} r \, \mathrm{d}r = \frac{\pi N_i}{NZ_i} (r_{i+1}^2 - r_i^2)$$
$$= \frac{\pi N_i}{NZ_i} (2r_i \Delta r - \Delta r^2)$$

and hence taking $Z_i = \pi (2r_i \Delta r + \Delta r^2)$, $\overline{\rho}_k$ gives the desired property; i.e.

$$\int_{A_i} \rho_k \, \mathrm{d}A \approx \int_{A_i} \overline{\rho}_k \, \mathrm{d}A.$$

4.5.2 Results

Figure 4.14 shows cross-sections of the two-dimensional pdf $\overline{\rho}_1(r, \phi)$ for $\phi = 0$. To achieve this we use the L = 25D simulation data of the previous section at a range of representitive Reynolds numbers. This data spans 10^4 time units where measurements of \hat{q}_1^j are taken 10 times per time unit, yielding a total of 10^5 samples.

These approximations to ρ are clearly noisy around the origin. In part this is due to the normalising technique employed above, since Z_i decreases linearly with *i* and around the origin, $1/Z_0 = 1/\Delta r \rightarrow \infty$ as $\Delta r \rightarrow 0$. More samples with zero Fourier amplitude are therefore needed in order to counter this. Unfortunately, for a fully turbulent flow (such as Re = 2,800), few samples possess this property as a proportion of the whole, and hence resolving the pdf correctly near zero is extremely expensive. However, the convergence properties away from the origin are relatively good with the small number of samples attained in this simulation.

At Re = 2,800, where the flow is uniformly turbulent, $\rho_1(r, 0)$ is sharply peaked and almost perfectly Gaussian, so that the most probable state of the system has no structure on the scale of the 25*D* pipe. This can be viewed as the disordered phase. As Re is reduced past Re_i, ρ_1 becomes non-Gaussian, with the peak reducing and the function spreading. This signifies the appearance of larger-scale structures; from examining the data and from figure 4.6, these take the form of puff-like structures forming within the pipe.

Finally, as Re reduces past 2,200, the peak clearly shifts to r > 0, signifying the presence of structures in the fluid on the scale of L; i.e. the appearence of equilibrium puffs. In



Figure 4.14: Above (a-c): Two-dimensional approximations of the complex pdf $\rho_1(r, \phi)$ of first-order Fourier modes. Below: Cross-sections of $\rho_1(r, \phi)$ at $\phi = 0$ for Re = 2,200 (bottom curve), 2,300, 2,400, 2,500 and 2,800 (top curve), as indicated by arrows and alternating line types.

addition, the most probable structure that can be observed is a puff, and the probability of finding uniform turbulence (r = 0) becomes 0. This can be seen as the ordered phase, whereby between Re_c and Re_i, the dynamics of the flow consist of a mixture of ordered and disordered phases. This demonstrates that the underlying dynamics form a reversible continuous transition between these transition points as Re is varied.

4.6 Conclusions

The main focus of this chapter has been the examination of the reverse transition process in pipe flow by use of large-scale direct numerical simulations in long, periodic pipes. Through the computational results presented here, we have demonstrated the existence of two transitions and critical points: the intensive-extensive transition which occurs at Re_c and the death of the intermittent regime occurring at $\text{Re}_i \approx 2,600$.

Key to the existence of these transitions are the highly fluctuating spatio-temporal dynamics highlighted in figure 4.6. Irregular turbulent patches within laminar pipe flow have been reported since the original experiments by Reynolds (1883b), and observations of puff splitting near transition go back many years (Rotta, 1956, Wygnanski *et al.*, 1975). These observations failed to capture the essential point, however, that in a well-defined range $\text{Re}_c \leq \text{Re} \leq \text{Re}_i$, intermittent turbulent-laminar flows are the intrinsic, asymptotic form of turbulence. Moreover, while near to Re_c irregularity takes the form of puff splittings and extinctions, intermittent states vary continuously with Re, and nearer to Re_i intermittency takes the form of laminar flashes in a turbulent background.

Taken in the context of what is known about other shear flows, especially plane Couette flow (Prigent *et al.*, 2002, Barkley & Tuckerman, 2005, Manneville, 2009), we see a generic and perhaps universal picture emerging for the route from turbulence to laminar flow in subcritical shear flows as Re is reduced to the smallest value that supports turbulence. At Re_{*i*} (whose value depends on the particular flow), uniform turbulence becomes unstable on long length scales and gives rise to an intermittent alternation of turbulent and laminar flow. At some lower value Re_{*c*}, this in turn gives rise to localised turbulence within a laminar background and such turbulence has a finite characteristic lifetime as shown in, for example, Avila *et al.* (2010). What remains to be investigated however is the precise value of the critical Reynolds number Re_c which denotes the transition between isolated turbulent puffs and more complex spatiotemporal intermittency. In the following chapter, we introduce the necessary statistical analysis used to determine the distributions of puff lifetimes, and using this run a series of numerical simulations to determine the nature of the splitting process. We then introduce the simple stochastic model of directed percolation which, as suggested by Pomeau (1986), shares many properties with pipe flow. By drawing upon the existing knowledge of this model, and through use of large quantites of experimental data, we show that puff splitting exhibits the same exponentially distributed behaviour seen in puff decay and calculate the characteristic splitting time. From this we are able to calculate the scaling of splitting lifetimes as a function of Re and thus determine the critical Reynolds number.

Chapter 5

Spreading turbulence in pipe flow

In the previous chapter, we considered the existence of a bifurcation point above which isolated turbulent regions (puffs) split and nucleate new pockets of turbulence in downstream laminar flow. Whilst the results were not conclusive, they indicate that such a bifurcation may exist, and furthermore, the spatial dynamics inherent to this process are extremely important in understanding the transition to turbulence. The purpose of this chapter is therefore to further investigate this process by drawing upon knowledge of other simple and universal models which possess similar properties and mechanisms.

Throughout this chapter we will consider turbulence as a *spreading* process, whereby areas of laminar flow are 'contaminated' by upstream turbulent pockets. Spreading processes are commonly found in many aspects of nature, such as in the transmission of diseases within populations or the dynamics of expanding fronts in forest fires. The analysis of both of these problems, whilst vastly complex in their own right, can be linked to other, far simpler statistical models such as coupled-map lattices and directed percolation (DP) where a control parameter dictates the balance between the rates of infection and recovery. When the rate of recovery is high, infection quickly dies out and in the infinite-time limit all sites become inactive. However, as the rate of infection increases and overwhelms recovery, the system undergoes a subcritical bifurcation and the limiting state consists of either fully active sites or an intermittent mixture of both active and inactive sites.

In order to consider pipe flow in a similar light, we will therefore examine the two processes which dictate the decay and spreading of turbulence. The first part of the chapter is dedicated to a summary of lifetime studies of the decay aspect of the problem. This is important both because only by knowing the decay timescale is it possible to know when growth overwhelms decay, and because similar methods will be used in spreading studies. Specifically we concentrate on puff decay, now a well-documented phenomenon, in which turbulent puffs re-laminarise following a memoryless process according to a timescale characterised by a mean lifetime. We first summarise the various studies which propose different scalings for this lifetime as a function of Re. We then introduce the statistical techniques and analysis of Avila *et al.* (2010), which strongly suggests that isolated puffs always have a non-zero probability of decay, and thus their mean lifetime diverges only as Re $\rightarrow \infty$.

The other process at work in pipe flow, as demonstrated in both figures 4.6 and 4.7 of the previous chapter, is that of puff *splitting*, in which active regions spread and give rise to spatio-temporal intermittency. The key difference between the approach taken in the previous chapter and the method adopted here is the use of the statistical analysis which is prevalent throughout decay studies. The remainder of this chapter is therefore devoted to applying these statistical methods to the splitting process and demonstrating its connections to directed percolation. By utilising data obtained both numerically and experimentally from a collaboration between Warwick (D. Moxey & D. Barkley) and Göttingen (K. Avila, A. de Lozar, M. Avila & B. Hof), and comparing the relavent timescales for splitting and decay, we determine the critical Reynolds number $\text{Re}_c = 2,040 \pm 10$. This critical point characterises a phase transition from states which are more likely to decay, to those which are more likely to spread. It is worth noting that the intensive-extensive transition described in the previous chapter does not therefore strictly exist, since every puff is not guaranteed to split under this new framework. However, it does offer a glimpse of the dynamics that one can typically expects as the Reynolds number is increased above Re_c .

5.1 Decay studies

Over the last decade, a topic of much interest in the hunt for the critical Reynolds number in a pipe has been the relaminarisation of turbulent flow below the suspected transition point. Faisst & Eckhardt (2004) performed the first rigorous study of this problem, by use of fully-spectral numerical simulations in a pipe of length 5*D*. In this work, simulations start with a flow field $u = u_{HP} + au_P$, where u_{HP} is the usual Hagen-Poiseuille flow. u_P is a perturbation generated by randomly exciting modes in a spectral expansion and *a* some amplitude lying between 0 and 1, so that after a short time the flow transitions to turbulence. A criterion is also defined based upon the total energy of the fluid, which is then used to determine when a given simulation relaminarises. The *lifetime* of a turbulent state is therefore the amount of time taken to relaminarise according to this criterion.

