

7	Final testing and data processing	100
7.1	Test names and flowrate variation	100
7.2	Alignment and optical considerations	103
7.3	Velocity profiles	107
7.3.1	Data pre-processing	107
7.3.2	Normalisation of the radial ordinate	108
7.3.3	Determination of the wall position	109
7.4	Timing measurements	110
7.4.1	Point of deviation	111
7.4.2	Method of measurement	114
7.4.3	Description of the phenomena	115
7.4.4	Prediction of the inception point	117
7.4.5	Variability	118
7.5	Flow meter correction and scaling of velocity profiles	121
7.5.1	Symmetry of the velocity profiles	121
7.5.2	Integration of the velocity profiles	122
7.5.3	Flowmeter calibration	124
7.5.4	Flowrate variation with time	126
7.5.5	Velocity profile normalisation	129
7.6	Derived quantities	129
7.6.1	Displacement thickness	129
7.6.2	Velocity profile shape parameter	132
7.6.3	Velocity profile curve fitting	132
7.7	Summary	134
8	Experimental results – velocity profiles	135

4.4.10	Pump & tank control	70
4.5	Overall design considerations	71
4.6	Summary	72
5	The acquisition & control software system	73
5.1	Basic design considerations	73
5.1.1	Object-orientation	74
5.1.2	User interface	75
5.2	The main control software	76
5.2.1	Command interface	76
5.2.2	The RUN command and the run file	77
5.2.3	Automatic data conversion	78
5.3	Slave software	79
5.4	Summary	81
6	Preliminary experimentation	82
6.1	Terminology	82
6.2	Flow control	84
6.2.1	Rig settings	84
6.2.2	Computer control	85
6.3	Preliminary testing	89
6.3.1	Data smoothing	90
6.3.2	Behaviour far downstream	90
6.3.3	Inlet behaviour	96
6.3.4	Entrance effects and the background environment	97
6.4	Summary	99

4	Experimental system	40
4.1	System overview	40
4.1.1	The flow rig	41
4.1.2	Velocity measurement components	41
4.1.3	Acquisition & control components	43
4.2	Rig sub-system	43
4.2.1	Header tank	44
4.2.2	Inlet contraction	46
4.2.3	Test section	49
4.2.4	Flowmeter & valve network	55
4.2.5	Catchment tank and temperature control	59
4.2.6	The pump	59
4.3	Velocity measurement sub-system	60
4.3.1	The main Laser Doppler Velocimeter	60
4.3.2	Remote laser head	61
4.3.3	Laser positioning	63
4.4	The acquisition & control hardware	64
4.4.1	LDV acquisition sub-system	66
4.4.2	Timing	66
4.4.3	Stepper-motor drive units	67
4.4.4	Linear actuator unit	67
4.4.5	Position control devices & limit switches	68
4.4.6	Measurement of small laser position	68
4.4.7	Analogue I/O system	69
4.4.8	Valve control	69
4.4.9	The flowmeter	70

2.2.3	Bypass mechanisms	14
2.3	Steady transition	13
2.3.1	The flat plate boundary layer	13
2.3.2	Pipe flow	17
2.4	Unsteady transition	19
2.4.1	Flows with steady flowrates	19
2.4.2	Non-periodic unsteady flows	20
II	EXPERIMENTATION	22
3	Laser Anemometry	24
3.1	Principles of operation	25
3.2	Aspects of the LDV method	27
3.2.1	Transmitting optics	27
3.2.2	Receiving optics	27
3.2.3	Seeding particles	28
3.2.4	Signal processing	29
3.2.5	Frequency shifting	30
3.2.6	Random time base	31
3.2.7	The multiple component system	31
3.3	Advantages of the LDV method	32
3.4	Data processing	33
3.4.1	Bias	33
3.4.2	Bias correction	35
3.4.3	The time-windowing technique	37
3.5	Summary	39

Contents

List of Figures	xlvii
List of Tables	lIII
Nomenclature	lv
I INTRODUCTION	1
1 Overview	2
1.1 Scope of the study	3
1.2 Objectives & proposed strategy	5
1.3 Structure and philosophy	6
1.3.1 Layout of the thesis	6
1.3.2 Presentation of raw data	8
1.3.3 Data processing and modelling tools	8
2 Literature overview	9
2.1 Scope and emphasis	9
2.2 General overview	10
2.2.1 Linear stability	10
2.2.2 Secondary instability and transition	11

To my now father-in-law, for his help in supplying me with powerful computing facilities for use in obtaining the arduous numerical solutions to the mathematical problems at hand.

To my family, for financial and moral support during hard times.

To my many friends, both here and abroad, whose favours and kindness were often not reciprocated due to my long preoccupation with this study.

Finally to my wife, who married me despite this PhD. Things can only get better!

Acknowledgements

The work presented here would certainly not have been possible without the support and encouragement of friends, family and colleagues. Below follows a by no means inclusive list of all the people to whom I owe much gratitude for either helping my endeavour, or putting up with my strange and antisocial behaviour during its coming together:

My supervisor Prof. E.A. Moss, who was instrumental in interest ing me in the field of fluid mechanics in the first place, and who managed always to encourage and guide me through difficult parts of the work. His pioneering work on the subject of pipe flow stability was invaluable for the production of this thesis.

Dr. ate da Silva, my precursor, who laid many of the foundations concerning the use of laser techniques in pipe flows, and was helpful in defining more precisely the realm within which my analysis occurred. Although some of his numerical predictions were in the light of current knowledge incorrect, the overwhelming thrust of his work was visionary: science progresses as much by refuting the wrong as by justifying the right.

To Prof. C. Wright, of the applied Mathematical department at the University. His gentle introduction to the spectral approach in general, and the operational philosophy in particular, deepened substantially my understanding of the general principle of numerical analysis. His insistence of the need to use only Matlab for the analysis quite possibly shortened the duration of the numerical work by a year, and certainly allowed me a far larger scope of analysis.

To my colleague G.H. Olivier, whose ideas, help and forbearance during the difficult and lengthy rig design were invaluable. The successful implementation of a tricky design was due in part to his intractable belief that it would in fact work.

The workshop staff, for their patience, and their help in the design of the many intricate components of the experimental apparatus.

To the staff of the School of Mechanical Engineering, and for the many invaluable discussions about various aspects of my work.

*Dedicated to my beloved wife
Caroline.*

temporally developing and linear, and the second spatially developing, steady and non-linear. These two sub-models predict velocity profiles possessing error function and Blasius shapes respectively, in the zero boundary layer limit. Non-parallel effects were formally shown to be small for the zero acceleration, steady entrance flow, and by inference for its accelerating counterparts.

The experimental trends confirmed the predictions by exhibiting an increase of minimum critical Reynolds number with exponential constant, and boundary layer thinning behaviour when viewed from the physical co-ordinate system. Furthermore, favourable quantitative comparisons were obtained between predicted and measured laminar flows on the basis of displacement thickness and selected velocity profiles. The experimental inception curve for each of the twelve tests was below its analytical linear stability counterpart, consistent with the presence of finite-amplitude disturbances. These disturbances increased with acceleration parameter from about 1.4% to 1.7%, corresponding to discrepancies between measured and predicted minimum critical Reynolds numbers which ranged from about 12% to 22%. The disturbance magnitudes were too low to initiate by-pass transition, although non-linear effects were certainly present at this disturbance level. On the basis of minimum critical Reynolds number, the current work shows a substantially closer correlation than previous pipe flow studies between linear stability predictions and measured deviations from the laminar state.

In conclusion the theoretical models developed in the present thesis describe all the important features of the measured system. Quantitative discrepancies between the two can be explained by finite amplitude effects whose origin was the experimental pipe inlet contraction.

Abstract

This thesis is an experimental, analytical and numerical investigation of pipe flows started impulsively from rest and thereafter accelerated exponentially with time. The broad objectives were to reconcile measured base flow and transition inception data with linear stability predictions.

A fully automated and computerised test facility was developed. The test section comprises a straight glass pipe of 48.2 mm diameter supplied with water via a smooth contraction from an axisymmetric pressure vessel. Flow rate control was achieved using a magnetic flow meter in conjunction with electro-pneumatic valves. Fluid velocities were measured using a laser Doppler velocimeter. Twelve different exponential variations were considered at a number of axial positions. All showed initial boundary layer growth, followed by a period of thinning in which deviations from laminar flow occurred within the boundary layer near the pipe inlet. The minimum critical Reynolds number in the pipe increased with increasing values of the exponential constant.

The unsteady, incompressible, axisymmetric Navier Stokes equations were scaled using instantaneous cross-sectional mean velocity, thus ensuring boundedness for the dependent variables. The resulting equations were used to derive (i) a high Reynolds number approximation to the primitive variable laminar flow system in which the pressure term was formally eliminated, and (ii) an unsteady extension of the Squire stability equation, involving no additional levels of assumption than those traditionally employed in the steady entrance flow formulation. Accurate numerical methods of solution were developed, utilising a mixed operational-Tau collocation method for the laminar flow equations and an operational-Tau method for the linear stability system. The operational-Tau spectral method was modified to avoid the inversion of badly scaled matrices, allowing the high order numerical solutions necessary in the current context. Solutions for the stability of plane Poiseuille flow and the Blasius velocity profile with its associated stability problem, were obtained for verification purposes, showing the method to be flexible, highly accurate and economical in terms of computer time. Laminar base flow predictions and corresponding linear stability results were obtained in dimensionless co-ordinates for a number of flows varying exponentially with time. At all times during the course of flow development the laminar flow model was shown to comprise two major elements (with an interface of inconsequential spatial extent between them); the first being parallel,

Declaration

I, Anthony Halley Abbot, declare that this thesis is my own unaided work. It is being submitted for the degree of Doctor of Philosophy in the University of the Witwatersrand, Johannesburg. It has not been submitted before for any degree or examination to any other University.



18th day of July 1995

The transition to turbulence in strongly accelerated pipe flows

Anthony Halley Abbot

A thesis submitted to the Faculty of Engineering, University of the Witwatersrand, Johannesburg, in fulfilment of the requirements for the degree of Doctor of Philosophy.

Johannesburg, June, 1995

E.2.2	Port control object	561
E.2.3	Timer I/O object	561
E.2.4	Digital I/O object	561
E.2.5	Analogue I/O object	561
E.2.6	Serial I/O object	562
E.2.7	Position I/O object	562
E.2.8	PC11 object	562
E.2.9	PC30 object	562
E.3	Transformation sub-system	562
E.4	The bench sub-system	563
E.4.1	Stepper-motor object	563
E.4.2	Actuator object	563
E.4.3	Limit object	564
E.4.4	Interactive bench control	564
E.4.5	Co-ordinate system & origin	564
E.5	Tank and pressure control	565
E.5.1	Pressure control	565
E.5.2	Tank control	566
E.6	The slave control, data transfer and processing sub-system	566
E.6.1	Slave control & data transfer	566
E.6.2	Data objects and data processing	567
E.7	Closed-loop flow control	567
E.8	Acquisition control sub system	568
E.8.1	Run-file execution object	569
E.8.2	Command execution	570
E.8.3	Safety & rig control	570

D.4.1	base2d.m	544
D.4.2	contint.m	547
D.4.3	contmat.m	547
D.4.4	kepint.m	548
D.4.5	linopint.m	548
D.4.6	makomati.m	548
D.4.7	nonlini.m	549
D.5	The Sexl stability equation	550
D.5.1	sex1.m	550
D.5.2	mako_opa.m	551
D.5.3	l1.opor.m	552
D.5.4	l2.opor.m	552
D.5.5	4.opor.m	552
D.5.6	makoua.m	553
D.5.7	ba.m	553
D.5.8	findro.m	554
D.5.9	getnoco.m	555
E	The main control software	557
E.1	General design principles	557
E.1.1	The elements of the figure	557
E.1.2	Containment	559
E.1.3	Inheritance	559
E.1.4	Reference	559
E.2	The input/output sub-system	560
E.2.1	I/O card object	560

D.1.1	chebvec.m	531
D.1.5	col2cheb	531
D.1.6	colldor	532
D.1.7	intmat.m	532
D.1.8	lobgrid	533
D.1.9	predmat	533
D.1.10	poly2d	534
D.1.11	coef2pts	534
D.1.12	transopa	535
D.1.13	convcheb	535
D.1.14	eta	536
D.1.15	mu	536
D.2	The Blasius solution	537
D.2.1	blasius.m	537
D.2.2	frechet.m	538
D.2.3	g_of_f.m	539
D.2.4	beblan.m	539
D.3	The one-dimensional base flow equation	540
D.3.1	baccid.m	540
D.3.2	otamu.m	542
D.3.3	knt.m	542
D.3.4	l_opor.m	542
D.3.5	omat.m	543
D.3.6	nmat.m	543
D.3.7	omegamat.m	543
D.4	The two dimensional base flow equation	544

19.3 Effects of non-axisymmetry and finite-amplitude disturbances . . .	462
20 Conclusions	466
21 Recommendations	471
V APPENDICES	473
A Error and sensitivity analysis	474
A.1 Analogue acquisition	474
A.2 Scatter in the flowmeter data	475
A.3 Systematic error in the profile wall slope due to uncertainty in the wall position	477
A.3.1 Error in the mass flow rate due to wall error	478
A.4 Errors in experimental velocity profiles	479
A.4.1 Velocity profile integration error	479
A.4.2 Effect of noise	480
A.4.3 Verification of the mean flow accuracy	482
B Base flow data	483
C Tabulated velocity profile and stability data	496
C.1 The far-downstream parallel developing flow	496
C.2 The steady pipe entrance flow	513
D Matlab m-files	520
D.1 The basic operators	529
D.1.1 muhat.m	529
D.1.2 otahat.m	530
D.1.3 chob2.m	530

17.1.4	Critical Reynolds number and wavenumber	407
17.1.5	Physical interpretation of stability results	411
17.2	The far-time entrance flow	413
17.2.1	Critical Reynolds number and wavenumber	414
17.2.2	The physical co-ordinate system	418
17.3	Comparison between the spatial and temporal results	422
17.4	Summary	424
IV	SYNTHESIS	425
18	Laminar base flow	427
18.1	Acceleration parameter	428
18.2	Dimensionless co-ordinate system	431
18.2.1	Asymptotic approximations to the full numerical solution	433
18.2.2	Flow development from a fixed physical position	434
18.2.3	Basis of comparison with experimental data	435
18.3	Flow development in time and space	435
18.3.1	Displacement thickness	436
18.3.2	The velocity profile shape parameter	440
18.4	Flow development with increasing acceleration	443
18.5	Velocity profiles	447
18.5.1	Parallel developing and steady entrance velocity profiles .	448
18.5.2	The effect of acceleration on the velocity profile	449
19	Stability and inception	451
19.1	Synthesis of measured data	453
19.2	Comparisons	457

15.4 Summary	376
16 The stability of steady pipe entrance flows	378
16.1 General overview	378
16.2 Analytical formulation of the unsteady Sesh equation	383
16.2.1 A singularity as a hindrance to spectral accuracy	383
16.2.2 Removal of the singularity	384
16.3 Types of linear instability	386
16.4 Numerical formulation	388
16.4.1 Basic equations	389
16.4.2 Boundary conditions	389
16.4.3 Method of solution	391
16.5 Comprehensive results	391
16.5.1 Displacement thickness Reynolds number	392
16.5.2 Wavenumber	394
16.5.3 Celerity	395
16.5.4 Behaviour with increasing x	396
16.5.5 The point of maximum instability in a pipe	397
16.5.6 Neutral stability curves	397
16.5.7 Eigenvalue spectrum	399
16.6 Summary	400
17 The stability of exponentially accelerating pipe flows	402
17.1 The far-downstream temporally developing flow	402
17.1.1 Stability of the limiting error function profile	403
17.1.2 Temporal variation of the velocity profile shape	405
17.1.3 Solution procedure	406

14.1.6 The nearly-parallel nature of the flow	325
14.5 Exclusion of the inlet plane	328
14.6 Summary	328
15 Exponentially accelerating pipe flows	331
15.1 Flowrate variation	332
15.2 The far-downstream flow	333
15.2.1 Centreline velocity and entrance time	334
15.2.2 Displacement thickness	336
15.2.3 Velocity profiles	338
15.2.4 Velocity profile shape parameter	339
15.2.5 Scaled shape parameter	340
15.2.6 Limiting behaviour as $\Omega \rightarrow \infty$	342
15.3 The far-time entrance flow	344
15.3.1 Method of solution	345
15.3.2 Centreline variation	346
15.3.3 Displacement thickness	347
15.3.4 Velocity profile shape parameter	348
15.3.5 Scaled shape parameter	349
15.3.6 Comparison between the limiting temporal and spatial re- sults	351
15.4 The limiting cases as approximations to the proper solution . . .	353
15.5 Mapping to the physical domain	356
15.5.1 A normalised co-ordinate system for the physical domain	359
15.5.2 Development of the flow for $\Omega = 20$	360
15.6 Physical framework for comparison to experiment	364
15.7 Comprehensive results	367

13.2.4	Initial conditions and solution procedure	280
13.2.5	Convergence and accuracy	281
13.2.6	Derived parameters	283
13.2.7	Limiting behaviour of S_{δ^*}	286
13.2.8	Tabulated results	287
13.3	The Blasius problem - base flow and stability	289
13.3.1	Analytical formulation	289
13.3.2	Numerical formulation	295
13.3.3	Solution procedure	297
13.3.4	Numerical results	298
13.4	Summary	302
14	Impulsive laminar pipe entrance flow	305
14.1	Analytical formulation	306
14.2	Numerical formulation	307
14.2.1	Axial discretisation	308
14.2.2	Non-linear terms	309
14.2.3	Operators	312
14.2.4	Discretisation	312
14.3	Method of solution	315
14.4	Numerical results	316
14.4.1	Comparisons with previous results	317
14.4.2	Centreline velocity	319
14.4.3	Displacement thickness	320
14.4.4	Velocity profile shape parameter	321
14.4.5	Scaled shape parameter	323

12.6.1	Legendre polynomials	241
12.6.2	Chebyshev polynomials in $[-1, 1]$	242
12.6.3	Chebyshev polynomials in $[0, 1]$	244
12.7	Direct formulation of operators	245
12.7.1	Transformed operators	246
12.7.2	Verification of operator transformation	248
12.8	Collocation	250
12.8.1	Collocation grid	250
12.8.2	General description	251
12.8.3	Differentiation	252
12.8.4	The collocation approach as an operational system	253
12.8.5	Advantages of the collocation method	255
12.9	Summary	256
13	Verification	257
13.1	The Orr-Sommerfeld stability problem	258
13.1.1	Analytical and numerical formulation	260
13.1.2	Treatment of general velocity profiles	261
13.1.3	Construction	263
13.1.4	Boundary conditions	264
13.1.5	Solution procedure	266
13.1.6	Numerical results for plane-Poiseuille flow	267
13.2	Impulsively started laminar pipe flow far from the inlet	275
13.2.1	Analytical formulation	275
13.2.2	Temporal discretisation	276
13.2.3	Numerical formulation	278

10.4.4	Applicability and limitations	209
10.5	Summary	211
11	Numerical methods	212
11.1	Historical development	213
11.2	The finite-difference approach	215
11.2.1	Boundary conditions	216
11.2.2	Solution procedure	217
11.2.3	Convergence & accuracy	217
11.3	The method of weighted residuals	219
11.3.1	Collocation	220
11.3.2	Galerkin	222
11.3.3	Tau	223
11.3.4	Accuracy	224
11.4	Commonality of solution procedure	225
11.5	Summary	228
12	The operational philosophy	229
12.1	Notational description	230
12.2	Operators	232
12.2.1	Differentiation & polynomial pre-multiplication	232
12.2.2	The limit operator	234
12.2.3	The definite integral	236
12.3	The inner problem	237
12.4	Boundary conditions	238
12.5	Solution procedure	239
12.6	Suitable sets of polynomials	240

9.1	Inception times and Reynolds numbers	173
9.2	General	178
9.3	Local and convected influences	178
9.3.1	The leeward side of the stability curve	180
9.3.2	Stability versus inception	181
9.4	The minimum critical Reynolds number	181
9.5	Summary	183
 III ANALYSIS		 184
10	The governing equations	186
10.1	The incompressible Navier-Stokes equations	187
10.2	Unsteady pipe flow	190
10.2.1	Co-ordinate transformation and equations of motion . . .	190
10.2.2	Boundary conditions	192
10.2.3	Scaling	194
10.2.4	Strongly varying flows	195
10.2.5	The acceleration parameter	198
10.2.6	The behaviour of the axial co-ordinate	199
10.3	Laminar base flow equations	201
10.3.1	The high Reynolds number approximation	201
10.3.2	Elimination of pressure term	202
10.4	The unsteady stability equation	203
10.4.1	Types of linear instability	204
10.4.2	Assumptions & methods of derivation	205
10.4.3	Derivation	205

8.1	Methodology of presentation	135
8.2	Typical evolution of the velocity profiles	136
8.2.1	The flow variation near the inlet	136
8.2.2	Profile development further downstream	138
8.2.3	Similarity between the parallel development regions	140
8.2	Inlet variability and inlet effects	141
8.3.1	The inlet hump	143
8.3.2	Flows at low accelerations	144
8.3.3	Non-receptivity to the flow waviness	147
8.4	The level of finite amplitude disturbances	148
8.4.1	Normalisation of the traces	149
8.4.2	Disturbance intensity	149
8.4.3	Power spectral density	152
8.5	Further data processing	153
8.5.1	Errors in the velocity derivative at the wall	154
8.5.2	Correction of systematic errors in S'	155
8.6	Typical behaviour of the derived parameters	156
8.6.1	Velocity profile shape parameter	156
8.6.2	Displacement thickness	158
8.7	Comprehensive variations of the derived parameters	159
8.7.1	The tests EXPE, EXPF, EXPG and HXPA	159
8.7.2	The tests FXPA, FXPB, FXPC and FXPD	165
8.7.3	The tests GXPA, GXPB, GXPC and GXPD	168
8.8	Summary	172
9	Experimental results - inception	173

14.9	The variation of the parallel parameter scaled with respect to displacement thickness κ_{δ^*} , with displacement thickness δ^* , in a steady-state pipe entrance flow, showing the Blasius limit at the inlet.	327
15.1	The variation of normalised flowrate Q/Q_0 with time in exponentially accelerating pipe flow, for various values of the acceleration parameter Ω	332
15.2	The variation of normalised centreline velocity u_c/U_m with time in parallel exponentially accelerating pipe flow, for various values of Ω . Also shown is the entrance-time T_e for each acceleration.	335
15.3	The variation of the entrance time T_e for parallel exponentially accelerating pipe flow, for various values of the acceleration parameter Ω	335
15.4	The variation of displacement thickness with time in parallel exponentially accelerating pipe flow, for various values of the acceleration parameter Ω	336
15.5	The variation of displacement thickness as $t \rightarrow \infty$, δ_{∞}^*/R , with acceleration parameter Ω in a parallel exponentially accelerating pipe flow.	337
15.6	The variation of velocity profiles at $t \rightarrow \infty$, with acceleration parameter Ω in parallel exponentially accelerating pipe flow.	338
15.7	The variation of shape parameter S with time t in parallel exponentially accelerating pipe flow, for varying values of the acceleration parameter Ω	339
15.8	The variation of scaled shape parameter S_{δ^*} with time t in parallel exponentially accelerating pipe flow, for varying values of the acceleration parameter Ω	340
15.9	The variation of scaled shape parameter S_{δ^*} with displacement thickness δ^*/R in parallel exponentially accelerating pipe flow, for varying values of the acceleration parameter Ω : $\infty \rightarrow \infty$, locus of limiting values as $t \rightarrow \infty$; $S_{\delta^*[\text{eff}]}$, parallel boundary-layer limit as $t \rightarrow 0$	341

13.6	The eigenspectrum for the even modes for plane-Poiseuille flow with $\alpha = 1$ and $Re = 50\,000$	274
13.7	The temporal development of the centreline velocity in an impulsively started pipe flow far from the inlet.	281
13.8	The variation of the displacement thickness δ^* with time in an impulsively started pipe flow far from the inlet.	283
13.9	The variation of the shape parameter S with time in an impulsively started pipe flow far from the inlet.	285
13.10	The temporal variation of the normalised wall shear stress S_{δ^*} in impulsively started pipe flow far from the inlet.	286
13.11	The variation of the normalised shape parameter S_{δ^*} with displacement thickness δ^* in an impulsively started pipe flow far from the inlet.	287
13.12	The development of the laminar boundary layer over a flat plate at zero incidence.	290
14.1	The normalised centreline velocity u_c/U_m , compared to the results of da Silva (1990); ϕ , da Silva; present.	310
14.2	The variation of the normalised centreline velocity u_c/U_m with axial station x in a steady-state pipe entrance flow, showing the entrance length Z_c	320
14.3	The variation of the displacement thickness δ^* with axial station x in a steady-state pipe entrance flow.	321
14.4	The variation of the shape parameter S with axial station x in a steady-state pipe entrance flow.	322
14.5	The variation of the scaled shape parameter S_{δ^*} with axial station x in a steady-state pipe entrance flow.	323
14.6	The variation of the scaled shape parameter S_{δ^*} with displacement thickness δ^* , for the two cases of parallel impulsive pipe flow (---), and steady-state pipe entrance flow (- - -).	324
14.7	A comparison between the parallel and Blasius boundary layers for identical displacement thickness δ^* : parallel boundary layer; - - -, Blasius profile.	325
14.8	The variation of the parallel parameter κ with displacement thickness δ^* in a steady-state pipe entrance flow.	326

9.15	The variation of the experimentally measured minimum critical Reynolds number with acceleration factor k : •, measured data; —, regression curve; ○, corrected value for EXP(4); □, theoretical linear minimum critical Reynolds number for steady pipe entrance flow [Abbot & Moss (1991)].	182
10.1	The radial co-ordinate system for pipe flow.	191
10.2	The simplified co-ordinate system for axisymmetric pipe flow.	191
10.3	The variation of a general function $w(x, y)$ in the range $x \in [0, k]$, $y \in [0, 1]$	196
10.4	The variation of a normalised general function $w(x, y)$ in the range $x \in [0, k]$, $y \in [0, 1]$, showing the encapsulation of the unbounded variation of the original function w in the term w_0	197
10.5	The effect of the variable mapping $x \mapsto \bar{x}$ on the perception of an impulsively started and exponentially accelerating pipe entrance velocity flow (indicated using displacement thickness δ^*) in the dimensionless co-ordinate system, from a position of constant x in the physical co-ordinate system. The dashed lines indicate lines of constant x . The physical variation of δ^* at an x station $x = k$ is indicated in the accompanying plot.	200
11.1	A typical one dimensional finite-difference grid.	216
11.2	The process of numerical formulation and numerical solution.	226
12.1	The Gauss-Lobatto grid points, shown for $N = 20$	250
13.1	The curves of constant temporal amplification for plane-Poiseuille flow [Shen (1954), after Lin (1955)].	259
13.2	The solution domain for the stability of flow between parallel plates.	259
13.3	A schematic representation of the neutral stability curve for temporal disturbances in plane-Poiseuille flow.	268
13.4	The variation of c_1 with Reynolds number for a particular value of the wavenumber α	269
13.5	The eigenspectrum for the even modes for plane-Poiseuille flow with $\alpha = 1$ and $Re = 10\,000$. Modes numbered according to the scheme of Mack (1976). □, A family; •, P family; □, S family.	272

9.1	The variation of inception time and Reynolds number with axial position for the test EXPE ($k = 0.03$, $Re_0 = 3018$).	174
9.2	The variation of inception time and Reynolds number with axial position for the test EXPF ($k = 0.03$, $Re_0 = 3018$).	174
9.3	The variation of inception time and Reynolds number with axial position for the test EXPG ($k = 0.07$, $Re_0 = 3018$).	174
9.4	The variation of inception time and Reynolds number with axial position for the test EXPA ($k = 0.02$, $Re_0 = 2659$).	175
9.5	The variation of inception time and Reynolds number with axial position for the test EXPA ($k = 0.04$, $Re_0 = 4143$).	175
9.6	The variation of inception time and Reynolds number with axial position for the test EXPB ($k = 0.05$, $Re_0 = 4143$).	175
9.7	The variation of inception time and Reynolds number with axial position for the test EXPC ($k = 0.07$, $Re_0 = 4143$).	176
9.8	The variation of inception time and Reynolds number with axial position for the test EXPD ($k = 0.03$, $Re_0 = 4143$).	176
9.9	The variation of inception time and Reynolds number with axial position for the test GXPA ($k = 0.05$, $Re_0 = 6510$).	176
9.10	The variation of inception time and Reynolds number with axial position for the test GXPB ($k = 0.07$, $Re_0 = 6510$).	177
9.11	The variation of inception time and Reynolds number with axial position for the test GXPC ($k = 0.04$, $Re_0 = 6510$).	177
9.12	The variation of inception time and Reynolds number with axial position for the test GXPD ($k = 0.03$, $Re_0 = 6510$).	177
9.13	The variation of experimental critical Reynolds number with axial position, for the test GXPA. Lines of mean flow 'convection' are also shown: (a) parallel development region; (b), entrance region; (c), minimum critical point.	179
9.14	The inferred shape of the stability curve and wash-down region in the pipe entrance: (a), region of linear or small amplitude instability; (b), convection and growth of disturbances; (c), 'front' at which turbulence occurs; (d), region destabilised by large disturbances from the inlet; (e), relatively undisturbed region of flow seen experimentally.	181

8.21	The variation in experimental shape parameter and displacement thickness, as a function of time, for the test HXPA ($k = 0.02$, $Re_0 = 2959$): o, HAG ($x/R = 113.57$); +, HAH ($x/R = 81.62$); x, HAI ($x/R = 51.56$); *, HAJ ($x/R = 32.99$); •, HAK ($x/R = 11.41$).	161
8.22	The variation in experimental shape parameter and displacement thickness, as a function of time, for the test FXPA ($k = 0.04$, $Re_0 = 4143$): o, FA0 ($x/R = 25.52$); +, FA1 ($x/R = 13.07$); x, FA3 ($x/R = 67.84$).	165
8.23	The variation in experimental shape parameter and displacement thickness, as a function of time, for the test FXPB ($k = 0.05$, $Re_0 = 4143$): o, FB3 ($x/R = 13.69$); +, FB4 ($x/R = 44.61$); x, FB5 ($x/R = 90.54$).	166
8.24	The variation in experimental shape parameter and displacement thickness, as a function of time, for the test FXPC' ($k = 0.07$, $Re_0 = 4143$): o, FC3 ($x/R = 38.38$); +, FC4 ($x/R = 79.96$).	166
8.25	The variation in experimental shape parameter and displacement thickness, as a function of time, for the test FXPD ($k = 0.03$, $Re_0 = 4143$): o, FD2 ($x/R = 79.54$); +, FD3 ($x/R = 36.31$).	167
8.26	The variation in experimental shape parameter and displacement thickness, as a function of time, for the test GXPA ($k = 0.05$, $Re_0 = 6510$): o, GAG ($x/R = 11.51$); +, GAA ($x/R = 44.40$).	169
8.27	The variation in experimental shape parameter and displacement thickness, as a function of time, for the test GXPB ($k = 0.07$, $Re_0 = 6510$): o, GBH ($x/R = 11.74$); +, GBI ($x/R = 41.91$); x, GBJ ($x/R = 80.79$).	169
8.28	The variation in experimental shape parameter and displacement thickness, as a function of time, for the test GXPC' ($k = 0.04$, $Re_0 = 6510$): o, GCE ($x/R = 11.49$); +, GCF ($x/R = 24.68$); x, GCG ($x/R = 41.28$).	170
8.29	The variation in experimental shape parameter and displacement thickness, as a function of time, for the test GXPD ($k = 0.03$, $Re_0 = 6510$): o, GDE ($x/R = 67.01$); x, GDG ($x/R = 30.94$).	170
8.30	The approximate variation of the final parallel flow shape parameter with acceleration factor k : o, 'E' series of tests; +, 'F' series of tests; x, 'G' series of tests; —, linear regression line; shaded region, region of uncertainty.	172

8.9	Normalised velocity profiles for the test IIAJ ($x/R = 113.57$), represented as a surface.	145
8.10	Velocity profiles for the test IIAJ ($x/R = 32.99$), represented as a surface.	147
8.11	The variation in disturbance intensity with time at centreline near the pipe inlet, for the test FXPD.	150
8.12	The variation in disturbance intensity with rate of exponential acceleration, for the flow at centreline near the pipe inlet.	151
8.13	The power spectrum for the centreline flow at the inlet, for the test FXPA. Data re-sampled at 100 Hz.	153
8.14	The uncorrected experimental shape parameter as a function of time, for the test FXPA, showing the systematic deviation between results from the different measuring stations: (1), FA0 ($x/R = 25.52$); (2), FA1 ($x/R = 13.07$); (3), FA3 ($x/R = 67.84$).	155
8.15	The corrected experimental shape parameter as a function of time, for the test FXPA: (1), FA0 ($x/R = 25.52$); (2), FA1 ($x/R = 13.07$); (3), FA3 ($x/R = 67.84$).	156
8.16	The variation of the corrected velocity profile shape parameter with time, for the test FXPA: o, FA0 ($x/R = 25.52$); •, FA1 ($x/R = 13.07$); x, FA3 ($x/R = 67.84$); —, mean fitted lines.	157
8.17	The variation in experimental displacement thickness as a function of time, for the test FXPA: o, FA0 ($x/R = 25.52$); •, FA1 ($x/R = 13.07$); x, FA3 ($x/R = 67.84$); —, mean fitted lines.	158
8.18	The variation in experimental shape parameter and displacement thickness, as a function of time, for the test EXPE ($k = 0.03$, $Re_D = 3018$): o, EF1 ($x/R = 12.24$); +, EF2 ($x/R = 29.25$); x, EF3 ($x/R = 64.52$); *, EF5 ($x/R = 320.50$).	160
8.19	The variation in experimental shape parameter and displacement thickness, as a function of time, for the test EXPF ($k = 0.05$, $Re_D = 3018$): o, EF1 ($x/R = 63.69$); +, EF2 ($x/R = 40.04$); x, EF3 ($x/R = 13.49$).	160
8.20	The variation in experimental shape parameter and displacement thickness, as a function of time, for the test EXPG ($k = 0.07$, $Re_D = 3018$): o, EF40 ($x/R = 13.49$); +, EF41 ($x/R = 64.52$); x, EF42 ($x/R = 105.06$); *, EF43 ($x/R = 84.81$).	161

7.16	The variation of the measured Reynolds number with time, for the tests: (1), EXPE; (2), EXPF; (3), EXPG; (4), HXPA; ---, expected variation; , actual variation.	126
7.17	The variation of the measured Reynolds number with time, for the tests: (1), FXPA; (2), FXPB; (3), FXPC; (4), FXPD; ---, expected variation; , actual variation.	127
7.18	The variation of the measured Reynolds number with time, for the tests: (1), GXPA; (2), GXPB; (3), GXPC; (4), GXPD; ---, expected variation; , actual variation.	127
7.19	The definition of the boundary layer displacement thickness. The flows represented by the two shaded areas $a \rightarrow b \rightarrow c \rightarrow d$ and $a \rightarrow b \rightarrow e \rightarrow f$ are equal.	130
7.20	The approximate points used to extract the wall derivative and centreline velocity from a normalised velocity profile: (a), points used to find the wall derivative; (b), point used to determine the centreline velocity.	133
8.1	Normalised velocity profiles for the test FB3 ($x/R = 13.69$), represented as a surface.	136
8.2	Superimposed normalised velocity traces for the test FB3 ($x/R = 13.69$), from selected radial stations: (1), $r/R = 0.975$; (2), $r/R = 0.932$; (3), $r/R = 0.866$; (4), $r/R = 0.838$; (5), $r/R = 0$ (centreline).	138
8.3	Normalised velocity profiles for the test FB4 ($x/R = 41.61$), represented as a surface.	139
8.4	Normalised velocity profiles for the test FB5 ($x/R = 90.54$), represented as a surface.	140
8.5	The temporal variation of dimensionless centreline velocity in time, for the test FXPB; $\circ$, FB3 ($x/R = 13.69$); $\bullet$, FB4 ($x/R = 41.61$); ∞ , FB5 ($x/R = 90.54$).	141
8.6	The velocity profile from the test FB5 ($x/R = 90.54$) at time 22 seconds: $\circ$, measured data; , interpolated curve.	142
8.7	Velocity profiles for the test FB3 ($x/R = 36.31$), represented as a surface. (1), the inlet bump.	143
8.8	The development of the entrance 'bump' with axial position, for the test EXPE. The dotted line shows the approximate locus of the inlet bump.	144

7.7	Scaled velocity time traces corresponding to the velocity profile FB5, illustrating the point of deviation from laminar flow in the pipe entrance. The traces correspond to the following radial stations: (1), $r = 0.973$; (2), $r = 0.952$; (3), $r = 0.930$; (4), $r = 0.907$; (5), $r = 0.885$; (6), $r = 0.863$; (7), $r = 0.841$; (8), $r = 0.819$; (9), $r = 0$.	113
7.8	A typical velocity - time trace in the boundary layer at the pipe inlet, showing the point of deviation from laminar flow. (1), velocity time trace; (2), velocity time trace normalised with respect to mean cross-sectional velocity; o, point of deviation; •, point of transition.	115
7.9	The nature and timing of the transition phenomena in the pipe entrance region. Data is selected from the test GXPA. Axial stations are: (1), $x/R = 11.5$; (2), $x/R = 17.6$; (3), $x/R = 23.0$; (4), $x/R = 28.8$; (5), $x/R = 33.0$; (6), $x/R = 41.3$; (7), $x/R = 66.2$; (8), $x/R = 82.0$; (9), $x/R = 112.7$; (10), $x/R = 122.9$; ~ ~ ~, inception (deviation from laminar flow); ~ ~ ~, transitional behaviour.	116
7.10	A velocity - time trace in the pipe centreline far downstream, showing the convected turbulent front. The data is EE5_0032, with axial position $x/R = 320.50$.	118
7.11	Multiple traces from the measuring station GAO, ($\bar{x}/R = 28.76$), illustrating the variability of the point of inception and the resultant perturbed velocity field: (1), laminar flow; (2), perturbed laminar flow; (3), transitional flow; (4), turbulent flow; (5), shut-down variability.	119
7.12	The distribution of the points of deviation about their mean, for the entire test GXPA. Data has been combined from all measurement stations.	120
7.13	Velocity profiles for the entire pipe diameter, showing the symmetry of the flow. Data from the far portion of the pipe have been reflected about the pipe origin to lie on the near-side data: (1), $t = 10$ s; (2), $t = 20$ s; o, near-side data; x, far-side data.	122
7.14	The numerical integration of the experimental velocity profile, in order to calculate the flowrate: (a), point of intersection of the fitted curves.	123
7.15	The calibration curve for the flowmeter: o, experimental data; ~ ~ ~, fitted curve.	125

6.7	A schematic representation of the development of the velocity profile in an accelerating pipe flow: (a), The physical system; (b), the normalised system. Profiles are labelled (1) to (5) in order of development. The right-hand side of the figure shows the profiles superimposed.	94
6.8	A normalised set of velocity-time traces from the axial station $x/R = 160.25$, from selected radial positions: (1), $r/R = 0.931$; (2), $r/R = 0.863$; (3), $r/R = 0.716$; (4), centreline; (i), inlet 'front'; (t), point of transition.	95
6.9	A normalised set of velocity-time traces from the axial station $x/R = 12.24$, from selected radial positions: (1), $r/R = 0.966$; (2), $r/R = 0.926$; (3), $r/R = 0.916$; (4), $r/R = 0.893$; (5), centreline.	96
6.10	The various types of impulsive starts, in order of their effect on the flow: (a), rounded start; (b), impulsive start; (c), under-damped impulsive start.	98
7.1	The principle whereby the LDV measuring volume was placed on a major diameter: (a), mis-aligned, beams reflecting down; (b), vertically aligned.	104
7.2	The beam distortion occurring through the various interfaces in the pipe: κ , free-air half-angle; κ' , actual half-angle; δ , beam displacement due to water-glass-water interface.	106
7.3	A three dimensional representation of a set of velocity-time traces, showing a particular velocity profile.	108
7.4	The method of determination of the wall position from a particular experimental velocity profile: $\circ$, experimental data; $\sim$, approximate velocity profile; $-$, low-order boundary layer curve fit; $\bullet$, position of the wall. The dotted rectangle contains the points selected for the boundary layer fit.	109
7.5	The general distribution of testing stations in the pipe entrance region.	111
7.6	Velocity time traces corresponding to the velocity profile FB5, illustrating the point of deviation from laminar flow in the pipe entrance. The traces correspond to the following radial stations: (1), $r = 0.973$; (2), $r = 0.952$; (3), $r = 0.930$; (4), $r = 0.907$; (5), $r = 0.885$; (6), $r = 0.863$; (7), $r = 0.841$; (8), $r = 0.819$; (9), $r = 0$	112

4.7	The pipe joining flanges.	51
4.8	The testing section of the rig.	52
4.9	The refraction of light through a water filled tube: (a), without a surrounding water jacket; (b), with a water jacket.	53
4.10	The flowmeter preceded by the stilling chamber. The direction of the flow is down. A cross-sectional schematic of the stilling chamber is shown on the right.	56
4.11	The valve network.	58
4.12	The LDV mounted on the traversing bench.	60
4.13	A schematic diagram of the remote laser head, showing the redirection of the vertical beams from the main unit.	62
4.14	The LDV and traversing bench, showing the vertical traversing assembly.	64
4.15	The layout of the data acquisition hardware.	65
4.16	The digital data lines between the polling computer, the signal counter and the host computer.	66
5.1	The user interface screen.	75
5.2	The data acquisition flowchart.	79
6.1	The macro to execute a single acquisition run on the acquisition system.	86
6.2	The macro to execute multiple acquisition runs on the acquisition system.	87
6.3	A run-file to control the flow to a particular exponential acceleration, and acquire LDV data therefrom.	88
6.4	An unprocessed centreline trace from the axial position $x/R = 160.25$ from the preliminary experimentation, showing the occurrence of a turbulent front.	91
6.5	The smoothed centreline trace from the axial position $x/R = 160.25$	91
6.6	The smoothed wall trace from the axial position $x/R = 160.25$. The radial station is $r/R = 0.931$	93

List of Figures

2.1	The general nature of transition in wall bounded shear flows [from Morkovin & Reshotko (1989)].	12
3.1	The beam crossing point in a dual beam LDV system, illustrating the measuring volume and fringes [from Fingerson <i>et al</i> (1989)].	26
3.2	The measured signal from a dual beam LDV.	26
3.3	A schematic representation of a dual beam, back-scatter LDV, showing the principal components [from Fingerson <i>et al</i> (1989)].	27
3.4	The signal from the anemometer photomultiplier shown in the frequency domain.	29
3.5	The progressive processing of the signal from the photodetector, resulting in a frequency measurement.	30
3.6	The temporal bias problem and its correction in LDV data processing.	36
4.1	The testing facility.	41
4.2	A schematic view of the apparatus, showing the various sub-systems: (a), flow rig sub-system; (b), acquisition & control sub-system; (c), measurement sub-system.	42
4.3	A cross-sectional view of the supply tank.	44
4.4	The header tank.	46
4.5	The shape of the inlet contraction.	48
4.6	Problems associated with cold joining of pipes: (a), step caused by mismatch in diameters; (b), lack of access does not allow a transition region to be blended; (c), flange allows access for blending of the pipe surface.	50

E.8.4	Analogue acquisition & interrupts	571
E.8.5	The abort sequence	571
E.8.6	General considerations	572
E.9	Auto traversing sub-system	572
E.10	Main command loop	573
E.10.1	Command,	573
E.10.2	Main command loop	574
E.10.3	Data display sub-system	575
E.10.4	Run-file compiler	575
E.11	Command syntax and definitions	575
E.11.1	Command-line and macro syntax	576
E.11.2	Run file commands	576
F	The LDV components	580
G	Raw data on CD-ROM	584
	References	585

List of Tables

1.1	The analogue connections.	69
7.1	The naming convention for the experimental runs, showing their abbreviations and file names. The terms x and y indicate arbitrary numeric values.	101
7.2	Description of the file suffixes used in data storage.	102
7.3	The values of the parameters Re_0 , k to generate the flowrate variations for the separate tests.	103
8.1	The axial stations, approximate acceleration constants and impulsive start-up Reynolds number for the ‘E’ series of tests. Included is the data IIXPA.	162
8.2	The adjustments for the shape parameter for the ‘E’ series of tests.	164
8.3	The axial stations, approximate acceleration constants and impulsive start-up Reynolds number for the ‘F’ series of tests. . . .	167
8.4	The adjustments for the shape parameter for the ‘F’ series of tests. . . .	168
8.5	The axial stations, approximate acceleration constants and impulsive start-up Reynolds number for the ‘G’ series of tests. . . .	171
8.6	The adjustments for the shape parameter for the ‘G’ series of tests.	171
11.1	Convergence results of da Silva (1990) for the Orr-Sommerfeld equation, applied to plane-Poiseuille flow, with $\alpha = 1$ and $R = 10\,000$, and compared to the ‘exact’ result of Orszag (1971a). . . .	219

B.32	Velocity profiles for the test HAG ($x/R = 113.57$), represented as a surface.	494
B.33	Velocity profiles for the test HAH ($x/R = 81.62$), represented as a surface.	494
B.34	Velocity profiles for the test HAI ($x/R = 54.56$), represented as a surface.	494
B.35	Velocity profiles for the test HAJ ($x/R = 32.99$), represented as a surface.	495
B.36	Velocity profiles for the test HAK ($x/R = 11.41$), represented as a surface.	495
E.1	The top-level object relationships in the system: $\leftrightarrow$, inheritance; $\bullet$, containment; $\circ$, reference.	558
E.2	The principle of controlling the tank pressure.	566
E.3	The data acquisition & control loop.	569
F.1	Layout of the model 9100-7 LHV system (from TSI System 9100-7 LHV instruction manual).	580

B.16 Velocity profiles for the test FB4 ($x/R = 44.61$), represented as a surface.	488
B.17 Velocity profiles for the test FB5 ($x/R = 90.54$), represented as a surface.	489
B.18 Velocity profiles for the test FC3 ($x/R = 38.38$), represented as a surface.	489
B.19 Velocity profiles for the test FC4 ($x/R = 79.96$), represented as a surface.	489
B.20 Velocity profiles for the test FD2 ($x/R = 79.54$), represented as a surface.	490
B.21 Velocity profiles for the test FD3 ($x/R = 36.31$), represented as a surface.	490
B.22 Velocity profiles for the test GAA ($x/R = 11.51$), represented as a surface.	490
B.23 Velocity profiles for the test GAA ($x/R = 44.40$), represented as a surface.	491
B.24 Velocity profiles for the test GBII ($x/R = 11.74$), represented as a surface.	491
B.25 Velocity profiles for the test GBI ($x/R = 41.91$), represented as a surface.	491
B.26 Velocity profiles for the test GBJ ($x/R = 99.85$), represented as a surface.	492
B.27 Velocity profiles for the test GCT ($x/R = 11.49$), represented as a surface.	492
B.28 Velocity profiles for the test GCT ($x/R = 24.69$), represented as a surface.	492
B.29 Velocity profiles for the test GCG ($x/R = 11.29$), represented as a surface.	493
B.30 Velocity profiles for the test GDF ($x/R = 67.01$), represented as a surface.	493
B.31 Velocity profiles for the test GDG ($x/R = 30.91$), represented as a surface.	493

A.5	Numerical velocity profile with noise of magnitude 1% added: ϕ , noisy data; ϕ_{exact} , exact curve.	481
B.1	Velocity profiles for the test EE1 ($x/R = 12.24$), represented as a surface.	483
B.2	Velocity profiles for the test EE2 ($x/R = 29.25$), represented as a surface.	484
B.3	Velocity profiles for the test EE3 ($x/R = 64.52$), represented as a surface.	484
B.4	Velocity profiles for the test EE5 ($x/R = 320.50$), represented as a surface.	484
B.5	Velocity profiles for the test EF1 ($x/R = 63.69$), represented as a surface.	485
B.6	Velocity profiles for the test EF2 ($x/R = 40.01$), represented as a surface.	485
B.7	Velocity profiles for the test EF3 ($x/R = 13.49$), represented as a surface.	485
B.8	Velocity profiles for the test EG0 ($x/R = 13.49$), represented as a surface.	486
B.9	Velocity profiles for the test EG1 ($x/R = 64.52$), represented as a surface.	486
B.10	Velocity profiles for the test EG2 ($x/R = 105.06$), represented as a surface.	486
B.11	Velocity profiles for the test EG3 ($x/R = 169.63$), represented as a surface.	487
B.12	Velocity profiles for the test FA0 ($x/R = 25.52$), represented as a surface.	487
B.13	Velocity profiles for the test FA1 ($x/R = 13.07$), represented as a surface.	487
B.14	Velocity profiles for the test FA3 ($x/R = 67.84$), represented as a surface.	488
B.15	Velocity profiles for the test FB3 ($x/R = 13.49$), represented as a surface.	488

19.6	The variation of critical Reynolds number with axial position, for various values of the acceleration parameter: (a), curves extracted ¹ from the fitted experimental surface: thin lines, the theoretical stability curves.	458
19.7	The variation of the minimum critical Reynolds number with acceleration parameter: (a), curve extracted from experimental data surface fit; (b), theoretical pipe entrance minimum critical curve; $\infty = \infty$, extrapolation to zero acceleration.	458
19.8	The variation of the minimum critical axial position with acceleration parameter: (a), curve extracted from experimental data surface fit; (b), theoretical pipe entrance minimum critical curve.	459
19.9	The relationship between the uncertainty in the position of a curve, and the uncertainty in the position of its minimum point: (a), fitted curve; (b), possible valid variation of Re_c	460
19.10	The variation of critical displacement thickness Reynolds number with axial position, for various values of the acceleration parameter: (a), curves extracted from the fitted experimental surface: thin lines, theoretical stability curves.	461
19.11	The variation of the RMS background disturbance magnitude with acceleration parameter, measured at centreline near the pipe inlet.	464
19.12	The reduction in experimental minimum critical Reynolds number with respect to the linear theory, with disturbance amplitude.	465
19.13	The expected reduction in the minimum critical Reynolds number with respect to the linear theory, as the disturbance amplitude approaches zero.	465
A.1	Unsmoothed flowmeter data from the test EXPE, showing the innate broadening in the data. The full range is 4 l/s.	475
A.2	The distribution of the flowmeter data about the mean. The horizontal scale is normalised with respect to the full range (in this case 4 l/s). A model Gaussian distribution is shown superimposed.	476
A.3	The effect of wall variability on the estimation of the velocity derivative at the pipe wall.	477
A.4	Numerical velocity profile, showing sampled points to be used for the estimation of the integration error.	480

18.9	The variation of dimensionless displacement thickness with dimensionless axial position, for the boundary layer thinning portion of the 'E' series of tests: •, EXPE ($\Omega \approx 20$); ○, EXPF ($\Omega \approx 33$); ◐, EXPG ($\Omega \approx 45$); ×, HXPA ($\Omega \approx 13$); —, theoretical curves.	444
18.10A	schematic representation of the portions of the flow from where velocity profiles were obtained: (a) profiles from the parallel and thinning flow; (b) profiles from the parallel flow region - different acceleration parameters.	447
18.11E	Experimental velocity profiles from the parallel and entrance regions of the flow, for the station EE2 ($x/H = 29.25$, $\Omega \approx 20$): ◐, $t = 7.685 \times 10^{-3}$ (parallel development); ×, $t = 2.152 \times 10^{-2}$ (entrance flow); the two solid lines represent the equivalent numerical parallel and entrance flow profiles.	448
18.12	The variation in the shape of the velocity profile with acceleration parameter in the parallel developing region for the 'E' series of tests: ◐, station EE2, $t = 7.685 \times 10^{-3}$, $\Omega \approx 20$; ×, station EG3, $t = 1.261 \times 10^{-2}$, $\Omega \approx 45$; the two solid lines represent the equivalent numerical velocity profiles for $\Omega = 20$ and $\Omega = 50$	450
19.1	The schematic depiction of the transition to turbulence in the entrance region of an exponentially accelerating pipe flow: (a), laminar flow; (b), local boundary layer breakdown; (c), transitional flow converts downstream, local boundary layer breakdown persists near the inlet; (d), local breakdown upstream, turbulent front converts downstream; (e) local breakdown upstream, pipe-wide turbulent flow downstream; - - -, region of inception. . . .	452
19.2	The axial variation of the experimental flow breakdown, from the test GXPA ($Re_0 = 6510$, $\Omega = 32.5$), as a function of time and Reynolds number. The error-bars represent one standard deviation.	453
19.3	The variation of the experimental critical Reynolds number with dimensionless axial position, for the test GXPA ($Re_0 = 6510$, $\Omega = 32.5$).	454
19.4	The distribution of the experimental data about the fitted surface.	456
19.5	The three dimensional fitted surface to the experimental inception data in (x, Ω, Re) space, showing the experimental points superimposed. The extent of the surface is indicated on the floor of the plot.	457

17.14	The variation of the critical displacement thickness Reynolds number with displacement thickness, for discrete values of the acceleration parameter: (1), parallel developing flow; (2), entrance flow; (i), Error-function stability limit, $Re_{\delta^*[\text{crit}]} = 1675$; (ii), Blasius stability limit, $Re_{\delta^*[\text{crit}]} = 519.06$; (a), $\Omega = 0$; (b), $\Omega = 10$; (c), $\Omega = 20$; (d), $\Omega = 30$; (e), $\Omega = 40$; (f), $\Omega = 50$; (g), $\Omega = 60$; (h), $\Omega = 70$	423
18.1	The flow variation with time, for the 'E' series of tests. Both the Reynolds number and acceleration parameter variation are shown: (a), EXPE; (b), EXPE; (c), EXPG; (d), EXPA; — — —, expected variation;, actual variation.	429
18.2	The flow variation with time, for the 'F' series of tests. Both the Reynolds number and acceleration parameter variation are shown: (a), FXPA; (b), FXPB; (c), FXPG; (d), FXPD; — — —, expected variation;, actual variation.	430
18.3	The flow variation with time, for the 'G' series of tests. Both the Reynolds number and acceleration parameter variation are shown: (a), GXPA; (b), GXPB; (c), GXPG; (d), GXPD; — — —, expected variation;, actual variation.	431
18.4	The temporal development of δ^*/R for the test EXPE: —, theoretical results; •, EE1 ($x/R = 12.24$); o, EE2 ($x/R = 29.25$); ◊, EE3 ($x/R = 61.52$); ×, EE5 ($x/R = 320.50$); A, B, C, chosen points of increasing time for EE3.	436
18.5	The spatial development of δ^*/R for the test EXPE: —, theoretical results; •, EE1 ($x/R = 12.24$); o, EE2 ($x/R = 29.25$); ◊, EE3 ($x/R = 61.52$); (a), (b), (c), inlet 'bump'; A', B', C', chosen points of increasing time for EE3.	437
18.6	The variation of the boundary layer shape parameter with dimensionless time, for the test EXPE: •, EE1 ($x/R = 12.24$); o, EE2 ($x/R = 29.25$); ◊, EE3 ($x/R = 61.52$); ×, EE5 ($x/R = 320.50$); —, theoretical curve for $\Omega = 20$	440
18.7	The variation of the boundary layer shape parameter with dimensionless axial position, for the test EXPE: •, EE1 ($x/R = 12.24$); o, EE2 ($x/R = 29.25$); ◊, EE3 ($x/R = 61.52$); —, theoretical curve for $\Omega = 20$	441
18.8	The variation of dimensionless displacement thickness with time, for the parallel portion of the 'E' series of tests: •, EXPE ($\Omega \approx 20$); o, EXPE ($\Omega \approx 33$); ◊, EXPG ($\Omega \approx 45$); ×, EXPA ($\Omega \approx 13$); —, theoretical curves.	443

17.5	The trajectories for various developing flows in parallel impulsively started pipe flow. (a), $Re = 10\,000$; (b), $Re = 20\,000$; (c), $Re = 30\,000$	411
17.6	The neutral stability curve for $\Omega = 70$, showing The trajectories for developing flows in parallel impulsively started accelerating pipe flow, for different values of Re_0	412
17.7	The temporal variation of the initial critical Reynolds number in parallel, temporally developing flow, for discrete values of the acceleration parameter; (a), $\Omega = 0$; (b), $\Omega = 10$; (c), $\Omega = 20$; (d), $\Omega = 30$; (e), $\Omega = 40$; (f), $\Omega = 50$; (g), $\Omega = 60$; (h), $\Omega = 70$	414
17.8	The spatial variation of the critical Reynolds number in pipe entrance flow, for discrete values of the acceleration parameter; (a), $\Omega = 0$; (b), $\Omega = 10$; (c), $\Omega = 20$; (d), $\Omega = 30$; (e), $\Omega = 40$; (f), $\Omega = 50$; (g), $\Omega = 60$; (h), $\Omega = 70$	415
17.9	The spatial variation of the critical boundary layer displacement thickness Reynolds number in pipe entrance flow, for discrete values of the acceleration parameter; •, Blasius boundary layer limit, $Re_{\delta^*crit} = 519.06$; (a), $\Omega = 0$; (b), $\Omega = 10$; (c), $\Omega = 20$; (d), $\Omega = 30$; (e), $\Omega = 40$; (f), $\Omega = 50$; (g), $\Omega = 60$; (h), $\Omega = 70$	416
17.10	The variation of the minimum critical Reynolds number with acceleration parameter in steady pipe entrance flow, compared to the data for the parallel downstream case: —, entrance; ---, downstream.	417
17.11	The variation of the critical wavenumber based on displacement thickness with acceleration parameter in steady pipe entrance flow; •, Blasius boundary layer limit, $\alpha_{\delta^*crit} = 0.303\,77$; (a), $\Omega = 0$; (b), $\Omega = 10$; (c), $\Omega = 20$; (d), $\Omega = 30$; (e), $\Omega = 40$; (f), $\Omega = 50$; (g), $\Omega = 60$; (h), $\Omega = 70$	418
17.12	The spatial variation of the critical Reynolds number in pipe entrance flow, for discrete values of the acceleration parameter; (a), $\Omega = 0$; (b), $\Omega = 10$; (c), $\Omega = 20$; (d), $\Omega = 30$; (e), $\Omega = 40$; (f), $\Omega = 50$; (g), $\Omega = 60$; (h), $\Omega = 70$; superimposed are lines of constant physical axial position $\hat{x} = \text{const}$	420
17.13	The variation of the critical Reynolds number in pipe entrance flow with respect to the physical co-ordinate system, for discrete values of the acceleration parameter; (a), $\Omega = 0$; (b), $\Omega = 10$; (c), $\Omega = 20$; (d), $\Omega = 30$; (e), $\Omega = 40$; (f), $\Omega = 50$; (g), $\Omega = 60$; (h), $\Omega = 70$; superimposed are lines of constant normalised axial position $\hat{x} = \text{const}$	421

16.4	The axial variation of the critical wavenumber, for axisymmetric disturbances in steady pipe entrance flow.	394
16.5	The axial variation of the critical wavenumber scaled with respect to displacement thickness, for axisymmetric disturbances in steady-state pipe entrance flow. •, the Blasius boundary layer limit, $\alpha_{\delta^*} = 0.303\ 77$	395
16.6	The axial variation of the product $\alpha_{\delta^*[\text{crit}]}Re_{\delta^*[\text{crit}]}$ with wavenumber $\alpha_{[\text{crit}]}$, for axisymmetric disturbances in steady-state pipe entrance flow.	397
16.7	Neutral stability curves in the $\alpha - Re$ plane, for selected velocity profiles: (i), $x = 0$ (Blasius boundary layer limit); (ii), $x = 0.000\ 197$; (iii), $x = 0.003\ 59$ (Minimum critical point); (iv), $x = 0.007\ 83$; (v), $x = 0.008\ 18$	398
16.8	Migration of the eigenvalues with axial distance in the range $0 \leq x \leq 0.008\ 18$ for the minimum critical Reynolds number at each position. Eigenvalues are numbered according to the scheme of Mack (1976). The continuous spectra are indicated as a series of lines. •, spectrum at $x = 0$ (Blasius limit); o, spectrum at $x = 0.008\ 18$; intermediate spectra correspond to $x = 0.000\ 071\ 1$, $0.000\ 631$, $0.001\ 51$, $0.002\ 72$, $0.004\ 55$, $0.005\ 94$, $0.007\ 06$	399
17.1	The temporal variation of the critical Reynolds number in parallel, temporally developing flow, for discrete values of the acceleration parameter: (a), $\Omega = 0$; (b), $\Omega = 10$; (c), $\Omega = 20$; (d), $\Omega = 30$; (e), $\Omega = 40$; (f), $\Omega = 50$; (g), $\Omega = 60$; (h), $\Omega = 70$	407
17.2	The variation of the minimum critical Reynolds number with acceleration parameter in parallel, temporally developing flow. . . .	408
17.3	The temporal variation of the critical displacement thickness Reynolds number in parallel, temporally developing flow, for discrete values of the acceleration parameter: (a), $\Omega = 0$; (b), $\Omega = 10$; (c), $\Omega = 20$; (d), $\Omega = 30$; (e), $\Omega = 40$; (f), $\Omega = 50$; (g), $\Omega = 60$; (h), $\Omega = 70$	409
17.4	The temporal variation of the critical displacement thickness wave number in parallel, temporally developing flow, for discrete values of the acceleration parameter: (a), $\Omega = 0$; (b), $\Omega = 10$; (c), $\Omega = 20$; (d), $\Omega = 30$; (e), $\Omega = 40$; (f), $\Omega = 50$; (g), $\Omega = 60$; (h), $\Omega = 70$	410

15.34	The variation of displacement thickness δ^*/R with normalised ordinate x , for various values of the dimensionless physical ordinate x^j ; $\Omega = 30$	371
15.35	The variation of displacement thickness δ^*/R with time t , for various values of the dimensionless physical ordinate x^j ; $\Omega = 40$	372
15.36	The variation of displacement thickness δ^*/R with normalised ordinate x , for various values of the dimensionless physical ordinate x^j ; $\Omega = 40$	372
15.37	The variation of displacement thickness δ^*/R with time t , for various values of the dimensionless physical ordinate x^j ; $\Omega = 50$	373
15.38	The variation of displacement thickness δ^*/R with normalised ordinate x , for various values of the dimensionless physical ordinate x^j ; $\Omega = 50$	373
15.39	The variation of displacement thickness δ^*/R with time t , for various values of the dimensionless physical ordinate x^j ; $\Omega = 60$	374
15.40	The variation of displacement thickness δ^*/R with normalised ordinate x , for various values of the dimensionless physical ordinate x^j ; $\Omega = 60$	374
15.41	The variation of displacement thickness δ^*/R with time t , for various values of the dimensionless physical ordinate x^j ; $\Omega = 70$	375
15.42	The variation of displacement thickness δ^*/R with normalised ordinate x , for various values of the dimensionless physical ordinate x^j ; $\Omega = 70$	375
16.1	The axial variation of the critical Reynolds number for axisymmetric disturbances in steady-state pipe entrance flow, showing the results of various Authors: •, Huang & Chen (1974); ×, Garg (1981); o, da Silva (1990); +, Sarpkaya (Experimental).	379
16.2	The axial variation of the critical Reynolds number for axisymmetric disturbances in steady-state pipe entrance flow: ×, Garg (1981); o, da Silva (1990); —, Abbot & Moss (1994); •, minimum critical Reynolds number.	381
16.3	The axial variation of the critical Reynolds number based on displacement thickness, for axisymmetric disturbances in steady-state pipe entrance flow: o, da Silva (1990); —, Abbot & Moss (1994); •, Blasius boundary layer limit $Re_{\text{critical}} = 519.06$	393

15.21	The variation of velocity profiles in time for $\Omega = 20$, at the physical station $x' = 0.002$.	362
15.22	The variation of velocity profiles in time for $\Omega = 20$, at the physical station $x' = 0.01$.	362
15.23	The variation of velocity profiles in time for $\Omega = 20$, at the physical station $x' = 0.02$.	363
15.24	The variation of displacement thickness δ^*/R in time for $\Omega = 20$; (1), $x' = 0.002$; (2), $x' = 0.01$; (3), $x' = 0.02$; (4), parallel ($x' = \infty$).	363
15.25	The simulated variation of experimentally 'measured' displacement thickness δ^*/R with normalised axial station x for $\Omega = 20$, corresponding to three dimensionless physical positions; (1), $x' = 0.002$; (2), $x' = 0.01$; (3), $x' = 0.02$; (4), steady-state entrance flow.	365
15.26	The actual variation, for $\Omega = 20$, of displacement thickness δ^*/R in x (as given in the previous figure), compared to data according to the parallel-temporal and steady-state spatial limits of the system; (1), (2), (3), parallel temporal flow, corresponding to $x' = 0.002$, $x' = 0.01$, and $x' = 0.02$ respectively; (4), steady-state entrance flow.	366
15.27	The variation of displacement thickness δ^*/R with time t , for various values of the dimensionless physical ordinate x' ; $\Omega = 0$.	368
15.28	The variation of displacement thickness δ^*/R with normalised ordinate x , for various values of the dimensionless physical ordinate x' ; $\Omega = 0$.	368
15.29	The variation of displacement thickness δ^*/R with time t , for various values of the dimensionless physical ordinate x' ; $\Omega = 10$.	369
15.30	The variation of displacement thickness δ^*/R with normalised ordinate x , for various values of the dimensionless physical ordinate x' ; $\Omega = 10$.	369
15.31	The variation of displacement thickness δ^*/R with time t , for various values of the dimensionless physical ordinate x' ; $\Omega = 20$.	370
15.32	The variation of displacement thickness δ^*/R with normalised ordinate x , for various values of the dimensionless physical ordinate x' ; $\Omega = 20$.	370
15.33	The variation of displacement thickness δ^*/R with time t , for various values of the dimensionless physical ordinate x' ; $\Omega = 30$.	371

15.10	The variation of scaled shape parameter S_{δ^*} with displacement thickness δ^*/R in parallel exponentially accelerating pipe flow, for varying values of the acceleration parameter Ω ; — — —, locus of limiting values as $t \rightarrow \infty$ of $(\delta^*/R, S_{\delta^*})$; $S_{\delta^* crq}$, parallel boundary-layer limit as $t \rightarrow 0$; lines of constant-time are shown superimposed.	343
15.11	The variation of normalised centreline velocity u_c/U_m with axial distance x , in the entrance region of a pipe, for various values of the acceleration parameter Ω	347
15.12	The variation of displacement thickness δ^*/R with axial position x in the entrance region of a pipe, for the chosen values of Ω	348
15.13	The variation of shape parameter S with axial position x in the entrance region of a pipe, for various values of the acceleration parameter Ω	349
15.14	The variation of scaled shape parameter S_{δ^*} with axial position x in the entrance region of a pipe, for various values of the acceleration parameter Ω	350
15.15	The variation of scaled shape parameter S_{δ^*} with displacement thickness δ^*/R in the entrance region of a pipe, for various values of the acceleration parameter Ω . The dashed line indicates the locus of limiting values for $x = \infty$	351
15.16	The variation of scaled shape parameter S_{δ^*} with displacement thickness δ^*/R in the entrance region of a pipe, for various values of the acceleration parameter Ω . The equivalent temporal development curves are superimposed, and are shown dashed. The dotted line indicates the locus of limiting values for $x = \infty$	352
15.17	The variation of displacement thickness δ^*/R with axial distance x in the developing entrance region of a pipe, for $\Omega = 20$. The curves represent the value of the parameter at various times.	354
15.18	The variation of displacement thickness δ^*/R with axial distance x in the developing entrance region of a pipe, for $\Omega = 20$ and $t = 7 \times 10^{-3}$; (1), steady-state entrance flow; (2), temporal developing flow; (3), the actual inlet flow.	356
15.19	The relationship between the dimensionless co-ordinate systems and the physical independent variables.	360
15.20	Lines of iso- x in the $t - x^d$ plane, for $\Omega = 20$	361

1.2 Objectives & proposed strategy

This study is specifically concerned with the various aspects of pipe entrance flows subjected to strong monotonic accelerations, including the base flow variation, instabilities and later transition, if any.

In particular it was decided to investigate exponentially accelerating flows within the pipe. This decision was based solely on the fact that the factor $(1/U')dU'/dt$, used by Shen (1961) in his limited analysis remained constant for an exponential flow. A dual strategy was adopted: experimental measurement of such flows was to be considered along with a detailed analytic and numerical analysis, hopefully concluding with some level of agreement between the two. The following specific objectives were envisaged:

- (i) To develop an experimental apparatus capable of accurately generating exponential and other arbitrary accelerations in a pipe. In particular a broad range of Reynolds numbers and varying rates of acceleration had to be attainable by the apparatus.
- (ii) In combination with the above goal, accurate acquisition of flow data was desired. Since it was desired to observe transitional flows, a sensitive, flexible and non-disruptive flow measurement was required. A laser Doppler velocimeter (LDV) was available for use. The applicability of this instrument for measurement of transitional flows was to be evaluated, and means of processing data developed.
- (iii) Since only the simple Shen (1961) criterion for accelerating flows was available, the development of an appropriate analytical basis was seen as a necessity. Although the structure of this system was initially unknown, its desired form was to be an exact theory and not asymptotic.
- (iv) A broad assortment of numerical analyses are usually used in the study of fluid mechanics problems, some of dubious applicability. A primary goal was thus to develop a uniform and accurate numerical methodology, possibly using the so-called 'spectral' approach, to tackle the equations encountered in the study of pipe entrance flows. It was decided to conduct a rigorous verification of the methods on various simple flow geometries, in order to verify their accuracy using known results for steady-state flows. Prototypal flows considered were: the stability of the Blasius boundary layer, and the still contentious issue of the linear stability of pipe entrance flows, among others.
- (v) Referring to point (iv) above, two particular numerical investigations were foreseen. The first concerned the development of the laminar $^{(1)}$ $^{(2)}$ flow under conditions of acceleration. Unlike in the impulsive case $^{(3)}$ where the general nature of the laminar pipe development was known [da Silva & Moss (1991)], in the current context even this had to be determined. Secondly the stability characteristics of the base flow were to be investigated.

Unsteady flows

It is however the more general types of flow that are the more interesting and in fact the more common in nature; for instance flows like the induction and combustion process within the internal combustion engine are by their nature highly unsteady, while wind does not blow steadily, but gusts. For this reason, steady approximations to general fluids phenomena will at best be quantitatively inaccurate and at worst completely wrong.

Several fundamental types of unsteady flow occur: flows with varying velocity fields but steady flow rates, flows with periodically varying flowrates that are steady *in the mean*, and truly unsteady flows showing unbounded mean flow variation in time. Of all these unsteady flows, it is the latter, non-periodic accelerating ones that are the least studied. A good comprehension is developing of the behaviour of the first type of unsteady flow. For instance in the context of pipe flow Moss & da Silva (1993), amongst others, have undertaken detailed studies of the impulsively started pipe entrance flow. Some headway, too, has been made in the case of the stability of periodic flows. In such flows the mean velocity fluctuates usually sinusoidally about a mean value of either zero or one offset from zero. Here parametric resonance effects can induce instability, and the principles of Floquet theory [Drazin & Reid (1981)] are applicable. However, in the generally accelerating case, the flow variation is unbounded, and besides the initial work by Shen (1961), there is little theoretical basis by which to understand the stability, let alone later non-linear transition of such flows.

The main dilemma centres around the unbounded nature of generally accelerating flows. In the case of the linear stability of such a flow, an instability might be growing in an absolute sense, but if the flowrate is increasing at a rate greater than the disturbance growth, then that disturbance is in reality diminishing with respect to the base flow. Uncertainty about the terms of definition of such flows are indicative of a poor understanding of the behaviour of the system.

The scope of the work is as a result of the basic lack of understanding, very broad. There is no general mathematical framework in which to analyse strongly accelerating flows; neither is there any clear understanding of what parameters govern instability and transition. The lack of research in this area in fact calls for a rigorous and fundamental investigation into the basics of accelerative flow. The fact that no one since Shen (1961)¹ has considered an accelerative analysis of these flows is indicative of the need for research in this field.

Discovery of an underlying principle in accelerating flows would be significant in that it would permit the extension of many of the current steady flow methods to describe the more general case. Instead of possessing any particular significance, steady flows would then be seen as simply a specific case of a more general theory. Although this is an idealistic aim, it serves to indicate the broad scope of any discovery in the field.

¹Except for the brief and simplified analysis by Moss (1991)

solutions to be obtained, while experimentally, the relative ease of obtaining steady experimental data makes it the obvious choice. Steadiness, although a useful mathematical abstraction, is never realised in real flows. Flows showing constant flow conditions in reality have to be first brought into motion from rest or changed from some other state. Although the dynamic nature of the flow is vitally important, this behaviour of flows has in general been avoided by most researchers in the field.

This thesis is an attempt to reveal some of the characteristics of the behaviour of truly unsteady flows. It addresses the flow in a particularly simple geometry - that within a smooth straight pipe of circular cross-section - in the hope that the fundamental effect of strong accelerations on flows of all types will be revealed within this simple framework.

1.1 Scope of the study

The accurate prediction of transition in fluid flows is not just an interesting exercise but affects almost every aspect of our lives. Turbulent flows show much higher heat-transfer rates than laminar flows, while the drag of a turbulent boundary layer is larger than one which is laminar. Aspects of both flows can be useful in practice. For instance, in the design of a high performance aircraft, control of boundary layer growth, prediction of transition and the reduction of turbulence are important considerations in the aerodynamic design. On the other hand strong turbulence is advantageous in promoting chemical mixing and even burning in the combustion process occurring in the same aeroplane's jet engine.

Steady flows

Study of steady flows over bluff bodies, and the development and control of boundary layers, is in the realm of 'conventional' fluid dynamics, and a multitude of commercially available numerical packages can be applied to these problems with good effect. The transition of steady boundary layers is also well understood in general, and good correlation between theory and experiment has been obtained for the entire transition process in the case of simple flows such as the Blasius boundary layer and parallel channel flows.

Considering turbulence: useful statistically based empirical models for strongly turbulent steady flows have been available for some time, and these give reasonable results for many flow systems. Despite their innate limitations (for instance the 'closure' problem of turbulence), successful results have been achieved using these methods in fairly complex domains, such as gas turbine engines.

Chapter 1

Overview

A central mystery of the universe still lies in the behaviour of fluids and how they flow. Water from a tap, the blood in animals' veins, and the interaction of nebulous intergalactic gas are vastly removed from each other in context and scale, yet all display a universality of behaviour that is remarkable to consider.

Fluids when flowing at low velocities, under calm conditions, maintain a structure that tends toward a simple equilibrium of the internal and external driving forces. Such flows exhibit a laminated or *laminar* structure, where subsequent layers of fluid slide smoothly past others, and where points in the fluid initially close to each other remain relatively close, or become distanced from each other in relative proportion to their separations. However, when a fluid is driven sufficiently far from equilibrium, a fundamental change in behaviour is observed within that fluid. At this critical point, the laminar flow of the fluid breaks down and a sinuosity of motion develops until a highly stochastic or turbulent flow regime persists. Unlike the laminar system, such turbulent flow shows chaotic motion of individual particles, and points in the fluid close to each other will quickly become distant. Such flows show no instantaneous uniformity in their motion as in laminar flows, but are only seen to ascribe to an average predictable behaviour.

The prediction of this point of transition between well-ordered laminar flow and the turbulent regime has been the subject of intensive research for many years. Despite the wealth of effort thrown at the problem, much of the fundamental behaviour of transition still remains unclear and only the general behaviour of a very few simple flows is well known.

A fundamental assumption made by most workers in the field of transition - both experimentalists and theoreticians - is one of steadiness. Studies are generally performed on flows whose mean values do not change in time (steady flow-fields), or where at the very least the mean flow variation at the boundaries is kept constant (steady boundary conditions). The steady assumption simplifies the mathematical description of these flows. Analytically it allows

Part I

INTRODUCTION

Greek symbols

α	wave number
	boundary layer thickness
δ^*	boundary layer displacement thickness
η	operational differentiation matrix
κ	parallel parameter (analysis)
	LDV beam half-angle (experimentation)
λ	wavelength of light
μ	dynamic viscosity
	operational multiplication matrix (analysis)
ν	kinematic viscosity
ψ	disturbance streamfunction
ρ	density
τ_w	wall shear stress
θ	azimuthal co-ordinate
Ω	acceleration parameter

Accents, super- and sub-scripts

The dummy variable z or Z is used in all cases below, to indicate the placement of the accent, superscript, or subscript. It does not denote any particular variable. The capital form is used in those cases where a capital letter is most expected.

z_{crit}	refers to conditions at the critical point on the stability curve
z_0	dimensional parameters
z	dimensionless variables,
	polynomial basis vector
$\mathcal{E}$	polynomial coefficient vector
$\hat{\mathcal{E}}$	collocation coefficient vector
$\hat{Z}$	matrix operators in the Chebyshev basis
z'	disturbance quantities
z_x, z_r	derivatives of z with respect to x, r respectively

Nomenclature

This nomenclature gives the meaning of the major symbols in this work. More specific symbology that is localised in scope, is properly introduced where it occurs, and is thus omitted here.