By performing an ensemble of 50 simulations, each of which with a different choice of a and u_P , Faisst & Eckhardt (2004) determined that the lifetimes of turbulent states are in fact distributed according to some probability distribution function; that is, the lifetime of a given turbulent state is dictated by a random variable and is therefore a *stochastic* process. As explained in chapter 4 the fact that random behaviour is observed in a deterministic system arises from the complexity of the phase space of the Navier-Stokes equations, whereby two similar but distinct initial conditions will behave in radically different ways.

At each Re therefore, one must construct probability distributions in order to determine the mean turbulent lifetime τ . The easiest and most intuitive approach is to calculate the survival function S(t) defined as

$$S(t) = \mathbb{P}($$
flow remains turbulent up to time t);

i.e. the function which describes the probability of an individual turbulent sample surviving up to time t.

Constructing S(t) is computationally very expensive, as we must obtain a large number of samples (and thus perform many expensive simulations) in order to guarantee convergence to the true distribution. We also wish for the correlation between each initial condition of each simulation to be very close to zero, so that the observed distribution is not biased. Faisst & Eckhardt (2004) determined that the survival function obeys an exponential distribution, so that

$$S(t) = \exp[-t/\tau(\text{Re})]$$

where τ (Re) denotes the average lifetime of a puff at a given Reynolds number. Puff decay is therefore a *memoryless* process, so that if the random variable *T* is distributed according

to the lifetime of turbulent puffs at some Re then for any s and t,

$$\mathbb{P}(T > t + s \mid T > t) = \mathbb{P}(T > s).$$

The essential issue here is determining the scaling of the lifetime τ as Re varies. By performing simulations at a variety of Re and obtaining a fit for the data, Faisst & Eckhardt (2004) determined that

$$\tau(\text{Re}) \propto \frac{1}{\text{Re} - \text{Re}_c}$$

where $\text{Re}_c \approx 2,250$, so that at Re_c turbulent lifetimes asymptotically diverge to infinity and turbulence becomes sustained.

Whilst this novel study demonstrated that pipe flow exhibits interesting statistical properties, there are several issues which required further investigation. From the previous chapter as well as numerous other studies, turbulence is known to take the form of puffs at the proposed critical Re. The length of the pipe used for these simulations is set at L = 5D to both allow for direct comparison with Eggels *et al.* (1994) and to reduce computational costs. This therefore precludes the possibility of observing puffs, which require a length of at least 25D to be accurately reproduced, the observed distribution of lifetimes may be biased. In addition, it is unclear whether both the type of perturbation used to trigger transition to turbulence and its amplitude affect S(t) and thus τ . Finally, the method used to determine τ from individual lifetime measurements requires further statistical analysis, since some states did not relaminarise before the end of simulations, thus biasing measurements of τ .

The experimental results of Peixinho & Mullin (2006) further validated the original work and found that τ observed a similar scaling, albeit at a much lower Re_c \approx 1,750. In this study, puffs were generated at a high Reynolds number using a perturbation at the pipe wall, and assuming the puffs survived for at least 100 diameters, Re was reduced to values below the critical point. However, the validity of this Re-reduction technique remained unjustified. Also, unlike the previous numerical work where the three-dimensional average kinetic energy was used as a criterion for decay, in this study puff decay was visually observed. However, more evidence supporting this scaling was presented by Willis & Kerswell (2007), who performed numerical simulations in a much longer L = 50D periodic pipe and determined that Re_c \approx 1,860. The conclusion of these studies conflicts with the experimental results of Hof *et al.* (2006) obtained using a gravity driven pipe, which suggested that $\tau^{-1} = \exp(a\text{Re} + b)$ for fixed constants *a*, *b*. Under this distribution, isolated puffs have finite lifetime at *any* given Re. Whilst this was contested, the greatly improved measurements of Hof *et al.* (2008) using many additional samples in a variety of pipes demonstrated that in fact a super-exponential scaling of the form $\tau = \exp[\exp(a\text{Re} + b)]$ provided a very good fit for the data. This fit was validated again by experimental observations of Kuik *et al.* (2010).

For the most part, the apparent disagreement between all of these studies – either in the distribution of lifetimes obtained, or the value of Re_c – can be attributed to statistical error from the relatively small number of samples observed. Numerically, the problem is extremely challenging due to the long domains required and the very large timescales that must be resolved. Indeed, most Re are inaccessible numerically due to the unfeasibly large amounts of computer time which would be required. In a similar fashion, experimental studies must use a pipe which is sufficiently long in order to resolve the timescales involved.

Finally, the techniques used to determine τ initially relied on extrapolating the median half-life of the data set. Statistically, a much better quantity to use is the maximum likelihood estimator (MLE), and from this we may readily obtain 95% confidence intervals which determine the accuracy of the estimate for τ given the observed data. In the following sections, we outline the relevant statistical knowledge needed to define the MLE and show how it is applied to decay studies.

5.1.1 Lifetime statistics

Lifetime problems are often encountered in a number of real-world applications such as manufacturing for quality control purposes or drug trials and medical studies. The statistical analysis of lifetimes is therefore a well-studied topic, and many techniques exist to obtain the best estimate of the mean lifetime and the accuracy of this estimate. In this section, we will briefly outline the relevant measurement techniques of Lawless (2003), which provides an extremely comprehensive reference outlining a wide variety of statistical methods which can be used to analyse lifetime problems.

We assume that the lifetime of individuals in a population can be described by a random

variable T which lies in $[0, \infty)$ and is distributed according to a probability density function f(t). Lifetime distributions are most obviously constructed via the survival function S(t).

Definition 5.1.1 (Survival function). The survival function f(t) is defined by

$$S(t) = \mathbb{P}(T \ge t) = \int_{t}^{\infty} f(x) \, \mathrm{d}x$$

In a general lifetime problem, it is often the case that by observing a finite number of lifetimes $\{t_i\}_{i=1}^n$ which are distributed according to f, we surmise that f belongs to a family of possible density functions $\{g(x; \theta) | \theta\}$. For example, in pipe flow each g represents an exponential distribution, and θ corresponds to τ . The problem therefore is to find:

- an approximation of the true value of θ from the given lifetime data, denoted by $\hat{\theta}$;
- how accurate this value is, given the sample size.

The best statistical approach for the first question is to apply the method of *maximum likelihood*, which provides an optimal estimate of the true value given any observed data.

Definition 5.1.2 (Maximum likelihood estimator). Given i.i.d. lifetimes $\{t_i\}_{i=1}^n$ distributed as $f(x; \theta)$, we define the *likelihood function* $J(\theta)$ as

$$J(\theta \mid t_1, \ldots, t_n) = \prod_{i=1}^n f(t_i; \theta).$$

If $\hat{\theta}$ maximizes J, then it signifies the most likely parameter and is known as the *maximum likelihood estimator*.

When the survival function is exponential so that $S(t; \theta) = \exp(-t/\theta)$ and $f(t; \theta) = \theta^{-1} \exp(-t/\theta)$, it can be shown that $\hat{\theta}$ is given by the usual sample mean

$$\hat{\theta} = \frac{1}{n} \sum_{i=1}^{n} t_i.$$

However, this analysis assumes that we are always able to observe the lifetimes of any given sample, which is often not the case. For example in pipe flow, puffs have lifetimes exceeding 10^6 time units for even moderately large Reynolds numbers which is often inaccessible through either experiment or simulation.

A natural way to avoid this problem is to censor the data by imposing some restrictions on the observation time. Lawless (2003) examines two common cases. Type 1 censoring imposes an upper limit on all observation times, so that $t_i \leq t_{max}$ for all *i*, in which case we define the number of uncensored lifetimes *r* to be the number of samples which have $t_i < t_{max}$ (i.e. those which do not exceed the upper limit on the observation time). In type 2 censoring, prior to the start of the simulation we determine the number of uncensored lifetimes *r* which will be observed, leaving n - r censored cases. Simulations are then performed until *r* lifetimes have been observed. Both methods have their merits; type 1 is often easier and more natural to implement, but is difficult to obtain estimates for the accuracy of the MLE. On the other hand, the accuracy of the MLE for type 2 may be derived exactly, but the total observation time remains unknown.

In either case, the MLE is obtained through a slight variation of the sample mean, with

$$\hat{\theta} = \frac{1}{r} \sum_{i=1}^{n} t_i.$$

The reliability of this MLE may be determined through the standard statistical method of confidence intervals. Given a confidence level $1 - \alpha$ with $0 < \alpha < 1$, we determine an interval [L, U] such that $\mathbb{P}(L < \hat{\theta} < U) = 1 - \alpha$; i.e. the range of values likely to contain $\hat{\theta}$ up to probability $1 - \alpha$. In the case of type 2 censoring, exact confidence intervals may be derived as

$$[L, U] = \hat{\theta} \times \left[\frac{2r}{\chi^2_{2r, \alpha/2}}, \frac{2r}{\chi^2_{2r, 1-\alpha/2}}\right]$$

where $\chi^2_{r,p}$ denotes the *p*-th quantile of the χ^2 distribution with *r* degrees of freedom; i.e. the value such that if $X \sim \chi^2_r$ then

$$\mathbb{P}(X \leq x) = p.$$

In the case of type 1 censoring however, exact intervals are impossible to determine and so must be estimated numerically. One approach is to appeal to the central limit theorem, in which case intervals may be estimated by $\hat{\theta} \times (1 \pm 0.96n^{-1/2})$. However, for small sample sizes this approximation is not accurate and is easily affected by non-symmetric distributions. Instead then, we use the relatively simple (and accurate) technique of *bootstrapping* as presented in DiCiccio & Efron (1996), amongst others.

- Generate a pseudorandom sample t₁^{*}, t₂^{*}, ..., t_n^{*} by sampling from the set of available lifetimes {t_i}_{i=1}ⁿ with replacement. For example when n = 3, a possible pseudorandom sample may be {t₁, t₁, t₂}.
- 2. Calculate the maximum likelihood estimator $\hat{\theta}_1^*$ of this data as

$$\hat{\theta}_1^* = \frac{1}{r_1^*} \sum_{k=1}^n t_k^*,$$

where r_1^* is the number of uncensored lifetimes.

- 3. Repeat the previous two steps *B* times to obtain bootstrap samples $\{\hat{\theta}_1^*, \hat{\theta}_2^*, \dots, \hat{\theta}_B^*\}$.
- 4. Sort this data so that $\hat{\theta}_b^* < \hat{\theta}_{b+1}^*$.

Assuming $B = 10^k$ where $k \ge 2$, confidence intervals at level $1 - \alpha$ may then be read off as $[\hat{\theta}^*_{B\cdot\alpha/2}, \hat{\theta}^*_{B\cdot(1-\alpha/2)}].$

5.1.2 Current state of lifetime statistics

The work presented in Avila *et al.* (2010) is the first study which applies the above statistical analysis on large sets of numerical and experimental data over a wide range of Re. The main aims of this paper are to describe how these techniques can be applied to pipe flow, and also to consolidate many existing studies in order to determine a scaling for τ . To accomplish the first task, we must incorporate the use of a *formation time* t_0 (Willis & Kerswell, 2007, Schneider & Eckhardt, 2008). The formation time represents the amount of time necessary for the flow to transition from a laminar state to a puff, through whatever experimental or numerical perturbation is applied to the flow. Equivalently, the value of t_0 dictates the time above which S(t) becomes exponential, and so we may therefore re-write the survival function as

$$S(t) = \exp[(t - t_0)/\tau].$$

 t_0 clearly depends non-trivially on Re and the perturbation technique, and thus is difficult to determine. In Avila *et al.* (2010), t_0 is found by examining the hazard function

$$h(t) = -\frac{\mathrm{d}}{\mathrm{d}t}\log(S(t))$$

which for the exponential distribution is the constant escape rate $\kappa = 1/\tau$. One approach to determine the formulation time is therefore to allow the escape rate to vary as a function of t_0 , so that $\kappa(t_0) = 1/\tau(t_0)$. For large enough values of t_0 , κ remains constant and thus determines the formation time. It should be noted however that this technique relies on large numbers of samples in order to be a reliable method.

The formation time must also be incorporated into the calculation of the MLE. For the exponential distribution, this is straightforward and gives the result

$$\hat{\tau} = \frac{1}{r} \left(\sum_{i=1}^{n} t_i - n t_0 \right).$$
 (5.1)

The formation time also allows us to validate that there is no statistical difference between initial conditions generated using a perturbation and those using an Re-reduction method. The conclusion from both Avila *et al.* (2010), and the results of de Lozar & Hof (2009), strongly indicate that technique used is unimportant, so long as it is sufficient enough to reliably generate puffs.

By applying these techniques to re-analyse the data of Hof *et al.* (2008) and Hof *et al.* (2006), as well as incorporating a substantial number of lifetimes obtained numerically, Avila *et al.* (2010) present very compelling evidence which strongly suggests that τ obeys the super-exponential scaling of Hof *et al.* (2008). Also included in this paper is a re-analysis of the numerical data of Willis & Kerswell (2007) which originally supported an asymptotic increase in τ at Re_c = 1,870. Using these new statisitical techniques, a super-exponential curve provides a much better fit to the data. All evidence then suggests that infinite lifetime turbulence cannot be attained by observing the decay process alone, and that another process is at work.

5.2 Puff splitting

The phenomenon of puff splitting is well-known and can be traced back to the studies of Reynolds (1883b) and Wygnanski & Champagne (1973). Figure 5.1 depicts a typical puff splitting in a L = 100D periodic pipe, where an isolated puff nucleates a new puff into downstream laminar flow. Whilst the detailed physical mechanism underlying splitting is



Figure 5.1: Spatio-temporal dynamics of puff splitting in a L = 100D domain at Re = 2,350. The turbulence intensity q(x,t) as defined in equation (4.1) is visualised in a frame of reference moving at speed $c = 0.914\overline{U}$. The white dashed line indicates the splitting time detected by the criterion of section 5.2.1.

at present unresolved, the splitting process has been observed in detail (Nishi *et al.*, 2008). Initially, a patch of vorticity escapes from the 'seed' puff into downstream laminar flow. Provided that the disturbance is large enough, and escapes sufficiently far downstream, it will evolve into a puff structure. This process, and the growth of the downstream puff, is highlighted in figure 5.3. Figure 5.2 demonstrates the space-time characteristics of the splitting process in an L = 200D simulation at Re = 2,350, and shows how repeated splits can lead to the spreading of turbulence throughout the entire domain as long as the flow is observed for a sufficient period of time.

Fundamentally, in much the same way as puff decay, splitting evolves according to a *stochastic* process, and therefore must be investigated through an ensemble of simulations. We must therefore determine the statistical distributions which underly the splitting process, and from this investigate the scaling of the mean lifetime of an isolated puff. Complicating this issue is that large timescales are inherent to the splitting process. By manually observing a small selection of samples, it is easy to ascertain that the probability of observing a split is much higher for larger Re than lower Re. Cases at very low Re close to 2,000, at which splitting events are extremely infrequent, cannot therefore be considered numerically due to the considerable computational costs involved. To overcome this problem, we shall later combine the numerical data with results obtained from a series of physical experiments.

5.2.1 Methodology

The general process for this series of numerical simulations is much the same as described in the previous section and Avila *et al.* (2010). We may define the survival function S(t) so that it represents the probability of a puff remaining isolated up until time t; i.e. it determines the probability that the puff does not split. These simulations therefore resolve the timescales which are inherent to the splitting process. In order to construct S(t), we must determine the method by which initial samples are generated and how splitting events are detected.

Simulations were performed using Semtex in a periodic pipe of length L = 50D; the minimum that can be used since we require that the pipe be capable of encompassing at least two puffs. The axial resolution remains fixed at $N_x = 768$ throughout, and the 36 element $K_{8,4}$ mesh similar to figure 4.1(c) is used at polynomial order P = 12 to correctly resolve the flow. All simulations impose periodic boundary conditions and use a constant mass flux







Figure 5.3: Snapshots of the streamwise vorticity ω_x through cross-sections of an L = 50D pipe at Re = 2,350 as a puff splits. Positive and negative vorticity are denoted by blue and red respectively; black areas represent $\omega_x = 0$. At t = 0 (bottom image), Re is changed from 2,100 to 2,350, and further snapshots are taken at t = 50, 138, 198 with the snapshot at t = 438 showing the final two-puff configuration.

of $\overline{U} = 1$ to drive the flow.

Prior to any splitting simulations, a 1,500 time unit data set was generated using an isolated puff at Re = 2,100 as an initial condition. The first 500 time units of this simulation provided sufficient time to confirm that the puff travelled at a relatively constant speed and did not abruptly relaminarise. After this initial period, snapshots of the flow field from the latter 1,000 time units were taken every 20 time units, providing 50 separate flow fields. Each of these form the initial condition for splitting simulations at some higher Re.

This Re-increase technique is similar to the Re-reduction technique employed in decay studies. Avila *et al.* (2010) determined that samples initialised with Re-reduced puffs are not statistically different from those generated using a perturbation-type approach, and only impact on the formation time t_0 . Whilst a similar investigation was not performed here, the later comparison of experimental results, obtained using puffs generated with a jet perturbation at the pipe wall, showed that the Re-increase technique is not statistically different. Primarily, as long as t_0 is properly determined for each distinct method, then the obtained distributions are statistically identical.