Roman symbols

c	celerity
f	frequency of LDV signal
k	exponential acceleration constant
p	pressure
r	radial co-ordinate
δr	offset of velocity profile origin from centre of pipe
Δr	radial displacement of LDV bench
$\Delta r'$	radial displacement of LDV control volume
R	pipe radius
Re	Reynolds number based on pipe radius and mean cross-sectional velocity
Re_δ	Reynolds number based on displacement thickness
s	axial displacement
S	velocity profile shape parameter
t	time
u	axial velocity
U, U_m	mean pipe cross-sectional velocity
U_c	pipe centreline velocity
v	radial velocity
x	axial co-ordinate
y	transverse co-ordinate (Cartesian system)

C.32	The correspondence between x and the other velocity profile parameters at discrete points in the pipe entrance, for the case $\Omega = 70$	528
E.1	The commands for the acquisition program.	578
E.2	The enhanced program control commands for macro use.	579
E.3	The run-file commands and their values.	579

C.20	The correspondence between x and the other velocity profile parameters at discrete points in the pipe entrance, for the case $\Omega = 10$	516
C.21	Variables and parameters at discrete points in the pipe entrance, for the case $\Omega = 20$. The underlined data are the values of the parameters at the minimum critical Reynolds number.	517
C.22	The correspondence between x and the other velocity profile parameters at discrete points in the pipe entrance, for the case $\Omega = 20$	518
C.23	Variables and parameters at discrete points in the pipe entrance, for the case $\Omega = 30$. The underlined data are the values of the parameters at the minimum critical Reynolds number.	519
C.24	The correspondence between x and the other velocity profile parameters at discrete points in the pipe entrance, for the case $\Omega = 30$	520
C.25	Variables and parameters at discrete points in the pipe entrance, for the case $\Omega = 40$. The underlined data are the values of the parameters at the minimum critical Reynolds number.	521
C.26	The correspondence between x and the other velocity profile parameters at discrete points in the pipe entrance, for the case $\Omega = 40$	522
C.27	Variables and parameters at discrete points in the pipe entrance, for the case $\Omega = 50$. The underlined data are the values of the parameters at the minimum critical Reynolds number.	523
C.28	The correspondence between x and the other velocity profile parameters at discrete points in the pipe entrance, for the case $\Omega = 50$	524
C.29	Variables and parameters at discrete points in the pipe entrance, for the case $\Omega = 60$. The underlined data are the values of the parameters at the minimum critical Reynolds number.	525
C.30	The correspondence between x and the other velocity profile parameters at discrete points in the pipe entrance, for the case $\Omega = 60$	526
C.31	Variables and parameters at discrete points in the pipe entrance, for the case $\Omega = 70$. The underlined data are the values of the parameters at the minimum critical Reynolds number.	527

C.9	Critical parameters at discrete times in the far-downstream impulsively started pipe flow, for the case $\Omega = 40$. The underlined data are the values of the parameters at the minimum critical Reynolds number.	505
C.10	The correspondence between t and the other velocity profile parameters at discrete times in the pipe entrance, for the case $\Omega = 40$	506
C.11	Critical parameters at discrete times in the far-downstream impulsively started pipe flow, for the case $\Omega = 50$. The underlined data are the values of the parameters at the minimum critical Reynolds number.	507
C.12	The correspondence between t and the other velocity profile parameters at discrete times in the pipe entrance, for the case $\Omega = 50$).	508
C.13	Critical parameters at discrete times in the far-downstream impulsively started pipe flow, for the case $\Omega = 60$. The underlined data are the values of the parameters at the minimum critical Reynolds number.	509
C.14	The correspondence between $\tilde{t}$ and the other velocity profile parameters at discrete times in the pipe entrance, for the case $\Omega = 60$	510
C.15	Critical parameters at discrete times in the far-downstream impulsively started pipe flow, for the case $\Omega = 70$. The underlined data are the values of the parameters at the minimum critical Reynolds number.	511
C.16	The correspondence between t and the other velocity profile parameters at discrete times in the pipe entrance, for the case $\Omega = 70$).	512
C.17	Variables and parameters at discrete points in the pipe entrance, for the zero acceleration case ($\Omega = 0$). The underlined data are the values of the parameters at the minimum critical Reynolds number.	513
C.18	The correspondence between x and the other velocity profile parameters at discrete points in the pipe entrance, for the case of zero acceleration ($\Omega = 0$).	514
C.19	Variables and parameters at discrete points in the pipe entrance, for the case $\Omega = 10$. The underlined data are the values of the parameters at the minimum critical Reynolds number.	515

17.1	The value for the minimum critical Reynolds number based on displacement thickness Re_{δ^*} for various orders of solution and domain truncation, for the parallel impulsively started boundary layer.	404
18.1	The values of the parameters Re_0 , Ω of the experimental flowrate variations for the separate tests.	429
C.1	Critical parameters at discrete times in the far-downstream impulsively started pipe flow, for the zero-acceleration case ($\Omega = 0$). The underlined data are the values of the parameters at the minimum critical Reynolds number.	497
C.2	The correspondence between f and the other velocity profile parameters at discrete times in the pipe entrance, for the zero acceleration ($\Omega = 0$) case.	498
C.3	Critical parameters at discrete times in the far-downstream impulsively started pipe flow, for the case $\Omega = 10$. The underlined data are the values of the parameters at the minimum critical Reynolds number.	499
C.4	The correspondence between f and the other velocity profile parameters at discrete times in the pipe entrance, for the case $\Omega = 10$	500
C.5	Critical parameters at discrete times in the far-downstream impulsively started pipe flow, for the case $\Omega = 20$. The underlined data are the values of the parameters at the minimum critical Reynolds number.	501
C.6	The correspondence between f and the other velocity profile parameters at discrete times in the pipe entrance, for the case $\Omega = 20$	502
C.7	Critical parameters at discrete times in the far-downstream impulsively started pipe flow, for the case $\Omega = 30$. The underlined data are the values of the parameters at the minimum critical Reynolds number.	503
C.8	The correspondence between f and the other velocity profile parameters at discrete times in the pipe entrance, for the case $\Omega = 30$	504

12.1	Variation most unstable eigenvalue with order in plane Poiseuille flow, for $\alpha = 1$, $Re = 3\,000$, compared to results generated by explicit transformation of operators into the orthogonal basis. Only the imaginary component is shown.	219
13.1	The convergence of the minimum critical Reynolds number with order for the present investigation.	270
13.2	Maximum accuracy attained for the most unstable eigenvalue at $\alpha = 1$, $Re = 10\,000$, as compared to the results of Orszag (1971a).	270
13.3	Comparison between present results for the most unstable eigenvalue with $\alpha = 1$, $Re = 10\,000$, and those from Orszag (1971a), with N varied to show the rate of convergence.	271
13.4	The variation of the most unstable eigenvalue for plane-Poiseuille flow with $\alpha = 1$ and increasing Re , compared to the results of Gersting (1981) and da Silva (1990).	273
13.5	The values for the constants a_1 for the the various temporal discretisation schemes.	277
13.6	Variation of centreline velocity in impulsively started pipe flow far from the inlet, showing also the values obtained using da Silva's (1990) finite-difference algorithm.	282
13.7	Variation of parameters in impulsively started pipe flow far from the inlet.	288
13.8	The truncation error for the Blasius solution, for various values of the domain size k and order N . *, to the level of the round-off error; (), not yet properly representative of the solution.	299
13.9	The value for the minimum critical Reynolds number based on displacement thickness Re_{δ^*} for various orders of solution and domain truncation, for the Blasius boundary layer.	301
14.1	The entrance length for a pipe inlet Z_c as determined by various authors.	317
15.1	The variation of the final displacement thickness δ_{∞}^*/R with acceleration parameter Ω , given to 9 significant figures.	338

2.3.2 Pipe flows

Far less research has been undertaken into pipe flows. They are seemingly analogous to flows between parallel plates (both will develop parabolic laminar profiles far downstream in the absence of transition), but whereas the planar flow has a well defined minimum critical Reynolds number for the parabolic profile, of 5772.2218 [Orszag (1971); A. Davey, unpublished - Drazin & Reid (1981)] and its eigenvalue spectra are well defined, the pipe-Poiseuille flow shows no such finite stability limit. The pipe inlet flow is boundary layer-like and behaves in a similar way to a flat plate boundary layer in a pressure gradient. Studies of the pipe system have thus resolved themselves into two branches: the studies of boundary layer-like behaviour within the pipe inlet, and research into the behaviour of the Poiseuille limit of the flow.

Poiseuille flow in a pipe.

Unfortunately much (with hindsight, misdirected) effort has been used to analyse the pipe-Poiseuille stability case, which has failed to show any instabilities to *any* small disturbances. The body of work directed to investigating this flow has been substantial [Sexl (1927), Pekeris (1948), Sexl & Spielberg (1958), Corcos & Sellars (1959), Gill (1965), Lessen, Sadler & Liu (1967), Davey & Drazin (1969), Salwen & Grosch (1972), Garg & Rouleau (1972), Burrige (1972), Metcalfe & Orszag (1973), Salwen, Cotton & Grosch (1980) and Schmid & Henningson (1994)], and except for a few flawed analyses [Graebel (1970), Munakata (1973, 1983)] which did find stability limits, all studies have proved categorically that Poiseuille flows are in fact unconditionally stable to infinitesimal perturbations. This conclusion has been confirmed experimentally by Leite (1958). Attention was initially focused on the subject quite possibly because of the temptingly simple description of the base flow - $u = 2U_m[1 - (r/R)^2]$, combined with the fact that the Cartesian counterpart of this flow is linearly unstable. The subsequent excessive effort is probably due to, as Morkovin & Reshotko (1989) put it, *... a mathematicians myth that instability and transition of interest take place in established parabolic mean flows* [in a pipe]. In fact turbulent flow far downstream in a steady pipe flow is directly attributable to transition in the inlet and its subsequent convection downstream.

Finite-amplitude studies of these flows have been made, but as with the linear studies, their relevance to the physical description of transition in these flows is questionable. Disagreement exists even here: studies by Sen *et al* (1985) gave finite-amplitude equilibria for this flow, while Patera & Orszag (1981) showed none. In reality, this profile needs the superposition of locally inflectional regions with magnitude of up to 7% [Morkovin & Reshotko (1989)] of the centreline velocity in order to become unstable.

result in any improvement. Subsequently a careful consideration of the existing evidence - both experimental and numerical - led to a better understanding of the various stages of the process, identifying the important elements. The basic steps of the transition process in boundary layer flows are summarised well by figure 2.1, and are:

- (i) Receptivity [Goldstein (1985), Itoh (1986), Tam (1986)] of the boundary layer to disturbances, and the subsequent slow linear growth of these disturbances as Tollmien-Schlichting (T-S) waves. This step comprises the primary instability of the flow
- (ii) The T-S waves develop three-dimensional characteristics at the start of a non-linear interaction with the flow, and deform onto the Λ vortices seen in both numerical work and experiment. This process is well summarised by Herbert (1985, 1988) and Morkovin & Redhotko (1989). A number of arrangements of these vortices have been observed - the staggered C' and II types, and the K type - and the agreement of the numerical work with experiment is remarkably good. This part of the flow development is usually termed secondary instability, but is still far from turbulence.
- (iii) Understanding of the *tertiary* flow breakdown has only begun recently, but concord has been reached between the experimental work of Klebanov *et al* (1962) and numerical simulations. Basically, high shear layers around the Λ vortices roll up into vortices themselves, resulting in a sustainable process that leads to early stages of a wall-bounded turbulent flow. Simulations of this process were first made by Gilbert (1988), and an excellent summary of the current understanding can be found in Kleiser & Zang (1991), while Sandham & Kleiser (1992) present results 'extracted' from Gilbert's (1988) numerical experimentation.

Many of these analyses have drawn upon the development of, not only powerful computers, but also numerical procedures such as the so-called spectral (Fourier) method [Canuto *et al* (1987)].

In general the major aspects of boundary layer transition are well understood. In summary, linear analyses are still valid for the early stages of transition, while excellent understanding of the later stages of transition have been attained by the use of direct methods rather than by the rather limited weakly non-linear theories. The early emphasis placed on non-parallel effects in high Reynolds number boundary layers has been shown to be inappropriate, these effects showing themselves to be wholly insignificant. This concludes the broader consideration of boundary layer research in general, and the following section deals with literature of relevance to steady pipe flows.

There are two general types of non-linear approach. The first type essentially involves expansions about the perturbation terms of the linear theory, allowing the solution of simplified equations to determine the threshold amplitudes and growth rates of non-linear disturbances. The second approach is simply direct numerical experimentation, simulating the flow using the full Navier-Stokes equations and 'observing' the results. The latter method has only recently become feasible, with the introduction of powerful computers. The former approach has however been followed for some time, and consists of a well developed theory.

Early non-linear studies were by Bohr (1909), Noether (1921) and Heisenberg (1924), although the fundamental concepts of non-linear analysis were only later laid by Landau (1944) [see Landau & Lifshitz (1959), Drazin & Reid (1981)]. The basic approaches were consolidated by Stuart (1960) and Watson (1960), and are generally termed *weakly non-linear* theories. These methods all are concerned with the expansion in amplitude of disturbance modes, described by the *Landau* equation. A detailed synopsis on this approach is contained in Drazin & Reid (1981).

The weakly non-linear theories are only approximations to the behaviour of the real flow, and thus must be used with circumspection³. Inappropriate choice of the Landau coefficients can lead to incorrect results, and thus care must be taken in the application of these methods. However, powerful insights have been gained by their use. For instance in plane-Poiseuille flow, substantial agreement has been obtained between non-linear studies and experiment. Both Pekeris & Shkoller (1967) and Reynolds & Potter (1967) detected sub-critical instability in this flow, and later both Chen & Joseph (1974) and Itoh (1974b) confirmed these results. The experimental results of Nishioka *et al* (1975) showed good agreement with the theoretical results above, both in terms of the amplitude and frequency of disturbance, and with the spatial growth rates of the linear [see Itoh (1974a)] and weakly non-linear modes.

Stronger non-linear effects can be accounted for in the above methodology by including extra terms in the expansions, governed by higher order Landau coefficients. There are complications in such approaches, but successful solutions have been obtained by many authors [Joseph & Saffinger (1972), Herbert (1980), Sen & Venkateswarlu (1983)].

Direct studies of the secondary instabilities leading to turbulence were initially entirely experimental [Klebanoff *et al* (1962), Kachanov & Levchenko (1984), Kozlov & Ramazonov (1984), Barros & Liu (1984), Saric & Thomas (1985)]. The weakly non-linear theories above could not conclusively describe the behaviour of the non-linear disturbances [see Herbert (1984)], and extensions to the basic concepts of Landau [e.g., Resonant wave interactions (Triuk (1985))] did not

³In comparison, the linear theory is exact in the limit of infinitesimal disturbances, and in fact does govern the real growth of measured disturbances until a fair amplitude is reached [see Kachanov (1987)].

Non-parallel effects

The discrepancies between the experimental and numerical results were initially addressed using non-parallel analyses. Many such calculations were carried out in the seventies and early eighties. Initial attempts were based on expansion procedures [for example Barry & Ross (1970), Volodin (1973), Ling & Reynolds (1973)], and were only successful in lowering the minimum critical Reynolds number to 500. Bouthier (1972, 1973) appeared to have resolved the discrepancy, when he presented results in substantial agreement with Ross' experimental work, not only in terms of the minimum critical Reynolds number but with regard to the alignment of the entire neutral stability curve. Gaster (1974), in a profound study, expressed caution about many of the methods used to infer instability in non-parallel methods. He presented a different approach and calculated a more conservative minimum critical value of 480. There the matter stood, and consequently a number of further approaches were attempted [Sarie & Nayfeh (1975), Smith (1979a), Van Stijn & Van der Vooren (1983), Bridges & Morris (1987)] which all yielded results scattered somewhere between the lower limit of Bouthier and the upper one of Gaster. In a fundamental study by Fasel & Konzelmann (1990), a numerical simulation of the problem was presented, vindicating Gaster's (1974) result² and discrediting the Bouthier result, showing it to be fortuitous. This paper presented sound reasons for the disagreements between the authors and concluded that non-parallel effects were in fact *not* important in the flat plate boundary layer flow. They ascribed the discrepancy with experiment to possible errors in the experimentation itself (pressure gradients, etc.) and called on further experimental work. This paper represented a landmark in the area of flat plate stability, and brought to a halt the perennial disagreements associated with the particular flow.

Finite amplitude effects and transition

The previous analyses were associated with the point of (parallel and non-parallel) linear instability in the pipe. The subsequent development of these small perturbations towards non-linear finite-amplitude states and finally to turbulence has been the subject of vigorous research.

The growth and interaction of the non-linear terms results in the richness of behaviour of finite amplitude flows. Non-linearity leads to unusual behaviour. For instance, flows unstable to the linear disturbances can become stable at some finite amplitude threshold (*super-critical stability*), or linearly stable flows may become unstable when subjected to a certain level of finite-amplitude disturbance (*sub-critical instability*).

² © 1990 Cambridge University Press. Printed in the United Kingdom. All rights reserved.

²It is interesting to consider the little-known addendum to Drazin & Reid (1981) [1991 reprint], which refutes some of their own claims that they make in that book, for the case of non-parallel effects being dominant in the flat plate boundary layer. They also vindicate Gaster's (1974) analysis.

2.3 Steady transition

In the case where bypass transition does not occur, and where environmental disturbances are small, initial growth of disturbances in a flow is initially most certainly due to the linear growth of small perturbations, as described by linear stability theory. Later development of these perturbations are non-linear. Other effects such as the non-parallel nature of the flow may have an influence on the stability behaviour of the system. This section considers the large body of literature concerning these different aspects of transition to turbulence in steady flows. The main categories of flows considered here are flat plate flows, and general pipe flows. Some of the literature for general channel flows that is of relevance is also discussed.

2.3.1 The flat plate boundary layer

The flat plate system is useful in that it is the asymptotic limit of pipe flows toward the inlet [da Silva & Moss (1994), Abbot & Moss (1994)]. The basic velocity profile has a well known analytic solution, as determined by Blasius (1908), and many solutions of the resulting non-linear equation have been obtained using various numerical techniques [Blasius (1908), Töpfer (1912), Dairstow (1925), Goldstein (1930), Howarth (1938)]. Experimental study of these boundary layers were carried out by Burgers (1925), van der Hege-Zijnen (1924), and later Nikuradse (1942). The latter author accounted for effects due to the leading edge of the plate or from the presence of small pressure gradients along the plate, and obtained measured velocity profiles in remarkable agreement with the theory.

The Blasius profiles are self-similar, and thus a linear stability analysis is simply concerned with a single velocity profile. Detailed experimental measurements of stability of these flows were made by Schubauer & Skramstad (1947), and later Ross *et al* (1970) carried out an extensive set of measurements which gave a minimum critical Reynolds number based on displacement thickness of about 400, and also defined the entire stability curve to a high level of detail.

Early calculations of the minimum critical Reynolds number of this flow were made by Tollmien (1929), Schlichting (1933), Liu (1945) and Shen (1954), and agreed with the experimental results of Ross. Unfortunately, more accurate calculations worsened the correlation: the solution of Jordinson (1970) raised the value to 520, while a definitive solution by A. Davey [unpublished, see Drazin & Reid (1981)] determined the exact value to be 519.060 for this limit.




Figure 2.1: The general nature of transition in wall bounded shear flows [from Morkovin & Reshotko (1989)].

In pipe and other similar wall bounded flows. By-pass transition is usually seen in environments where disturbances are large and the 'traditional' routes are skipped. This route to turbulence is possibly important downstream in the current flow, where destabilisation by large convected disturbances might occur.

The above sections were concerned with a global overview of instability and transition, including the history of the subject. The remainder of this chapter is concerned with detailed literature reviews in specific areas of interest for the thesis.

equations are perturbed by a general quantity and then linearised, were first applied to fluid systems by Stokes, Kelvin and Rayleigh. These analyses were difficult, thus investigations concerning the behaviour of simpler classes of flows were undertaken. The study of the instability of ideal (inviscid) fluids was investigated by Helmholtz (1868), Kelvin (1871) and Rayleigh (1880). Here destabilisation of the flow was by purely inertial effects. A fundamental and sufficient condition for instability was discovered by Rayleigh (of which Reynolds was apparently unaware). He showed that a point of inflection was a sufficient criterion to induce normal-mode or linear instability in a flow.

It was only much later that any significant gains were made in the understanding of the viscous problem. Orr (1907) and Sommerfeld (1908) developed the acclaimed viscous stability equation - the Orr-Sommerfeld equation, which fully describes the linear stability behaviour of parallel shear flows. This remarkable equation has been the subject of investigation until the present day. Heuristic methods of approximation were used to obtain approximate solutions to this fourth order differential eigenvalue equation, but with the advent of numerical computers, accurate solutions for certain types of flow systems have become available. Liu (1955) contains an excellent summary and description of linear stability.

2.2.2 Secondary instability and transition

The linear instability mechanism discussed above is concerned with the starting point of a series of changes of state, or transitions, of the flow, possibly leading to turbulence. Very recently, success has been achieved in the numerical simulation of most of the transition process. The mechanisms and processes involved have been discussed by, among others Bayly *et al* (1988), and Sandham & Kleiser (1992). The latter authors have presented detailed descriptions of the complicated events leading finally to turbulence in the case of channel flows. Thus in general the mechanisms of the full transition process to turbulence from infinitesimal disturbances are quite well understood. Herbert (1988) presents a thorough review on the subject which is still relevant today.

2.2.3 Bypass mechanisms

A lot of work has been done on what are termed bypass mechanisms in fluid flows, where the gradual sequential breakdown of the flow is not seen, but rather transition to turbulence occurs suddenly. A summary of the general transition process, is given in figure 2.1.

This figure serves as a good description of the current understanding of transition as a whole. The classical gradual route described in section 2.2.2 is shown in the middle of the figure. Consideration of the left-hand path in the diagram is necessary in the current context, because it is a possible route to turbulence

the non-linear and finite amplitude studies in this chapter; not because they are unimportant¹ but rather because the emphasis of this thesis is on the fundamental nature of accelerating flow. Although a finite-amplitude study of strongly accelerating flows would be of great benefit to the subject, it is a topic for later work; for now it is more important to establish the *basic* behaviour (mathematical framework, base flow variation and linear stability) of these unknown flows before tackling the more complicated phenomena such as the later stages of transition.

For this reason the discussion on finite amplitude effects given here is more of an overview of the subject, discussing the phenomena and mechanisms rather than the methods of solution, and references are given to other more detailed descriptions. Although an analysis will not be carried out, the experimentation will certainly be subject to finite-amplitude effects; thus knowledge of the mechanisms is invaluable for the experimental part of the study.

2.2 General overview

Initial investigations into the nature of transition were made in the 19th century, notably by Kelvin, Rayleigh and Reynolds. Reynolds (1883), in the now famous series of pipe transition experiments, studied the nature of the transition process in great detail, and many of the features noted by him are still the source of investigation today. Reynolds' contribution to the field was acknowledged by naming the law of similarity of fluid flow he discovered after him. This law states that the behaviour of fluids is governed by the ratio of the viscosity to the mean velocity and characteristic dimension (i.e. $(Ua)/\nu$) alone, and that fluids at the same Reynolds number would behave in the same way. He discovered the existence of a *critical* Reynolds number above which transition to turbulence occurred in the pipe, and that this value increased as the level of disturbance in the inlet flow decreased. Reynolds measured a critical Reynolds number of 2300, based on diameter. Many studies [cf. Barnes & Coker (1905), Schiller (1922)] have since verified the existence and behaviour of the critical Reynolds number. In the case of pipe flows Ekman (1910) maintained laminar flow up to a value of 40 000 based on diameter, using very smooth pipes and carefully tailored inlet conditions. On the other hand numerous experiments (see Schlichting (1979)) making use of roughened pipes and disturbed inlet conditions have resulted in the minimum value of this parameter being determined to be about 2 000.

2.2.1 Linear stability

Methods for analysing the onset of instability have been available since before Reynolds' famous experiments. The method of *normal modes*, where the flow

¹These mechanisms are indisputably and fundamentally important in any transition process by definition.

Chapter 2

Literature overview

Consideration of all the studies in the literature *directly* involved with the subject of non-periodic unsteady transition would yield only a small number of works. While these are of crucial importance to the current study, the consideration of the broader aspects of transition in wall bounded flows is needed, since any fundamental study of accelerating flows is an extension of, and will contain aspects of, the behaviour of already known and widely studied systems.

This chapter first presents a general introduction, where literature concerning the history of the study of fluid mechanics is briefly cited. Thereafter studies into related subjects concerning instability and transition in various important flows are given. Finally literature of direct interest is cited, and its important aspects discussed. The literature survey presented here is in effect a broad overview of the entire field of study, and detailed literature reviews are given throughout this work, in a more local context.

2.1 Scope and emphasis

This chapter presents literature of importance to the study of transition in flows in general. As this represents the background against which the current study developed, the relative emphasis of the topics in this overview reflect the emphasis of the study. It is primarily concerned with the *fluid mechanical* aspects of the study. Detailed surveys of numerical techniques or LDV measuring methods are restricted to the chapters in which they appear, to avoid cluttering.

Strong emphasis is placed on the factors affecting the linear stability of flows in general, since a major portion of the study concerns instability. The various results for different types of flow system are important, because they may represent limiting behaviours of the current study. Another important reason for a broad survey *around* the region of interest is to highlight the acute lack of knowledge about accelerating flows. Only superficial attention is given to

gained from the analysis, and comparison between the two are made. The conclusions [chapter 20] and recommendations [chapter 21] conclude this component of the work.

1.3.2 Presentation of raw data

Modern developments in data acquisition and the tremendous increase in computer power in recent years led to huge quantities (≈ 500 megabytes) of raw data being gathered from the current study. With such quantities of data it is inappropriate to present the basic raw information, because it would span many volumes. Thus, in keeping with modern developments, this data is presented here in digital form, as a CD-ROM included in appendix G. This allows full 'closure' of the thesis while not sacrificing compactness. Thus the so-called 'raw data' that is presented graphically in the appendix is in fact semi-processed information, much reduced in quantity from the actual raw data; and is given for the sake of convenience and quick reference.

1.3.3 Data processing and modelling tools

It is customary for the applied mathematician to use standard programming languages, such as FORTRAN or C, to develop the numerical simulation code of their study. Conversely, in the realm of data processing, use is usually made of spreadsheet packages (Lotus 123, Microsoft Excel), or of stand-alone code designed specifically for the purpose. However, in the present study, a single tool was available that embraced all the above facets of experimentation and analysis, and as such it deserves mention.

Both the solution of numerical problems, and the processing of the experimental data, was achieved using a fourth generation matrix language called **Matlab**. This product is a symbolic matrix tool with multifarious uses. Its efficient and powerful matrix manipulation engine allows simple coding of, and solution to, the numerical problems in this work. On the experimental processing side, large sets of data could be accommodated with ease, while the extensive numerical tool kit simplified data processing considerably. Its platform independence allowed processing to continue on a desktop PC, or by using a powerful UNIX computer. Matlab is available from The Mathworks, Natick, MA (USA), or they can be contacted via e-mail on the internet at support@mathworks.com.

The introduction

This is the current part of this work, including this chapter and the general literature survey of the following chapter. It is concerned with setting the context and justifying the study.

The experimentation

This section starts with a discussion of the LDV measurement method [chapter 3]. Thereafter, the development of the experimental apparatus is given, and design and construction of both the physical (flow rig) and control elements of the system is presented [chapters 4 and 5].

The preliminary experimentation follows in chapters 6 and 7. Final experimentation and the resulting data are then given in the final two chapters [chapters 8 and 9]. The results of the final chapter of this part are interesting, and leading from this it is natural to seek a mathematical explanation for their appearance. This leads to the analysis.

The analysis

The first chapter [chapter 10] of this part is concerned with addressing the mathematical problem of the unsteady flow. This is necessary due to the lack of a unified basis for analysing the problem of unsteady fluid behaviour. The next 2 chapters [chapters 11 and 12] are concerned with the development of the numerical methodology needed to solve the derived relations. The method is verified by applying it to various prototypal fluid dynamics problems in chapter 13, allowing comparisons to be drawn to existing results. Finally, the solutions to the main problems at hand are presented in chapters 14, 15, 16 and 17.

The overriding philosophy of approach in this work was to verify as much as possible every tool used to understand the problem. Because of the many unknown factors in the research, testing of the tools on known problems was taken as being very important. Confirmation of the accuracy and applicability of the chosen methods allowed confidence in the results obtained when these tools were applied to the unknown problems.

The synthesis

The two preceding parts - experimentation and analysis - were essentially parallel investigations of the same phenomena. The first two chapters of this part [chapter 18 and 19] re-examine the experimental data in the light of insights

- (vi) Agreement of the experimentation with theory is important as a verification of the veracity of both. Disparities would indicate deficiencies in either the measurements or analysis, thus prompting further research in search of the cause. Since there were two elements to the numerical study, there were two points of comparison: the base flow variation and the occurrence of instability in each system. It was decided to design the basic experimentation in order to allow extraction of the two 'kinds' of data, in order to allow easy comparison.
- (vii) In the case of a lack of substantial agreement between theory and experiment, a likely cause needed to be found. A lack of agreement would indicate other unknown mechanisms to those considered in the analysis, thus indicating the way forward for further study.
- (viii) In general, it was hoped that the accelerating results would be extensions of the simpler impulsive pipe flow case. Various transitional events were identified in these flows by both Moss (1985) and da Silva (1990), and the occurrence of similar phenomena in the accelerating flow was to be investigated.
- (ix) The verification of the shape parameter as a means of uniquely quantifying the velocity profile [da Silva (1990)] and stability behaviour in the extended case of accelerating flows was to be investigated. The use of other parameter (e.g. displacement thickness) for flow quantification was to be considered.

1.3 Structure and philosophy

The general layout of this thesis fell naturally into four parts: the introduction, the experimentation, the analysis, and finally the synthesis. The arrangement and general content of these parts are discussed here, outlining the reasons for their chosen arrangement. Thereafter, specific issues of interest concerning the structure are discussed.

1.3.1 Layout of the thesis

Both the experimentation and numerical studies showed phenomena which, in isolation were fairly novel, and the ordering of this thesis represents an optimum way of explaining these occurrences. Broadly speaking, the parts of this work are as follows:

3.2.4 Signal processing

The signal from the photomultiplier, as shown in figure 3.2, needs extensive processing to extract a frequency value proportional to the velocity. As it stands in the figure this signal consists of a low frequency pedestal combined with a high frequency superimposed wave. Figure 3.4 shows the Fourier components of the signal from the former figure.




Figure 3.4: The signal from the anemometer photomultiplier shown in the frequency domain.

Two basic methods of processing this signal to attain the frequency are available: measuring fringe spacing in the time domain or position of a signal peak in the frequency domain. The ability to process the LDV signals using the latter principle is a recent innovation in the field. Only units based on the former principle were available, and as a result only the former method of signal processing will be discussed below.

Simply put, the incoming signal (figure 3.2) is subjected to a high-pass filter, which removes the pedestal component from the signal. Next a low pass-filter is applied which removes any high frequency noise from the signal. This sequence of events is demonstrated in figure 3.5. A threshold level is applied to the resulting signal, and the frequency at which the signal exceeds this level for a set number of cycles is registered. This frequency can then be converted into a velocity value using equation (3.1).

In the current system this process was carried out by a counter supplied with the velocimeter equipment. The filter and level settings could be adjusted via dials on the front panel, and the frequency output was available via a digital connector on the front face.

In order to measure strong signals from low power laser beams, the receiving optics can be placed on the far side of the flow, aligned (on-axis) or misaligned (off-axis) with the transmitting optics. It is preferable to have the receiving unit slightly misaligned from the primary beams² to increase the signal-to-noise ratio. This arrangement is termed a *forward scatter* system.

An alternative to this is the *back scatter* arrangement, where the receiving unit is placed on the same side of the flow as the transmitting components. The most common layout is for the receiving components to be incorporated into the transmitting optics. Here, reflected light is passed back through the transmitting lens, reflected off a mirror (the transmitted beams pass this mirror on either side) and focused onto the photomultiplier. This method requires higher power transmitted beams for the same signal power as the forward scatter method, but results in an integral arrangement.

A major advantage of the on-axis back-scatter unit is that the reflected light follows much the same optical path as the incident beams, thus optical refraction of the beams through fluid-wall interfaces is 'cancelled out' during the return path of the refracted light, and thus the receiving unit remains constantly focused on the measuring volume. This is not generally the case with a forward scatter unit: movement of the transmitting optics usually results in the misalignment of the receiving unit.

3.2.3 Seeding particles

The generation of a Doppler signal requires that a seeding particle pass through the measuring volume. Seeding particles in the flow are either present naturally (dust), or must be added to the fluid. The prerequisites for a seeding particle are:

- (i) They should be highly reflective, refracting or reflecting the incident light strongly.
- (ii) They should follow the underlying fluid motions as closely as is possible. Thus they should possess a similar density to the fluid.
- (iii) Their size needs to be equivalent to the fringe spacing in the measuring volume. Pingerson *et al* (1989) contains an excellent discussion on particle sizing.

In practice, other effects, such as toxicity need to be considered in the choice of particles. In the current system, which used water as the working fluid, evenly sized (0.5 μm) silicon carbide particles gave excellent signal properties.

²The primary beams passing out through the back of the flow do not carry the velocity information, and are usually allowed to disperse on a far surface. Their interaction with the photomultiplier will simply swamp the signal carried by refracted light.

3.2 Aspects of the LDV method

The above description of the principle whereby a dual beam LDV operates was by its nature brief, and the user is referred to any of the extensive number of reference works available on the subject [e.g. Dürst *et al* (1976), Drain (1977), Buchhave *et al* (1979) Fingerson *et al* (1989)]. In the sections that follow below, specific aspects of the measurement instrument and procedure are considered, especially those pertinent to the current experimentation. A general schematic layout of a typical dual beam instrument is shown in figure 3.3, illustrating the various principal components.




Figure 3.3: A schematic representation of a dual beam, back-scatter LIDV, showing the principal components [from Fingerson *et al* (1989)].

3.2.1 Transmitting optics

The function of this sub-assembly is simply to split a laser beam into two equal strength components and then to allow them to cross. Beam splitting is usually achieved using a prism arrangement, and a precision ground lens is used to focus the beams at a point. Other components, such as beam expanders, polarising filters and the like, can be incorporated, but are not essential to understanding the basic operation of the system.

3.2.2 Receiving optics

Scattered light from particles in the beam crossing point or measuring volume is focused through a lens arrangement, and onto the aperture of a photo-detector, where the optical signal is converted into an electric signal. Since light is scattered randomly by passing particles, the direction from which this light is gathered can be changed at will. The point from which the light is gathered has a strong effect on the strength of the measured signal.




Figure 3.1: The beam crossing point in a dual beam LDV system, illustrating the measuring volume and fringes [from Fingerson *et al* (1989)].

which the particle passes through the bright fringes is directly proportional to the vector component of velocity perpendicular to the plane of the fringes. If the refracted signal is measured using a photodetector, then a wave pattern like that shown in figure 3.2 will be measured. Here the frequency of the wave pattern in this signal is directly proportional to velocity, and is given by the relation

$$U = \frac{f\lambda}{2\sin(\kappa)}, \quad (3.1)$$




Figure 3.2: The measured signal from a dual beam LDV.

Here f is the frequency of the signal in figure 3.2, λ is the wavelength of the laser light, κ is the beam half-angle (see figure 3.1) and U is the vector component of velocity through the measuring volume, perpendicular to the plane of the fringes.

The LDV method on the other hand merely intersects two beams of laser light within the flow of interest, simply requiring optical access to this flow, thus satisfying the requirement of non-intrusiveness¹ where the hot-wire probe fails. The temporal and spatial resolution of the LDV are equivalent to the hot-wire probe [Teitgen (1975)], although Gaster (1989) expressed doubts of its ability to resolve fluctuations below the order of 1% of the mean velocity. Further to this, it is a completely linear device and is thus superior to the hot-wire system in that it requires no calibration. The historical development and principle of operation of the LDV are both interesting and elucidating for understanding its strengths and limitations, and are briefly given below.

3.1 Principles of operation

The laser velocimetry method was initially developed by Cummins, Kuablie & Reh (1964), and fast became popular because of its accuracy and non-intrusive nature. The early systems were what is known as *reference beam* systems: laser light was split into two unequal strength beams, the more powerful of which was directed through the flow while the second weaker beam was directly incident on a photo-detector. Light from the primary beam, refracted from small particles moving with the flow, was also incident on the photo-detector, and the small difference in frequency due to the Doppler effect was measured as optical 'beating' at the detector, the frequency of the beat being directly proportional to velocity.

The reference beam set-up was unwieldy to use, and this resulted in the development of the *dual beam* system. Here, a single laser source is split into two equal strength laser beams, and these are intersected in the flow at the point of interest. At this intersection, fringes or standing waves are set up, due to the constructive and destructive interference of the coherent beams of light. These bright and dark fringes take the form of planes equally spaced throughout the laser crossing point or measuring volume. Although the dual beam set-up does not employ the Doppler effect, the name LDV (or LDA - the 'A' for anemometer) has remained. Figure 3.1 illustrates the beam crossing and the appearance of the measuring volume.

The spacing of the fringes shown in the figure is dependent only on the beam intersection angle and the wavelength of the light. Small particles in the flow, of the size of the fringe spacing or smaller, will, as they pass through this measuring volume, refract light whose intensity varies from bright to dark according to the particle being positioned within a bright or dark fringe. Clearly the *rate* at

¹Interestingly enough, a series of experiments by Cooper *et al* (1985) on instabilities in unstably Taylor-Couette flow showed that the intrusion of the laser caused local transition nearby. The most probable explanation for this was that, because this was a closed system, a ring of fluid was continually being heated by energy from the laser in the time before transition occurred, thus modifying the fluid properties locally and initiating the transition process. In a high velocity, open system (pipe flow), no repeated heating occurs and the effect is negligible.

Chapter 3

Laser Anemometry

The primary flow measurement device used in the experimentation was the Laser Doppler Velocimeter or LDV, selected for use of the innate advantages it holds for the accurate measurement of high, variable flows. In order to best understand the advantages and limitations of the LDV, it is necessary to give a brief description of its principle of operation. Specific points relating to the use of the method relevant to the current study are also discussed.