A criterion for determining the splitting time of any given sample is obviously very important in terms of generating correct statistics. In decay studies, one may simply observe some



Figure 5.4: Schematic drawings depicting criterion used to determine splitting times. Shaded rectangles indicate the region around the primary peak of the signal to be ignored, of width $d_{\min} = 20D$.

global quantity such as the total kinetic energy to determine the moment of decay. When decay occurs, the change in such quantities is abrupt and thus the uncertainty in determining the moment of decay is small. On the other hand, in order to observe a split we must determine a signal which is sensitive to the process, such that it differentiates a single puff from multiple puffs. One of the most convenient signals to process is the pressure field p, which for puffs has a distinctive shape and peaks around the area of maximal turbulence intensity. An added advantage of the pressure field is that it remains accessible experimentally at the pipe wall.

Let $p_k = p(x_k, 0, 0, t)$ denote the pressure field along the pipe axis at position $x_k = k\Delta x$. The criterion by which a splitting is determined is as follows:

• Roll data so that $\operatorname{argmax}_{0 \le k < N_x} p_k = 0$.

- Given a minimum interaction distance $d_{\min} = 20D$, determine the corresponding number of grid spaces $N_{\min} = \lceil d_{\min}/\Delta x \rceil$.
- If the condition

$$\max_{N_{\min} \leqslant k \leqslant N_x - N_{\min}} p_k \geqslant K \max_k p_k$$

is satisfied, then we say a split has occured at time t. K is a threshold parameter, in this case fixed at $K = \frac{2}{3}$.

Figure 5.4 visually depicts how this splitting criterion works. The top half of the figure shows a signal where the nucleation of a puff has started to occur from a 'parent' puff which is located using the maximum of the pressure signal. At this stage, the strength of the nucleated puff may not be sufficient in order to guarantee its survival. This is highlighted in figure 5.1, which shows the detected splitting time for a simulation in an L = 100D pipe. The splitting process begins approximately 100 time units before the observed time, with small patches of turbulence escaping downstream of the puff. These patches are not guaranteed to develop into puffs, and so the use of K and d_{min} allows us to distinguish small regions which quickly decay from those which survive and eventually form puffs.

The first half of the figure therefore demonstrates the purpose of d_{\min} which is chosen to correspond with the approximate length of a puff, so that we ignore small patches until they are sufficiently far downstream of the leading edge of the original puff. The second half of the figure shows how the nucleated puff may evolve, and depicts the moment in time that we classify a splitting. In this case, the puff has grown sufficiently in strength (determined by K) and established a large enough distance from the parent puff that it is highly likely to survive, and thus may be classified as a successful split.

The criterion was tested by both visually comparing splitting times to the space-time data of samples, and also by varying $\frac{1}{2} \leq K \leq \frac{9}{10}$, where virtually no change in the splitting time was detected. The additional advantage of using the pressure signal is that it is accessible experimentally at the pipe wall without disturbing the fluid, and thus allows for direct comparison between experimental and numerical data.



Figure 5.5: Initial condition number *i* against detected splitting time t_i for samples obtained when Re increased from 2,150 to 2,350.

5.2.2 Results

To overcome the problem of large timescales and to validate that numerical results are sufficiently accurate, a second series of simulations using the hybrid finite-difference–Fourier pseudo-spectral code of Willis & Kerswell (2009) was performed by Marc Avila. This code solves the Navier-Stokes equations in cylindrical co-ordinates (x, r, θ) , with 56 radial points, 64 azimuthal Fourier collocation points and 2048 axial Fourier collocation points to resolve the $L = 32\pi D \approx 100D$ domain. Periodic boundary conditions and constant mass flux were also imposed. Results from Semtex are denoted as DNS1, and from the alternative code by DNS2.

The aim of the numerical simulations is to encompass as large a range of Re as computationally possible. DNS1 simulations were therefore performed at 2,275 \leq Re \leq 2,375 using intervals of Δ Re = 25. Further simulations were performed using DNS2 for 2,250 \leq Re \leq 2,350 with Δ Re = 50 to widen the range of observed Reynolds numbers and provide a larger set of samples at Re = 2,300 and Re = 2,350 in order to provide better statistical averages.

Figure 5.5 demonstrates the stochastic nature of the splitting process. Each initial condition



Figure 5.6: Survival functions S(t) for Re = 2,300 (top) and Re = 2,350 (bottom). (*n*, *r*) = (90, 51) and (95, 90) samples are obtained respectively using Semtex (DNS1) and the code of Willis & Kerswell, 2009 (DNS2, courtesy of Marc Avila). Dashed line represents the curve $S(t) = \exp[-(t - t_0)/\hat{\tau}(\text{Re})]$, and shaded areas represent the 95% confidence interval associated with each simulation.

is numbered according to an index k corresponding to the snapshot taken at time 500 + 20k, and its detected splitting time t_k recorded. The lack of correlation between the splitting time of each sample further indicates that the Re-increase technique employed to generate initial conditions is valid.

The survival functions S(t) obtained through simulation data are visualised in figure 5.6

at Re = 2,300 and Re = 2,350. Samples obtained using both DNS codes are included to improve statistical accuracy, and the results from each are indicated by the marker colour. In both cases, a clear exponential fit is found and, like decay studies, S(t) may therefore be determined as

$$S(t) = \exp[-(t - t_0)/\hat{\tau}(\text{Re})],$$

where $\hat{\tau}$ is the MLE corresponding to the true mean lifetime τ . A maximum simulation time of $t_{\text{max}} = 1,500$ time units was imposed on each sample, so that the lifetime observations are subject to type 1 censoring and thus $\hat{\tau}$ may be determined through equation (5.1). The solid lines in each figure show this exponential fit. The shaded regions in each figure show the 95% confidence intervals as determined by the bootstrapping process of section 5.1.1. The vast majority of samples lie within these bounds, indicating the validity of the exponential fit. This exponentially decaying survival function is also found at every Re tested numerically.

In the decay studies of Avila *et al.* (2010), the hazard function was be used to determine a sufficiently large formation time t_0 . This approach depends on a large number of samples, and is not possible in this setting since higher Re demands higher domain resolution and thus more computational resources. Throughout these simulations then, we estimate $(t_0)_{true}$ by

$$t_0 = \max\{150, \min\{t_i\}\}.$$

The lower bound of 150 was determined from experimental studies, where again exact values for t_0 are unknown but are estimated as $50 \le t_0 \le 150$. It is therefore likely that this provides an over-estimate of $(t_0)_{true}$, especially at low Re for which splitting events are rare and thus min{ t_i } is large. However, this is much more beneficial than an under-estimate, since for exponential distributions a larger t_0 will not affect $\hat{\tau}$.

Table 5.1 summarises the results of all experimental data obtained through DNS1, including the number of samples *n* used at each Re, number of uncensored lifetimes *r* and the estimated confidence intervals $\pm \Delta \hat{\tau}$ for the MLE $\hat{\tau}$.

The results of figure 5.6 clearly demonstrate the stochastic nature of the splitting process (where again the word 'stochastic' refers to the sensitivity to initial conditions in the Navier-Stokes equations). In particular, the exponential probability distributions obtained after the initial formation time show that splitting is a memoryless process, and the probability for a

Re	τ	r	п	t_1	<i>t</i> _{max}	$+\Delta \tau$	$-\Delta \tau$
2275	1795	15	30	281	1500	1349	664
2300	1042	22	30	177	1500	550	344
2325	563	27	30	233	1500	247	184
2350	429	67	70	170	1500	110	92
2375	454	29	30	147	1500	153	116

Table 5.1: Data obtained from DNS1 of the characteristic splitting time τ as a function of Reynolds number, where each simulation observes *r* splittings from *n* initial samples. t_1 denotes the first observed splitting time, and t_{max} the maximum simulation time. Upper and lower confidence intervals are denoted by $\pm \Delta \tau$.

puff to split thus does not depend its prior evolution or the method which was used to generate it. The fact that these properties are also shared with the puff decay process indicates that these two processes are closely linked together. Indeed, in many simple stochastic models of thermodynamic systems, the critical value of the parameter which separates different phases may be classified by examining the relative timescales of two competing processes. In the next section, we will briefly outline the model of directed percolation, which may be used to represent simple epidemiological problems and in which the rates of infection and recovery determine the final evolution of the system. We then return to pipe flow, and by use of the numerical data and additional experimental evidence, obtain the scaling of τ as a function of Re. Drawing upon inspiration from directed percolation, the timescales of decay and splitting are then used to determine the critical Reynolds number Re_c.

5.3 Directed percolation

The scientific examination of diseases and how they contaminate a population has been a major driving force in determining new ways of modelling real-world problems. The simplest of these models partition the population into healthy and infected individuals. Initially only a small group of individuals will be infected. As time progresses the disease spreads to others whilst infected individuals eventually recover, although perhaps still remaining susceptible to further infection. The rates at which these different processes occur dictate whether the disease will perpetuate indefinitely or eventually reach the absorbing state where it dies out and no further infection is possible. Any model for this classic epidemiological problem

must therefore incorporate this behaviour, so that as the infection rate increases the system undergoes a phase transition.