An ideal analytical flow measurement device is one that has the following broad characteristics:

- (i) It must cause the minimum possible disturbance to the flow being measured.
- (ii) The spatial extent of the measurement device must be small enough so that it can resolve the small-scale spatial structure of the flow.
- (iii) The inertia of the unit must be low enough so that the device can track accurately the temporal variation of the flow, i.e. it must possess a high frequency response.
- (iv) The device should show stability in its readings, so that repeated calibration is not necessary.
- (v) The signal-to-noise ratio of the device should be high enough so that small scale fluctuations relative to the mean velocity can be resolved.

Devices commonly used in the past, such as the pitot tube, do not satisfy all the points listed above. For instance such a device is by its nature intrusive. The sensitive measurements required for transitional flows have usually necessitated the use of the hot-wire anemometer. This device has high temporal and spatial resolution and good signal-to-noise ratio, but is also intrusive and requires complicated calibration procedures because of the strong non-linearity of its output signal.

Preamble

This part of the thesis presents the experimental investigation into the general behaviour of strongly accelerating flows. The initial chapters describe both the development of the necessary infrastructure, and the discussion of the method of measurement. In particular, these concentrate extensively on both the physical test-rig development, and the acquisition and control software system. The preliminary testing is then described, leading to some initial results. By considering these preliminary data, the development of the data acquisition and processing methodology is shown. The measured results are presented in the final chapters of this part.

Part II

EXPERIMENTATION

Boathier (1972) gave an asymptotic method, applicable only in the limit of very low accelerations.

Moss (1991) presented an integral method for the study of the variation of laminar flow in pipe flow far from the inlet. He defined an acceleration parameter $\alpha = \frac{1}{U} \frac{dU}{dx}$ which arose naturally from his analysis and was identical to the stability parameter of Shen (1961). He discovered that for increasing acceleration parameter as he termed it, the velocity profiles did not approach the parabolic limit but retained a lower centreline velocity and higher wall shear stress. The crude nature of his integral approach, combined with use of the boundary layer thickness δ as a parameter, did however introduce an error into his results, leading him to notice a discontinuity at the acceleration where the boundary layer merged. Nevertheless this important paper introduces some clear indicators of the expected trends in accelerating laminar flows.

Experimentally, van de Sande *et al* (1980) measured velocity profiles and points of transition in a pipe flow driven by a constant pressure gradient. Although not a highly unsteady flow, it nevertheless strictly belongs in this category of flows, rather than those of section 2.1.1. These results showed an increase of the transition Reynolds number with acceleration.

Lefebvre & White (1989, 1991) presented experiments on constant acceleration pipe flow, and noted global transition everywhere in their pipe, similar to the initial transition of Moss (1985). They noted a monotonic increase of the critical Reynolds number with acceleration. They only took measurements at a single far-downstream portion of the pipe, and only presented timing correlations for their data. Their accelerations were substantial, and their highest observed transition Reynolds number was 1.1×10^6 . They achieved good correlation between the results from their two sets of tests, each conducted using a different diameter test section, indicating a lack of rig-independence of their data.

were identical in structure. da Silva noticed an increase in his measured critical Reynolds number with acceleration. He was not interested in generally accelerating flows, and the occurrence of acceleration was incidental, caused by the poor response of his rig at the higher flow rates. He did try to explain the effect of the acceleration, but the sparse nature of his data in this limit allowed for no more than hypothesising.

Moss and da Silva (1993) gave a summary of the regions of stability and instability within an impulsively started pipe. They showed that for early times the entire pipe was momentarily unstable, while in the steady-state limit only the region near the inlet remained linearly unstable.

2.4.2 Non-periodic unsteady flows

In these flows, the mean flow rate of the system changes in an unbounded manner, thus these systems could be described as *Reynolds number varying* flows.

The most profound study to date for such flows was possibly that by Shen (1961). He presented a theoretical description of the effect of temporal acceleration on an *inviscid* fluid with self-similar velocity profiles, of the form $U = T(t/t_0)U(y)$. He reasoned that a good description of the flow stability would be given by defining the concept of *momentary stability*. For instance an accelerating flow with growing disturbances is momentarily stable if the rate of increase of the base flow is greater than that of the disturbance. He quantified this concept by using the ratio of the kinetic energy of the disturbance to that of the base flow; thus momentary instability occurred when the kinetic energy growth of disturbance exceeded that of the base flow.

For applicable flows, Shen found that the criterion for instability was not whether the imaginary component of the disturbance celerity was greater than zero, but rather that it exceeded in value the term

$$\frac{1}{U_m} \frac{dU_m}{dt} \quad (2.1)$$

where U_m is the mean velocity. Thus in general he theorised that flows undergoing acceleration would have higher critical Reynolds numbers than steady ones, because of the 'shrinking' of the neutral stability curve due to the change of the stability definition as described above.

Couard & Criminale (1965a,b) used an energy method to understand the stability behaviour of generally unsteady flows. However they only gave sufficient conditions and bounds on parameters such as the Reynolds number. Unfortunately, this approach can lead to the estimation of very wide bounds of sufficiency, because of the broad nature of the energy method.

2.4 Unsteady transition

A good description of the comprehension of accelerating flows is given by Drazin & Reid (1981). They make the statement that ... *at this stage satisfactory methods of finding the stability characteristics of general unsteady basic flows have still to be found.* Since 1981 little has been added to the understanding of this topic. Consideration below is given specifically to non-oscillatory flows: systems with periodicity have unique mathematical properties, and analyses of such systems take advantage of their special symmetry. Unfortunately, these methods of analysis are not applicable to more general unsteady systems.

2.4.1 Flows with steady flowrates

Flows in this category have unsteady flow fields, but are steady in the mean. This means that the driving or boundary conditions remain constant, and the unsteadiness is associated with the development of the internal velocity field.

The impulsively started pipe flows of Maruyama Kato & Mizushima (1977), Moss (1985) and da Silva (1990) were of this category. In all cases flow was started from rest to a chosen Reynolds number, and the unsteadiness was associated with the development of the boundary layer in time, while the mean flow rate remained constant. Maruyama *et al* used a highly disturbed inlet, and monitored wall shear stress and velocity profile shape simultaneously. Moss (1985), in a predominantly experimental study, used only a single wall mounted probe to measure the wall shear stress at a point several hundred diameters down a smooth pipe with disturbed inlet conditions. He measured two distinct transitional events. At early times the flow became turbulent, this being possibly ascribable to local instability of the developing profile. After a reversion to laminar flow, another turbulent transition occurred, due to convection of disturbed flow from the inlet. The severe limitations of these results (a single measuring station) was unfortunate and did not allow any calculation of for instance the rate or direction of movement of the first transition event. Moss also introduced a convenient parameter to quantify the boundary layer - the velocity profile shape parameter - which was simply a dimensionless wall shear stress.

da Silva (1990) experimented with impulsively started flows, but he utilised a smooth inlet and a shorter aspect-ratio pipe than Moss. He did not see the early event, possibly because of the reduced length of his pipe, and the second event of Moss was absent because of the lack of inlet disturbance. He did however observe the occurrence of a centreline 'spike' within the developing pipe flow, which evolved into a turbulent slug as it traversed downstream. Strangely he did not see any initial boundary layer disturbances before the development of this centreline spike. The shape parameter of Moss (1985) was used to quantify both profiles in the pipe inlet, and those existing far downstream where the flow was still parallel, but he erroneously concluded that the two types of profiles

Pipe entrance flows

Pipe flows are sensitive to destabilisation by infinitesimal disturbances near the inlet only, and finite-amplitude effects are required to induce instability elsewhere [Abbot & Moss (1994)]. Many studies have been carried out to determine the stability behaviour of the pipe entrance to both axisymmetric and non-axisymmetric small disturbances [Tatsumi (1952), Huang & Chen (1973, 1974), Garg (1981), Gupta & Garg (1981), da Silva (1990), da Silva & Moss (1994), Abbot & Moss (1994)], and neutral stability curves for the inlet were presented in most cases. There has been until recently no substantial quantitative agreement between authors. The earlier works [Tatsumi, Huang & Chen] were based on approximate pipe entrance base flow models, which were known to be inaccurate. Garg (1981) used an accurate base-flow approximation, but the later and apparently more accurate calculations by da Silva [da Silva (1990), da Silva & Moss (1994)] disagreed substantially with his data. However the results of Abbot & Moss (1994), obtained using a spectral approach for both the base-flow and stability analysis (da Silva used a finite-difference technique), showed almost perfect alignment with the data of Garg. This finally cleared up the controversy surrounding the matter, vindicating the Garg results, and showing da Silva's analysis to be flawed.

The only experimental study to date is that by Sarpkaya (1975), who studied the effect of both axisymmetric and non-axisymmetric small disturbances. The theoretical and experimental results converge towards the inlet, but the agreement wanes downstream, with the two data showing different trends in this limit [see da Silva & Moss (1994)]. In the light of these results it seems that the description of pipe flow stability is lacking. The differences are not accounted for by non-parallel effects [Abbot & Moss (1994); also see the earlier discussion of section 2.3.1], which in fact diminish rapidly away from the inlet. The only viable culprit is finite-amplitude effects, but conclusions cannot be based on a single set of experiments and far more experimental work needs to be undertaken to resolve the dilemma.

While no work has been done on the finite-amplitude response of pipe inlet flow, there have been a number of studies of large scale structures within pipe flow. The appearance and growth of puffs and slugs within the pipe have been investigated in some detail [Lindgren (1969), Wygnanski & Champagne (1973), Wygnanski *et al* (1975), Rubin *et al* (1979), Teitgen (1979)] while related studies [Wygnanski *et al* (1979), Glezer *et al* (1989), Katz *et al* (1990)] exist for general boundary layers. These phenomena in pipe flow are possibly associated with by-pass mechanisms [Morkovin & Reshotko (1989)]. These events are associated with downstream regions of the pipe, the convection of the initially small disturbances from the inlet catalysing their development.




Figure 4.1: The testing facility.

4.1.1 The flow rig

The physical components used to contain and then generate the exponential and other flows in the system were grouped under the category of the rig. Basically the flow rig consisted of a pressurised header tank, a measuring section (a long glass pipe), and a downstream storage tank. The system was optimum for the current accelerating flows in that its operation was not continuous, but rather of a 'blow-down' nature. In order to generate a flow the upstream tank had first to be charged with fluid. After a suitable stilling period a test could be run by allowing water to flow from the pressure tank into the catchment tank. Control of the flowrate during testing was by means of valves on the discharge end of the measuring section. Flow into the test section was kept undisturbed by a carefully tailored inlet contraction. A pump was used to recharge the header tank. The rig sub-system is indicated in figure 4.2

4.1.2 Velocity measurement components

The rig was specifically designed for use with a LDV. There was a large two channel TSI unit available at the university, capable of simultaneously acquiring two orthogonal velocity components within a flow. The LDV method was discussed in detail in chapter 3.

A strong advantage of the LDV method is its high spatial resolution; the dimension of the measuring volume in the current system was of the order of 0.1 mm by 0.6 mm. In order to position this measuring volume in the flow, a traversing

Chapter 4

Experimental system

The requirement of flexible dynamic control of flow in a pipe required the design and manufacture of a specialised testing facility, comprising a long straight pipe (the system to be measured), reticulation system and control valves. An integrated computer acquisition and control system was necessitated by the highly dynamic nature of the experiments to be conducted. The apparatus was specifically designed for use with a LDV system.

This chapter describes the apparatus and its constituent parts, with specific emphasis on the design philosophy that dictated their structure. A detailed explanation of the acquisition and control software, written specifically for the current purpose is given in the following chapter - chapter 5.

4.1 System overview

The need to generate highly repeatable exponential accelerations meant that an unusual and highly flexible experimental apparatus had to be devised. Whereas generation of a steady flow requires a passive test rig, with settings to valves, etc. being made once and then left alone, generation of a strongly variable flow requires constant dynamic control of the test rig using feedback mechanisms. Due to this requirement of dynamism, the experimental flow apparatus, data acquisition system and control components had to be tightly integrated. Figure 4.1 shows an overall view of the testing apparatus.

The entire experimental system consisted of three major elements: the physical flow rig itself, the measuring equipment, and the data acquisition and flow control system. These three components are illustrated schematically in figure 4.2 below. General brief descriptions of these three sub-systems are given here in order to supply a broad understanding of the basic structure and working of the apparatus. In the major sections that follow this overview, a more detailed description of the constituent parts and operation of each sub-system is given.

velocity of this portion of the flow) but the investigation was not concerned with turbulent flow.

The above mentioned data smoothing process was in the current context almost automatic, and was in fact applied uniformly to all velocity traces, with the exception of those used to estimate disturbance intensity. In a sense the output from this smoothing technique was treated as the *raw* data for the remainder of this work. For this reason the above process was referred to as data *pre*-processing.

3.5 Summary

The LDV method was shown to be highly advantageous for the non-invasive and accurate measurement of flow variation. The basic principles of operation of the method were outlined, as were its limitations. The specific nature of the data processing needed for unsteady flows was presented, and sources of error were highlighted. The following chapter is concerned with the functional description of the entire testing apparatus, and a detailed description is given of the actual LDV that was used.

parameters that concerned this technique: the time window, the time increment by which that window was stepped, and the number of standard deviations for the rejection criterion. A command-line utility (SMOCTYPEX) as well as an interactive scheme (LDVPROC.EXE) was written to process the LDV data. In these programs the three parameters could be set and the data was then processed automatically as follows:

- (i) The window was moved until it contained the starting portion of the data.
- (ii) The temporal bias calculation as described earlier was applied to the data; data outside the acceptance range was rejected, and a mean value was calculated. This data value was placed at the centre of the time window.
- (iii) The window was moved by an increment Δt forward in time. This increment was independent of the window size; thus overlapping or staggered window configurations could be generated. Step (ii) was then repeated.

The resulting processed data was aligned at times $i \times \Delta t$ and consisted of a series of mean points. While the time increment defined the 'sampling rate' of the data, it was the time window that affected the frequency cut-off of the scheme. Thus, unlike traditional sampling, where the sampling rate determines the maximum frequency of the data, here the additional factor of the time window had to be considered. Certainly the resultant data sampled at a frequency of $f = 1/\Delta t$ would accurately represent only those frequencies below $f/2$, according to the well-known Nyquist criterion. However, the frequency cut-off would also be limited by the time window at a value of $f_w \approx 1/t_w$. Thus the total frequency cut-off limit of the scheme was estimated to be

$$f_{\max} \approx \min \left(\frac{1}{2\Delta t}, \frac{1}{t_w} \right). \quad (3.4)$$

The time increment was set at $\Delta t = 0.1$ s for the entire test. The time window was varied, between values of $t_w = 0.2$ s for the data used to infer transition, to about $t_w = 0.5$ s for traces used to construct base-flow profiles. Thus in the former the frequency limit was about 5 Hz whereas in the latter it was reduced to around 2 Hz. The above values were chosen after careful consideration of the raw data. For the velocity profile data, there was generally no high frequency components in the flow; the criterion for choosing these filter settings was to ensure a high data population per window. The initial stages of transition were characterised by a long wavelength deviation, thus the stringent filtering implied by these filter settings did not alias the important features of the data. Later turbulent flow was severely aliased (the method successfully calculated the mean

⁵The windowing technique, because it measured gradient, was not subject to the same criterion as a normal sample-and-hold process. It was estimated (by numerical experiment) that the technique started to alias frequencies only above approximately $1/t_w$.

strong temporal gradient, the distribution of the data is not represented by the horizontal pair of lines but rather by the angled lines parallel to $u \rightarrow u$. The bias is a result of incorrect distribution of data about the mean. For instance the point (1) which is obviously far from the true mean lies very close to the mean in the biased calculation, whereas point (2) which is clearly close to the true mean is far away in the biased mean calculation.

The extra broadening of the data caused by temporal gradient bias has been considered by Greenblatt (1990a, b), in this case for turbulent flows. He recorded errors as high as 70% in turbulence levels. The technique employed to remove the bias was similar to that used by da Silva (1990). The method can be best explained by reference to figure 3.6:

- (i) For a small time window t_w shown in the figure, a least-squares fit is made through the data. The experimental points are flattened by removing this slope from the data, while maintaining the same mean value⁴.
- (ii) The weighted mean and standard deviation calculation is applied to the flattened data. The data rejection criterion of n standard deviations (in all cases $n = 3$ was used in this study) was then used to reject points outside these values. This was equivalent to rejecting data outside the sloped lines.
- (iii) Process (i) was then repeated, but for the reduced data set. The central mean point of the fitted mean line - (t_c, U_c) in the figure - was taken as the mean value for the time window. Its temporal position was taken as central to the window. Thus by progressing the window by successive increments Δt through the data, mean processed values were calculated at evenly placed points in the data Δt apart.

3.4.3 The time-windowing technique

A method of calculating a moving average was hinted at in the above section. This section describes the algorithm employed. As the variation in the mean of a strongly variable flow system was of primary concern, it was necessary to calculate the mean of small, successive portions of data, with the time 'window' of data considered small enough so that the velocity variation within that region was close to linear. This moving average method was used by da Silva (1990).

The basic method chosen was to apply the temporal correction scheme to successive windows in the data, also employing the above rejection criterion to remove 'drop-outs' and other anomalies in the data. There were three basic

⁴The calculations to remove the velocity bias effect require that the data be corrected by a factor containing its magnitude. Thus on average the magnitudes of the velocity points are kept constant, in order to maintain the correct velocity bias correction. This factor was ignored by da Silva (1990).

must be made, and the correct equations applied. In the current context of laminar and transitional (initially disturbed laminar) flows with high mean values and small relative fluctuating component, comparison of the mean as calculated above with the normal mean calculation showed discrepancies in only the 4th significant figure at *worst*. Thus although the corrected equations were used, in the current context this bias was negligible.

Fringe and filter bias

These two effects were not really present in the flow system. The use of the frequency shifting technique eliminated to a great extent the possibility of the filter bias, by reducing the relative variation of the velocity to the zero datum [see section 3.2.5 above for an explanation]. Fringe bias was absent because of the high velocity, one-dimensional nature of the flow, aided again by use of frequency shifting.

Temporal gradient bias

A predominant new effect arises in temporally varying flows; that of *temporal* gradient bias. This is similar to the spatial broadening problem outlined above, but concerns the artificial broadening of the data in a small time window, due to the mean trend of the data increasing noticeably within that window. Figure 3.6 illustrates this effect.




Figure 3.6: The temporal bias problem and its correction in LDV data processing.

In this figure a small portion of data is shown, within the time window of width t_w . This data possesses a temporal increase in value, which can be approximated by the linear mean variation $a \sim a$ within this region. Because of the

3.4.2 Bias correction

The primary aim of data processing in unsteady LDV measurements is usually to remove the bias and broadening, and to produce a resultant set of points sampled at *regular* time intervals, and representing the mean variation of the flow. In order to achieve this a windowing technique was employed, whereby small portions of the data at regular times were processed in order to produce a mean value for the interval. Corrections to the data had to be applied to avoid any biasing of the results.

Corrections for the velocity bias

There are many methods that can be employed for correcting the data for this effect, for instance see Dimotakis (1976). In most cases the mean and standard deviation calculations are weighted according to some correction factor in order to remove this bias.

A particularly simple scheme can be employed if the flow is one dimensional, namely the technique of McLaughlin & Tiederman (1973). Here the mean and mean-square fluctuation calculations are simply weighted by the inverse of the velocity magnitude, as

$$\bar{U} = \frac{\sum_{i=1}^N |u_i|^{-1} u_i}{\sum_{i=1}^N |u_i|^{-1}}, \quad (3.2)$$

$$\overline{u^2} = \frac{\sum_{i=1}^N |u_i|^{-1} [u_i - \bar{U}]^2}{\sum_{i=1}^N |u_i|^{-1}}, \quad (3.3)$$

The use of these calculations will give good results as long as the flow remains largely one-dimensional or steady [Buchave *et al* (1979)]. In the present case the velocities were unsteady but highly one dimensional. This fact, combined with the use of frequency shifting, ensured that the use of these simple relations yielded corrected data.

For flows with significant multi-dimensional velocity fields, the more correct approach is to use a residence-time weighting technique [Buchave *et al* (1979)], which requires extra information from the LDV such as particle residence time. This data was available in the current system but was not stored because of the applicability of the simpler scheme.

In practice, substantial corrections to data using (3.2) and (3.3) only become significant when the data variation becomes large with respect to the mean. Thus in turbulent flow measurements, careful consideration of the bias effect

ious causes of bias occurring during processing can be readily identified and removed. McLaughlin & Tiederman (1973) give a detailed explanation of the biases occurring on LDV measurements. The main traditional culprits of bias are:-

- (i) **Velocity bias:** If the seeding particles are assumed to be evenly distributed throughout the flow, then measurement of portions of the flow with high velocities will show a greater data rate than portions with lower velocities, in direct ratio to the ratio of the velocities. Thus because of this bias it would be more likely for a high velocity reading be taken than one with low velocity. Consequently the simple mean of any group of velocity values will be biased towards high velocities, simply because more particles with high velocities are likely to be measured.
- (ii) **Spatial gradient bias:** If a strong spatial velocity gradient exists across the measuring volume (for instance in the boundary layer), then a wide band of data will be seen, the 'slow' side of the measuring volume measuring a much lower velocity than the 'fast' side. Because of this a laminar velocity time plot from a region with strong spatial gradient can be mistaken for a turbulent velocity profile. Furthermore, a simple mean of the data will be biased towards the higher velocities as in point (i) above.
- (iii) **Filter bias:** The high and low cut-off filter used for removing the high frequency noise and pedestal from the signal respectively, can also result in the introduction of bias. If the range of velocity variation exceeds the acceptance band of the filters, then all velocities outside the filter range will be removed from the signal, leaving only velocities within this range. This bias can be removed only by the correct choice of filter settings.
- (iv) **Fringe bias:** The LDV signal counter requires that the particle traverse a certain number of fringes before the signal is accepted as valid. This is achieved simply by counting the number of cycles in the input signal that exceed the acceptance threshold, accepting the reading only if the chosen number is exceeded. For particles possessing high velocities at right angles to the direction of measurement, only a small number of fringes will be counted before the particle again leaves the region. This is because the high lateral velocity component causes the particle to traverse the measuring volume by only a small distance in the normal direction. The upshot of this effect is that only particles moving in certain directions will be measured. The angle over which readings will be registered is called the acceptance angle. Use of frequency shifting dramatically alters the angle of acceptance because of the fringes moving counter to the direction of particle motion, allowing more fringes to be traversed. Thus correct use of frequency shifting can alleviate this bias. In the current context, the velocity naturally has a very narrow angle of incidence on the measuring volume, and this bias is thus of no consequence in the current system.

- (iii) The dependence of the measured velocity on only the beam intersection angle and frequency of the light results in no calibration being needed. This allows accurate measurement from creeping flows, right up to ones possessing supersonic velocities. No calibration is required when changing working fluids, or in fluids where the fluid properties vary. Errors can be introduced by variable fluid properties; for instance if the refractive index of the fluid is not isotropic, in which case the beam angles would vary. However, anisotropy would most probably result in the beams not intersecting at all, thus not lowering accuracy but eliminating the signal altogether.
- (iv) The small spatial extent of the measuring volume, combined with the high dynamic response of the method makes it ideal for measuring strongly varying flows. The dynamic response is not limited by the instrument at all, rather by the seeding particles, the criterion being that these particles follow the flow variation accurately.

Despite these strong advantages, as with any instrument, care must be taken in its use. In particular the LDV method can yield very poor readings if the operator shows a lack of understanding of the basic principles and pitfalls inherent in the method. In general the LDV can be treated as a 'black-box' measuring device, but only after careful setting up of the system (beam alignment, frequency shifting, filter settings and so on) to match the current flow configuration. Poor initial set-up will require constant adjustment to maintain readings, but careful initial adjustment will allow a 'point and shoot' methodology to be adopted. This is especially so with the more modern equipment, where higher powered lasers combined with devices such as the frequency shifters alleviate many of the problems associated with simpler devices.

The following sections are concerned with the specific nature of the data gathered from an LDV, and the care that must be taken in processing this data to avoid serious errors.

3.4 Data processing

Individual measurements from a LDV are highly accurate indications of the velocity of the particle that was measured. This is because of the linearity of the method. Problems arise when the mean *fluid* variation in the system has to be extracted from large groups of these individual measurements.

3.4.1 Bias

The difficulty of processing LDV data was hinted at earlier. The main concern with data processing is the introduction of spurious bias into the data. Var-

simply by utilising the principle whereby light beams of different frequencies do not interfere with one another. If a laser that emits light at many discrete frequencies is used as the power source for the apparatus, then individual component frequencies of this laser beam (split using a prism) can be superimposed without interfering. Thus effectively two separate velocimeters using different light frequencies can be superimposed co-axially, with the mutual plane of their beams at right angles to one another. Suitable filters can be employed on two separate photomultipliers such that measurement of the signals from each component can be completely separated.

In order to measure all three velocity components a third velocimeter, operating using a third frequency of light, can be aligned with the measuring volume from an acute angle. Three component systems suffer from the disadvantage that the third component cannot be co-axial with the other two. As a result a very bulky optical bench is needed, coupled with a need for a second optical path into the flow. Furthermore, complications to do with optical misalignment in complex flow geometries can occur, and such systems are thus only really useful in refractive index-matched systems, or for measurements in free air. New developments in the field have recently resulted in tiny LDV units incorporating fibre optic technology, allowing much simpler design of these multi component systems, and thus eliminating some of the reservations concerning three component systems expressed above. The current system was a two component device, although only single component axial flow measurements were made during the final testing.

3.3 Advantages of the LDV method

As mentioned in the discussion above, the dual beam laser velocimeter holds a number of advantages as a flow measurement device. To summarise, these are as follows:

- (i) The system measures vector components of velocity. Use of a multi-component LDV allows the direct extraction of the separate velocity components. In comparison a hot-wire device is a scalar measuring instrument³, and much processing must be done to extract the velocity components.
- (ii) The system is non invasive, requiring only that the working fluid and containing walls be transparent. This holds a number of advantages, namely: no disruption of the fluid being measured, and the ability to be used to measure velocities in hostile environments (combusting or corrosive flows).

³A single 'hot wire' will measure the scalar value of the velocity in the plane perpendicular to its axis; the combination of two vector velocities. To discriminate the velocity components in such a system, three wires must be positioned in a mutually orthogonal manner, and the component velocities can be extracted from the correlation between the three separate readings.

3.2.6 Random time base

Usually a flow device such as a hot-wire anemometer or pitot tube measures the flow variation continually. Sampling of the device from a computer or other such instrument is then done regularly, at a set frequency. As a result, the time increment between the sampled readings is fixed, allowing simple storage (Δt needs to be stored only once) and simple processing; for instance use of the Fast Fourier Transform (FFT) can be made to extract frequency domain information, numerical differentiation, integration or filtering can proceed using any number of standard algorithms.

The laser anemometer, on the other hand, only registers a signal when a particle traverses the measuring volume, thus the interval between readings is not fixed but random. For high seeding density and reasonably high velocities, this rate of particle traversal approaches a more or less fixed average in the mean, and thus the time increment remains randomly distributed about some mean value. For low particle densities or low velocities however, the time of arrival of the next particle becomes completely unpredictable.

The random time base causes a number of problems, both in acquisition and in processing. Filtering of the resulting discrete velocity data to remove high frequency components requires specially developed algorithms (see section 3.4), while frequency domain information necessitates the use of slow and cumbersome algorithms such as that by Gaster & Roberts (1975, 1977). However, if the data rate is high enough, then the time base can be successfully approximated by its mean value, and the more common FFT type algorithms can be employed [for instance, see McDonnell & Fitzpatrick (1992)] to obtain frequency information. A more subtle problem arises with respect to the acquisition of data from an LDV, because on the lack of knowledge concerning the arrival of the next particle. Whereas a hot-wire probe can be sampled regularly by computer, leaving the computer free to conduct other tasks between readings, the acquisition of LDV data requires the computer to be attentive at all times, in case a reading is made. This is especially a problem in the case of a multi-channel system. Here a signal may be ignored by the computer on one channel because it is busy sampling another channel. In this case the first channel sampled will always be favoured, causing a bias in the data towards this channel. The possibility of this occurring in the present apparatus was avoided by utilising a separate computer for each LDV channel, thus avoiding the favouring of any particular channel.

3.2.7 The multiple component system

The preceding description concerned a single channel laser Doppler system, comprising a single pair of incident beams intersecting in the flow, combined with a single set of receiving optics. It is possible to construct a multiple component system for measuring more than one component of velocity at a given point,




Figure 3.5: The progressive processing of the signal from the photodetector, resulting in a frequency measurement.

3.2.5 Frequency shifting

A strong disadvantage of the simple dual beam method is that only the magnitude of the velocity vector is measured, not its sign. Thus in reversing flows there is no way to ascertain in which direction the flow is moving. To overcome this, the principle of frequency shifting was devised.

Frequency shifting entails changing the frequency of one of the laser beams slightly, by means of a Bragg cell. When the shifted and unshifted beams cross, the fringes no longer remain stationary but move across the measuring volume at a rate proportional to the difference in the light beam frequencies. The resulting frequency registered by the photodetector is thus not relative to the measuring volume itself but to the moving fringes. In order to calculate the actual fluid velocity, the fringe velocity must be subtracted from the obtained velocity value.

This shifting of the zero datum allows flows with negative velocities to be measured, as long as those velocities do not exceed those of the fringes. In the context of flows with a large variation in velocity, frequency shifting reduces the effective range of the velocity relative to the zero datum, and thus allows high quality readings to be taken for the entire flow range without varying the filter settings on the counter. This can be explained as follows: suppose that the range of flow variation in terms of frequency is from 0.1 to 1 kHz, the variation covering 90% of the range. Without frequency shifting, wide filter settings would be needed, admitting noise and degrading the signal. If a frequency shift of 2 kHz is added, then the measured range of frequencies at the photodetector would be 2.1 to 3 kHz, only 30% of the range, thus allowing far tighter filtering of the input signal. Another way of viewing the effect of frequency shifting in this context is to consider it as a means whereby the frequency of the pedestal and signal can be separated in value, thus allowing more accurate filtering of the signal.

Optical access

Fluid within a round tube produces a strong lensing effect, due to the curved pipe interface in combination with the different refractive indices of water, glass and air. A light ray entering the pipe off the major diameter will be sharply refracted as shown in figure 4.9, (a). In the case of the LDV, the lens effect will tend to focus the beams themselves due to their finite diameter. Thus a round laser beam, even if entering the pipe on the major axis will be distorted with a resultant decrease in signal quality. Furthermore, the reflected light is gathered not back down the beam paths but over a conical region; here the lensing will cause most of the returning light not to fall on the photodetector, causing an extremely poor signal.




Figure 4.9: The refraction of light through a water filled tube: (a), without a surrounding water jacket; (b), with a water jacket.

The simplest way of reducing this effect is to place the pipe in a bath of liquid with the same refractive index as the working fluid, and with a planar interface between the outer jacket and the air. This arrangement is shown as (b) in figure 4.9. With this arrangement the outer planar wall will distort all the LDV beams equally¹, and the only effect of this interface would be a change in the beam intersection angle. The pipe wall is now simply a water-glass-water interface. For a thin pipe wall, the curvature of the glass becomes insignificant, and the passage of the laser beams through this interface - so long as they remain near the major diameter - will result in a small translation only. The substantial reduction in the lensing effect also permits most of the returning light to remain incident on the photodetector. The resulting substantial increase in signal quality fully justifies the use of an outer water jacket.

The ideal way to eliminate refractive problems is to utilise a fluid with identical refractive index to the fiber glass pipe. In this case the pipe will 'disappear', indicating that no refraction is occurring, and thus LDV readings will be possible anywhere within the pipe. However, most refractive index matched fluids are either extremely corrosive, toxic, or both. For this reason it was decided to use water as the working fluid² and take advantage of the already impressive gains in signal quality introduced by the water jacket. Edward & Dyblis (1984) give a good description of refractive index matching in LDV work.

¹This will be so as long as the axis of the LDV is kept perpendicular to the planar face.

onto the machined bore. The pipes were kept parallel during the curing of the epoxy by using an assembly jig. The gap left between the pipe and flange (after excess epoxy was removed) was filled with car body putty, and then smoothed and burnished to a shine. This filled region allowed careful matching of pipe and flange, thus avoiding any steps.

The join between the flanges was finished before attachment of the pipes; thus care had to be taken not to damage this interface after the pipes had been glued, because access was obviously restricted after fixing the pipes. The edge was masked off during the smoothing of the transition fillet to avoid any damage. The result of the above process was a join between the sections of pipe that was as smooth and regular as possible, the major variation being due to the variation in dimension of the pipes themselves.

In total eight units were assembled, with a resulting length of 12.7 m, and a length to diameter ratio of $l/D = 263$. However, only the first 9.6 m were accessible by laser ($l/D = 200$). The figure 4.8 shows the pipe segments assembled on the rig. The header tank is in the background, and the pipe leads through a hole in the far wall (not shown), to the valve system.




Figure 4.8: The testing section of the rig.

Construction of the flanges

Flanges were manufactured out of polypropylene, because of its inertness and good structural stability. A detailed plan of the flanges is shown in figure 4.7.




Figure 4.7: The pipe joining flange.

Each flange was manufactured out of a solid block of polypropylene, of dimension $180 \times 180 \times 80$ mm. Sealing was by means of an 'o'-ring seal mounted in the front face, in a recessed groove in the one flange. The matching flange had a raised ring which sealed onto the 'o'-ring within this groove. By this means the sealing function of the flange was moved from the abutting edge, allowing a flush fit between the units. The matching surface between the seal and the bore was slightly raised, by 0.2 mm, to ensure an interference fit at the surface of the bore when the flange was bolted together. This eliminated the possibility of any gap at this join.

The entire assembly was held together by eight bolts. A tapered dowel was fitted to ensure perfect alignment between the units. The final machining of the bore was done with the two units bolted together; in this way the perfect match of the two flanges was ensured. Flanges were in matched pairs, and could not be interchanged.

Each flange was specifically manufactured to match a chosen section of glass pipe, and the final internal diameter of the flange bore was matched to the adjoining sections of glass. The sections of glass in turn were arranged so that they showed the smallest difference between the matching ends. The glass tube was faced off at the end to give a sharp interface, which was then bonded into the recessed portion of the flange using a high quality epoxy resin [Araldite AV129 & HV997 hardener]. A taper prevented the pipe from facing directly

Joining of pipes

In glass test rigs previously assembled at the university [da Silva (1990), Greenblatt (1990b)], sections of glass tubing were joined together by fusing the ends of the pipe sections together. This process caused rather large ripples to be formed on the inside of the pipe near the join. These were unavoidable and largely uncontrollable in the manufacturing process. To avoid the problem of distortion it was decided to resort to a cold-joining process, using bolted flanges.

Initially it was thought that the glass tubes could be simply faced off against each other and held together using a collar. However measurement of the available piping showed that slight ovality and discrepancies in the diameter (of the order of 0.1 mm ovality and 0.1 mm variation in diameter), and direct butting of pipes together would have resulted in steps. Figure 4.6 illustrates this problem.




Figure 4.6: Problems associated with cold joining of pipes: (a), step caused by mismatch in diameters; (b), lack of access does not allow a transition region to be blended, (c), flange allows access for blending of the pipe surface.

A possible solution of leaving a region between the pipes - (b) in the figure - that can be blended to avoid any steps immediately raises the problem of access: if the pipes are bonded with a transition region, then smoothing of the region to the pipes needs to be carried out only after the pipes have been bonded, and this requires access down a 1.5 m length of pipe. Furthermore no inspection and repair of the joins can be carried out once assembly of the entire pipe is complete, as access then would be even more restricted.

Using a flanged end to each section of pipe - (c) in figure 4.6 - allowed the pipes to be blended to each other while allowing access both for manufacture and later repair. The flanges were manufactured as matched pairs, and their internal bore was machined together to allow for a perfect match on the abutting surfaces. Using this approach, two regions had to be blended - the pipe-to-flange interfaces for each side of the flange.

region. The inlet was bonded to a polypropylene flange (described later) and the inner matched surface between flange and inlet was beveled and filled to avoid ripples or ridges. The method of joint construction is discussed in detail in the next section.

4.2.3 Test section

One of the major problems in the way of constructing a long smooth pipe is the availability of high quality single lengths of tube. Seamless stainless steel pipes can be purchased with l/D ratios of a few hundred, but in the case of glass tubing the maximum manufactured length is 1.5 m. For the current study, access was required everywhere within the length of the pipe, thus the entire pipe system had to be optically transparent.

Selection of pipe

It was decided to use high quality glass tubing for the test section. The mean internal diameter of the pipe needed was 48 mm - this diameter was decided upon because it offered a good compromise between radial resolution and the length of pipe (as a ratio of diameter) that could fit within the laboratory, and because smaller diameter tubing possessed too thick a wall in relation to the diameter. The need for the thinnest possible wall was dictated by optical considerations (discussed later). The closest available diameter of glass tubing had a nominal internal diameter of 48.2 ± 1 mm and a wall thickness of 1.6 ± 0.2 mm. The glass compound was duran borosilicate - a type of Pyrex glass. This glass has a high level of transparency to light in the visible range, and was thus suitable for LDV measurements. It possessed a refractive index of 1.474 for the sodium D line, and a transmission factor of 92% for a 2 mm wall thickness [Bibby, technical data sheet].

The fairly high variability of diameter shown above - 2% - was reduced by selecting particular pipes from the supplier with a variation of only ± 0.1 mm in diameter, reducing this variability to 0.2%. The presence of ovality was also considered, and pipes with ovality of 0.1 mm or less were chosen.

The use of a perspex tube was considered, because of their availability in greater lengths, but the optical quality of perspex was found to give extremely poor readings near the wall. This is mainly due to perspex's susceptibility to scratching³. Because of the above choice, some means of accurately joining lengths of glass pipe had to be considered.

³In measurements taken using frequency shifting on other perspex rigs, extremely high data rate, zero velocity readings were obtained near the wall. It was found that the LDV was simply 'reading' the velocity of the scratches. With glass walls the LDV measuring volume could be positioned partially within the wall, with no such problems.




Figure 4.5: The shape of the inlet contraction.

to provide more uniform conditions than ambient to the second stage. In the current case the ambient tank conditions were very uniform. Moreover, it was reasoned that for unsteady flows, a large as possible turn-down ratio in a short length would allow the spatial pressure gradients to dominate the flow system, thus allowing conditions in the inlet to remain fairly steady.

The inlet was manufactured by turning a wooden die on a lathe, and then hand finishing this mould using sanding sealer and then later wax polish to achieve a very smooth and regular shape. The inlet was then constructed on this die using fibreglass with waterproof epoxy resin as the binding compound - the die surface was first coated with a waterproof gel-coat (for a smooth finish) before layers of woven fibreglass mat were applied. Flaws in the inner surface of the inlet (some bubbles were formed under the gel-coat) were filled using automotive body-patty² and the surface was finished off with fine grade water-paper, burnishing compound and then polish. The result of the above process was an extremely smooth and uniform inlet contraction. The use of a lathe to turn the die (rather than the inlet directly) allowed better control of the general surface shape, and a high quality finish on the die was simple to achieve, by smoothing it while mounted in the lathe chuck. The inlet incorporated two diameters of straight pipe section after the end of the contracting cone, to remove the inevitable join as far as possible from the very thin boundary layer at the end of this contracting

²Of all the expensive and exotic waterproof fillers and bonding agents used in the construction of the rig, this inexpensive compound proved the most stable and useful as a filling agent.

(1944), Lighthill (1945), Goldstein (1951), Whitehead *et al* (1951), Libbey & Reiss (1951), to mention a few). Although many of these methods result in different shaped inlet contractions, it is unclear which design is physically the most optimum, and under what range of conditions remains so.

In the context of *unsteady* flows of interest here, no clear indication is available at all as to what design is the most correct. However, any well designed inlet contraction must produce the following effects in the inlet flow; it must ensure an undisturbed inlet velocity profile of uniform velocity and with a very thin boundary layer. Concerning the unsteadiness of the flow, the avoidance of separation due to its dynamic nature, for instance due to instantaneous adverse pressure gradients within the contraction, is an important consideration.

The present flows were subjected to strong positive accelerations at all times¹; thus the restrictions concerning separation were far less important than would be the case in a decelerating flow, where negative temporal pressure gradients might overcome the positive spatial gradients within the inlet.

As a result of the plethora of design procedures, and the lack of clarity about the response of an inlet under unsteady conditions, it was decided to opt for a design that was known to have been successfully used for the generation of round jets in free air, by Popiel & Trass (1991). This contraction design was determined using the formula by Witoszynski (1924), based on a uniform pressure variation through the contraction. The formula applied was

$$r = R_2 \left\{ 1 - \left[1 - \left(\frac{R_2}{R_1} \right)^2 \right] \frac{(1 - 3z^2/a^2)^2}{(1 - z^2/a^2)^3} \right\}^{-\frac{1}{2}}. \quad (4.1)$$

Here R_1 is the larger radius and R_2 the smaller. The factor a is a chosen multiple of R_2 , in this case it was taken as $a = 10R_2$. With suitable choice of terms, equation (4.1) then gives a variation of the radius r with axial position z . The contraction shape is given in figure 4.5 for the current values of R_1 , R_2 and a .

As can be seen in the figure, the contraction was very abrupt (the figure is presented in approximately correct aspect ratio). In fact from a performance consideration, abruptness was advantageous, because the spatial pressure gradient increased as the length of inlet reduced, and thus possibly improved performance. However ease of manufacture was also a constraint, and the current solution was a compromise between these two factors.

It was decided to use a single stage contraction [as opposed to the dual-stage unit of da Silva (1990)], because of the symmetry in the tank offering uniform inflow conditions. One of the functions of the first stage in a two-stage contraction is

¹The experimental results showed an anomaly, possibly due to separation at start-up within the inlet. It was thought to be related to overshoot of the flowrate variation at the impulsive start of the flow.




Figure 4.4: The header tank.

Advantages of the design

The pressurised concept allowed an extremely compact discharge tank that could supply water at high pressure, safely up to 2 bar (it was the strength of the test section that limited the working pressure). An equivalent free surface system would need about a 20 m head to achieve the same results. The pressure was not limited by the strength of the internal air-water interface; in fact the conditions across the diaphragm always remained constant, because the absolute pressure in both cavities rose equally. Although the forces on the piston were quite severe, they were only static, and remained the same regardless of the supply pressure.

The system was completely axisymmetric, to the point that water was pumped into the tank symmetrically. The tank was initially filled through a pipe in the front face, next to the drain, but a large swirling eddy was induced which persisted in the tank for *days*. Pumping the fluid back up the test section was introduced to successfully eliminate this effect. It seems the symmetric circulation induced by the latter approach decayed much quicker. The procedure for further elimination of the swirl during testing is discussed in section 6.3. To further reduce the possibility of circulation an aluminium honeycomb section (21) [depth 150 mm with 10 mm width honeycomb] was placed within the tank before the inlet.

4.2.2 Inlet contraction

Inlet contraction design is a ‘black art’; most references to design procedures are not recent, and are based on simple potential flow theory [e.g. Smith & Wang

the minimum pressure into the water was slightly higher than the air pressure, and keeping the diaphragm taut. An inspection hatch (9) allowed access to the linkage and springs in the air cavity.

As well as the pressure differential problem described above, the presence of the water-air interface within the invaginated portion of the diaphragm (10) produced a strong buoyancy effect, forcing the piston upward. To avoid putting undue strain on the linkage, large lead weights (11) were placed inside the piston. Their mass was selected to counter the maximum buoyancy force; the force was a minimum when the piston was retracted (as shown in the figure), and a maximum when fully forward. However the flexibility of the linkage was at a minimum in the retracted position, thus the piston was weighted to be approximately neutrally buoyant in the fully extended position (not shown).

Access to the water cavity was through the front face of the tank (12). This face was securely bolted to the tank, and through its centre a large inlet contraction (13) constricted the flow into the test section. The test section attached to the end of the contraction by a flanged join (14), which was carefully matched to the downstream pipes. The details of the flange design are discussed in conjunction with the test-section below.

The air compartment could be pressurised via the electrically actuated solenoid valves [ASCO model 8210 - 2 inch bore] (15) and (16), while the air pressure was roughly maintained by a pressure regulator (17) [Norgren Olympian, 2 inch supply pipes]. The supply pipe diameter to the regulator and valves was large (2 inches), in order to allow unhindered supply of large volumes of air. Fine control of the air pressure was by means of a small pressure bleed solenoid [ASCO $\frac{1}{2}$ inch] (18) which could be throttled to allow slow escape of air. High pressure (≈ 6 to 8 bar) air was supplied to the regulator. Valves (15) and (16) were connected electrically so that only one of the two were open at a time. Thus air was either admitted from valve (16) to the tank, or air from the tank could be released using valve (15). Both valves (15) and (16) could be throttled by means of manual gate valves in series with them; thus the rate of tank filling and discharge could be regulated manually. All three solenoid valves were controlled from the central computer.

Water was both discharged and admitted to the tank through the bell-mouth contraction - the tank was recharged back up through the test section. Air could be bled from the tank from a bleed pipe (19), and it could be drained completely through the drain (20). A rheostat attached to the linkage and a water level sensor (not shown) measured the piston position and the presence of any water leakage into the air compartment respectively. A photograph of the tank is shown in figure 4.1. The approximate capacity of this unit was 500 l.

4.2.1 Header tank

The function of this component was to hold a quantity of fluid, and to supply it on demand to the test section. A detailed cross-sectional view of the unit is given in figure 4.3.




Figure 4.3: A cross-sectional view of the supply tank.

One of the most important considerations in the design of the test rig was the requirement of axisymmetry. The flows measured by da Silva (1990) showed some lack of symmetry, and this was thought to be due to the lack of symmetry in the header tank he used (it was a vertical tank with a free surface). Another major requirement was the need to maintain a large head of water in the tank, in order to generate the large Reynolds number flows needed for the experimentation. These two requirements dictated a rather unusual design for the tank, which is described below. Reference is made to figure 4.3 in this description.

The tank was essentially a large symmetrical pipe (1), with diameter 600 mm and length 1.4 m, coaxial to the test section. Within the tank a water compartment (2) and air chamber (3) were separated by a large diaphragm (4) supported by a rigid piston (5). The diaphragm was made of a waterproof knitted vinyl material, identical to that used in the manufacture of inflatable dinghies; thus it could fold but was not elastic. The piston was supported on a 'scissors-like' mechanism (6), which constrained it to move backwards and forwards only. The piston was not driven by this linkage, but merely supported. The joints in this linkage incorporated sealed roller-bearings, to reduce friction to a minimum.

Cavity (2) was full of water, while (3) was air-filled. The hydrostatic gradient down the front face (7) of the piston resulted in the pressure at the top of the water cavity being much lower than the (constant) pressure in the air tank. As a result, there was a tendency for the diaphragm to collapse downward into the water cavity, fouling the piston. To counteract this, strong springs (8) were added to the scissors linkage, to offset the hydrostatic effect, thus ensuring that

4.1.3 Acquisition & control components

The rig described in section 4.1.1 above contained elements, such as valves, solenoids and pumps, the control of which was the concern of a central computer. Further to this various data had to be acquired from the apparatus, both as measurements (for instance the LDV data was such an element) and for the purposes of control (flowmeter data). The responsibility of both acquisition and control was vested in a single computer system, which is shown schematically in figure 4.2.

The combination of the acquisition and control was because some elements of the rig generated data necessary for both processes. For instance the flowmeter data was required by the flow control system, and as a measurement. More importantly, automation of the entire testing procedure was desirable for the sake of repeatability. In order to achieve this, centralisation of both acquisition and control was necessary. The combination of all elements into a single system allowed great generality in control of the experimentation which is not possible when using separate and inflexible acquisition and control systems.

While it was the main computer's task to control all elements of the rig, the function of real-time acquisition of the LDV data itself was delegated to two smaller computers, each under direct control of the main unit via a network constructed specifically for the task. Two of these acquisition computers were required for the current system, one per LDV channel. Separate systems were used to allow high data acquisition rates, but more importantly to avoid favouring an acquisition channel. A discussion on this biasing effect was contained in chapter 3, in section 3.2.6. The main computer ran a dedicated software program, concerned with all the above aspects of measurement and control. Because of the general requirements of the experimentation - the generation of arbitrary accelerating flows - the user interface had to be flexible. To achieve this, a command-line control package was developed, utilising a macro command language, and allowing easy modification of all the rig parameters during execution, without having to resort to alteration of source code.

This overview is by its nature, brief. It introduced the general elements of the experimental apparatus, which are described in full detail in the sections that follow.

4.2 Rig sub-system

This unit comprised the components that contained and physically controlled the flow: the pressurised header tank, pipe test-section, valve network, catchment tank, pumping, flow measurement and pressure mechanisms. These elements are listed and described in detail below.




Figure 4.2: A schematic view of the apparatus, showing the various sub-systems: (a), flow rig sub-system; (b), acquisition & control sub-system; (c), measurement sub-system.

bench was required with the ability to accurately position the LDV in three orthogonal spatial directions. The entire laser velocimeter had to be moved, a mass of approximately 150 kg; thus there was the daunting task of positioning a large mass accurately in three dimensional space.

An existing horizontal traversing table (originally a textile cutting bench) was modified to allow automated vertical movement. This bench already possessed stepper-motor control in the horizontal plane, and was used by da Silva (1990) in his measurements of impulsively started flows. Accurate position measurement devices were attached to all three axes of this bench in order to acquire three dimensional position data. The resulting system could position the LDV under computer control anywhere within the flow system. The figure 4.2 shows the measurement components of the system schematically.

In addition to the large device described above, a small portable laser velocimeter was available, in order to acquire data from portions of the flow outside the range of the large traversing bed. This unit which was attached to the main unit by means of a fibre optic link, was light (≈ 5 kg) and thus could be positioned freely in order to measure the flow variation anywhere in the rig. This small unit is shown in the schematic in figure 4.2. It had a one dimensional manual traversing mechanism.

Apart from the LDV and its traversing equipment, various other measurement devices were positioned on the rig: an electromagnetic flowmeter was mounted in the flow, a temperature probe was positioned in the catchment tank, and a pressure transducer monitored the discharge tank pressure. While flowmeter data could be acquired during testing, the other two sensors were mainly required for rig control, and are thus really part of the acquisition sub-system.

thereafter the hardware configuration (computers and other control equipment) is given. The controlling software is presented in chapter 5

The system was primarily designed to replace the Hewlett-Packard 9836-based acquisition system which existed at the university [see da Silva (1990)]. This machine ran an interpreted basic operating system, and allowed a single LDV channel to be polled at about 100 Hz, and also permitted several lines of digital communication. Because of the interpretative nature of the operating system it was incredibly simple to use for basic acquisition (programs could be modified 'on the fly') but was very slow and inflexible, and thus unsuited to the high performance multiple control nature of the current rig.

It was decided to utilise the IBM PC compatible based infrastructure, basically because of its inexpensiveness, widespread usage, and extreme flexibility. Unlike the HP system above, the IBM based systems allowed incorporation of hardware components (acquisition cards, etc.) from diverse sources. Furthermore the most advanced programming languages were available for these machines. It was decided to remain with the Microsoft MS-DOS operating system for these machines. The Microsoft Windows V3.0 operating system was available at commencement of the project, but both the package and design tools for this system were still not fully developed. The low-level power of DOS also allowed more direct control of the hardware components than did Windows. The software design philosophy is presented in section 5.1.

The basic structure of the hardware was discussed in the overview. A more detailed schematic of the layout is shown in figure 4.15.




Figure 4.15: The layout of the data acquisition hardware.

The central computer was the cornerstone of the system. It controlled a number of sub-systems, both external and internal, which are described below:

a 4 m travel in the direction axial to the pipe and a 1.4 m range horizontally at right angles to the pipe axis.

Both horizontal co-ordinates were driven by MAE HY200 stepper-motors, controlled by Digiplan CD30 motor drive units. The interface between the stepper driver and the computer was by a RS-232C serial cable link. The bench offered no vertical movement, so a vertical traversal mechanism had to be designed. A fork-lift type arrangement was decided upon. The LDV bench was supported on a cradle, which could move vertically on two IKO LWA15 linear bearings. This mounting arrangement allowed for an extremely rigid design. The vertical traversing structure is shown in figure 4.14, and allowed 0.9 m of vertical movement in total.




Figure 4.14: The LDV and traversing bench, showing the vertical traversing assembly.

A Warner-Electric Electrac 100 linear actuator was used to raise and lower the rig vertically. This was a sealed unit, incorporating a DC motor, gearbox and recirculating ball worm drive. This device was powerful enough to raise the entire bench vertically with no counterbalancing. As safety features it incorporated limit switches and a friction clutch.

4.4 The acquisition & control hardware

The acquisition and control system was comprised of the main computer, the two LDV acquisition computers, and related hardware components. Control was implemented by means of a software programs on the main machine and LDV acquisition boxes. In this section the general philosophy of design is discussed;

The receiving optics were far simpler than for the larger unit. The unit worked in the back-scatter mode, and thus the receiving components were combined with the transmitting optics. No field-stop was incorporated. This was because the compact nature of the beams from the fibre-optic unit, combined with the incorporation of the optical box, use of glass pipes and the general cleanliness of the rig, resulted in no noticeable flare off the pipe walls⁶. The basic components of the receiving assembly were:

TSI 940 Photodetector mount. A single channel mount was used to direct the returned light to the photodetector. This unit incorporated a focusing lens to allow the reflected light to be made properly incident with the photodetector pin-hole.

TSI 9160 Photodetector. The second photodetector from the large rig was mounted on the photodetector mount.

The small unit did not incorporate a frequency shift unit; however the one laser beam had already been frequency shifted on the main unit, before being directed down the optic fibre. This meant that frequency shifting behaviour was available on this small unit with no increase in bulk. Because of the remote position of the frequency shift control box (on the main unit), a long signal cable (≈ 6 m) had to be routed from the remote photodetector, to the frequency shift box for down-mixing. It was feared that some reduction in signal quality would be introduced by this extra length of cable, but none was noticeable. In fact, as mentioned, the unit produced remarkably good signals despite the optical fibre link and long signal cables.

The high quality of signals was partly ascribable to the optical link. The accurate focusing arrangement of the fibre optic unit 9263 allowed for an extremely small beam diameter to be produced, of the order of 0.1 mm. With the beam expander, the measuring volume width was approximately 0.044 mm and its length only 0.5 mm. The small beam diameter and the good distribution of light within each beam allowed for a high-precision measuring volume and thus excellent signals. Of course the quality and cleanliness of the optical path aided this.

4.3.3 Laser positioning

Although the laser traversing bench was not strictly part of the LDV measuring system, its function was to position the large LDV, and thus its description is included here. The bench was a large Aristomat textile cutting bench, providing

⁶The signal-to-noise (SNR) ratio was so high that the noise in the burst signal could not be resolved on the oscilloscope on the same scale of $\approx 10^{-6}$ as the signal itself. With such clean signals, the presence of the field-stop merely resulted in degradation of signal strength, and thus actually decreased the SNR.




Figure 4.13: A schematic diagram of the remote laser head, showing the redirection of the vertical beams from the main unit.

TSI 9263 optical fibre unit. This unit, described above, converted a pair of light beams from within the fibre-optic cables to free-air laser beams. The unit had fine adjustment screws for beam direction and collimation, thus there was no need for steering prisms or a collimator in the small unit.

TSI 932 optics bench. This was a small one-dimensional traversing bench, with a total pitch of about 50 mm. The traversing mechanism was hand adjustable by means of a small knob attached to a threaded rod. The unit had adjustable feet and a standard TSI mounting rack for the optical components.

TSI 9113-0 beam spacer. This component brought the beams closer together prior to the beam expander.

TSI 9188 beam expander. The function of this component was, as with the larger unit 9189 on the main LDV, to reduce beam diameter at the measuring volume; in this case by a theoretical factor of 2.27, thus possibly increasing the signal-to-noise ratio about 5 times.

TSI 917A lens. This lens had a focal distance of 121.2 mm with a beam spacing of 21.4 mm, resulting in a free-air half-angle of 5.045° . Because of its very short focal distances, the effective solid angle over which scattered light was gathered was the same as for the big laser with its large diameter lens.

also allowed optimum collection of scattered light, thus ensuring a good signal-to-noise ratio. The unit gave a theoretical increase in the signal-to-noise ratio of 50 and a decrease of measuring volume diameter of 3.75. [TSI 9100-7 LDV manual (1987)].

- (iv) The TSI 9169-150 front lens with a focal length of 480 mm was used, giving a free air beam half-angle of 4.911° .

Basic arrangement of components

Light from the laser passed through a prism which split the beam into its component frequencies. The two strongest light frequencies - green at 514.5 nm and blue at 488 nm - were directed by mirrors down the transmitting optical barrel. Each beam was then split into two equal strength components, one pair in the horizontal plane and the other in the vertical. Each pair then passed through a Bragg cell, and on through the focusing optics, finally leaving the front lens to focus at a point in the flow.

Light refracted from particles returned along the same path as the transmitted beams. This light was then focused through the field-stop aperture (the transmitted beams passed on the sides of this unit), to eliminate light from outside of the control volume. It then was reflected to the side by an angled mirror, where it was focused onto the photo-detectors. Light was split between the detectors by a semi-silvered mirror. Each detector was preceded by a narrow band filter, each filter to eliminate light from the other channel.

The LDV as described above was a standard 9100-7 unit. Although this unit was designed to be used as a two-component instrument, the modular nature of the TSI equipment allowed other components to be included in its assembly.

In the current testing, the vertical component beams were required to power the remote laser head; thus a model 9259 focusing unit followed by a 9262 fibre-optic attachment were placed in the LDV assembly directly after the frequency shifters. These units allowed the horizontal beams to pass unhindered on their normal path. The two vertical component beams were then directed down a pair of fibre-optic cables to the remote system. The modified arrangement and remote head are shown schematically in figure 4.13 below.

4.3.2 Remote laser head.

This unit was assembled from separate TSI components. The source of light was via a fibre optic link from the main LDV unit: the two beams within the fibre optic cables were converted back into free air beams by the TSI unit 9263. They were then directed along a simplified set of transmitting and receiving optics. The components of this small head appear in figure 4.13 and are listed below:

used, because pre-filtered water was used (from the laser coolant drain) to top up the system. Water fouling was minimal because of the extensive use of plastics in the system, the deliberate choice of which is discussed below.

4.3 Velocity measurement sub-system

4.3.1 The main Laser Doppler Velocimeter

The unit used for testing was a TSI 9100-7 series laser: a two component, back-scatter unit, which operated in conjunction with a Coherent Imnova 90, 5 W multi-band argon-ion laser. A photograph of the system mounted on the traversing bed is given in figure 4.12. A full list of the components comprising this unit, including their arrangement, are given in appendix F, but a brief summary of the functionally important items are given below:




Figure 4.12: The LDV mounted on the traversing bench.

- (i) The LDV had a pair of frequency shift units, one for each channel. These each consisted of a TSI model 9182 Bragg cell mounted in line with the optics, and a TSI model 9180 electronic control box.
- (ii) A model 9143 field-stop system was included, to ensure elimination of erroneous reflected signals when measuring near to the pipe wall. This component focused the returning light through a small pin-hole, thus only admitting signals from the measuring volume.
- (iii) A large beam expander, model 9189 was included to ensure accurate beam alignment and a small measuring volume. The large aperture of this unit

4.2.5 Catchment tank and temperature control

The catchment tank was partially shown in figure 4.11. Its major function was to contain the flow from upstream. The tank had a capacity of 1 000 l, and was thus large enough to contain water from the entire rig (important during cleaning, or during a malfunction). Other than simply holding water, it was also concerned with the regulation of temperature in the rig.

Temperature control unit

Two large immersion heaters [Hottube, power 5 kW each] extended into the lower portion of the tank, one from each end. A temperature probe was immersed in the water, and a small pump recirculated water from the bottom of the tank to the top. The probe and heaters were controlled by an accurate temperature control unit [TCU, TCR-1], to 0.1°C . This unit actuated two large 30 W contact breakers, which controlled the supply of 380 V, 3-phase current to the heaters. The temperature was regulated to 25°C for the entire test.

High capacity heating units were used to allow rapid correction of temperature disparities. The system was completely self-contained and did not need any computer intervention. Temperature was regulated in the bottom tank because this was the most permanent place that the water resided. During testing, water was pumped repeatedly from this tank upstream; the initial circulation of fluid each day before testing brought the entire apparatus to equilibrium. The repeated movement of water during testing kept it that way. Another important reason why the temperature elements were not mounted in the upstream tank (apart from the complicated internals of the tank) is that this would have induced strong convection currents in this tank, and thus disturbed the symmetry of the settling flow.

4.2.6 The pump

There was a need to return water from the lower tank into the pressure reservoir. To do this a simple swimming pool pump was used. This unit had all plastic working parts, to prevent fouling of the water by corrosion. As mentioned, this unit returned the water via the test section, which it did through a one-way valve feed into the pipe between the flowmeter and valve network. This valve stopped fluid flowing back through the pump when the rig was under pressure, but admitted fluid from the pump into the pipe when the rig was not pressurised.

A filter (of 1 or 5 μm) was mounted after the pump in a secondary branch of the pipe, and water could be diverted through this filter by closing two valves. Its function was to clean the water (when new water had been added to the system) or remove seeding particles and sediment. However it was not often




Figure 4.11: The valve network.

Large electro-pneumatic valve. This device (6) - a Mittech 1 inch unit - was similar to the unit (4) above, but larger. It had a 2 inch bore supply, and operated on the same principle as the smaller unit. It was fitted with oversized air supply pipes in a bid to increase its response; the response time of the unit was about 0.3 s closed-to-open. The branch of the valve network containing this valve also contained a diaphragm valve (7) identical to the component (2), and a ball valve (8), the same as unit (3).

Flow left the network by way of two pipes (9), from where it passed into the catchment tank (10), shown partially here. The ends of the pipes were submerged in the tank to avoid air getting into the valve system during conditions of no flow.

The solenoid and two flow valves were controlled from the main computer. Basically these three units could be controlled in any desired manner simultaneously. Practically however, only one of the valves was used at any time. The throttling valves on each branch allowed a manual setting of the flow range, especially in the case of the solenoid. Setting of the range was also important for the pneumatic units because it allowed the valves to operate over a fuller range of positions, thus allowing maximum sensitivity of control. Air was supplied to both valves from a regulator (not shown).

Flowmeter

The flowmeter was a commercially available electromagnetic device [Endress & Hauser Speedmag - 40 mm throat] with a maximum flow-range of $12.6 \times 10^{-3} \text{ m}^3/\text{s}$ and a dynamic response time constant of 0.01 s. It is a non intrusive device, inferring the flow by measuring the electric current caused by the change in magnetic flux of the liquid passing through a strong magnetic field. The signal that was transmitted from the flow unit shown here was a proprietary serial format. This was interpreted by a remote unit (not shown) which produced an analogue signal representing the flowrate. The range of the unit could be modified from the remote head.

To increase the quality of the reading from the meter, two earthed brass rings were mounted upstream and down of the unit. The function of these was to remove any charge from the liquid that might have affected the readings.

From the flowmeter the pipe continued downwards, through a one metre length of 10 mm wall thickness high pressure rubber piping (to absorb any pressure pulses due to water-hammer in valve shutdown) and into the controlling valve network.

Valve network

A photograph of the valve network is shown in figure 4.11. The flow enters the system from the right (the down tube from the flowmeter is just off the right side of the photograph), and leaves through the two vertical tubes on the left.

The valve network was the primary flow control mechanism for the entire rig, and consisted of three branches:

Solenoid valve. This is indicated by (1) on the figure. This valve was a fast opening ($\approx 35 \text{ ms}$ specification), electrically operated solenoid valve with 2 inch threaded fittings. On the solenoid valve branch was a diaphragm valve (2) and ball valve (3) to allow manual throttling of the flow.

Small electro-pneumatic valve. This (4) was a Mittech $\frac{1}{2}$ inch throat pneumatically actuated unit. The device was electronically controlled via a 4–20 mA signal from the computer. It had a response time of $\approx 0.2 \text{ s}$ from fully closed to fully open with a supplied air pressure of 6 bar. Its branch of the network contained a ball valve (5) and shared the diaphragm valve (2) with the solenoid branch. This unit was added in the middle of testing, to allow measurement of the lower flowrates with better control.




Figure 4.10: The flowmeter preceded by the stilling chamber. The direction of the flow is down. A cross-sectional schematic of the stilling chamber is shown on the right.

Stilling chamber

The chamber was introduced to ensure as uniform conditions as were possible through the flowmeter section. Basically it consisted of an expansion (1), a pleenum (2) containing flow straightening cocktail-straws, and a smooth contraction (3) leading into the flowmeter. From the cavity (2) led an outlet (4) to a pop-off valve. This unit was incorporated as a safety measure to allow the bleeding off of excess pressure in the system.

The exit of the stilling chamber incorporated an integral flange to clamp the flowmeter in place, in conjunction with a matching flange on the far side of the meter, which directed the flow back into a section of PVC piping.

The stilling device was considered necessary because of the proximity of the flowmeter to a bend. It was not intended that the device produce laminar flow; on the contrary a turbulent profile was probably more desirable. The expansion was abrupt enough to induce separation in the flow. This, along with the straws and contraction was hoped to have removed the asymmetry from the flow.

was made of the temperature of the effluent liquid, but there was found to be no appreciable temperature variation in the fluid, even after substantial settling times. What was occurring was that the thermal *momentum* of the water jacket was sufficient to maintain approximately isothermal conditions in the pipe. The system had to initially be brought into equilibrium, by pumping the working fluid repeatedly through the pipe; conduction of heat through the glass then brought the two water bodies into thermal balance. Thereafter the flow in the pipe maintained this balance. The initial pumping was necessary anyway, to disperse the seeding evenly in the flow. Thus effective temperature control was attained within the system without the use of the complicated earlier system.

Rig support

Both the tank and the eight pipe units were supported on a beam arrangement, using adjustable pedestals. The tank was mounted on a frame that could be slid backwards and forwards, or adjusted to make it level. This structure is visible in figure 4.1 above.

The supports for the individual pipes can be seen attached to the main beam in figure 4.8. Each pipe unit was held at the flange by an adjustable support. Vertical adjustment was by means of a threaded rod, while crude lateral adjustment could be achieved by shifting the support sideways on the beam before clamping it in position. Fine adjustment of lateral position was through grub-screws mounted in the flange cradle.

The rig was aligned by assembling the pipe sections from the tank end. Each unit was adjusted to be horizontal and parallel to the last before tightening the flange bolts. The tank had very limited vertical or lateral adjustment, and thus the pipe had to be adjusted to match the tank.

4.2.4 Flowmeter & valve network

From the end of the test section, the pipe passed through a hole in the far wall. Here the final flange joined to a section of standard high-pressure PVC tubing, of outer diameter 50 mm (inner diameter was ≈ 45 mm). The diameter was reduced gradually, over about 200 mm, by chamfering the inner surface of the PVC piping. This was to lower the frictional resistance of the system, and remove any sources of disturbance from near the test section.

The horizontal pipe led into a gradual 90° bend, which turned the pipe downward. The pipe then led into a flow straightening chamber before passing through the flowmeter. The stilling chamber and flowmeter are shown in figure 4.10.

The flanges shown in figure 4.7 had provision for the attachment of a square box. Long thin glass panels were cut from 6 mm untempered float glass. The thickness was chosen because this was closest to the natural float thickness during manufacture⁵, thus minimising optical distortions. For the same reason the untempered glass was chosen - the heat of the tempering process introduces ripples in the glass. Float glass was chosen over Pyrex because the optical quality of the surface was higher in the former, and because of the price premium of Pyrex.

Initially the entire pipe was enclosed with four pieces of glass. However, the inability to clean this configuration and the difficulty in sealing the joints resulted in the lid being removed. Each section of pipe thus had a glass box surrounding it, which could be simply filled from above. The pipe sections can be seen clearly in figure 4.8.

Temperature control of the pipe

The fluid in the optical boxes was initially to serve a dual purpose: as well as reducing optical distortion it was to be used to control the temperature of the fluid in the pipe and thus retain isothermal conditions in the system. It was decided to do this because measurement on a previous rig, of the effluent fluid from the pipe that had been left to stand for several minutes (necessary settling time for reducing disturbances), showed a variation of several degrees, dependent on the ambient temperature.

Initially, fluid was pumped through the boxes from a temperature controlled reservoir, and returned to the reservoir afterwards. However, a number of problems arose with this system:

- (i) The driving pressure needed to circulate the water caused problems with the sealing of the boxes. The large surface area of these boxes caused the faces to expand, thus requiring clamping.
- (ii) The water jacket surfaces became fouled with algae and water sediment, and cleaning was impossible without opening the boxes.
- (iii) Air bubbles from pumping, and from dissolved gases in the water formed on the glass surfaces, disrupting optical access.

Due to these problems, it was decided to abandon the recirculation approach, and simply open the boxes. This allowed easy cleaning and removal of bubbles. However the primary goal of temperature control was defeated. In order to estimate the temperature variation with the water jackets open, measurement

⁵To make thinner panes, float glass is stretched from this initial thickness, with resulting distortion.

explicitly de-pressurised. Thus the seemingly simple **PRESSURE** operation in reality required continual control. This was only possible because of the nature of the design.

5.2.2 The RUN command and the run file

Most of the commands shown in the tables are self explanatory. The most fundamental of these was the **RUN** command, which conducted the testing. During use of the other macro commands, the rig remained essentially 'dumb'. For instance the **RETRACT** command would retain control until the tank was full, or the user pressed a key. The rig operation under macro control was thus simple and not time-critical. The state of the rig after each command was static: if an error occurred in macro execution the rig would never be in a dangerous dynamic state. The **RUN** command used a pre-compiled *run file* to control elements of the rig during the highly dynamic testing phase. To execute a test, this command was passed the name of the run file as an argument, for example **RUN:thistfile** would run the experimental test defined by **thistfile.mov**.

The syntax of the run file is given in section E.11. Simply put, the experimental test was defined by a text file that stated the timing and sequence of tasks required during testing. Each line of this file required three parameters: the first was the time of execution, the second the command, and the last was the command parameter. This file had to be compiled before it could be executed. The compiler checked for syntax errors and produced a binary file that allowed quick execution³. This compiled file was loaded from disk when the **RUN** command was invoked. Subsequent lines had to have identical or increasing values of time to the preceding line. Any exceptions to this were stripped out by the compiler. The command part of the argument could be any one of the three-letter character strings given in table E.3, and the nature of the third parameter argument is also shown in that table. The detailed structure of the testing subsystem is described in appendix E. In essence, each command in the run-file performed a basic task at a certain time. Some of these were simple actions, for instance opening or closing a valve, or starting acquisition from the LDV. Other, more complex tasks, such as the dynamic control of the flow, could be fully controlled from the run-file. The flow was allowed to vary in three ways: polynomial, sinusoidal and exponential. These are shown below in equations (5.1).

$$\begin{aligned}
 Q(t)_{\text{poly}} &= Q_{\text{max}} [k_0 + k_1 \times t + k_2 \times t^2], \\
 Q(t)_{\text{trig}} &= Q_{\text{max}} [k_0 + k_1 \sin(k_2 \times t)], \\
 Q(t)_{\text{exp}} &= Q_{\text{max}} [k_0 + k_1 \exp(k_2 \times t)].
 \end{aligned}
 \tag{5.1}$$

³Interpreting the test file during testing would have been both slow and dangerous.

Further to this it was specified that the commands be readily combinable into short programs or 'macros', thus allowing automated execution of sequences of commands. Underlying the interpretative front-end, time critical control and acquisition was to occur unseen by the user.

5.2 The main control software

Due to the complex nature of the main software, space will not allow a detailed description of all its component parts, and the reader is referred to appendix E for further detail. This section contains a brief functional description of the elements of the software important to the control and the acquisition.

5.2.1 Command interface

From the user's point of view, the main software consisted of a command terminal at which control commands could be entered. Below this terminal there was a status area, where the state of the rig, and other important parameters were presented. The terminal window had a prompt at which the control commands could be typed. Previous commands scrolled up the screen (much like a UNIX terminal or DOS screen) and a buffer stored the last 30 commands, allowing the user to recall common instructions quickly using the arrow keys. The commands were all entered using a simple syntax, and the full command set is given in table E.1 of the appendix. Each command was concerned with a particular task; for instance the command **PRESSURE** would cause the tank to be pressurised. Every component of the rig could be controlled from the command line, and further commands were concerned with altering parameters such as conversion factors and counter settings. All commands were however *static* in nature; that is each performed its action and then returned control to the user. The highly *dynamic* control of the acquisition was the responsibility of a single function - the **RUN** command, which is described in detail later.

The commands could be combined into short programs or macros. Essentially a macro was a text file containing lines of commands. The macro environment also allowed the use of some extended commands concerned with program control and looping. The basic macro syntax is given in section E.11, and the extended command set is shown in table E.2. The looping and other program control commands were very basic. Essentially there was a crude **FOR** loop, a **GOTO** statement and a user input **DECISION** command, but this was sufficient to allow fairly complex control of the apparatus. Command execution occurred at a high level. All the while low level monitoring of the rig ensured safety and flexibility which would not have been possible had the commands been for instance separate programs. For example the **PRESSURE** command caused the rig to pressurise and then returned user control. However the low level control elements continued to monitor and control the pressure until the system was

The above description of OOP is very brief, and there are detailed books available on the subject [Rumbaugh *et al* (1991), Jacobson *et al* (1992), Booch (1994)]. This digression on the OOP philosophy was necessary to explain how the complicated data acquisition system could be written within a year without any major flaws and problems. Basically the philosophy allows a very robust and generic system to be developed quickly, if the correct object abstractions of the real world are applied. A gross measure of the success of the approach is the number of lines of code (excluding comments) generated. In the present case over 22 000 lines were written. The resulting software is extremely robust, and the only places where problems did arise were explicitly due to short-cuts taken to save time, not due to unknown programming errors.

5.1.2 User interface

The previous Hewlett-Packard 9836 system was briefly mentioned in the previous chapter. The operating system of this machine was essentially a BASIC interpreter, and the acquisition software was a simple set of BASIC programs. This system, while very slow, held some innate advantages. For instance a BASIC command to operate the pump could be issued directly from the command prompt, as well as being included in some executable program. In the current design it was decided to avoid a more usual inflexible menu and windowing approach commonly seen in commercial software packages, simply because the rig control and acquisition was inherently a command based process. In order to retain the flexibility of the HP paradigm, while allowing high speed real-time control, a simple command processor front-end to the package was envisioned; effectively an emulation of the old HP system. The user screen is shown in figure 5.1 below.




Figure 5.1: The user interface screen.

The design specification was to have a textual user interface at which the operator could type in appropriate commands (i.e. **RETRACT** to retract the piston).

5.1.1 Object-orientation

At the outset of the project, the Object Oriented Programming (OOP) technique was developing commercially, and an object oriented version of Borland Pascal (version 5.5) became available. The OOP philosophy was adopted in its entirety for development of the system. In order to understand the merits of object orientation in software design, a short explanation of its principles needs to be given.

Fundamentally, the OOP technique allows the concept of *encapsulation* in software design. Instead of the common variable-passing, or global variables of say FORTRAN, the OOP philosophy works with objects. Data and code are no longer separate, but are combined into entities. In a sense the object 'knows'² about its own behaviour and how to implement it. For instance, instead of having a function that, say displays data passed as an argument to it (traditional approach), with OOP the data knows how to display itself.

Objects usually represent elements of the real world; for instance an object would be created for the stepper-motor drive. The internals of the object would concern itself with the inner workings of the elements for which it was responsible (i.e. writing specific commands to the stepper control box), while the interface - the commands used by other code invoking the object - would be generic. For instance in the stepper-motor case, the interface would consist of commands such as `steppers.Move(200)`; thus the details of the stepper control are hidden and unavailable for direct manipulation. The user of an object simply calls high level, generic and common-sense commands. This is termed *encapsulation*. It allows sections of the code to be debugged in isolation because of the modularity of design the method enforces.

Another even more powerful mechanism of the OOP methodology is what is termed *Inheritance*. A generic *ancestor* object can be set up that represents the behaviour of, for instance, all movement devices (stepper-motors, linear actuators, etc.), but does nothing - it is effectively a 'place-holder' for all objects with similar behaviour, and merely defines a set of functions that all movement devices have in common. Devices such as the stepper-motors or the linear actuator share the basic behaviour of this abstract object, but add implementation-specific internal operations to the object (for instance the `Move` command must actually move the stepper motor). Thus these specific objects 'inherit' from the base object, but extend its behaviour. The power of the approach lies in the way the inheritance is used. Software using the object need not be concerned with what type of object it is dealing with, so long as it uses the generic ancestor interfaces. To continue the example, the program can simply tell a movement device to move (`AParticularDevice.Move(...)`). If that device is a stepper-motor, specific code will move the stepper, while if it is the linear actuator it will also move.

²The anthropomorphism seems strange but is necessary to explain the basic principles of the method simply.

Chapter 5

The acquisition & control software system

The list of controllable or measuring hardware components given in the previous chapter was extensive. These items were all connected to the main control computer, and thus the high level control of the entire rig was the exclusive responsibility of a single software program. This chapter outlines the general philosophy behind the design of this program, and then briefly describes its user interface and operation. The description of the internal software structure is cursory, but a detailed top-level design is given in appendix E.