Directed percolation (DP) is perhaps the simplest class of models which exhibits this bifurcation. The name arises from the study of fluid flow through a porous media along a given direction – for instance, groundwater flow may be driven vertically downwards by gravity. In DP, the permeability of the substrate determines whether the flow fully percolates through it or if the material is ultimately impermeable. However, DP forms a universality class, which encompasses simple epidemiological models such as the one described above. There are many models for directed percolation; the simplest are the standard contact process (Harris, 1974) and directed bond percolation (Broadbent & Hammersley, 1957). In this section, we shall investigate the latter process following the guidance of Hinrichsen (2000), and demonstrate that it shares many similarities with pipe flow.

5.3.1 Bond percolation

In directed bond percolation, the population or substrate is represented by a lattice. Each site on the lattice may be either active (infected) or inactive (healthy). In general, the lattice is (n+1)-dimensional; that is, *n* spatial dimensions represent the population and a single dimension describes the percolation direction. Figure 5.7 shows the simplest example of bond percolation in (1+1) dimensions, where the process is therefore represented by a rotated two-dimensional square lattice. In this figure, the direction of percolation is represented as the evolution of time, where spatial locations are updated according to a difference relation.

Each site on the lattice is enumerated by s_i^t , where $i \in \{1, ..., N\}$ and $t \in \{0, 1, 2, ...\}$ represent the spatial and temporal co-ordinates respectively. Any given simulation must start with an initial configuration $s^0 = \{s_1^0, ..., s_N^0\}$. Edges of the lattice represent bonds which connect sites together. Each bond may be either open or closed, and fluid (or infection) is only allowed to flow through open bonds.

At any given time t, the dynamics evolve according to a simple stochastic rule. For each pair of bonds connecting the points $(i \pm 1, t)$ to (i, t + 1) we generate two uniformly distributed random numbers $z_i^{\pm} \in [0, 1]$. Then a bond is open if $z_i^{\pm} < p$, where p is the control parameter dictating the rate of infection; otherwise it is deemed to be closed. Sites



Figure 5.7: Directed bond percolation on a (1+1)-dimensional square lattice utilising periodic boundary conditions, where the percolation direction (time) is vertically upwards. Active and inactive sites are denoted by black and white dots respectively. Open bonds are represented by solid lines, and closed bonds by dashed lines. Numbers along the bottom correspond to spatial locations *i*.

remain active so long as they possess at least one open bond, and therefore dynamics evolve according the difference equation

$$s_i^{t+1} = \begin{cases} 1, & s_{i-1}^t = 1 \text{ and } z_i^- < p, \\ 1, & s_{i+1}^t = 1 \text{ and } z_i^+ < p, \\ 0, & \text{otherwise.} \end{cases}$$

At each end of the lattice where $i \in \{1, N\}$ if t is even we must apply a boundary condition. In the examples considered here, a periodic boundary condition is used so that $s_1^{2t} = s_N^{2t}$.

Figure 5.8 shows space-time visualisations for three representative values of p. In each case, simulations start from a random initial configuration (bottom) or an isolated active site (top). In general, directed percolation systems undergo a continuous phase transition as the parameter p is altered. Below some critical value p_c – part (a) of the figure – the absorbing state, where all sites remain inactive and there is zero probability of returning from, is always reached. Around p_c (b), intermittent dynamics are prevalent, where activity is widespread and persistent but grows slowly. Sites alter sporadically between active and inactive states, and where it occurs active sites are restricted to narrow bands. When $p > p_c$ as in (c), even a single active site causes infection to break out and quickly contaminate the entire domain.


Figure 5.8: Examples of critical behaviour in directed bond percolation using N = 500 sites over t = 2,000 timesteps. Figures along the top row start with a configuration possessing a single active site, and those below use a random initial configuration.

It is interesting to note that even in this simple two-dimensional geometry, an exact value for p_c is yet to be analytically derived; Essam *et al.* (1986) estimate $p_c \approx 0.644701$.

5.3.2 Universality and connections to pipe flow

Universality is an important concept in dynamical systems and statistical mechanics, whereby properties of sometimes vastly differing systems can be linked by examining the order of phase transitions occurring within the systems. Directed percolation is a particularly interesting example of a universality class which encompasses many problems in physics, chemistry and statistics. Whilst there have been many suggestions that turbulent transition in shear flows may be connected to the DP universality class, the spatial-temporal dynamics of pipe flow in particular are not well studied. This can be explained by both the difficulty of observing space-time data experimentally due to the strong advection of the flow, and the computational requirements needed to perform large-scale pipe flow simulations.

The idea that fluid dynamics may undergo stochastic phase transitions dates back to the work of Pomeau (1986), who suggested that directed percolation may share certain characteristics with fluid dynamics in a general fashion. Further studies of Couette flow presented in Bottin & Chaté (1998) and Bottin *et al.* (1998) revealed that intermittent dynamics have various ties to directed percolation and other similar systems such as coupled map lattices. More recent work by Manneville (2009) also suggests that Couette flow has very close ties with the DP universality class, and the very recent DP simulations of Sipos & Goldenfeld (2011) in a (3+1)-dimensional pipe-shaped domain may indicate that pipe flow observes scalings relevant to DP.

From a completely qualitative perspective, some of the behaviour seen in figure 5.8 can be readily reproduced in pipe flow. Figure 5.9 demonstrates some representative dynamics in pipe flow by use of a long L = 200D pipe, which resolves the pipe using $N_x = 3,072$ collocation points and the $K_{8,4}$ mesh similar to figure 4.1(c) at polynomial order P = 12. As in figure 4.6, the quantity $q(x,t) = \sqrt{v^2 + w^2}$ where (v, w) represent the transverse velocity components is visualised in a reference frame moving at the bulk velocity \overline{U} . The initial condition for each simulation is an isolated puff state obtained from a simulation at Re = 2,000, which is instantaneously changed to 1,700 (bottom) and 2,700 (top).

The bottom half of the figure corresponds to the case where $p < p_c$ and active sites – i.e. puffs – eventually decay to background laminar flow. The top half of the figure demonstrates how turbulence rapidly spreads throughout the domain at high Re and shows the link between pipe flow and DP when $p > p_c$. In this regime the structure seen is known as a *slug*, depicted diagramatically in figure 4.5. The evolution of slugs is radically different to that of puffs, and is not something which will be further considered here. However, recent work by Duguet *et al.* (2010) suggests that the underlying mechanism behind this process may involve vortex shedding near the pipe wall via a Kelvin-Helmholtz mechanism.

Both of these figures clearly correspond visually to the dynamics of DP seen in figure 5.8, and with other studies suggesting that fluid flow in general has close ties with DP, it is apparent that the spatio-temporal intermittency in the flow is crucial to the nature of turbulence. In particular, the decay studies seen earlier only describe half of the true dynamics of the system where $p < p_c$ or equivalently Re $< \text{Re}_c$. In order to determine the onset of turbulence, we must therefore contrast the characteristic timescales of decay to those of splitting. In the next



Figure 5.9: Space-time figures demonstrating a relationship between directed percolation and pipe flow. In each case, the initial condition is an isolated puff, and Re is instantaneously set to 2,700 (top) and 1,700 (bottom).

section, we describe how the numerical results of the previous section can be incorporated into experimental data in order to find a close approximation to Re_c .

5.4 Critical point

The results of the previous two sections clearly demonstrate the stochastic and memoryless nature of the splitting process beyond the formation time, and thus the scaling of the mean lifetime τ of an isolated puff must be calculated for each process. The numerical results of section 5.2 are confined to a narrow range of Re due to the large computational requirements at low Re, and thus to determine an accurate scaling we must obtain data points as close to the suspected critical Reynolds number as possible. Since the data from table 5.1 clearly demonstrates the quick growth of τ as Re decreases, it is not computationally feasible to obtain these numerically.

To this end, as part of the Warwick-Göttingen collaboration we combine the numerical data obtained from DNS1 and DNS2 with the experimental data of Kerstin Avila, Alberto de Lozar and Björn Hof in a very long pipe of length L = 3,750D. The setup of the experiment is detailed thoroughly in de Lozar & Hof (2009) and Avila *et al.* (2011), and allows for up to 60,000 samples to be obtained at a single Re if required. Splitting events are detected using the pressure signal at the pipe wall, and thus the numerical results presented here therefore provide an immediate validation for the experimental results and vice versa.

In figure 5.10 (which is adapted from Avila *et al.*, 2011), we visualise the characteristic puff lifetime τ as a function of Re for both decaying (left) and spreading (right) turbulence, using all of the available experimental (coloured points) and numerical data (solid black triangles). Decay lifetimes are presented using the experimental and numerical data from Hof *et al.* (2008), Kuik *et al.* (2010) and Avila *et al.* (2010).