The much simpler acquisition program resident on the acquisition computers is also discussed. The slave LDV machine - both hardware and software - in reality simply constituted a peripheral component of the main computer; however because of its importance, and because its resident software program was developed together with the main system, its description is included here.

5.1 Basic design considerations

The hardware choice for the system was an IBM compatible system as described in section 4.1, using the MS-DOS operating system. Many software development tools were available for this platform, and Borland Turbo Pascal was chosen because of its power, ease of use, and the familiarity of the author with it¹. The compiler used was initially version 5.5, but version 7 was available by the time the final code was compiled.

¹The "C" language was considered, but at that time the Borland product was not stable; other manufacturer's compilers (Microsoft) were too expensive, and the Borland implementation of Pascal was not any more limited than C was.

4.6 Summary

The design of the LDV supporting apparatus, rig and control hardware was described in this chapter. The design was very general, permitting a broad range of experimentation to be conducted. However, all the equipment described here was incapable of any complicated acquisition or control without the presence of some controlling software on the main computer. To this end a dedicated computer program was developed to control the rig, and this is described in the following chapter.

4.5 Overall design considerations

Avoidance of corrosion

Except for the front wall of the pressure tank, some small pipe fittings and the flow valves, the remainder of the rig was entirely non-corrosive. This was done to avoid poor LDV signals caused by fouling of the water by corrosion. In previous test rigs at the university, the use of a cast iron pump and many metal components resulted in poor and noisy LDV signals after a day of operation. The present system could operate for a fortnight without substantial degrading of the LDV signal, except for settling of seeding particles which resulted in a decrease in the data rate, which had to be counteracted simply by adding more.

Avoidance of electrical interference

Control cables to the valves carried a 2-10 V analogue signal, while the solenoid was actuated with a 220 V AC supply. Moreover the temperature control cables were carrying high current at 380 V. Thus in general there was a lot of electrical noise sources in the laboratory to interfere with sensitive signals such as those to the valves. To specifically avoid this problem, all low power analogue lines were shielded from sources of noise, by running them in separate earthed conduits.

Safety

Due to the fact that the system operated at pressure, safety was of concern. Various design decisions were made to minimise any danger, while not interfering with the working of the rig. An unavoidable risk was that of the pipe breaking. There was no way in the current design to avoid all the water from the tank discharging onto the floor, because no valve could be placed upstream of the glass pipe. However measures were taken to allow quick venting of the tank pressure should any such eventuality occur while the rig was under pressure. These were implemented in the computer control system, described in the following chapter. Concerning the possibility of power failure, all valves were chosen to revert to the safest possible position under such conditions; for instance the air solenoids on the back of the tank released air if power was removed, and all the downstream valves closed. As mentioned earlier, the lower tank could contain the entire rig's water, thus any discharge of the rig due to malfunction could be contained. All other safety measures were implemented via software, for instance the program would immediately abort any test if any error conditions were detected, for instance over-range on the piston position, or zero flow on the flowmeter when there should have been flow. These safety features are discussed in detail in the next chapter.

The solenoid valve required a 220 V AC signal to open. This was supplied from a four channel relay box that was driven by digital lines from the PC36 card. The signal from the card was amplified using a transistor bridge (the relay box had its own 12 V power supply) allowing the 220 V input to the box to be directed out via four plugs to various controllable devices. The solenoid was directly supplied with current from this unit. The other three channels were used for other components.

4.4.9 The flowmeter

The flowmeter unit was described earlier, in section 4.2.4. The remote measuring unit communicated to the control unit (mounted near the computer) via a serial link. A variety of signals were available from this unit, among them a 4–20 mA signal, proportional to the flowrate in the pre-set range. As with the valves, it was trivial to convert this signal to voltage. Here calibration was important, since the actual value of the voltage indicated flowrate. The AD card was factory calibrated, but it was uncertain whether this calibration was correct. It was decided to leave out direct calibration, rather to use LDV data to calibrate the unit later. This process, carried out in section 7.5.3 on page 124 showed that there was indeed an error in the flowmeter, quite probably due in part to the AD card. Calibration by this means was adopted because it revealed the *total* error of all the components of the measurement.

4.4.10 Pump & tank control

The downstream pump - a single phase, 220 V device - was simply controlled by a channel from the relay box described earlier. Various elements of the pressure tank shown in figure 4.3 were either controllable or could be measured. The supply and relief solenoid on the back of the tank (wired in parallel) could be opened or closed simultaneously by supplying a 220 V signal. The small bleed valve was also actuated by 220 V. The other devices on the tank were a pressure transducer [Watson Smith model 53701307, range 0–3.5 bar, 4–20 mA signal] and a rheostat for sensing the piston position. Both these units returned an analogue signal, and the acquisition channels used for these devices are shown in table 4.1. The rheostat, being a resistive device, had to be supplied with a voltage. A twelve volt line from the AD card was brought to the rheostat, and via a simple circuit, the output was brought into the range 0–10 V, for the entire range of rotation of the device. The position of the piston was inferred from the support linkage, using a pulley on the rheostat attached with fishing gnt to the linkage. This position information was used for safety purposes and to ascertain whether the tank was full, thus the imprecision of this simple arrangement was not important.

4.4.7 Analogue I/O system

Many of the peripheral control and acquisition devices required or returned an analogue voltage signal - the linear actuator above required a 0–10 V, while the flowmeter returned a voltage in this range. An Eagle electric PC30B analogue acquisition card was used to read in and put out analogue data.

This card was mounted in the main computer, and had twelve analogue channels: four for output and eight for input. The resolution of the input channels was 12 bits, in the range 0–10 V, while only two of the output channels operated at 12 bits, the remaining two only supplied 8 bit resolution. Interface to the card was via a 50 pin 'D' connector on the back of the machine. A shielded cable was made up, to take the signal connections from the back of the computer to a connector box. This connector box was properly shielded and earthed, and on its front face had twelve BNC bayonet connectors, to allow the separate analogue channels to be distributed via co-axial cable. The inputs and outputs were connected as shown in table 4.1 below:

Table 4.1: The analogue connections.

Direction	Resolution	Channel	Use
input	12	0	flowmeter
	12	1	tank pressure transducer
	12	2	piston position
	12	3–7	unused
output	12	0	large electro-pneumatic valve
	12	1	small electro-pneumatic valve
	8	2	linear actuator
	8	3	unused

The A/D card also contained 24 digital lines (another 8255 I/O chip on-board), some of which were used for control of the other elements of the rig.

4.4.8 Valve control

The two electro-pneumatic valves were controlled using the analogue output channels shown in table 4.1. This signal was 0–10V, while the valves expected a 4–20 mA signal. The one was easily converted into the other by the insertion of a 500 Ω resistor in series with the signal. This was done at the valve. High absolute accuracy of the conversion was not important because absolute valve opening had no meaning; rather relative positioning by the control software was the prerequisite. The resultant voltage range to operate the valve was thus 2–10 V, the lower signal just unseating the device.

4.4.5 Position control devices & limit switches

Stepper-motors are accurate 'feed-forward' devices, in that they can be moved to a chosen destination without monitoring their motion or position. However, in the current context, highly accurate information could not be gathered from their prospective positions, firstly because of backlash in the gearing system of the bench, and secondly because the motors were prone to missing steps from time to time. Considering the linear actuator, this device returned extremely crude position information, from a simple rheostat mounted on the gearbox.

In order to acquire accurate bench position information, linear position encoders [Tekel ROM50C/100 series] were attached to the transverse and vertical axes, while a rotary encoder [Tekel TK661] was fixed to the axial co-ordinate of the bench. The rotary device was used to reduce costs (a 4.5m axial encoder was prohibitively expensive) and because this axis did not require high resolution.

All three devices were capable of *resolving* to 10 μm , although temperature variations and structural flexibility of the rig reduced the accuracy of this figure to about 50 μm in reality. Signals from all three encoders were relayed directly to a custom-built acquisition card in the main controlling computer. This device simply had three registers that could be read by the computer, each containing the current measurement from consecutive axes.

Limit switches were attached to the traversing bed in order to determine the maximum travel of the bench, and thus to allow interruption of bench movement. There were six switches in total - two per axis - and the signal from these elements was fed directly to the computer via six input lines on a PC36 digital I/O board. This board was effectively identical to half of the PC14 board described earlier, incorporating only a single 8255 I/O chip. Data from these lines informed the controlling program to cease movement upon receiving the appropriate limit signal.

4.4.6 Measurement of small laser position

The small traversing bench only had one degree of freedom. This meant it had to be positioned accurately near the flow and aligned with the pipe. The alignment procedure is described later, in section 7.2 of chapter 6. Neither the movement or position of this bench were governed by computer; rather manually, by adjustment of a knob. To determine the radial position accurately, a dial gauge was attached to the small bench and measured its relative movement in the radial position to within 0.01 mm. As with the large unit, this measurement was relative to some arbitrary bench starting position, and not to the pipe. The details of calibrating this distance measurement to coincide with the pipe origin was part of the data processing discussed in section 7.2 on page 103.

used to measure time by allowing its counter to decrement from value 65535 to 0 continually. The controlling computer program stored the number of times the counter 'clocked over' to 65535, and thus, knowing the counter rate, the time in seconds of a specific event could be ascertained. The chip could use a number of data sources for a time base: the computer's clock (inaccurate), an internal oscillator, or an external data line. To achieve a uniform time base between all three computers, this timing chip in each slave computer was synchronised externally with a signal from an identical chip on the main computer's PC14 card. Thus any need for synchronisation of time bases between computers was avoided and only the main computer clock needed to be calibrated.

The primary function of the slave machine was simply to act as a 'memory-box', with the ability of storing large quantities of LDV data when told to do so. In fact the use of a computer in this role was only decided upon when it was realised that the assembly of a simple memory buffer would be too expensive. The economy of scales resulted in an entire computer costing substantially less than the development of a dedicated memory box. The absence of hard-drives and other high specification peripherals was simply to keep the price low.

4.4.3 Stepper-motor drive units

The two horizontal co-ordinates of the traversing bed were powered by stepper-motors. These were controlled by a Digiplan CD30 two channel stepper-motor drive, which connected to the computer via a serial connection. Commands were issued to the box using simple ASCII strings, such as X,100, meaning "move the x-co-ordinate motor by 100 steps". The unit could move the motors at variable rates, simultaneously, as well as accelerating and decelerating them in a chosen manner. Internal buffering of commands was permitted, thus the box could be instructed in an asynchronous manner.

4.4.4 Linear actuator unit

In order to control the vertical DC actuator accurately, a DC motor controller, operating on the pulse-width principle⁷, was used. This unit accepted as an input signal a voltage in the range 0 - 10 V, with 5 V the 'zero' position. This voltage signal could be provided directly from the analogue/digital card (described later).

Both stepper-motors and actuator had a safety interlock that was connected to a latched relay control. The interlock had to specifically be defeated every time the system was powered on, and could be actuated by a button on the front face of the computer, thus immediately stopping any movement.

⁷Simply put, instead of reducing voltage to slow the motor, the unit supplies a square wave at high voltage, and simply varies the width of the power pulse.

4.4.1 LDV acquisition sub-system

This system as mentioned consisted of the two identical acquisition computers. These machines were each IBM compatible AT 286's running at 20 MHz. They had no hard drive - only a single floppy - but were fitted with 4 megabytes (Mb) of memory. They were each fitted with an Eagle Electric PC14 digital input/output (I/O) card, to allow communication with the LDV signal counter and to send data to the main computer. The PC14 cards used the *de facto* I/O chip; the Intel 8055. This chip allowed up to 24 lines of high speed buffered digital I/O. These cards allowed sample-and-hold acquisition to be carried out with the signal processors: the card could poll data independently of the computer CPU, which simplified the acquisition software design considerably. Figure 4.16 shows the data interconnections between the signal counter, polling computer and main computer.




Figure 4.16: The digital data lines between the polling computer, the signal counter and the host computer.

The PC14 card also supported a card-to-card, dual handshaking communication. In this mode, two PC14 cards in separate computers could be made to communicate with one another. This mechanism was used to allow transfer of data between the slave machine and the main computer. This connection is shown in figure 4.16, between CN2-A on the slave acquisition machine (on the left) and CN1-A on the host computer (on the right).

4.4.2 Timing

Time base information on the counters was provided by a timer chip on the PC14 card; the Intel 8253, a high speed 16 bit count-down timer. This chip could be

where $Q_{\max}$ was the maximum range of the flowmeter. In the current case, the flowmeter was set to have a maximum range of 2 l/s.

The commencement of the LDV acquisition is indicated by the line 0 slave 1; next the valve is pre-opened to position 1200 (the total range of operation being $800 \leq \text{valve} \leq 4096$), and closed loop control (0 PID 1) is instantiated. All three of these commands are executed sequentially at time $t = 0$. In this particular case the test was allowed to run for 75 seconds, at which time the closed loop control was stopped, as was the LDV data acquisition. The system automatically released pressure on the header tank, and ensured zero flow in the system at the end of the run-file; thus the last two lines of this file were in fact redundant.

The run-file given in figure 6.3 is only a particular instance of an exponential run-file; to achieve the various accelerations used during testing, the various flow constants were varied suitably, as was the duration of the testing. However, all the run-files and macros used in both the preliminary testing and the final experimentation were based on the form of those shown in figures 6.1, 6.2 and 6.3 above.

As mentioned earlier, the simple selection of flow constants did not guarantee that a particular flow could be attained in practice. An optimum combination of the control values and rig settings were needed to achieve a chosen flow variation. Having managed to attain a general method for generating exponential flow, a preliminary investigation using the LDV flow measurement device was conducted, in order to obtain a general picture of the flow phenomena in the pipe. The following section is concerned with these preliminary investigations, and the emergent broad conclusions.

6.3 Preliminary testing

The automation of the entire flow apparatus allowed a quick initial investigation of the flow; the control macro could be run as often as desired, with the LDV measuring volume positioned appropriately anywhere in the flow. The behaviour of the flow was initially completely unknown; to obtain an initial picture of the flow, a particular exponential flow variation was chosen⁴ and the LDV measuring volume was positioned far downstream in the centreline of the pipe.

Data sets resulting from these tests were firstly the flowmeter data by which comparisons could be made to the expected flow variation, and secondly an LDV velocity-time trace from the point of measurement. Firstly, the following section discusses the generic pre-processing necessary to make the LDV data more useful.

⁴The variation in fact corresponded to that shown for the test EXPT6 in the table 7.1 given in the following chapter.

Run-file control

The command `RUN:EXPF` present in both figures 6.1 and 6.2 is central to the actual running of the particular test. This command executes the compiled *run-file* named `EXPF`. The information contained in this run-file pertains to the actual length of testing, and the nature of the flow variation in time. The separation of normal control (interpreted macro commands) and run-time control (the compiled run-file) was described in detail in section 5.2, and was necessary to allow quick, safe and flexible data acquisition and real-time rig control. The run-file `EXPF` used in the macros above is shown in figure 6.3.

```
% Acceleration Movement file, Exponential E
%
% Values of the parameters for exponential ...
% pk1=0.051, pk2=0.03, pk0=0;
% Q = Q_max * (PK0 + PK1 * exp (PK2 *t))
%
0 PK2      0.03
0 PK0      0.0
0 PK1      0.051
0 PTY      2

0 slave    1
0 Valve    1200
0 PID      1
75 PID     0
75 Slave   0
```

Figure 6.3: A run-file to control the flow to a particular exponential acceleration, and acquire IDV data therefrom.

The `%` symbol indicates that the rest of the line is a comment. Each line in the run-file consists of three parts: the first number being the time of execution in seconds (zero being the commencement of execution of the run-file command `0 PID 1`), the second - mnemonic - argument indicating what action to take, and the final parameter the value the action must take. The run-file shown in figure 6.3 consists of two parts. The first four commands set the values of the control constants; the mnemonic `PTY` was set to value 2, indicating that an exponential acceleration was desired³. The constants `PK0`, `PK1` and `PK2` were set to the indicated values, in order to generate a flowrate variation of the form

$$Q = (Q_{\max} * [PK0 + PK1 \exp(PK2 * t)]). \quad (6.2)$$

³A value of 0 indicated a linear flow variation, while 1 produced a sinusoidal one. The other three constants had different interpretations according to the type of flow desired.

the actual testing, and is described below.

TRANSFER: Retrieve the data from the LDV acquisition computers and store it locally in memory.

VIEW: Execute the interactive graphical viewing shell, allowing the user to view the LDV and flowmeter data.

INPUT: VARS, ... Prompt the user for a file name and store it in the variable **VARs**.

SAVE: VARS Save the LDV and other data using the file name entered above.

The macro given in figure 6.1 needed user interaction to view the data, and then to enter a file name. In some cases where it was necessary to repeat an identical run a number of times, a single macro could be implemented to achieve this. Figure 6.2 shows a typical such macro.

INFO: Generic acceleration macro - multi-shot

RETRACT:

SETPRESS:2

CHANNELS:2

SETTLE:60

PRESSURE:

SETTLE:2

RUN:EXPF

TRANSFER:

SAVE:EFO_0000

SETTLE:60

PRESSURE:

SETTLE:2

RUN:EXPF

TRANSFER:

SAVE:EFO_0001

RETRACT:

...

END:

Figure 6.2: The macro to execute multiple acquisition runs on the acquisition system.

In this macro, the sequence of events of the interactive macro are repeated, but the interactive commands are omitted. The experiment is run multiple times by repeating the relevant commands (there were only very crude looping commands devised for the command interpreter). In the current implementation, the test was run twice before recharging the header tank. Data was explicitly saved to particular file names. By repeating the macro fragment, the test could be repeated an arbitrary number of times.

up from the bottom catchment tank. The system was then allowed to settle for a time; thereafter the run was initiated. At the end of the test, the macro transferred LDV data from the acquisition computers to main controller, displayed this information, and then saved it to file. The flowmeter data, which was acquired locally on the control computer, was also saved to file at this stage. Figure 6.1 shows a macro used in the current testing.

```
INFO: Generic acceleration macro - 1 shot
RETRACT:
SETPRESS:2
CHANNELS:2
SETTLE:60
PRESSURE:
SETTLE:2
RUN:EXPF
TRANSFER:
VIEW:
INPUT:VARS,Enter file name to save (MAX_XXXX):
SAVE:VARS
END:
```

Figure 6.1: The macro to execute a single acquisition run on the acquisition system.

The sequence initiated by this macro is as follows:

RETRACT: The header tank is recharged by pumping water up from the lower storage tank.

SETPRESS:2 The header tank control pressure is set to 2 bar. The actual pressure was governed primarily by a manual supply regulator, the control being implemented by means of a bleed from the tank. In this case the regulator was set to 1 bar. Setting the control pressure to 2 bar effectively overrode the computer's pressure control mechanism. This was to allow the pressure to increase continually during the early stages of the test², without being limited at 1 bar. This command did not cause the tank to pressurise; it merely set the control level.

CHANNELS:2 Use only the second LDV channel for polling. The argument was changed to 3 (logical 1 OR 2) in the case where both LDV channels were to be polled.

SETTLE:60 Wait for 60 seconds before continuing.

PRESSURE: Pressurise the header tank.

SETTLE:2 Wait for two seconds before continuing, to allow the pressure to develop to a suitable level.

RUN:EXPF Execute the run-file (named in this case EXPF). This run-file conducts

²The supply to the regulator was in fact throttled, such that its set pressure of one bar was only reached after about 10 seconds. In this way it was ensured that the pressure was lower at the beginning of the test, thus effectively allowing the valve to control more accurately here.

quick response, especially for the impulsive start⁴. It was decided to restrict the working pressure for all the tests to 1 bar, for the reason that a 1.5 bar working pressure was considered too close to the design maximum for the system. The chosen pressure delivered adequate response for the present purpose.

Valve choice was initially restricted by availability (the small control valve was not present at the commencement of testing), and so the early tests were conducted using the larger valve only. It was realised that the big valve was not optimum for the lower accelerations, and thus a smaller one was installed for the final few tests. As well as exhibiting better response at the lower flowrates, the new valve allowed lower initial impulsive Reynolds numbers to be achieved.

The choice of static valve settings was somewhat arbitrary. Crudely, these valves were adjusted so that the full working range of the control valve encompassed the full range of the test. From this initial estimate, the settings were further optimised by observing the resulting flow and making fine adjustments accordingly. In particular, the part of the flow most sensitive to these adjustments was the impulsive start and the flow immediately thereafter. Unfortunately, the behaviour of this part of the flow was coupled to the way the pressure increased at the start of the test [see point 1 above] and thus an iterative solution was embarked upon. The criteria were that the resultant flow had to show no overshoot or other anomalies at inception, while still achieving the chosen maximum flow rate later in the test.

6.2.2 Computer control

The operation of the rig and acquisition of data during a run was controlled completely by the main computer. As shown in chapter 4, the rig could be controlled by means of typed commands issued to the main control computer. Groups of commands could be combined in a file called a macro, allowing automation of any sequence of such commands. In the present case, control of the exponential flowrate and acquisition of the LDV data was implemented via a macro and run-file [See section 5.2 for an introduction to these control files]; the macro controlling the rig during priming and settling, while the run-file controlled the actual testing.

Macro control

The general form of the exponential control macro was initially evident, and all subsequent macros used in the study were based on the same macro template. In short, the macro ensured that the header tank was full by pumping water

⁴A high pressure driving head with flow restricted, compared to a low pressure driving head with lower restriction, are *not* equivalent, even if they result in the same steady flowrate. The dynamic response of the system (i.e. dQ/dt) is directly proportional to the driving force, all other things being equal.

6.2 Flow control

Perhaps the most important aspect of the experimentation was the accurate attainment of an exponential flowrate. The dynamic nature of this requirement necessitated careful consideration of a number of factors concerning both the flow rig and control system. These elements are discussed below.

6.2.1 Rig settings

As discussed in the previous chapter, the rig was fully automated. The user had control over only a limited number of parameters in order to achieve a desired flow variation. The desired temporal flow variation was exponential; the control software had the ability to control the flow exponentially as

$$Q = Q_0 \exp kt, \quad (6.1)$$

and the user was capable of selecting the constants Q_0 and k at will. However, any arbitrary selection of values would not necessarily result in a flow that could be realised. Many other factors contributed to a chosen flow being possible. These were:

- (i) **Pressure in the header tank:** This had a direct effect on the maximum flowrate achievable, and by that means, it affected the range of operation. Thus, pressure played a major part in the selection of the optimal configuration. More subtly, it was possible to allow the pressure to vary somewhat during the test, being lower at first, and reaching a fixed value well into the test. This was implemented by throttling the pressure supply line to reduce the flowrate of the supplied air, alleviating to some extent the valve range of movement.
- (ii) **Choice of valve:** There were two different size electro-pneumatic control valves to choose from: the smaller one was more suited to the lower flow ranges whereas the larger one worked better for larger flowrates. Furthermore the response and thus maximum attainable flow increase differed between the valves.
- (iii) **Static valve settings:** The valve network had a number of diaphragm and ball valves in series with the control valves, [see section 4.2.4] each of which could be set to modify the range of the control valves.

The rig pressure could be varied in the range 0 to 1.5 bar above atmospheric. The maximum Reynolds number achievable at the high pressure (with no downstream restriction) was well in excess of 80 000. Preliminary considerations dictated that the supply pressure should be as high as possible, in order to ensure a

invasiveness and linearity of the tool offset this disadvantage somewhat. In order to clarify specific issues concerning this method of measurement, a nomenclature is introduced to discriminate the various elements of the testing. Below is listed a brief summary of the naming conventions to be used in this study.

A clear distinction needs to be made between the various aspects of the experimental testing to avoid confusion. In general the entire study consists of many different measurements of a few flow variations. Terminology must be introduced to clarify what aspect of the experimentation is being referred to. To do this the following terms are defined:

RUN: A single measurement taken. A run requires that the rig undergo one complete measurement cycle, from priming to settling, and finally to the realisation of a particular flow variation in time within the pipe. A single set of LDV and other related data are usually acquired per run.

TEST: A particular variation of the flowrate in time. A test encompasses all runs that possess the same mean flow variation in time.

STATION: An axial position in the pipe at which several runs may be conducted. Usually a set number of runs will be conducted at a particular axial station for a particular test.

PROFILE: A set of runs from a single test, each at the same axial station in the pipe, and offset radially one from another such that the radial variation in velocity across the pipe is accurately described by their combination.

TIMING: This term refers to a set of runs at a station used to infer the inception of transition in the pipe at that station.

Thus in the defined terminology, the entire study consisted of a number of tests. Each test consisted of many runs, each conducted at a particular axial station in the pipe. At some of these stations, the runs defined a profile, while the runs from the remainder of the stations were used to infer timing information. There were two aspects to consider when conducting the tests. The primary one was adjusting the rig settings and control parameters to achieve a pre-defined variation in the flowrate with time; subsequent to this the positioning of the LDV and method of data acquisition had to be planned. The chapter concerns the preliminary flow control and data acquisition, while chapter 7 that follows describes the the final LDV measurements.

Chapter 6

Preliminary experimentation

The overall goal of the experimentation was to capture the behaviour of the flow in an exponentially accelerating pipe entrance flow. Two broad aims were initially apparent: firstly, the radial variation of the laminar velocity profile in time at various points in the pipe was of interest, and secondly the initial manifestation of deviation from this laminar state within the pipe entrance was important.

The general nature of the flow behaviour was initially unknown. Preliminary experimentation was conducted to ascertain the approximate behaviour of the system. Leading from this a methodical approach was adopted, in order to capture a wide range of the behaviours of the system. This chapter describes the methods employed to measure the flow; furthermore the acquisition of auxiliary information such as the flowmeter data is discussed. The major sections that follow each discuss a separate aspect of the experimental method from a general perspective. Emerging from the investigation in this chapter, an experimental methodology was decided upon, and this is discussed in chapter 7 that follows.

6.1 Terminology

The LDV is a point measurement device, only capable of making a measurement at a single place in the pipe at any one time. Because of this limitation, in order to investigate multifarious aspects of exponentially increasing flows for any *particular* flow variation, an identical experiment needs to be run repeatedly, with the LDV measuring volume positioned at different points in the pipe for each run. The composite picture of the developing flow system can then be assembled from these individual data. Because of this, the method is sensitive to anomalies in the flow; non-repeatable occurrences will be shown at only some measuring positions, and naive interpretation of the resulting composite data can lead to false conclusions. However, as discussed in the chapter 3, the non-

0.5 seconds, thus the time resolution was about 10 μ s. The use of a further two bytes for time measurement allowed test times of up to 9 hours to be measured.

The expanded memory of the computer was used to store the resultant data. This was accessible as 64 Kbyte pages, thus after 65 536 data bytes had been acquired, a time-consuming 'page-swap' had to be implemented. This resulted in small gaps (≈ 10 ms) after every block of data. In the current experimentation (data rate < 500 Hz) the gap was, however, so small as to be unnoticeable in the acquired data. Although the data rate ceiling was never achieved in the present readings, the speed of the loop was necessary to ensure that the data was not aliased. In the previous IIP system the input port was polled at about 120 Hz, and the time of acquisition of the data was taken as the time of sampling, even though the data may have been latched on the port for far longer. This system was thus prone to temporal aliasing. In the current system, aliasing would have occurred at about 10 kHz. As mentioned, the maximum data rate during the test was only about 500 Hz. The current system showed a marked increase over its predecessor, both in data rate ($\approx 200\times$ improvement) and in its ability to poll several channels independently.

5.4 Summary

An extremely general data acquisition and control system was developed using the object-oriented approach, which allowed the creation of a complex and integrated acquisition and control system, essential for the dynamic nature of the testing. Although the design was costly in terms of development time, this was necessary in order to generate and measure complex flow variations with ease. A number of advantages were apparent:

- (i) The flexibility of the interface and the resultant ease with which macros and run-files could be assembled and used, allowed quick experimentation with the various rig parameters (control constants, accelerations, etc.) that a more rigid system would have permitted.
- (ii) Automation ensured repeatability by removing human error and bias, and thus allowed for accurate testing.
- (iii) The system freed the user to concentrate on the more important aspects of experimentation, such as the data itself.

However, a far stronger case can possibly be made for the design; that the experimentation would not have been possible without the development of the rig and controlling software. The detailed description of the rig and controlling software was necessary to describe how the complex exponential flow variations were achieved. The following chapter now returns to matters more closely associated with fluid mechanics; the testing itself.

Purge memory. The amount of data stored on the slave computer was indicated by an index into an array of data. As data was acquired, this number was incremented by one for each point measured. Invoking the purge memory function set this counter to zero, effectively allowing data in the computer memory to be overwritten.

Set test time. The polling procedure, when initiated, continued for a set duration before terminating. This time (in seconds) was set using this command.

Poll data. This, the most important operation of those listed above, was concerned with the actual acquisition of data from the counter. When this command was issued, acquisition was immediately initiated, and continued for the time set by the previous command. Running a test removed control from the main loop for the duration of the test, and the acquisition could thus not be stopped. This inflexibility was necessary to allow the acquisition loop to poll data at the maximum permissible rate, without having to check the status of the communication port.

Get number of points. After polling, the number of points could be obtained from the computer. Upon receiving this command the acquisition computer transmitted this number down the network to the main computer. This was to allow for memory allocation on the main machine prior to transmission of the data.

Transmit data. The transmit command caused the computer to make the polled data available on the data link. After issuing the command, the main computer made itself ready for receipt of a block of data from the slave machine. The size of the block sent was restricted to 64 kilobytes or less. After the transmission of each block the main machine verified receipt, upon which the slave machine transmitted the size of the next block. This continued until the slave transmitted a zero block size, indicating completion. Each data point consisted of 7 bytes: three for the velocity from the LDV, and four for the time value.

The acquisition loop was concerned with polling the counter during a test, and transferred any readings to memory. It also measured the time at which the counter was read, and transferred this value to memory as well. The time data consisted of 4 bytes of information - 2 from the timer chip, and two indicating the number of times this timer had 'clocked' over. The velocity data was simply the three bytes of information from the signal counter. The loop terminated when the maximum time period was reached. By optimising the loop it became possible to poll the data at rates as high as 20 kHz. The timer clocked over every

5.3 Slave software

The job of the slave computers was to acquire data from the signal counter via the data link and store it in memory. When required, these data had to be sent to the main computer, which would process them further. Each action was controlled completely by commands sent from the central computer down this link to the slave machine.

The slave computers had no hard-disk, thus they booted the MS-DOS operating system from a floppy disk. This floppy disk also contained the controlling program, which was executed automatically after booting was complete; thus no user intervention was required upon start-up of these units. Basically, the program consisted of a loop, the software continually polling the communication line to the main computer for a command. If such a command was sent by the controlling computer then the acquisition software would carry it out. For instance, a command could be sent to tell the acquisition unit to purge its memory of old data. Upon receipt, the acquisition computer would perform the task and then return to the waiting state. A simple schematic representation of the program operation is shown in the flowchart in figure 5.2.




Figure 5.2: The data acquisition flowchart.

Referring to this figure, the program starts when the machine is powered up, and waits for a command. When a command is received the relevant portion of code is executed depending on the value of that command. When the command has finished executing, the program control returns to the command polling loop. This cycle continues until the exit command is issued, at which point the program stops executing. The six commands shown below are the more important functions of the system.

Reset counter. This command set the signal counter into an idle state, by stopping data polling of the counter. It served to allow the counter to read data independent of the acquisition process, which was useful because it allowed the data rate counter on the unit's front face to poll the current data rate, a feature necessary during adjustment of the LDV.

Here $Q_{\max}$ is the maximum range of the flowmeter and t was the time in seconds. The type of variation and the various constants in the equations could be defined from within the run-file. After defining the type of variation, the closed loop control could be started or stopped by the PID command. Different sequences of the equations could be used over different portions of the test, allowing the assembly of a general flow variation in time. For the current testing only the exponential variation was used. The run command also polled analogue data from the flowmeter (or other selected sources) during testing which was stored locally in memory. All events and data acquisition were synchronised off a single timer chip, which operated to 10 μ s resolution. Upon completion of the run-file or abnormal termination (key pressed, tank empty, etc.), control returned to the command-line interpreter and the data could be recovered both from the LJV machines and locally from the analogue sub-system at leisure, using further macro commands.

In conclusion the dividing of program control into two separate types: non-critical rig administration and time-critical acquisition and control, was the optimum solution to allow accommodation of the peculiarities of both

5.2.3 Automatic data conversion

The rig was fully automated, and the processing of the raw data from the slave computer was done according to conversion factors and parameters either calculated by the computer, or loaded from a configuration file.

The configuration file contained parameters such as the light frequency and free-air beam angle of the laser for each channel. Also contained here were parameters such as the default frequency shift and cycles per burst. The frequency shift, cycles per burst and other parameters could also be changed using macro commands. The free-air beam angle was not the actual angle of intersection of the beam within the flow, due to optical distortions caused by the mutual glass, air and water interfaces. A module of the program was concerned with determining the position and angle of the control volume using the specified geometry in combination with the principles of Snell's law of refraction. Using geometrical information retrieved from the configuration file in combination with information obtained from the bench position sensors, the actual position and intersection angle were calculated. This sub-system is discussed in further detail in appendix E. The velocity data was thus corrected according to the change in beam angle within the pipe.

Unfortunately, the small LJV head had no position encoder, and this mechanism could not be used in conjunction with measurements from it. However, measurements made in the horizontal direction on the major diameter were subject to a *constant* beam angle deviation; thus a constant correction could be made to the data later.

Furthermore, an entire test was referred to by a four letter name, the two letter prefix being an abbreviation of that name. Table 7.1 gives the full names and abbreviated prefixes of the tests.

Table 7.1: The naming convention for the experimental runs, showing their abbreviations and file names. The terms x and y indicate arbitrary numeric values.

Test name	Abbreviated suffix	File name template
EXPE	EE	EE x $yyyy$
EXPF	EF	EF x $yyyy$
EXPG	EG	EG x $yyyy$
FXPA	FA	FA x $yyyy$
FXPB	FB	FB x $yyyy$
FXPC	FC	FC x $yyyy$
FXPD	FD	FD x $yyyy$
GXPA	GA	GA x $yyyy$
GXPB	GB	GB x $yyyy$
GXPC	GC	GC x $yyyy$
GXPD	GD	GD x $yyyy$
HXPA	HA	HA x $yyyy$

The component x of the file name in this table was used to denote the axial station of the run. These stations were simply enumerated according to the order of testing. This enumeration was from zero to nine, thereafter the alphabet was used, beginning at 'A', and continuing until all the stations had been labelled. The components $yyyy$ referred to the particular run at a station. Repeated runs at a station were again enumerated using this component, starting from 0000, and continuing as 0001, 0002, etc. Thus for instance a series of 50 runs for test HXPA and station number 'C' defining a velocity profile would be labelled HAC'0000 through to HAC'0049. Where runs were conducted at varying radial stations the radial position of each run was stored separately in a log book; complications arising from trying to encode radial position into the file name were thus avoided. The unusual naming convention for the prefix (i.e. FXPB) resulted from the convention used in the preliminary testing. The very first preliminary test was named EXPA, where the 'EXP' stood for exponential and the 'A' denoted this as the first test of the series. At this point it was decided to introduce a further level of distinction: all the 'EXP' tests had the same starting Reynolds number, the following series of tests, with a different impulsive Reynolds number were thus denoted 'FXP' and so on. Differentiation between data from different sources (LDV, flowmeter, etc.) was by means of the file extension or suffix. The MS DOS operating system allows files to have up to three alphanumeric characters in this suffix. Table 7.2 lists the suffixes that were used, and the data type to which they applied.

The data acquired during a run was automatically saved to computer hard disk according to the naming convention given by table 7.2. The acquisition channels

Chapter 7

Final testing and data processing

The preliminary readings of the previous chapter indicated that the instability occurred near the pipe inlet, in the boundary layer. Using these observations, a methodology was decided upon, in order that the maximum information concerning the flow - both the laminar variation and the transition phenomenon - be captured efficiently. This chapter discusses the methodology developed to both capture and describe the data. A naming convention for the testing is introduced to simplify reference to the data later in the work; thereafter the chosen range of flowrate variations is given. Further sections are concerned with the description of the two types of measurement; namely base flow and stability. Finally, the flowmeter data is calibrated, and the actual flowrate variation is given for each test.

7.1 Test names and flowrate variation

A number of different exponential flows were investigated, in order to capture a representative range of accelerations. This section discusses the naming convention used to denote the different tests, and then presents the expected flowrate variation with time for each of these tests.

Test naming convention

Twelve different tests (comprising twelve different flowrate variations) were conducted in total. These separate tests were assigned distinct names, both for the purpose of data storage (file names were required), and to simplify reference to the data. Each test was given a unique two letter prefix; all runs retrieved from a particular test were stored using file names beginning with that prefix.

short length, to be optimum for strongly accelerating flows. It is thought that the instantaneous *negative* pressure gradient occurring during the overshoot might have induced a small amount of separation at the inlet at start-up, thus originating the bump that converts downstream. In the present set-up the contraction was opaque, and thus velocity readings could not be conducted within this region to confirm or refute this. The behaviour of inlet contractions under highly unsteady conditions is unstudied at present, and warrants further research. In the light of the sensitivity to the start-up, great care was taken to eliminate the occurrence of overshoot. The presence of overshoot was affected by amongst other factors the tank pressure at start-up, valve settings, and by pre-priming of the control valve. The control valve could be manually opened before commencement of the closed loop control by issuing the command `x VALVE y` in the run-file, where `x` was the time of execution and `y` the value to which the valve should be opened. By varying these two parameters, the valve could be controlled so that it was open by the correct amount when closed-loop control commenced. Thus in addition to the iterative set-up needed to produce the correct flow, the rig settings had to simultaneously ensure a proper start-up.

6.4 Summary

This chapter discussed broadly the initial experimental investigations and presented some early results. A brief consideration of the nature of the flow development in the pipe allowed a general understanding of the behaviour of exponentially accelerating pipe entrance flows. In particular, flow breakdown in these flows was seen to originate at the boundary layer near the pipe inlet. Arising from this preliminary study, a methodology of testing was developed, in order to optimally capture the early laminar variation, and later turbulent breakdown, of a broad range of exponentially varying flows. This methodology is explained in the following chapter, where the specific aspects of the final testing and data processing are described.

such bump. The dotted line in the figure represents the approximate behaviour of the flow when the bump was absent. The mass-flow compensation of the flow at other radial stations, due to the presence of the bump in the boundary layer, seemed to be slight. It was not possible to accurately measure the details of this bump, due to its erratic nature coupled with the point-measurement limitation of the LDV. An instantaneous global measurement of the velocity profile shape would be required to accurately determine its extent and nature.