For both decay and splitting processes, super-exponential curves of the form

$$\tau(\text{Re}) = \exp[\exp(a\text{Re} + b)]$$

provide a very good fit to the data, where for decay we obtain (a, b) = (0.005556, 8.499)and for splitting (a, b) = (-0.003115, 9.161) using standard regression techniques. The numerical data, whilst limited in scope to a small range of Re, provides excellent agreement with the experimental data, with the super-exponential fit lying within the 95% confidence intervals at each of the measured data points.

One important consequence of the data presented here is that there is no particular Re beyond which spatial proliferation abruptly sets in, as suggested in the results of the previous chapter and other studies (e.g. Wygnanski & Champagne, 1973, Nishi *et al.*, 2008). Indeed, like decay, it appears that the tendency to split is intrinsic to all puffs, regardless of the Reynolds number. Therefore 'equilibrium' puffs, described by Wygnanski *et al.* (1975) as those which do not grow in length and advect at a constant speed, do not exist: after a time governed by a random variable, any localised puff will either decay or split.

This point is especially emphasised in the experimental samples, for which much larger data sets can be obtained and thus far lower Re examined. Here, the measured samples extend as low as Re = 2,032. At this Reynolds number, splits are extremely infrequent, but still occur: out of 57,823 initial conditions only 7 were observed, so that $\tau \approx 2.7 \times 10^7$. Further details of the experimental data can be found in the supplementary materials of Avila *et al.* (2011).

The intersection of the super-exponential fits at $\text{Re} \approx 2,040$ therefore marks a tipping point between the two competing processes of splitting and decay. As Re increases, pipe flow undergoes a statistical phase transition so that the probability of splitting outweigh that of decay.

A typical trait of many systems which undergo stochastic phase transitions is that the tipping point between competing processes often under- or over-estimates the true critical point of the system. A classical example of this is in the standard contact process, in which the infection rate must outweight the rate of recovery by a factor of just over three (see Hinrichsen, 2000, Harris, 1974) before the absorbing state yields to intermittent or active states which are fully active throughout the domain.

However, for pipe flow, the super-exponential scalings of τ mean that the intersection point must be almost indistinguishable from the critical point. At Re = 2,050, 10 above the intersection point, the splitting rate already outweighs the decay rate by a factor of four. Indeed, reduced models of pipe flow which exhibit the same stochastic properties seen here indicate that the actual critical point is within around 0.3% of the value stated here (Barkley, 2011). Therefore we may surmise that the actual critical point lies within the



Figure 5.10: Characteristic lifetime τ as a function of Reynolds number for an isolated puff. Coloured points represent experimental data (courtesy of Kerstin Avila, Alberto de Lozar and Björn Hof), and solid black triangles those obtained from Semtex (DNS1) and a hybrid spectral finite-difference code (DNS2). The dashed line denotes the super-exponentially increasing fit for mean decay times, and the solid line the decreasing fit for mean splitting times. The crossover point at Re_c = 2,040 ± 10, determines the transition between transient and sustained turbulence in pipe flow.

interval $\text{Re}_c = 2,040 \pm 10$, and this marks the onset of sustained turbulence in pipe flow.

One further point requiring clarification is the definition of 'sustained turbulence'. In the original experiments of Reynolds (1883a), b Re_c is defined to be the number above which turbulence persists indefinitely, and the results presented here clearly indicate an exponentially small probability of decay above Re_c. However in this context, a more natural definition of 'sustained turbulence' is instead to consider the thermodynamic limit (i.e. an average over space, time and ensemble of initial conditions, where space and time tend to infinity) of the turbulence fraction or intermittency factor γ . Below Re_c, $\gamma = 0$ as decay is the more likely scenario; above Re_c we obtain sustained turbulence where $\gamma > 0$.

5.5 Conclusion and Discussion

In this chapter, we have further examined the interesting and complex spatio-temporal dynamics of pipe flow found around transition. In contrast with the classical view of Landau & Lifshitz (1959) and Ruelle & Takens (1971) which postulate that temporal aspects of the flow determine the onset of turbulence, we have shown that pipe flow exhibits many properties native to simple stochastic systems, such as models of directed percolation. A common trait of these systems is that the critical point is determined by the balance of rates of spreading and decay of active sites, and therefore complex spatio-temporal coupling plays a fundamental role in understanding how the dynamics alter throughout the domain. By drawing upon the experimental and numerical evidence obtained through lifetime studies of puffs, and taking inspiration from these stochastic models, we have ascertained that the spatial aspects puff splitting process prove to be the decisive factor in identifying the critical Reynolds number, which is recorded as $Re_c = 2,040 \pm 10$.

One aspect of the comparison to directed percolation which has not been addressed here is the measurement of scaling laws which characterise DP. Models belong to the DP universality class if they possess the same critical exponent β and both spatial and temporal length scales ξ_{\perp} and ξ_{\parallel} , which depend on the spatial dimension *d*. These scalings have not been measured for pipe flow due to the large time and length scales which must be resolved, both of which require significant computational effort when combined with the number of simulations that must be performed to ensure accurate averaging. A possible line of future research is therefore to calculate these scalings and place pipe flow firmly in the DP universality class.

Whilst the question of transition in a pipe is of obvious interest to the academic community, the complexity of the dynamics found within the transition is mirrored in many other shear flows — particularly Couette flow and various types of channel and duct flows. The intermittently alternating laminar and turbulent regions encountered in pipe flow just above criticality are intrinsic to the problem, and such bands are known to exist in these flows, as demonstrated in Tuckerman & Barkley (2011) for Couette flow. Since the technique adopted here, in which decay and spreading processes are considered separately, does not rely on knowledge of how the relevant mechanisms work, this approach should provide a method which can be adopted to determine the critical Reynolds number in other geometries.

Additionally, there are many other interesting, important and unanswered questions that can be posed regarding puffs and slugs in pipe flow and the role they play in transition. From the perspective of dynamical systems, it is not immediately clear why solutions involving localised pockets of turbulence should exist at all. The incredibly challenging nature of this question inevitably leads to a very broad scope of possible solutions. Possible avenues of interest may involve further examination of the travelling wave solutions and edge states (Faisst & Eckhardt, 2003, Wedin & Kerswell, 2004, Pringle & Kerswell, 2007), examining pipe flow in a setting which reduces the number of degrees of freedom of the system (Willis & Kerswell, 2009), or deriving models which mimic many of the essential properties of pipe flow (Barkley, 2011).

In physical terms, the mechanism which governs puff decay and splitting is of great practical and theoretical significance. In particular, examining splitting will inevitably prove useful in terms of understanding the nature of fluid turbulence, and being able to delay or encourage the onset of turbulent flow. From a computational perspective however, the results of this chapter and other studies (Avila *et al.*, 2010, Hof *et al.*, 2010) demonstrate that large-scale fluid problems are now well within reach of numerical methods. As such, fields and quantities which are difficult to obtain experimentally (e.g. vorticity) can be captured throughout the entirety of the domain and thus influence the direction of further studies.

Chapter 6

Summary

Starting with the work of Pomeau (1986), many studies over the last thirty years have emphasised how the transition to turbulence, and understanding the turbulent attractor generally, hinges on considering not only the temporal aspects of the flow, but simultaneously on both the spatial and the temporal aspects. The work in this thesis further emphasises this point, and shows that the spatio-temporal dynamics in pipe flow are fundamental to determining the onset of turbulence and understanding the transition process.

The study of the reverse transition in chapter 4 demonstrates the natural form of transitional dynamics in a pipe, and examines the two critical points key to the transition process. Above $\text{Re}_i \approx 2,600$, flow remains uniformly turbulent. As the Reynolds number is reduced below this value, the flow undergoes a transition to intermittency, which initially appears as small pockets of laminar flow residing in a turbulent background. As Re is reduced further, intermittency in pipe flow takes the form of unsteady laminar-turbulent bands which constantly decay and split. Finally, as the Reynolds number is reduced below Re_c , the decay process dictates the dynamics of the flow, causing the bands to dissipate and the flow to relaminarise. The appearance of these laminar regions in otherwise uniformly turbulent fluid means that many laws of fluid dynamics, such as the 5/3 power law for the Kolmogorov spectrum or the law of the wall do not hold in this setting, and therefore non-standard techniques are required to investigate the transition process.

To this end, chapter 5 reveals that the key to determining Re_c , and thus the onset of sustained turbulence, is to examine the process by which localised pockets of turbulence split. The

fundamental observation is that splitting, like decay, is a memoryless stochastic process so that the lifetimes of isolated puffs are governed by some characteristic time τ which scales with Re. Indeed, we further validate the observations of Pomeau (1986) by demonstrating that pipe flow exhibits the same properties found in simple stochastically-driven systems such as directed percolation. The critical Reynolds number does not therefore control the onset of turbulence in a fluid in the same way as, for example, a parameter for a system of ODEs which dictates that a particular solution remain stable above and unstable below some critical value. Instead, it denotes a phase transition so that for Re < Re_c turbulence is more likely to decay and for Re > Re_c turbulence is more likely to spread. Thus by comparing the relevant timescales for decay and splitting, we obtain an estimate for the onset of turbulence in pipe flow of Re_c = 2,040 ± 10. Since the analysis we present here is independent of the mechanisms which govern the splitting process, the approach taken in this chapter should be equally applicable to determining Re_c in other flow geometries.