Although the bump was worrying, its presence or absence had no discernible effect on the later transition. The bump simply grew slowly as it convected downstream, and the flow continued to behave in the same manner as before when the bump had passed. Because of its presence, it was decided to perform multiple experimental runs at each measurement station, in order that at least a single run *without* the bump would be captured. This was especially so when the data was to be used for the determination of velocity profiles, where non-repeatability of a phenomenon would result in spurious trends being shown in the processed data. The frequency of occurrence and magnitude of the bump was aggravated by the nature of the initial impulsive start-up. A large bump could be repeatably reproduced if any overshoot was permitted in the flowrate at start-up. Conversely, a 'rounded' start-up was far less inclined to produce any bump. Figure 6.10 shows schematically the different types of flow start-up, and their resultant effects.




Figure 6.10: The various types of impulsive starts, in order of their effect on the flow: (a), rounded start; (b), impulsive start; (c), under-damped impulsive start.

In this figure, (a) represents the ideal start to minimise or eliminate the presence of the inlet bump. The approximate time-scale of the rounded portion of the flow necessary was of the order of under a second. In (b), an 'ideal', impulsive start is shown. Such a start was not physically possible due to the momentum of the fluid and rig components. In some of the preliminary tests, flows which seemed to show proper impulsive starts were generated, but these always produced inlet bumps. The third type of start - (c) - generated large and repeatable effects. While the figure shows the overshoot exaggerated, the smallest overshoot was sufficient to cause problems. In fact the perfectly 'impulsive' starts seen in some of the preliminary results were thought to rather be of the form (c), but with a very small overshoot. Because the theoretically 'ideal' start - start (b) - could not be attained, it was decided to begin each test with a rounded start (a).

The cause of the phenomenon was unknown, although it was thought to originate in the inlet contraction. The inlet bell-mouth was specifically designed with a

(2)] which lead to a strong oscillatory structure by the time 60s is reached. As before, the data here is smoothed, but the raw data traces did not show any turbulent artefacts in this disturbed region - all oscillations were of a laminar nature. The centreline trace (5) remained laminar for the entire duration of the test. Traces (1) to (4) were restricted to a very small portion of the pipe - the region $1 < r/R \leq 0.893$. The traces from the outer part of this layer, (3) and (4), showed deviations, but these could be inferred to be due to mass flow conservation effects, caused by the variations occurring closer to the wall. Thus it seemed as if the boundary layer oscillations were restricted to a very small portion of the inlet profile, near the wall.

6.3.4 Entrance effects and the background environment

Background environment

Close observation of the traces in the boundary layer thinning region reveals the presence of small oscillations, whereas in the early boundary layer 'thickening' flow the traces are very smooth and regular. Fluid in the earlier region originated from within the pipe, whereas the latter fluid had entered the pipe during the experiment from the bell-mouth. Thus it seemed that some factor at the inlet was inducing a background finite-amplitude 'noise' in the flow. The origin of this noise was not the fluid in the tank itself: setting of the fluid overnight had the same effect on the level of disturbance as a normal two minute settling time. These perturbations were thus thought to originate either from the inlet join, or from imperfections in the contraction itself. Regardless of the source, these perturbations were present in all the tests conducted. It should be noted that this background environment is virtually constant with time in the *normalised* traces, that is the ratio u'/U_m (u' being the disturbance amplitude) remains fixed. This invariant behaviour is ideal for characterising the resultant flow breakdown; finite-amplitude analyses [see Herbert (1988)] commonly assume a fixed level of background disturbance relative to the mean flow rate. In the light of the importance of the background disturbance on the stability of the system, it is important to quantify the level of disturbance present. This topic is addressed in detail in the following chapter.

The inlet bump

Another inlet effect in the flow is shown in trace (1) of figure 6.9, where a laminar 'bump' can be seen at the transition point between the boundary layer thickening and thinning regions (at about $t = 15$ s). This bump is only evident in this particular trace in the figure - its presence or absence during testing was entirely random - but its magnitude seemed to be governed by the details of the impulsive flow start. Thus the shape of the 'wobble' between time 10 and 20 seconds in trace (1) was not repeatable; another run at the identical radial station might show no

6.3.3 Inlet behaviour

It was shown above that the flow exhibited various stages of development, passing through the early parallel stage, through an entrance flow stage, and resulting in a transitional stage. The origin of the flow breakdown - the time (t) in figure 6.8 - within the pipe was now sought, because the turbulent breakdown of the flow is intimately involved with the instabilities of the flow.

A hard turbulent front was seen far downstream, and was assumed to originate somewhere in the pipe entrance, having been convected downstream. It was however unclear what would be visible right at the pipe inlet. To clarify this, the LDV was positioned as far upstream as possible⁸ and readings were taken both at centreline and at various positions near the wall. Velocity time traces from this station, normalised again with respect to the centreline velocity, are given in figure 6.9.




Figure 6.9: A normalised set of velocity-time traces from the axial station $x/R = 12.24$, from selected radial positions: (1), $r/R = 0.966$; (2), $r/R = 0.926$; (3), $r/R = 0.916$; (4), $r/R = 0.893$; (5), centreline.

As before, the units of the vertical scale are arbitrary and unimportant. Considering the early laminar development of the flow ($2 < t < 50$ seconds), it is immediately evident that boundary layer thinning is occurring near the inlet too. However the thinning portion of the flow is far more extensive here than at the downstream station. At the current measuring station, the boundary layer thickening ends at time $t \approx 8$ s; again the reversion to the thinning type system is associated with the passing of the inlet flow. Laminar development in the boundary layer thinning region continues for an extensive time. At about $t = 45$ s perturbations develop in the boundary layer [specifically traces (1) and

⁸The large LDV could not be moved all the way to the pipe inlet, so the small LDV head,

linked to the laser by fibre-optic cables was used.




Figure 6.8: A normalised set of velocity-time traces from the axial station $x/R = 100.25$, from selected radial positions: (1), $r/R = 0.931$; (2), $r/R = 0.863$; (3), $r/R = 0.716$; (4), centreline; (i), inlet 'front'; (t), point of transition.

of about 25 s, as given earlier by equation 6.6, but that estimate was a simple first approximation. The phenomenon occurring at this time is of a global, pipe-wide, nature; the entire flow development reverts smoothly to one with a thinning boundary layer. This region is restricted to the time $36 < t < 55$ seconds.

- (iii) Starting initially in trace (2) at the time (t), flow breakdown begins. The instability spreads radially and grows temporally, and a pipe-wide transition to turbulence ensues. The current traces are smoothed, thus eliminating the transition spike shown in figure 6.4, but spikes are evident in the raw data in the later stages of transition. The period of breakdown is approximately $55 < t < 60$ s, after which the flow becomes fully turbulent.

The reversal of the boundary layer development was evident at every measuring station in the accelerating pipe. The point at which flow development reversal occurred was always consistent with the arrival of the inlet front. Thus at axial stations closer to the inlet the reversion to a thinner boundary layer occurred earlier than at the further downstream positions. For the sake of clarity, the region (i) - (t) in figure 6.8 was designated as the boundary layer thinning or simply the thinning region of the flow. The breakdown of the laminar flow to a turbulent state always occurred in the boundary layer thinning part of the flow, and the timing and position of this phenomenon is discussed in the following section.




Figure 6.7: A schematic representation of the development of the velocity profile in an accelerating pipe flow: (a), The physical system; (b), the normalised system. Profiles are labelled (1) to (5) in order of development. The right-hand side of the figure shows the profiles superimposed.

Using the concept of normalising the velocity profiles with respect to the mean flow rate, the traces in figures 6.5 and 6.6 can also be normalised. This is attained simply by dividing the velocity trace by the mean flow rate as obtained from the electromagnetic flowmeter. Figure 6.8 shows the data of these previous traces thus represented, but with some traces from intermediate radial stations also given.

The vertical scale in this figure is arbitrary and unimportant, since the interest is in the variation of the traces with respect to each other. The flow development shown here clearly consists of three regions:-

- (i) In the first region encompassing $2 < t < 30$ seconds, the flow shows classic developing behaviour with boundary layer growth indicated by the decrease in the normalised boundary layer velocity and commensurate increase at centreline. As can be seen, the flow develops rapidly away from the initial 'top-hat' profile at flow inception. In this region, the normalised development is remarkably similar to that of an impulsively started flow⁷.
- (ii) At the point (i), the flow behaviour changes rather abruptly, with the boundary layer development *reversing* its trend and actually reducing in thickness. The time (i) is consistent with the convection of the inlet front downstream; that is, for times after (i), the flow is an entrance flow. The timing of this point - $t \approx 35$ s - is longer than the estimated inlet time

⁷An impulsive flow can also be 'normalised' with respect to its mean velocity, but in that case this velocity is a constant divisor on the velocity; thus a normalised trace in the impulsive case is simply the trace scaled by the mean velocity.




Figure 6.6: The smoothed wall trace from the axial position $x/R = 160.25$. The radial station is $r/R = 0.931$.

the turbulent front is of major importance; the preliminary data above indicated a phenomenon originating within the pipe some time after flow inception. For this reason the exact nature of the initial laminar flow is equally important since it is this flow that initiates the instabilities that lead to turbulence.

The figures 6.5 and 6.6 indicate individually the development at the pipe centreline and wall respectively. In order to obtain a fuller picture a composite representation of the flow behaviour must be obtained. What is occurring in this particular flow is the development of the velocity profile away from the 'top-hat' form at flow inception, through a range of other shapes as time progresses. Simultaneously the mean flow is increasing rapidly; thus the magnitude of the profile is also increasing. What is of primary interest however, is the change in shape alone, the change in magnitude being simply exponential and described exactly by the mean flow variation in time. Thus, a better representation for the data would be scaled according to the mean velocity. Figure 6.7 shows schematically the development of the velocity profile, both in the physical framework, and normalised with respect to the mean velocity.

In this figure an arbitrary velocity profile development has been shown, from a 'top-hat' profile (1), through some intermediate profiles, to a parabolic final profile (5). In the physical reference system (a), it is obvious that the profile shape is changing, but it is not exactly clear to what extent the subsequent profiles differ from one another. In (b), the same profiles as (a) are shown, but each is normalised such that its mean flow rate is unity. Without the strong scaling effect, the boundary layer development is immediately apparent. The profiles in this figure are however simply fabrications; more interesting would be the behaviour of the actual experimental velocity profiles.

combined with knowledge of the exponential flow variation, allowed the calculation of the approximate time at which fluid from the inlet at flow inception would arrive at the current measuring station. Using the flow mean velocity as the approximate convection velocity, employing the relation

$$s = \int_0^t U(t) dt \quad (6.3)$$

would allow simple calculation of the mean fluid displacement with time in the inlet. If the mean flow variation with time is

$$U(t) = U_0 \exp(kt), \quad (6.4)$$

then the mean displacement of fluid downstream for a given time t is given by

$$s = \frac{U_0}{k} \exp(kt). \quad (6.5)$$

For the prototypical flow employed here a value of $k = 0.03 \text{ s}^{-1}$ was used, the value of the constant U_0 was 0.1118 m/s , and the axial displacement from the inlet was $s = 7.724 \text{ m}$. Rearrangement of (6.5) resulted in the equation

$$t = \frac{1}{k} \ln \left(s \frac{k}{U_0} \right) \quad (6.6)$$

and solution in terms of t gave a value of $t = 24.3$ seconds.

The simple calculation above demonstrated that the timing of this front did not coincide with the convected inlet fluid; that is, this front did not originate at the inlet at start-up but at a much later time somewhere in the entrance region. Measurements at the wall for the same axial station showed simply the convected front too: figure 6.6 shows a smoothed trace from the same axial station, but from radial position $r/R = 0.931$ - a position well inside the boundary layer.

Here, after the initial impulsive start, the boundary layer starts to thicken with a result that the velocity at the wall reduces. However in the present case the mean flow rate is increasing exponentially; it is unclear whether the boundary layer velocity is increasing or decreasing relative to the centreline value. However at about 40 seconds, the velocity in the boundary layer starts to increase at an accelerated pace, until the occurrence of the turbulent front at 58 seconds. After the passing of the front, the mean turbulent velocity is much higher than the laminar one before. This is consistent with the behaviour of a turbulent boundary layer, which possesses a high rate of shear at the wall. The origin of

of the inlet contraction was unknown. By means of a simple timing argument, the downstream measurement would indicate whether a perturbation originated from the inlet bell-mouth at inception, or whether it originated as a boundary layer phenomenon elsewhere in the pipe. Figure 6.4 shows a typical such trace from the axial station at a distance $x/R = 160.25$ from the inlet plane, while figure 6.5 shows the smoothed result.




Figure 6.4: An unprocessed centreline trace from the axial position $x/R = 160.25$ from the preliminary experimentation, showing the occurrence of a turbulent front.




Figure 6.5: The smoothed centreline trace from the axial position $x/R = 160.25$.

The unprocessed data - figure 6.4 - clearly shows a convected turbulent front passing the measuring station at time $t \approx 58$ s. Prior to this no convected phenomena were evident. Consideration of the axial station of the measurement,

6.3.1 Data smoothing

Data obtained from an LDV is innately problematic, primarily because of its random time base; that is, the time between readings is random. This occurs because the LDV measures discrete events, namely the traversal of the measuring volume by a seeding particle. Furthermore, measurements in the boundary layer are subject to broadening due to spatial gradients in the flow⁵. The smoothing algorithm defined in section 3.4.2 was used on all LDV data obtained, in order to place the data on a regular time-base, and remove broadening.

The smoothing procedure was optimum for the current data, because the laminar variation of the flow was of interest. The procedure effectively acted as a low-pass filter on the data, removing the high frequency component. As mentioned, this resulted in the turbulent parts of the flow being averaged out; however, optimum representation of the laminar and early transitory parts of the flow were of primary importance. As will be seen later, the deviation from the simple laminar state was itself laminar with no high frequency components, thus smoothing of the data was acceptable. On the practical side, the smoothing process reduced the data size considerably, allowing easy manipulation of the reduced set. The raw data had an average data rate of between 150 and 500 Hz; the output data from this pre-processing process was sampled at 10 Hz, which gave a more than adequate data density for subsequent observation and processing.

It should be recalled that the processing algorithm in section 3.4 required three parameters: a time window, a time increment and a standard deviation cut-off. The time increment was taken as 0.1 seconds, resulting in the 10 Hz sampling rate described above. The time window effectively controlled the band-width of smoothed data and was set to a value between 0.5 and 1.2 seconds⁶. The standard deviation cut-off limit was set at three standard deviations, in common with da Silva (1990). Of course the weighted mean and weighted standard deviation algorithms [McLaughlin & Tiederman (1973)] were employed as is appropriate for LDV data.

6.3.2 Behaviour far downstream

The most downstream station was chosen because it was the position most likely to eventually exhibit a deviation from the laminar state. This was because any transition occurring at or near the pipe inlet would most certainly be convected downstream. It was reasoned that once an event was identified, it could be tracked upstream to its point of inception. Furthermore the behaviour

⁵All the effects contributing to the nature of the LDV signal were discussed in detail in chapter 3, and will not be elaborated upon here.

⁶This window parameter was varied according to the data density of the input signal. It was attempted to ensure that each time-slice of the data contained upwards of 25 data points, in order to ensure a statistically valid population.




Figure 7.7: Scaled velocity time traces corresponding to the velocity profile FB5, illustrating the point of deviation from laminar flow in the pipe entrance. The traces correspond to the following radial stations: (1), $r = 0.973$; (2), $r = 0.952$; (3), $r = 0.930$; (4), $r = 0.907$; (5), $r = 0.885$; (6), $r = 0.863$; (7), $r = 0.841$; (8), $r = 0.819$; (9), $r = 0$.

The early deviation in traces (1) and (5) was not apparent in most of the other profiles measured, but the point of strong departure (at 27 s) always was. Section 6.3 earlier in the chapter discussed the variability associated with the thinning portion of the flow, and the early deviations in these two traces were thought to be part of that flow phenomenon, *not* early flow breakdown. That this was so, was corroborated by the fact that the early deviation only occurred intermittently. The latter deviations at ≈ 27 s were thus taken as the point of deviation. Again, the confusion and ambiguity from limitations of the LDV, namely its point-measurement capability.

The latter departure ($t \approx 27$ s) occurred virtually simultaneously throughout the entire inner boundary layer [traces (1) to (6) in figure 7.7]. The variability in inception time between traces was comparable to the variation in timing between traces from the same radial station. In other words, the flow breakdown seemed to be occurring simultaneously throughout a large portion of the boundary layer at one time⁴. The large radial extent of the initial deviation implied that its timing could be inferred from a single carefully placed measuring station within the boundary layer.

⁴It should be recalled that the process of flow breakdown was subject to variability; thus the composite representation shown in figure 7.7 is not a ‘snapshot’ of a single event, but rather the superimposition of different events.




Figure 7.6: Velocity time traces corresponding to the velocity profile FB5, illustrating the point of deviation from laminar flow in the pipe entrance. The traces correspond to the following radial stations: (1), $r = 0.973$; (2), $r = 0.952$; (3), $r = 0.930$; (4), $r = 0.907$; (5), $r = 0.885$; (6), $r = 0.863$; (7), $r = 0.841$; (8), $r = 0.819$; (9), $r = 0$.

toward the boundary layer and the centre of the pipe. Strong laminar ‘wiggles’ are seen in traces (3) to (6), while the wall traces (1) and (2) show a characteristic drop-off in velocity. The centreline trace (9) remains completely unaffected until about $t = 32$ s, at which point it exhibits a decline in accordance with the start of turbulent flow [as indicated by the sharp increase in traces (1) and (2)] in the boundary layer. After about 36 s, the entire boundary layer is turbulent (the traces have been smoothed and thus only show the mean component), although the centreline flow remains mostly laminar. As before, the representation in figure 7.6 is not optimum in that the traces scale in time, due to the exponential nature of the flow. A more optimum representation is obtained by dividing the trace velocity by the mean cross-sectional velocity as obtained from the flowmeter. The result of the scaling of the above figure is shown in figure 7.7.

Only the relative behaviour of the normalised velocity traces is important. This is because the study of the flow breakdown in the current context is unavoidably qualitative. Since the magnitude of the trace was of no consequence in this part of the study, it was not necessary to scale the data correctly. The LJV velocity was simply divided by the flowrate, and the vertical scale in the figure is thus arbitrary. The full behaviour of the flow is far more apparent in this figure: the centreline velocity diminishes from about $t = 12$ s, indicating the start of the thinning flow regime. The deviation point shown in traces (4) and (5) can be seen to originate at an earlier time - at about 21 s - but is still clearly in the thinning portion of the flow. More marked deviations occur in a number of the traces at 26 - 27 s; however, turbulence only starts developing much later, at 32 s.

per test. Due to the small expense of the timing measurements, many such measurements were interspersed evenly between the widely placed profile stations. Figure 7.5 shows the general arrangement of testing stations in the pipe entrance region.




Figure 7.5: The general distribution of testing stations in the pipe entrance region.

7.4.1 Point of deviation

The applicability of choosing a single radial position at which to measure the laminar deviation is perhaps best illustrated by considering data from a full velocity profile. A suitable candidate to illustrate this was the measurement FB5, a velocity profile taken for the test FXPB, at the axial station $x/R = 90.5$. As with the other experimentally measured profiles, pairs of readings were taken at varying radial positions, starting at the wall, and moving in small steps toward the centreline. In the current context only the inception time and the subsequent nature of the flow deviation were of importance.

Figure 7.6 shows smoothed data from this set of measurements for selected radial stations, mostly clustered near the wall. The corresponding centreline velocity trace is shown for comparison. Only one of the pairs of measurements for each radial station is shown here. The trace showing the most marked flow deviation was chosen in each case. The validity of the composite representation as a description of a single coherent transition occurrence in the pipe was as before doubtful, because of the traces being measurements from different runs. As will be seen in section 7.4.1 below, the time of deviation was subject to a statistical variability, however the *general* pattern of behaviour of the velocity traces was repeatable. Thus the data gives a good qualitative description of the flow deviation in the system.

The nature of the initial laminar development is unclear in figure 7.6, due to the strong scaling effect of the mean cross-sectional velocity. The boundary layer thickening for early times and the later thinning behaviour for the entrance flow is not obvious in this figure, although the approach of some of the boundary layer traces to the centreline velocity for $t > 20$ s indicates a reduction of the boundary layer dimension for this region. At about $t = 25$ s, the traces near the boundary begin to deviate from the laminar state. Trace (4) shows an increase in velocity, and (5) a commensurate decrease. Later, the influence spreads both

- (ii) For each profile, about 5 points encompassing the boundary layer were selected - the points shown inside the dotted rectangle in figure 7.1 - and a low order least squares fit was made through these data. This fit was deemed good if it passed accurately through all the data. Only a portion of the boundary layer was selected. A higher order fit would have been necessary to fit properly the entire boundary layer, whereas only the wall region was of importance. Using more of the boundary layer and a higher order fit resulted in far greater variance in the determined wall position. The fit was used as an extrapolation to find the intersection with the radial axis - the inferred position of the wall.
- (iii) The above process was repeated for between 10 and 20 profiles from different times, in order to arrive at a statistically valid population of wall positions. The resulting positions were combined, their mean was taken as the wall position and their standard deviation as an indication of the error. The factor δr was set to this mean.
- (iv) The data for all the profiles could then be shifted radially by the amount δr , resulting in data properly aligned in the range $0 \leq r \leq 1$. The data in the above figure had already been shifted by the displacement δr , hence the coincidence of the wall position with $r = 1$.

Once the wall had been determined, an artificial velocity value of zero was appended to all the extracted velocity profiles, at $r = 1$. The wall values determined by the above process showed a surprisingly small standard deviation. In most cases this value was in the range $[0.001, 0.003]$ - a percentage variation of between 0.1 and 0.3%. The order of the fit used was in all cases quadratic. The so-called 'centreline' trace was only approximately aligned with the pipe centreline (because it was impossible to judge *ab initio* where the wall was and hence the centreline position). Because the core flow was in all cases uniform, no loss in accuracy was caused by this trace being artificially moved to the position $r = 0$. The result of this entire process was a set of velocity traces normalised in the range $[0, 1]$, and with data points at each end of this range.

7.4 Timing measurements

The second type of measurement taken was the timing measurement. Whereas the velocity profile measurements above encompassed the entire radial flow domain, it was decided to limit timing measurements to a single point in the pipe boundary layer at each axial station, and rather to repeat the run many times while measuring the velocity variation at this point. Repetition at an identical point in the pipe was decided upon in order to gather statistical information about the inception of the transition process. The profile measurements described in the earlier sections were expensive in terms of time and data storage; for this reason alone it was decided to limit their number to about three to five

$r = 0$ being the centre of the pipe, and $r = 1$ the wall. Importantly, this equation allows direct conversion of the *bench* co-ordinate system to the ordinate r : If the radius R is set to $R = 17.76$ mm, then y can be taken directly as the bench co-ordinate.

7.3.3 Determination of the wall position

The normalisation carried out above did not necessarily yield the correct radial ordinate r for the data. This was because the wall position had yet to be accurately determined². The actual normalising equation was thus an extension of (7.5), namely

$$r = \frac{R - y}{R} + \delta r, \quad (7.6)$$

with δr the wall offset to be determined. Determination of this offset is described in point form below, and is illustrated in figure 7.4:




Figure 7.4: The method of determination of the wall position from a particular experimental velocity profile: $\circ$, experimental data; $—$, approximate velocity profile; $—$, low-order boundary layer curve fit; $\bullet$, position of the wall. The dotted rectangle contains the points selected for the boundary layer fit.

- (i) The radial co-ordinate of the data was non-dimensionalised according to equation (7.5); thereafter profiles were extracted at various times. This data was then actually over the range $-\delta r \leq r \leq 1 + \delta r$.

²It should be remembered that the zero radial station y_0 was simply set by positioning the measuring volume by eye somewhere in the pipe wall.




Figure 7.3: A three dimensional representation of a set of velocity-time traces, showing a particular velocity profile.

The sampling of the raw data at regular intervals was important in that it allowed the assembly of all the data for a particular station into a three-dimensional lattice, and then quick extraction of that data by simply taking a slice from that lattice. In the raw data the individual points did not align to any particular time; in that case extraction of a profile would have required more extensive processing. Further to this, the order of magnitude larger quantity of raw data would not have allowed consideration of all the traces from a particular station in parallel, which was necessary for the extraction of a profile.

7.3.2 Normalisation of the radial ordinate

The units of measurement of the velocity profile in the radial direction were millimetres in the range 0 mm to 24.1 mm, with zero at the near pipe wall. This system of axes was specific to the current study. A more generic means of giving the radial co-ordinate would be relative to the pipe centreline as a ratio of the pipe radius. That is, by representing the radial position as

$$r = \frac{r}{R} = \frac{R - y}{R}. \quad (7.5)$$

Here R is the pipe radius, r is the radial position as measured from the pipe centre, and y is the position as measured from the near wall of the pipe. Equation (7.5) yields a dimensionless ordinate r which varies in the range $0 \leq r \leq 1$.

profile would verify the assumption of *symmetry* innate in only measuring flow in half the pipe - a discussion on symmetry is contained in section 7.5.1. The following two sections each contain a discussion concerning the two types of measurements: profile and timing. Specific processing and interpretation issues relating to each type of measurement are discussed there.

7.3 Velocity profiles

In order to measure velocity profiles at a particular axial station in the flow, the laser measuring volume was positioned at the near wall according to the procedure outlined above, and the measuring volume was moved towards the centre of the pipe in sufficiently small steps to resolve the flow variation. The initial step-size chosen was the same as above, namely 0.4 mm, and the LDV was moved inwards through the boundary layer in these small increments. At each radial position the run was repeated twice to verify repeatability. Observation was made of the resulting velocity-time traces by plotting them in composite on the computer. When the difference between traces at different radial stations started to diminish, the step-size was increased gradually, first to 0.5 mm, then 0.6 mm and so on, until by the time the measuring volume was approaching the centre of the pipe, the increments were approximately 2 mm. It was possible to do this because all the velocity profiles that were measured consisted of a thin boundary layer combined with a core flow. The philosophy of increasing the step-size resulted in a far more economical capture of data. If the small step-size was used uniformly, then about 60 measurements would have been required per radius, whereas the current scheme only required about 25-30 different radial stations over the same radius. This represented a two-fold saving in both time and data storage space.

7.3.1 Data pre-processing

Various levels of data processing were required throughout the study; initial processing such as the smoothing process described earlier (see section 6.3.1) was of an almost automatic nature, and was concerned mainly with reducing substantially the vast quantities of raw data acquired and removing both the velocity bias and gradient broadening. In the current context this processing also served another important function, namely to sample the data at regular time intervals. The importance of sampling the data at a regular time interval was important for the following reason: the data acquired at each radial station for a profile were velocity-time data, each of these traces corresponding to a particular radial position. Thus in a time - radial position - velocity space, the full set of data would be represented as a group of velocity-time data, each data positioned at a particular radius. Velocity *profiles* would be represented in this framework as slices of constant time. Figure 7.3 shows schematically data represented in this manner.




Figure 7.2: The beam distortion occurring through the various interfaces in the pipe; κ , free-air half-angle; κ' , actual half-angle; δ , beam displacement due to water-glass-water interface.

More subtly, the change in half angle affected the *rate* of motion of the measuring volume within the fluid relative to the bench motion: if the LDV head was moved externally by an amount Δr radially, then the movement of the measuring volume within the pipe would be a different amount $\Delta r'$ according to

$$\frac{\Delta r}{\Delta r'} = \frac{\tan \kappa'}{\tan \kappa}, \quad (7.3)$$

Movement of the large laser compensated for this effect via the control software; the small head had no such compensation because its position was measured manually. The calculation of the distortion gave a value for the movement ratio $\Delta r/\Delta r'$ of

$$\frac{\Delta r}{\Delta r'} = 0.7369, \quad (7.4)$$

Thus for a certain movement of the LDV head, the measuring volume inside the pipe moved by a greater amount in the ratio $> 1/0.7369$. The actual nominal radius of the pipe was 24.1 mm; optical distortion resulted in the laser head only having to move through 17.76 mm to traverse this distance. The above distortion ratio (7.4) was arrived at using geometrical considerations. As mentioned, in some velocity profiles the test was conducted from wall to wall, in order to verify this. By acquiring an entire velocity profile, various components of the experimentation could be verified. Firstly the profile would give the *perceived* diameter of the pipe, thus directly giving the distortion ratio. Secondly such a

intersection point at all, since it was defined merely as the mutual intersection of a pair of laser beams. It was, however fairly easy to identify the crossing point by wearing laser goggles. These glasses filtered out entirely the frequencies of the incident beams, thus making them invisible. Slight phosphorescence of the glass yielded the laser beams visible at their intersection with the wall of the pipe. This allowed the cross-over point to be positioned anywhere *within* the wall simply using visual alignment. It was then decided to avoid attempting to align the system in an absolute sense; rather the in-wall position as set above was designated the origin and all measurements were taken according to this datum.

The first measuring station within the flow was found by running the test while traversing the control volume in small steps towards the flow. The increment used was 0.1 mm, approximately half the length of the measuring volume itself. The first station was taken as that point at which coherent signals were first obtained¹. The exact position of the wall was found subsequently by extrapolating the data during processing. This is described in section 7.3.

The far wall and optical distortion

A further complication introduced was optical distortion caused by the various components of the pipe possessing different indices of refraction. The laser beams had to pass from air, through the outer planar glass wall, into the water jacket, then through the glass pipe, and finally back into the water within the pipe. The beams were travelling in the horizontal plane, thus the curvature of the pipe was of no consequence. The crossing angle of the two beams was changed by these series of intersections; effectively only the air-glass-water interface contributed to the angle change because the other interface only resulted in a translation of the beams. Figure 7.2 illustrates the beam distortion through these series of interfaces.

The resultant distortion of the beams could be interpreted best as a change in their half-angles. This angle change affected the calibration of the LDV according to the relation between velocity and frequency

$$U = \frac{f\lambda}{2\sin\kappa'} \quad (7.2)$$

where λ is the wavelength of the light, f is the frequency of the signal received from the LDV photo-multiplier, and κ' is the actual beam half-angle. The acquisition software correctly calculated the actual half-angle and the conversion (7.2) based on geometrical factors contained in a start-up file.

¹With the measuring volume in the wall, the data rate was zero. The data rate increased quite abruptly at the point where most of the measuring volume was within the flow.

Vertical alignment

This alignment meant ensuring that the plane of the two laser beams was co-incident with a plane through the centreline of the pipe. In other words, this alignment allowed LDV traversal to occur along a major diameter of the pipe. To simplify matters further, this plane was set as horizontal. This allowed the LDV to be set initially using a spirit level. Once the LDV was horizontal, the entire head was adjusted vertically (using adjustable feet) until its plane was coincident with the centre-plane of the pipe. This was achieved very simply and highly accurately as follows: the LDV was left on during adjustment, and the reflected beams from the far pipe wall were observed. When the reflected beams became co-planar to the incident ones, they were known to be intersecting that wall perpendicularly; thus by inference the plane of the beams was on a major diameter of the pipe. Figure 7.1 illustrates this principle.




Figure 7.1: The principle whereby the LDV measuring volume was placed on a major diameter: (a), mis-aligned, beams reflecting down; (b), vertically aligned.

In practice, perfect vertical alignment resulted in reflected light entering the photo-multiplier aperture, thus resulting in a low signal-to-noise ratio. The optimum solution was to mis-align the system by a tiny amount - usually no more than 0.1 mm - eliminating the direct intersection of the reflected light while still remaining on the major diameter to a high level of approximation.

Alignment with the wall

Whereas the vertical alignment described above was relatively easy to achieve, fixing the position of the measuring volume relative to the pipe horizontally was far more difficult. Firstly, the dimension of the measuring volume in the radial direction was large (≈ 0.7 mm), introducing uncertainties as to what should be aligned to the wall. Secondly, even if a point of alignment with the measuring volume could be ascertained, there was the daunting task of positioning the small

Table 7.3: The values of the parameters Re_0 , k to generate the flowrate variations for the separate tests.

Test	Re_0	$k(s^{-1})$	Test	Re_0	$k(s^{-1})$
EE	3018	0.03	FD	4143	0.03
EF	3018	0.05	GA	6510	0.05
EG	3018	0.07	GB	6510	0.07
FA	4143	0.01	GC	6510	0.01
FB	4143	0.05	GD	6510	0.03
FC	4143	0.07	HA	2959	0.02

This table presents the expected flowrate variation in time for the twelve tests. Section 7.5.4 later in the chapter discusses the actual flowrate variation achieved in the system.

7.2 Alignment and optical considerations

All velocity measurements in the pipe system were conducted using the LDV described in chapter 3. This device was mounted on an automated traversing bed, so that the LDV measuring volume could be positioned by the acquisition computer anywhere along the length of the pipe, and anywhere within its cross-section. The movement of the bench was however restricted to axial stations greater than about three metres downstream of the inlet plane. In order to overcome this limitation, a separate - portable - LDV head was used, connected to the main LDV by a fibre-optic link. This small unit could be traversed in the radial direction manually, and its displacement measured using an attached dial gauge.

Unfortunately, the nature of the flow dictated that most of the LDV measurements were to be made near the inlet plane, due to the transition originating in the boundary layer there. As most of the data was acquired through the small LDV head, the repositioning of the measuring volume during macro control was manual. Its axial position was varied by moving the entire assembly along a bench placed next to the pipe. To vary the radial station, a micrometer dial on the small mounting bench could be turned; relative radial position was measured using a dial gauge mounted on the small head.

For both measurement types, the alignment of the LDV control volume relative to the pipe system was of major importance. Whereas the LDV measuring volume could be moved highly accurately relative to its current position, the absolute positioning of this point was subject to some uncertainty. Thus it was important to align the LDV to the pipe at some reference point; thereafter accurate movement could be achieved.

Table 7.2: Description of the file suffixes used in data storage.

Suffix	Data type
.RA1	Velocity - time data from the LDV - channel one
.RA2	Velocity - time data from the LDV - channel two
.CHI0	Voltage - time data from the flowmeter (this was assigned to the analogue channel zero)
.CHLr	Voltage - time data for any other analogue acquisition channel
.IN1	Information concerning the .RA1 file (Beam angle, light wavelength, counter mode, etc.)
.IN2	As above, but for .RA2

were configured such that channel 1 received data from the large, downstream LDV, and channel 2 acquired its data from the small LDV positioned near the inlet. In some tests, data was gathered from both channels simultaneously. In these cases the same test names referred to *different* axial and radial stations, and differentiation between the two data was made according to the file suffix.

Expected flowrate variation

Each test name given in table 7.1 referred to a different flowrate variation with time. The objective of the study was to observe the behaviour of exponentially accelerating pipe entrance flows and to achieve this, a range of exponentially increasing flowrates was chosen.

The experimental test rig described in chapter 4 was capable of generating flows with flowrates varying generally in time, over a large Reynolds number range (approximately $3\,000 \leq Re \leq 80\,000$). Thus within the inherent limitations of the rig, it was possible to generate a broad range of different exponentially increasing flows. The entire study centred around the determination of the velocity variation within a representative selection of such flows. Specifically, in all cases the flow was started from rest impulsively to an initial Reynolds number; thereafter the flow was accelerated exponentially in a prescribed manner, until the maximum possible flowrate was reached. Thus all the flows were of the general type

$$Re = \begin{cases} 0, & t \leq 0 \\ Re_0 \exp(kt), & t > 0 \end{cases} \quad (7.1)$$

where $Re = Re(t)$ is the prescribed Reynolds number variation with time, and Re_0 is the initial impulsively initiated Reynolds number in the pipe. The twelve separate tests were each conducted with a distinct pair of parameters (Re_0, k) . The values of the constants for these tests are listed in table 7.3 below.

Also indicated is the fifth order polynomial least-squares curve fitted to this data, taken as the calibration curve for the flowmeter.




Figure 7.15: The calibration curve for the flowmeter: $\circ$, experimental data; —, fitted curve.

The most striking aspect of figure 7.15 is the tight correlation between data from different tests. The circles represented flowrate values from all twelve tests, with about 30 to 40 points per test. The scatter shows a standard deviation about the fitted line of only 2% (the left-hand side of the calibration curve - the low flow end - had a standard deviation of only 1.5%). The analysis of the flowmeter scatter in appendix A.2 showed its variability to be proportional to its range. Since the flowrate variability was not determined from individual points; rather by fitting a smooth curve through the entire data, the variability indicated in that section was not relevant as an indicator of errors in the flowmeter⁹. The scatter shown in the figure was consistent with the error in profile integration, and was quite possibly primarily due to this effect. Due to the smooth nature of the flow variation, the more appropriate measure of variability here is the standard error of the mean [Brinkworth (1973)], defined as

$$\sigma = \frac{s'}{\sqrt{N}}, \quad (7.8)$$

where N is the number of data points and s' is the standard deviation of the data from the mean. The flowmeter data was sampled at 100 Hz and the test

⁹If the flowmeter were being used to track a high frequency flow variation - with duration of the same order as the sampling frequency - then the scatter in the flowmeter data would become significant, because in that case small populations of flowmeter data would be used to infer actual flow events. In the current case the variation is smooth, effectively the entire data set is used to fit a single smooth polynomial.

would be possible in a constant flow system⁶. An error analysis in appendix A, in section A.1, attempted to deduce the uncertainty in integration, using numerical data with random noise superimposed. This concluded that, for well placed velocity points throughout the boundary layer, the innate error in integration was about 0.02%.

Once the flowmeter was calibrated, a fairly good independent error estimate became available in that the flow rate from individual integrated velocity profiles could be compared directly to the calibrated data from the flowmeter⁷. The analysis in the appendix considered only velocity profile data from the initial parallel portion of the flow, and concluded that the actual error in integration was better than 1% of the mean cross-sectional velocity.

7.5.3 Flowmeter calibration

Initial comparison between the flowrates as determined by the two different methods showed that the flowmeter data was as much as 15% in error, especially at the lower flowrates. However, comparison between flow data from different tests showed that the deviation of the flowmeter data was highly systematic and thus repeatable, there being no noticeable dependence of the error on the flow acceleration, only on the flowrate itself. Due to this fact, it was decided to calibrate the flowmeter using *all* the experimental data. The reasoning behind this was as follows: the LDV-determined flowrate, while potentially accurate in the mean, showed fair scatter in individual cases, related to the specific nature of the data from the particular measuring station⁸. The flowmeter data on the other hand, while inaccurate, was repeatable; proper calibration of the unit would allow accurate determination of the flowrate directly from this data alone, without resorting to profile integration. Use of a large data population (integrated flowrate data from all the tests) would allow sufficiently accurate calibration of the flowmeter.