The conclusions of these chapters highlight the important role that numerical methods are now taking in understanding hydrodynamic stability problems. Fully-resolved simulations of pipe flow in particular have only proven in reach of computational power within the last twenty years (and of puffs and slugs within the last ten), but have already demonstrated their worth. Indeed, even a highly complex algorithm such as the spectral element method of chapter 2 has proven to be a reliable and highly accurate numerical method capable of resolving both large timescales and spatial domains. When coupled with an appropriate joint discretisation and the optimisations which exploit the underlying architecture of the computing facilities, such as the parallel transpose algorithm of chapter 3, it becomes an efficient method to study a wide variety of fluid problems in a variety of complex domains.

Potentially, the most important topic to investigate from this work is understanding the mechanism governing the splitting process, and determining why the process is stochastic. The study of both of these topics may lead to control mechanisms for the onset of turbulence and thus have direct application in many engineering problems. Another important aspect of the transition process, which has only been briefly touched upon here in chapter 4, are the more complicated states found in the middle of the transitional regime when Re $\gtrsim 2,450$. Here the dynamics are characterised by turbulent regions which merge and annihilate in quick succession. Determining how these states evolve will prove key in further understanding the nature of turbulence and the transition process, and sets a target for future studies.

Bibliography

- K. Avila, D. Moxey, A. de Lozar, M. Avila, D. Barkley and B. Hof (2011). *The onset of sustained turbulence in pipe flow*. Science **333** (6039), 192–196.
- M. Avila, A. P. Willis and B. Hof (2010). *On the transient nature of localized pipe flow turbulence*. J. Fluid Mech. **646**, 127–136.
- P. R. Bandyopadhyay (1986). *Aspects of the equilibrium puff in transitional pipe flow*. J. Fluid Mech. **163**, 439–458.
- D. Barkley (2011). Simplifying the complexity of pipe flow. Phys. Rev. E 84, 016309.
- D. Barkley and L. S. Tuckerman (2005). *Computational study of turbulent laminar patterns in Couette flow*. Phys. Rev. Lett. **94** (1), 014502.
- D. Barkley and L. S. Tuckerman (2007). *Mean flow of turbulent-laminar patterns in plane Couette flow*. J. Fluid Mech. **576**, 109–137.
- H. M. Blackburn and S. J. Sherwin (2004). Formulation of a Galerkin spectral element-Fourier method for three-dimensional incompressible flows in cylindrical geometries. J. Comput. Phys. 197 (2), 759–778.
- H. Blasius (1913). Das Änlichkeitsgesetz bei Reibungsvorgängen in Flüssigkeiten. Forsch. Arb. Ing 134.
- S. Bottin and H. Chaté (1998). *Statistical analysis of the transition to turbulence in plane Couette flow*. Eur. J. Mech. B-Fluids 6 (1), 143–155.
- S. Bottin, F. Daviaud, P. Manneville and O. Dauchot (1998). *Discontinuous transition to spatiotemporal intermittency in plane Couette flow*. Europhys. Lett. **43** (2), 171–176.

- S. R. Broadbent and J. M. Hammersley (1957). *Percolation processes*. In: Math. Proc. Cambridge. Vol. 53. Cambridge Univ. Press, pp. 629–641.
- C. D. Cantwell (2009). *Transient growth of separated flows*. Ph.D. thesis, University of Warwick.
- C. Canuto, M. Y. Hussaini, A. Quarteroni and T. A. Zang (1988). *Spectral methods in fluid dynamics*. Springer-Verlag.
- D. Chu, R. Henderson and G. E. Karniadakis (1992). Parallel spectral-element-Fourier simulation of turbulent flow over riblet-mounted surfaces. Theoret. Comput. Fluid Dynamics 3 (4), 219–229.
- H. J. Crowder and C. Dalton (1971). *On the stability of Poiseuille flow in a pipe*. J. Comput. Phys. **7** (1), 12–31.
- A. G. Darbyshire and T. Mullin (1995). *Transition to turbulence in constant-mass-flux pipe flow*. J. Fluid Mech. **289**, 83–114.
- A. de Lozar and B. Hof (2009). *An experimental study of the decay of turbulent puffs in pipe flow*. Philos. T. R. Soc. A. **367** (1888), 589–599.
- T. J. DiCiccio and B. Efron (1996). Bootstrap confidence intervals. Stat. Sci. 11 (3), 189–212.
- Y. Duguet, A. P. Willis and R. R. Kerswell (2010). *Slug genesis in cylindrical pipe flow*. J. Fluid Mech. **663**, 180–208.
- J. G. M. Eggels, F. Unger, M. H. Weiss, J. Westerweel, R. J. Adrian, R. Friedrich and F. T. M. Nieuwstadt (1994). Fully developed turbulent pipe flow: a comparison between direct numerical simulation and experiment. J. Fluid Mech. 268, 175–209.
- J. G. M. Eggels, J. Westerweel, F. T. M. Nieustadt and R. J. Adrian (1993). *Direct Numerical Simulation of Turbulent Pipe Flow*. Appl. Sci. Res. **51** (1), 319–324.
- J. W. Essam, K. De'Bell, J. Adler and F. M. Bhatti (1986). *Analysis of extended series for bond percolation on the directed square lattice*. Phys. Rev. B **33** (3), 1982–1986.
- H. Faisst and B. Eckhardt (2003). *Traveling waves in pipe flow*. Phys. Rev. Lett. **91** (22), 224502.

- H. Faisst and B. Eckhardt (2004). *Sensitive dependence on initial conditions in transition to turbulence in pipe flow.* J. Fluid Mech. **504**, 343–352.
- C. W. Gear (1971). *Numerical initial value problems in ordinary differential equations*. Prentice Hall.
- G. Grötzbach (1983). Spatial resolution requirements for direct numerical simulation of the Rayleigh-Bénard convection. J. Comput. Phys. **49** (2), 241–264.
- G. Hagen (1839). Ueber die Bewegung des Wassers in engen cylindrischen Röhren. Ann. Phys. (Berlin) 122 (3), 423–442.
- T. E. Harris (1974). Contact Interaction on a Lattice. Ann. Prob. 2 (6), 969–988.
- I. H. Herron (1991). *Observations on the role of vorticity in the stability theory of wall bounded flows.* Stud. Appl. Math. **85** (3), 269–286.
- H. Hinrichsen (2000). Non-equilibrium critical phenomena and phase transitions into absorbing states. Adv. Phys. **49**, 815–958.
- B. Hof, A. de Lozar, M. Avila, X. Tu and T. M. Schneider (2010). *Eliminating turbulence in spatially intermittent flows*. Science **327** (5972), 1491.
- B. Hof, A. de Lozar, D. J. Kuik and J. Westerweel (2008). *Repeller or Attractor? Selecting the Dynamical Model for the Onset of Turbulence in Pipe Flow*. Phys. Rev. Lett. **101** (21), 214501.
- B. Hof, A. Juel and T. Mullin (2003). Scaling of the turbulence transition threshold in a pipe.Phys. Rev. Lett. 91 (24), 244502.
- B. Hof, C. W. H. van Doorne, J. Westerweel, F. Nieuwstadt, H. Faisst, B. Eckhardt, H. Wedin,
 R. R. Kerswell and F. Waleffe (2004). *Experimental observation of nonlinear traveling* waves in turbulent pipe flow. Science **305** (5690), 1594.
- B. Hof, J. Westerweel, T. M. Schneider and B. Eckhardt (2006). *Finite lifetime of turbulence in shear flows*. Nature **443** (7107), 59–62.
- D. D. Joseph and S. Carmi (1969). *Stability of Poiseuille flow in pipes, annuli, and channels*.Quart. Appl. Math 26 (575), 21.

- G. E. Karniadakis, M. Israeli and S. A. Orszag (1991). *High-order splitting methods for the incompressible Navier-Stokes equations*. J. Comput. Phys. **97** (2), 414–443.
- G. E. Karniadakis and S. Sherwin (2005). *Spectral/hp Element Methods for Computational Fluid Dynamics*, 2nd Edition. OUP.
- R. R. Kerswell and O. R. Tutty (2007). *Recurrence of travelling waves in transitional pipe flow*. J. Fluid Mech. **584**, 69–102.
- T. K. Kim (2002). A Modified Smagorinsky Subgrid Scale Model for the Large Eddy Simulation of Turbulent Flow. arXiv:math/0209377.
- M. A. Kopera (2011). *Direct Numerical Simulation of Turbulent Flow over a Backward-Facing Step.* Ph.D. thesis, University of Warwick.
- D. J. Kuik, C. Poelma and J. Westerweel (2010). *Quantitative measurement of the lifetime of localized turbulence in pipe flow*. J. Fluid Mech. **645**, 529–539.
- L. D. Landau and E. Lifshitz (1959). Fluid Mechanics. Vol. 225. Oxford: Pergamon.
- J. F. Lawless (2003). *Statistical models and methods for lifetime data*, 2nd Edition. Wiley Series in Probability and Statistics.
- P. D. Lax and A. N. Milgram (1954). *Parabolic equations*. In: Contributions to the theory of partial differential equations. Vol. 33. Princeton University Press, pp. 167–190.
- T. S. Lund, X. Wu and K. D. Squires (1998). Generation of Turbulent Inflow Data for Spatially-Developing Boundary Layer Simulations. J. Comput. Phys. 140 (2), 233–258.
- P. Manneville (2009). *Spatiotemporal perspective on the decay of turbulence in wall-bounded flows*. Phys. Rev. E **79** (2), 025301.
- P. S. Marcus (1984). Simulation of Taylor-Couette flow. Part 1. Numerical methods and comparison with experiment. J. Fluid Mech. 146, 45–64.
- D. M. McIver, H. M. Blackburn and G. J. Nathan (2000). Spectral element-Fourier methods applied to turbulent pipe flow. ANZIAM J. 42 (E), C954–C977.
- B. J. McKeon, C. J. Swanson, M. V. Zagarola, R. J. Donnelly and A. J. Smits (2004). *Friction factors for smooth pipe flow*. J. Fluid Mech. **511**, 41–44.

- B. J. McKeon, M. V. Zagarola and A. J. Smits (2005). *A new friction factor relationship for fully developed pipe flow*. J. Fluid Mech. **538**, 429–443.
- A. Meseguer and L. N. Trefethen (2003). *Linearized pipe flow to Reynolds number* 10⁷. J.
 Comput. Phys. 186 (1), 178–197.
- D. Moxey and D. Barkley (2010). Distinct large-scale turbulent-laminar states in transitional pipe flow. Proc. Nat. Acad. Sci. 107 (18), 8091–8096.
- R. Narasimha and K. R. Sreenivasan (1979). *Relaminarization of fluid flows*. Adv. Appl. Mech. 19, 221–309.
- M. Nishi, B. Ünsal, F. Durst and G. Biswas (2008). *Laminar-to-turbulent transition of pipe flows through puffs and slugs*. J. Fluid Mech. **614**, 425–446.
- P. Orlandi and F. Fatica (1997). *Direct simulations of turbulent flow in a pipe rotating about its axis.* J. Fluid Mech. **343**, 43–72.
- S. A. Orszag, M. Israeli and M. O. Deville (1986). *Boundary conditions for incompressible flows*. J. Sci. Comp. **1** (1), 75–111.
- J. Peixinho and T. Mullin (2006). *Decay of turbulence in pipe flow*. Phys. Rev. Lett. **96** (9), 094501.
- J. Peixinho and T. Mullin (2007). *Finite-amplitude thresholds for transition in pipe flow*. J. Fluid Mech. **582**, 169–178.
- W. Pfenniger (1961). *Transition in the inlet length of tubes at high Reynolds numbers*. In: G.V. Lachman (Ed.), Boundary layer and flow control. pp. 970–980.
- J. L. M. Poiseuille (1840). *Recherches expérimentales sur le mouvement des liquides dans les tubes de très-petits diamètres*. Séanc. Acad. Sci. Paris **11**, 961–967.
- Y. Pomeau (1986). *Front motion, metastability and subcritical bifurcations in hydrodynamics.* Physica D **23**, 3–11.
- S. B. Pope (2000). Turbulent Flows, 1st Edition. CUP.
- A. Prigent, G. Grégoire, H. Chaté and O. Dauchot (2003). Long-wavelength modulation of turbulent shear flows. Physica D 174 (1-4), 100–113.

- A. Prigent, G. Grégoire, H. Chaté, O. Dauchot and W. van Saarloos (2002). Large-scale finite-wavelength modulation within turbulent shear flows. Phys. Rev. Lett. 89 (1), 014501.
- C. C. T. Pringle, Y. Duguet and R. R. Kerswell (2009). *Highly symmetric travelling waves in pipe flow*. Philos. T. R. Soc. A. **367** (1888), 457.
- C. C. T. Pringle and R. R. Kerswell (2007). *Asymmetric, Helical, and Mirror-Symmetric Traveling Waves in Pipe Flow.* Phys. Rev. Lett. **99** (7), 074502.
- O. Reynolds (1883a). An Experimental Investigation of the Circumstances which Determine whether the Motion of Water shall be Direct or Sinuous, and of the Law of Resistance in Parallel Channels. Proc. R. Soc. **35**, 84–99.
- O. Reynolds (1883b). An Experimental Investigation of the Circumstances which Determine whether the Motion of Water shall be Direct or Sinuous, and of the Law of Resistance in Parallel Channels. Philos. T. R. Soc. A. **174**, 935–982.
- J. Rotta (1956). *Experimental contributions to the development of turbulent flow in a pipe*. Ing. Arch. **24**, 258.
- D. Ruelle and F. Takens (1971). *On the nature of turbulence*. Commun. Math. Phys. **20** (3), 167–192.
- D. Samanta, A. de Lozar and B. Hof (2010). *Experimental investigation of laminar turbulent intermittency in pipe flow.* arXiv:1008.2294v1.
- T. M. Schneider and B. Eckhardt (2008). *Lifetime statistics in transitional pipe flow*. Phys. Rev. E **78** (4), 046310.
- T. Sexl (1927). Zur stabilitätsfrage der Poiseuilleschen und Couetteschen strömung. Ann. Phys. (Berlin) 388 (14), 835–848.
- H. Shan, B. Ma, Z. Zhang and F. T. M. Nieuwstadt (1999). *Direct numerical simulation of a puff and a slug in transitional cylindrical pipe flow.* J. Fluid Mech. **387**, 39–60.
- M. Sipos and N. Goldenfeld (2011). *Directed percolation describes lifetime and growth of turbulent puffs and slugs*. arXiv:1101.5835.
- S. Tavoularis and S. Corrsin (1981). *Experiments in nearly homogenous turbulent shear flow* with a uniform mean temperature gradient. Part 1. J. Fluid Mech. **104**, 311–347.

- R. Temam (1969). Sur l'approximation de la solution des equations de Navier-Stokes par la methode des fractionnaires II. Arch. Rational Mech. Anal. **33**, 377.
- R. Temam (1979). Navier-Stokes Equations: Theory and Numerical Analysis. AMS.
- C. Temperton (1992). A Generalized Prime Factor FFT Algorithm for any $N = 2^p 3^q 5^r$. SIAM Journal on Scientific and Statistical Computing **13**, 676.
- T. Tsukahara, Y. Seki, H. Kawamura and D. Tochio (2005). DNS of Turbulent Channel Flow at Very Low Reynolds Numbers. In: Proc. 4th Int. Symp. on Turbulence and Shear Flow Phenomena. pp. 935–940.
- L. S. Tuckerman and D. Barkley (2011). *Patterns and dynamics in transitional plane Couette flow*. Phys. Fluids **23**, 041301.
- L. S. Tuckerman, D. Barkley and O. Dauchot (2008). *Statistical analysis of the transition to turbulent-laminar banded patterns in plane Couette flow*. J. Phys. Conf. Ser. **137**, 012029.
- M. P. B. Veenman (2004). *Statistical analysis of turbulent pipe flow: A numerical approach*. Ph.D. thesis, Technische Universiteit Eindhoven.
- H. Wedin and R. R. Kerswell (2004). *Exact coherent structures in pipe flow: travelling wave solutions*. J. Fluid Mech. **508**, 333–371.
- A. P. Willis and R. R. Kerswell (2007). *Critical behavior in the relaminarization of localized turbulence in pipe flow*. Phys. Rev. Lett. **98** (1), 014501.
- A. P. Willis and R. R. Kerswell (2008). Coherent Structures in Localized and Global Pipe Turbulence. Phys. Rev. Lett. 100 (12), 124501.
- A. P. Willis and R. R. Kerswell (2009). *Turbulent dynamics of pipe flow captured in a reduced model: puff relaminarization and localized edge states*. J. Fluid Mech. **619**, 213–233.
- I. Wygnanski, M. Sokolov and D. Friedman (1975). *On transition in a pipe. Part 2. The equilibrium puff. J. Fluid Mech.* **69** (02), 283–304.
- I. J. Wygnanski and F. H. Champagne (1973). *On transition in a pipe. Part 1. The origin of puffs and slugs and the flow in a turbulent slug.* J. Fluid Mech. **59** (2), 281–335.
- N. N. Yanenko (1971). The method of fractional steps. Springer-Verlag.