Flowrate-time variation was thus extracted for all velocity profiles by integration, and the results used to calibrate the electro-magnetic flowmeter. The result of this calibration is shown in figure 7.15, where the ratio of expected (LDV) to measured (flowmeter) mean cross-sectional velocity (flowrate divided by the cross-sectional area) is plotted against the LDV mean cross-sectional velocity.

⁶In a flow system where the flowrate remained constant, the integrated information from profiles at different times could be compared one to another, in order to compare their flowrates - which should of course remain constant if the flow itself remained so. In the current context the flowrate varied strongly: the only point of comparison - the flowmeter data - was subject also to error and thus could not be used directly for this purpose.

⁷The process of inferring the error was potentially circular, integrated data being used to find the calibration curve, the calibrated data then being used to infer the error in the integrated velocity profiles. However, since the calibrated data was obtained using *all* the velocity profile data from the entire study, the comparison was valid.

⁸In some instances a velocity-time trace was 'lost' due to corruption of the backup data, thus lowering substantially the accuracy to which the flowrate could be determined in these cases.

describes the relationship between LDV frequency signal and velocity]. As these quantities were known or could be accurately determined, velocity profiles could be accurately constructed and the only major source of error in the extraction of the flowrate was the numerical integration of the profiles themselves.

The mean cross-sectional velocity was extracted from the velocity profiles as follows: a low-order (2nd order) fit was made to the core flow in the pipe, while a fourth or fifth order fit was implemented in the boundary layer. These two fitted curves were forced to intersect at a chosen radial station. The curves were then integrated analytically to extract the mean flow according to

$$U_m = 2 \int_0^1 ru(r)dr. \quad (7.7)$$

Figure 7.14 shows the two curves fitted to a particular velocity profile. The point of intersection is shown. Since only integral information was required from the profiles, it was not important that the two curves match in their derivatives at the point of intersection; neither were any boundary conditions enforced at the centre of the pipe. As a curve was fitted to the entire boundary layer, the velocity profile slope at the wall was not well indicated by this procedure⁷.




Figure 7.14: The numerical integration of the experimental velocity profile, in order to calculate the flowrate: (a), point of intersection of the fitted curves.

Due to the fact that the flowrate varied exponentially with time, it was not directly possible to calculate the error introduced in the integration process as

⁷A curve fit to the outer edge of the boundary layer was specifically used to infer derivative information at the wall. A general boundary layer fit as used here was found to give highly variable derivative information, despite giving accurate integral results. The derivative curve-fitting procedure is described in section 7.6.3




Figure 7.13: Velocity profiles for the entire pipe diameter, showing the symmetry of the flow. Data from the far portion of the pipe have been reflected about the pipe origin to lie on the near-side data: (1), $t = 10$ s; (2), $t = 20$ s; $\circ$, near-side data; $\times$, far-side data.

Many closely spaced readings were taken from the near side of the flow, while only a few coarsely spaced measurements were made on the far portion. The far data have been superimposed on the near-side points to show their mutual correlation. Two velocity profiles are shown, in order to verify the symmetry for different portions of the flow development. As can be seen, the data from either side of the centreline describe identical velocity profile shapes to within the uncertainty of the LDV data. Thus although the far side data was coarsely sampled, the fact that each point falls exactly on the near-side profile is indicative of the high degree of symmetry of the flow. Of course the current measurements were not capable of inferring the *vertical* symmetry of the system, but the good agreement in the horizontal plane implied that the general symmetry of the flow was excellent. Considering the flow variability and the inlet ‘bump’, as before it was impossible to deduce whether these phenomena were in fact symmetrical, due to the single point measurement capability of the current LDV; however, as demonstrated above, the undisturbed laminar development is certainly highly axisymmetric.

7.5.2 Integration of the velocity profiles

Due to the symmetry of the system as mentioned above, the flowrate variation could be obtained directly from the measured velocity profiles, by simple integration. As a result of the linearity of the LDV measurement, the value of the velocity traces acquired by this means were guaranteed to be directly proportional to the actual velocity, whatever the velocity range. In order to accurately convert the LDV data to give the correct velocity values, only the laser wavelength and beam intersection half-angle needed to be known [equation (7.2)

Turbulent coherence

The turbulent region '4' in figure 7.11 also deserves mention. Here, the strong variability evident throughout the preceding transitional regions '2' and '3' suddenly disappears, and the traces once again become highly repeatable. This behaviour would be in keeping with the flow developing a truly turbulent boundary layer at this time. The presence of a well defined eddy viscosity and mixing length in a turbulent boundary layer would ensure steady behaviour in the mean; the correlation of the data in region '4' indicates that a fully turbulent boundary layer has in fact developed here.

The final region - region '5' is included for completeness. It is present in some tests, and is due simply to non-repeatability in the shutdown of the flow system; thus these apparent phenomena are not due to the exponential flow system, and should be ignored.

7.5 Flow meter correction and scaling of velocity profiles

The sections above concerned the measurement of the velocity profiles, and the determination of the point of flow breakdown at different axial stations. For the twelve tests considered, the flowrate variation in time was only known approximately, as given in table 7.3. This section is concerned with the determination of the actual variation of the flowrate with time, obtained by integration of the velocity profiles and calibration of the flowmeter. The process of velocity normalisation is also described, whereby the velocity profiles are divided by the corrected flowrate, in order to yield their magnitude of order one.

7.5.1 Symmetry of the velocity profiles

The experimental velocity profile measurements were concerned with measurements along the horizontal diameter of the pipe, from the near pipe wall to the centre line. Tacit in taking such limited readings was the assumption of axisymmetry. Symmetry in an experimental pipe system allows substantial simplification of the velocity measurements, in that the velocity field throughout the pipe can be inferred from a series of measurements along a single radius of the pipe. In order to verify the assumption of axisymmetry in the current context, readings in some tests were extended past the centre point of the pipe, to the far pipe wall. The resulting velocity data allowed verification of the flow symmetry, at least in the horizontal plane. Data from such a 'full' velocity profile is shown below in figure 7.13, where two profiles from the preliminary test EXPC are given ($x/R = 289.8$, flow variation same as EXPE), corresponding to the times 10 s and 20 s respectively.




Figure 7.12: The distribution of the points of deviation about their mean, for the entire test GXPA. Data has been combined from all measurement stations.

Early laminar variability

With reference to figure 7.11, apart from the laminar perturbations that develop after about 20 s, there are earlier deviations evident, notably the ‘bump’ shown by a single test at time 10 s, and the bifurcation of the traces at about 15 s into an upper and lower branch. The presence of this inlet bump was in some tests variable (for different traces it occurred to a greater or lesser degree). In all cases where present, it was confined to near the wall, and remained laminar and small even when far downstream. There was no correlation between its presence and the later flow deviation. Its nature is in marked contrast to the later flow deviation which led locally to a turbulent flow.

The later bifurcation (in some tests there was more than two modes) seemed to be a manifestation of the inlet variability described in section 6.3. The appearance of this early variability seemed truly random: there was no correlation between it and any external conditions except the settling time between tests. However, the current settling time was optimum in reducing this disturbance; decreasing the time increased the variability, but extending the time did *not* reduce the randomness further than the present level. A significant observation was that the later point of inception seemed unaffected by this bifurcation. This allowed the use of all data, regardless of the nature of these earlier deviations.




Figure 7.11: Multiple traces from the measuring station GA0, ($x/R = 28.76$), illustrating the variability of the point of inception and the resultant perturbed velocity field: (1), laminar flow; (2), perturbed laminar flow; (3), transitional flow; (4), turbulent flow; (5), shutdown variability.

and obtain an indication of the variability. The entire test, data from different axial stations were combined after first extracting the mean from each test. The resultant data thus had a mean velocity, u , with a sufficient population to draw broad statistical conclusions.

For the test GXPA of which the data in figure 7.11 is a subset, the standard deviation of the entire data was 1.6236 s; and the number of data points was 186. Figure 7.12 shows a histogram of the above data; evident here is the bell-shaped structure of the histogram, indicating that the data is distributed normally about the mean.

The fact that the inception data was close to normal, suggested that the flow variation had a simple statistical nature. Some of the scatter was possibly due to errors in estimation. However, the innate variability was so large compared to the error in selection of the point of transition - about 0.1 s uncertainty - that the latter effect on the data could be safely ignored. More worrying than simple-alignment errors, was the possibility that visual processing was introducing a bias into the data; for instance by misinterpreting phenomena. Unfortunately the nature of the data made it difficult to foresee any other means of processing; furthermore the normal distribution tended to indicate that systematic biasing was not occurring. There was however, no firm way of ruling out the presence of a systematic bias.




Figure 7.10: A velocity - time trace in the pipe centreline far downstream, showing the convected turbulent front. The data is EE5.0032, with axial position $x/R = 320.50$.

The data was first smoothed using the standard scheme of section 3.4.2. Traces in which the point was ambiguous or unclear were ignored. The resulting deviation time for each axial station was then obtained by taking the mean of the valid times thus obtained. Thus for each axial station a single value was obtained corresponding to the mean time of deviation there. Furthermore, the standard deviation of the data at each axial station could be determined.

Selection of this point was by the nature of the problem, manual. The data were loaded into Matlab, displayed on the computer screen, and the point of deviation was chosen by selecting it with the computer mouse cursor. To prevent any bias, the individual traces were displayed without a scale; furthermore, no reference was made to the emerging results until all the inception points were extracted.

7.4.5 Variability

As mentioned, a single inception point was obtained by taking the mean of twelve measurements at the same station. In all the tests, there was a fairly large variability between individual data, of the time at which the flow deviation was first noted. Figure 7.11 shows the twelve runs from station GA0, showing the disparity between tests. As can be seen, there is variability innate to the flow system. The standard deviation of the resulting points (in terms of time) was $\delta_t = 1.940$ s. This value translated into a Reynolds number variation of approximately $\delta_{Re} = 2.404$. The statistical distribution of the current points was of interest. In order to increase the data population above twelve values,

Although there is misalignment between the individual traces and the curves (due to the unpredictability and non-repeatability of the deviation process - the velocity traces are individual realisations whereas the curves represent means), the series shows the development of the flow system to good effect. The following features are evident in this figure:

- (i) Near the inlet [curve (1), (2)], the laminar breakdown is followed by an extended period of disturbed laminar flow. Turbulent puffs may be seen, but reversion to the disturbed laminar state usually occurs.
- (ii) The deviations near to the inlet are localised to the thin boundary layer; no core flow perturbations are evident here. Velocity traces from the pipe centreline thus show no deviations for the entire test duration.
- (iii) As the measurement station progresses downstream, [curves (3) - (6)], the point of deviation occurs slightly earlier. The consequent laminar structures are similar throughout this region.
- (iv) The point at which the boundary layer flow reverts to a turbulent state occurs at earlier times when progressing from curve (1) to curve (6). The dimensionless centreline velocity of the velocity profile does not change much in this region, indicating that this turbulent phenomenon is of a local nature, not effecting the shape of the entire velocity profile. In other words, mass-flow conservation effects are restricted to the boundary layer.
- (v) For the more downstream profiles [curves (7) - (10)], the point of flow deviation occurs again at a later time. The initial flow deviation is smaller, but possesses a larger radial extent.
- (vi) The turbulent front occurs closer to the deviation point as the point of measurement progresses downstream. The onset of turbulence in this region is more radially extensive, as evidenced by the large increase in boundary layer velocity following the transition. The central core flow still remains largely laminar in this region, but the boundary layer becomes completely turbulent.
- (vii) Very far downstream, the flow ceases to show any initial laminar deviation. Turbulence ensues suddenly as the measuring station is overcome by a pipe-wide turbulent front. An example of this is not shown on figure 7.9; but such a trace is given by figure 7.10. This figure shows the *centreline* deviation for the measuring station EE5, for test EXPE. The well defined, convected turbulent front [da Silva (1990)] is evident here.

7.4.4 Prediction of the inception point

The identification of the point of departure from the laminar state was made by inspecting each trace, and graphically selecting the start of the phenomenon.

The two events occurring in figure 7.8 - laminar deviation and transition - were typically seen in most of the traces; their relative prominence depended on the location of the trace within the inlet region. Towards the inlet plane, the laminar deviation was well defined and persisted for an extended time. Turbulent transition was either absent completely, or a gradual transition occurred to the turbulent state. As the measuring station was moved downstream, the turbulent transition began to occur more and more as a 'hard front', indicating that the transitional phenomenon was taking on the form of a convected slug [Wygnanski & Champagne (1973)]. Even further downstream, the deviation phenomenon disappeared completely, and the corresponding traces showed only a convected turbulent front. Figure 7.9 shows schematically the entire spectrum of transitional events for a single test.




Figure 7.9: The nature and timing of the transition phenomena in the pipe entrance region. Data is selected from the test G4XPA. Axial stations are: (1), $x/R = 11.5$; (2), $x/R = 17.6$; (3), $x/R = 23.0$; (4), $x/R = 28.8$; (5), $x/R = 33.0$; (6), $x/R = 41.3$; (7), $x/R = 66.2$; (8), $x/R = 82.0$; (9), $x/R = 112.7$; (10), $x/R = 122.9$; — — —, inception (deviation from laminar flow); — — —, transitional behaviour.

Traces from different axial stations are superimposed on the same set of axes for the test G4XPA. The traces are offset from one another in proportion to their axial position; their mean is positioned at their axial position and their magnitude is arbitrary. Thus the horizontal scale represents distance downstream from the exit. Superimposed on the figure is the stability curve extracted from the full data for this test; the solid line represents an approximate fit through this data, and the dotted line the occurrence of the wash-down front. These curves were obtained by fitting an approximate line through the mean deviation and transition points for each axial station. These mean values were calculated by averaging the twelve deviation/transition times at each axial position.

differed from run to run. The inability to capture the entire phenomenon with the LDV method justified the simplistic measurement approach of a single radial measurement station.

7.4.3 Description of the phenomena

Although the deviation phenomenon varies from run to run, and from station to station, most of the traces that were captured showed similar behaviour. A typical timing trace is shown in figure 7.8 [trace (1)]. This is the data GB2.0000, and shows clearly a number of common trends. Superimposed on the above figure is the same data, but with the trace normalised with respect to the mean cross-sectional velocity [trace (2)].




Figure 7.8: A typical velocity - time trace in the boundary layer at the pipe inlet, showing the point of deviation from laminar flow. (1), velocity time trace; (2), velocity time trace normalised with respect to mean cross-sectional velocity; o, point of deviation; •, point of transition.

In this figure, the initial deviation from the simple laminar development is clearly shown by the sudden laminar drop in velocity, marked with an open circle. In both the ordinary and corrected figures a smooth, laminar decrease in velocity was noted. After the velocity has decreased for a small time, transitional behaviour begins to occur. The entire boundary layer starts to increase rapidly in velocity - in both the physical and normalised representations - while classical intermittence is noted in the trace. The point of initiation of this transition effect is indicated by the solid circle in the figure. As before, the transition begins within the thinning region of the flow. This is indicated by the increase in velocity with time in the normalised representation of the undisturbed velocity trace.

7.4.2 Method of measurement

The object of the timing experimentation was ultimately to capture the inception of the flow deviation phenomenon within the axial extent of the flow. The above section demonstrated that a single well chosen radial measuring station was adequate to infer the timing information at an axial position. As mentioned, the use of a point-measurement device, in conjunction with the innate variability of the flow breakdown, precluded obtaining any further levels of information (nature and shape of the disturbance, etc.) from the experimental system. In order to capture detailed information on the flow departure, the following experimental procedure was used:

- (i) An axial station was chosen and the LDV measuring volume was moved from the wall into the pipe. A run was conducted and the results plotted. The control volume was moved radially, and many such traces were plotted.
- (ii) With reference to the plotted results, a radial position was visually selected that showed the greatest magnitude of perturbations. The measuring volume was then positioned at this station. Because the exact radial position of the measuring volume was uncertain¹, and in the light of the weak dependence within the boundary layer of inception time on radial position, this position information was not measured.
- (iii) With the measuring volume appropriately positioned, the apparatus was run twelve times at the same radial position. All twelve traces were stored. The number twelve was rather arbitrary; a significant population of events was simply required. After completion of the twelve runs the LDV was then moved to another axial position, and the above process was repeated.
- (iv) The resulting data from each timing station was smoothed using the standard procedure, in order to remove broadening, thus revealing any small laminar flow changes. The data was *not* corrected for optical distortion; this was because only the timing information was of interest, not the numerical value of the velocity. The data was then normalised with respect to the flowrate (measured using the flowmeter). Again, this normalisation was with respect to an arbitrary datum, the magnitude of the resultant plot being unimportant.
- (v) From each smoothed and normalised trace, the first point of deviation was identified using a graphical selection procedure. Thus for each timing station, twelve timing values were identified.

Twelve separate runs were conducted per station because the inception point showed an innate scatter. In general the specific shape of the perturbation

¹It was required that the boundary layer be accurately traversed in order to determine properly the wall position: see section 7.3.

Parallel development

The test was initiated at time $t = 2$ seconds (to allow sufficient time for the tank pressure to build up); thus all data is offset by this amount from zero. For early times ($2 < t < 5$ seconds) the profiles develop away from the top-hat condition at inception; that is the boundary layer dimension increases with a corresponding increase in dimensionless centreline velocity. The nature of this development is much like that seen in impulsively started pipe entrance flows. At about 5 seconds, boundary layer development halts. This time corresponds to that at which the inlet front reaches the measuring station, or in other words the time at which non-parallel effects begin to occur. The change to the non-parallel regime occurs early at this particular station due to its proximity to the pipe inlet, and because of this, the boundary layer remains relatively thin.

Entrance flow

Between 5 and about 30 seconds the flow shows the unusual trend of boundary layer thinning. Although slight in this instance, the boundary layer traces show an increase in velocity. Not obvious in this figure, but observable, is a commensurate decrease in dimensionless centreline velocity, dictated by flow conservation arguments. The trend toward a thinner boundary layer is most clearly demonstrated by the trace labelled (1): at 5 seconds this trace is clearly within the boundary layer region, whereas by 30 seconds its velocity value has virtually attained the centreline value again. Figure 8.2 shows selected traces from figure 8.1 superimposed to show these trends - this figure effectively represents a view of the data in figure 8.1 in a direction parallel to the radial (r/R) axis, with some traces omitted for clarity.

The decrease in centreline velocity (and simultaneously the increase in the boundary layer velocity) is clearly evident here. This reversion is gradual and uniform between the times 5 and 30 seconds. Some waviness is evident in the two boundary layer traces in this range, and this issue is addressed in section 8.3.

Flow breakdown

The velocity data have been sampled very coarsely - at between 0.5 and 1 Hz - in order to show the profile development clearly. Thus they are not suitable for demonstrating the nature or timing of the flow breakdown. However, the increase in amplitude of the fluctuations in the boundary layer traces [(1) and (2)] in figure 8.2 for the region after 30 seconds, indicates the presence of this breakdown. The flow showing this transition is indicated by the area labelled (2) in figure 8.1. The corresponding centreline flow at this time remains unperturbed. The following chapter will focus specifically on the transition, and to this end will employ optimally processed data.

appendix B. The chapter first discusses in detail data from a representative test; data most suited to this was that from the test FXPB. Three profile measurements were conducted for this test; one near the inlet at $x/R = 13.69$, one intermediate station at $x/R = 44.61$, while the furthestmost downstream measurement station was at $x/R = 90.54$. General characteristics of the data are discussed, thereafter, variability associated with the inlet is mentioned, drawing on specific instances of the data in order to illustrate the phenomena. Finally the limiting behaviour of the experimentation at low accelerations is presented.

8.2 Typical evolution of the velocity profiles

The data presented in the following section is, as mentioned, from the test FXPB, and embodies the generic patterns of flow behaviour seen experimentally. This section thus serves as a forum to introduce the general flow behaviour; more exceptional and complex behaviour is discussed later in the chapter.

8.2.1 The flow variation near the inlet

The station FB3 was the furthestmost upstream profile for this test; its axial position being $x/R = 13.69$. Velocity profiles were extracted at each second for early times, the increment increasing to two seconds towards the end of the test. Figure 8.1 shows these profiles plotted as a surface.




Figure 8.1: Normalised velocity profiles for the test FXPB ($x/R = 13.69$), represented as a surface.

Chapter 8

Experimental results - velocity profiles

The experimental velocity profile results obtained are presented in this chapter. Only the profile data acquired from the experiment are given here, processed in accordance with the methodology described in the previous chapter. The stability and inception results are discussed in the following chapter - chapter 9.

8.1 Methodology of presentation

The means of acquiring the data for all twelve tests was described in the previous chapter (see figure 7.5). For each test, approximately three profiles were measured, one as close to the inlet as possible, and the others evenly staggered in the downstream direction. All the profiles were normalised both in domain and range; the independent variable was $x \in [0, 1]$, while the velocity was normalised with respect to the mean velocity.

Velocity profiles were extracted at one or two second intervals at each station. From these profiles various derived quantities were inferred; for instance displacement thickness and velocity profile shape parameter¹ could be extracted from these velocity profiles. The Reynolds number was available directly from the flow meter, hence the profiles could be normalised without the necessity to extract their actual mean velocity. All the profile data from all the tests showed very similar behaviour. For this reason only representative profiles are given here; in particular data that depicts most clearly common or important features in the system. Conversely, stations showing unique behaviour are also presented in detail. Where necessary, reference is made to the full data set contained in

© 2007 The Author. Journal compilation © 2007 Blackwell Publishing Ltd

¹This term is referred to hereafter simply as the shape parameter, but is not to be confused with the identically named parameter arising from the von Karman-Pohlhausen analysis of the boundary layer.

thickness. The extracted values were in the correct dimensionless form to calculate δ^*/R and S . Firstly the centreline velocity (7.11) was already in the correct form to be used to define the displacement thickness - equation (7.11); secondly the shape parameter could be directly evaluated from the extracted slope as

$$S = -2 \times \left[\frac{du}{dr} \right]_{r=1}. \quad (7.16)$$

7.7 Summary

The testing methodology and method of data processing were outlined in detail in this chapter. The various exponential flowrates of the tests were defined and the means of alignment of the LDV with the rig was discussed. Following this, the method of obtaining both the velocity profile data, and the point of flow breakdown, were given. Finally, the flowmeter data was calibrated, and accurate estimation of the flowrate variation for the tests was thus obtained. The following two chapters respectively give the results for the velocity profile measurement and subsequent flow breakdown; first in chapter 8 the profile data are shown, and then the inception information are given in chapter 9.

This fitting procedure was optimum for the extraction of the mean flow rate (integral quantity), but yielded poor results when used to infer the wall shear stress (derivative quantity). For this reason a procedure was developed specifically to extract this parameter, along with the centreline velocity of the profile. To obtain an accurate approximation to the wall shear stress, a low order (quadratic) polynomial was fitted to only the wall region of the boundary layer. Furthermore, the fit was *not* forced to pass through the inferred position of the wall. The low order combined with no enforcement of the boundary conditions resulted in a curve whose slope accurately followed that of the experimental data. The centreline point was evaluated in much the same way; a linear least-squares fit was made through a limited region of the core flow, and the value at the intersection of this line with the vertical (velocity) axis was taken as the centreline velocity. Figure 7.20 shows an example of two sets of points used to fit the boundary layer and centreline approximations to a velocity profile.




Figure 7.20: The approximate points used to extract the wall derivative and centreline velocity from a normalised velocity profile: (a), points used to find the wall derivative; (b), point used to determine the centreline velocity.

As the profiles had been normalised both in radius and velocity, the slope at the wall and centreline values extracted above were also non-dimensional. The centreline velocity value was simply the dimensionless value

$$u_c = \frac{u_r}{U} \quad (7.14)$$

while the velocity derivative at the wall was

$$\left[\frac{du}{dr} \right]_1 = \frac{R}{t^*} \left[\frac{du}{dr} \right]_R \quad (7.15)$$

The derivative at the wall was required to calculate the *shape parameter*, while the centreline value was needed to evaluate the boundary layer displacement

7.6.2 Velocity profile shape parameter

While the displacement thickness is an integral classification, the shape parameter of Moss (1985) is simply the dimensionless derivative of the velocity profile at the pipe wall. It can be defined as

$$S = 2 \frac{\tau_w R}{\mu U} \quad (7.12)$$

with τ_w the wall shear stress and μ the dynamic viscosity. The wall shear stress is in itself given by $\tau_w = -\mu \left[\frac{du}{dr} \right]_{r=R}$, thus (7.12) can be simplified to yield

$$S = -2 \frac{R}{U} \left[\frac{du}{dr} \right]_{r=R}. \quad (7.13)$$

In the present context, where the velocity profiles are normalised radially to the range $[0, 1]$, and where the velocity has been made non-dimensional according to the mean cross-sectional velocity U , the shape parameter is simply $S = -2 \times \left[\frac{du}{dr} \right]_1$, with $u = u/U$ and $r = r/R$.

The shape parameter varies from value infinity where the boundary layer approaches zero thickness, to a value of eight corresponding to a parabolic velocity profile. It can also be used successfully to quantify a whole range of profiles with shapes more ‘pointed’ than parabolic, including profiles with negative pressure gradients. Such profiles can have shape parameters anywhere in the range $8 > S > -\infty$. Profiles with $S < 0$ in fact show flow reversal at the wall. Both displacement thickness and shape parameter were readily available from the experimental data, notably from the curve fitting procedure to be outlined in section 7.6.3 below. As all experimental profiles were normalised, both parameters could be extracted directly, with no reference to fluid viscosity or other parameters. This effectively eliminated sources of error from the extraction process, pertaining to the determination of such fluid properties. Interestingly, the only quantity directly dependent on the fluid viscosity was the Reynolds number.

7.6.3 Velocity profile curve fitting

Velocity profiles had been normalised with respect to their cross-sectional mean velocities in section 7.5.5. Thus the profiles themselves possessed a mean cross-sectional velocity of unity. The mean cross-sectional velocity was extracted from the velocity profiles in order to calibrate the flowmeter. The curve fitting procedure employed there used a polynomial fit to the entire boundary layer, matched to a lower order polynomial for the core flow.

of (7.10) is expanded and the assumption is made that powers in δ^* are negligible compared to the pipe radius¹², then the linear form becomes

$$\frac{\delta^*}{R} \approx \frac{1}{2} \left(1 + \frac{U}{u_c} \right). \quad (7.11)$$

The equation (7.11) is equivalent to saying that the boundary layer in the pipe - when thin - can be 'unrolled' and approximated by a flat plate boundary layer. Although equation (7.11) is not physically exact when the boundary layer thickens, this presents no problems in the current context for a number of reasons. Firstly, the displacement thickness is simply a convenience parameter anyway, thus an approximation to the exact parameter in the present context would not hold any disadvantages, so long as the approximate parameter was single valued for all pipe boundary layers, and the approximate parameter certainly is. Secondly, any parameter that quantified the boundary layer in an integral sense would have a one equally well here: the displacement parameter was chosen initially only because of its common usage. Finally, the chosen approximation does become exact in the limit as the boundary layer disappears, which allows direct comparison to flat-plate boundary layer results in this limit. Another interesting feature of the displacement thickness (both exact and approximate), is that it is dependent only on the dimensionless centreline velocity. Thus the centreline velocity normalised with respect to the mean flow rate, is in fact an equally valid means of quantifying the boundary layer in a pipe. However, for the present purposes the displacement thickness as given by (7.11) will be used instead, primarily because it relates directly to the boundary layer thickness itself, but also because it allows direct comparison to boundary layer results from other contexts.

The centreline flow is only really analogous to the free stream flow when the boundary layer is thin; that is when the flow still possesses a central core region. When the boundary effects pervade the entire pipe¹³, the centreline velocity is no longer the analogue of the free-stream velocity, and thus both exact and approximate definitions of the displacement thickness cease to have any *physical* meaning. In the light of the discussion above - concerning the physical correctness of this parameter - such a physical invalidity is of no consequence. Indeed a parabolic velocity profile possesses a displacement thickness of $\delta^*/R \approx 0.25$, but this number is simply a parametric classification of this profile. The entire difficulty is semantic, relating to the word *displacement thickness* itself.

¹²Such an assumption is tantamount to saying that the displacement thickness is much smaller than the pipe radius.

¹³A good criterion by which to judge the disappearance of a core flow is to consider when the second and higher derivatives of the velocity profile become large at the centreline. When this happens, it implies that the profile is no longer *flat* near the centre.




Figure 7.19: The definition of the boundary layer displacement thickness. The flows represented by the two shaded areas $a - b - c - d$ and $a - b - c - f$ are equal.

In the context of a planar boundary layer, U is the free-stream velocity, and the displacement thickness δ^* is defined such that the two regions $a - b - c - d$ and $a - b - c - f$ have equal areas, implying that the flow rate in the boundary layer across the length $a - f$ is equal to the equivalent mass flow rate at free-stream velocity across the interface $a - d$. In the boundary layer occurring in a pipe however, the definition must be modified to take into account the axisymmetry. In this latter context the mass flow rate in the boundary layer is not simply proportional to the area under the curve but is rather the quantity

$$Q_{af} = 2\pi \int_a^f ru(r)dr, \quad (7.9)$$

where r is defined as $r \equiv R - y$ and $u(r)$ is the velocity variation in the pipe. The quantity equivalent to the free stream velocity in a pipe is the centreline velocity. This is however only strictly physically valid when the boundary layer is thin; that is, when the flow consists of a boundary layer plus well-defined core region. Using the above concepts along with equation (7.9), the displacement thickness in a pipe can be defined exactly as

$$\left(1 - \frac{\delta^*}{R}\right)^2 = \frac{U}{u_c} \quad (7.10)$$

with U the mean cross-sectional velocity given by $U = \frac{2}{R^2} \int_0^R ru(r)dr$, u_c the pipe centreline velocity, and R the pipe radius. Equation (7.10) is quadratic in δ^* , not allowing direct evaluation of the quantity. An approximation can be made to the equation - valid as $\delta^* \rightarrow 0$ - linearising it and thus allowing direct evaluation of the displacement thickness. If the quadratic on the left-hand side

7.5.5 Velocity profile normalisation

Although the velocity profiles were now non-dimensional in radius, the velocity was still dimensional. It was decided to normalise each experimental profile using the mean cross-sectional velocity obtained from the calibrated flowmeter data. As a result all the profiles had a mean cross-sectional velocity of unity, and the flowrate information was contained only in the corrected data from the flowmeter. Thus, although all the flowrates obtained by integration were used for calibrating the flowmeter, the success of this calibration obviated the need to use the integrated data directly to determine the flowrate or Reynolds number. The normalisation made direct comparisons between different velocity profiles possible; profiles differing only in magnitude would be identical in the normalised framework. All the experimental velocity profiles were reduced to this normalised form, allowing general comparison between test data. The full velocity profile and other results are presented in the following chapter - chapter 8.

7.6 Derived quantities

While the consideration of the full velocity profile variation in time within the pipe is elucidating, the extraction of quantitative parameters for accurate determination of the measured flow variation is perhaps more useful. Two derived parameters were considered in the current context: firstly the dimensionless boundary layer displacement thickness δ^*/R , whose definition can be found in any elementary fluid mechanics textbook [e.g. Schlichting (1979)]; secondly, the dimensionless wall shear stress or *velocity profile shape parameter*, as defined by Moss (1985). Both these terms give a clue to the structure of the boundary layer. The displacement thickness is an *integral-like* quantification of the boundary layer thickness, whereas the shape parameter yields *differential* information from near the wall. Because of the different natures of these two factors, two independent characterisations of the boundary layer should be provided by their use.

7.6.1 Displacement thickness

The displacement thickness is a quantification of the deficit in the mass flow caused by the presence of the boundary layer. Simply speaking, if the boundary layer were to hypothetically disappear, and the entire velocity field was to assume the value of the far stream flow, then the mass flow would increase due to the excess flow adjacent to the wall. The displacement thickness would then be the thickness of the fluid layer at the wall that would have to be excluded from the hypothetical flow to ensure the same mass flow rate as when the boundary layer was present. This quantity is shown diagrammatically in figure 7.19.

Data from the flowmeter was sampled at 100 Hz. The raw data from the flowmeter was smoothed using the same procedure as used for the LDV traces, except the normal (and not weighted) mean and standard deviation algorithms were employed; the smoothed data was sampled at 10 Hz. Smoothing was necessary because the flowmeter data showed some scatter, especially for the lower flowrates. The distribution of the data about the mean was normal, revealing the scatter to be true random noise, thus averaging or smoothing of the data was entirely valid. The appendix A contains an error analysis of the flowmeter data.

Due to experimental uncertainties and non-linearity in the valve system, the flow variation attained did not exactly match the expected variation shown in the figures. The closest matching curves are those given in figure 7.16. In this figure the experimental data is virtually indistinguishable from the predicted flow variation. The only noticeable discrepancy in this figure was seen at the upper end of the flow curves. The deviations at the upper end were caused by a strong non-linearity in the control valve, resulting in the control system under-compensating at that point. In the following two figures (figures 7.17 and 7.18), deviations were slightly more marked; in both sets of tests the initial experimental Reynolds number did not match the predicted value, but the data converged thereafter. There was a problem maintaining an accurate impulsive start while still avoiding overshoot. For the higher Reynolds numbers that these two figures portray, undershoot was unavoidable. A further aggravating factor was the large range of Reynolds numbers covered by the test. Feedback control parameters that were optimum for the mid-range of the flow resulted in sluggish behaviour at start-up. If ever optimising parameters for start-up resulted in oscillations at higher Reynolds numbers.

Although the actual flow curves did deviate from the expected values, their variation throughout the range was smooth. This allowed these data to be approximated as low order (≤ 7) polynomials, thus allowing straightforward extraction of derivative and other information¹¹. Furthermore, ignoring history effects, it was not extremely important that the actual flow variation follow the prescribed flow exactly, so long as the actual variation be captured correctly. Using the curve fitting approach, the smoothing of the data described above was in fact not necessary. A curve fit to the unsmoothed data yielded results indistinguishable from the curve fit to the smoothed flow variation. A general polynomial fitting routine was chosen over, for instance, an exponential one, firstly because the former was readily available in Matlab - the environment in which the data was processed - and secondly, to avoid biasing the data in any preconceived manner.

¹¹Knowledge of the flowrate derivative dQ/dt will be important in the later analysis of the flow.




Figure 7.17: The variation of the measured Reynolds number with time, for the tests: (1), FXPA; (2), FXPB; (3), FXPC; (4), FXPD; ---, expected variation; —, actual variation.




Figure 7.18: The variation of the measured Reynolds number with time, for the tests: (1), GXPA; (2), GXPB; (3), GXPC; (4), GXPD; ---, expected variation; —, actual variation.

duration was of the order of 40 s. Thus with $N \approx 4000$ and the maximum deviation $u'/u = 2\%$, the standard error was insignificant at about 0.01%. Error in the flow measurement was thus negligible and the variability of 2% which applied to the entire calibration process was thus taken as the maximum error in the flowrate. Using the fitted data, all the flow rate readings for all the tests were corrected appropriately. The experimental data in figures 7.16 to 7.18 was already calibrated according to the data above.

7.5.4 Flowrate variation with time

The closed loop flow control of the rig generated flows which closely approximated those given by equation (7.1). Dependent on the values of the constants Re_0 and k , the flow control was able to match to a greater or lesser degree the chosen flow variation. Distinct values of these constants within the range of maximum accuracy were chosen¹⁰, such that this range was adequately covered. The expected variation of the Reynolds number with time for each test is shown in figures 7.16, 7.17 and 7.18 below. Superimposed on these figures is shown the actual flow variation. Each figure represents a set of tests with the same initial Reynolds number Re_0 .




Figure 7.16: The variation of the measured Reynolds number with time, for the tests: (1), EXPI2; (2), EXPIF; (3), EXPG; (4), HXPA; ---, expected variation; —, actual variation.

¹⁰As mentioned earlier, the response of the rig was determined by multifarious factors, such as the control valve response, system momentum, and flowmeter response. The system thus had an inherent range of variation in which the systems' response was sufficiently close to the input signal (the chosen flow variation). This range was directly ascribable to these minute limitations of the individual components.

Further evidence for the non-receptivity of the flow to the low frequency waves is revealed by considering the amplitude of these waves and the subsequent nature of the transition. In the present test, transition occurs more or less instantaneously everywhere in the entrance at about $t = 70$ seconds, but the environment (in terms of the waviness) is vastly different between the far-upstream stations and those far downstream. The real originators of the transition process - whether linear waves amplified in the entrance, finite-amplitude disturbances converted in through the inlet, or a combination of both - are likely to be much smaller than the wiggles seen in the flow, but at the correct, receptive frequency to induce transition. The determination of the spectral content of the flow and the magnitude of the background disturbance is considered in detail in the following section.

8.4 The level of finite amplitude disturbances

All the data used until now consisted of or was extracted from smoothed velocity-time traces. The smoothing process removed the fluctuating, high-frequency component from the traces, resulting in smooth variations of the velocity with time being presented. The values of the smoothing parameters were set such that the high frequency cut-off of the filter was approximately 5 Hz, that is, frequencies above 5 Hz were removed from the signal.

It is well known that the presence of finite-amplitude disturbances in a flow has a profound effect on that flow's stability behaviour [cf. Herbert (1988)]. To this end it was decided to investigate the level and nature of the finite-amplitude perturbations present in the pipe flow. The LDV is unfortunately not an ideal device for measuring perturbation magnitude, because it is subject to many forms of signal broadening which can imitate 'false' turbulence. This nature of the LDV signal was discussed in detail in chapter 3. As a result of these limitations, it was decided to eliminate as much as possible the sources of these effects in the spectral measurements.

To remove the above limitation, and to simplify the comparison between tests, it was decided to measure the disturbance intensity only at the centreline position near the inlet. Measurement of the spectral characteristics here was made primarily because this point was the best indication of the innate level of disturbance converted through the inlet. Two types of measurement were made of the small disturbances: firstly the classical *disturbance intensity* was extracted as a measure of the gross disturbance magnitude, secondly the *power spectrum* was measured in order to give an indication of the frequency range of the perturbations. Unsmoothed and unfiltered LDV data were used for the measurements in all cases. Twelve velocity-time traces were investigated in total, corresponding to the most upstream centreline measurement for the twelve different tests conducted.




Figure 8.10: Velocity profiles for the test HAJ ($x/R = 32.99$), represented as a surface.

8.3.3 Non-receptivity to the flow waviness

The presence of large oscillations in the base flow could be taken as eventual cause for transition, by virtue of the ‘finite-amplitude’ environment they imply. However the experimental evidence points to their not affecting the transition of the flow to any large extent. While the eventual transition almost certainly takes place in a finite-amplitude environment, the current waves seem to be not of the right frequency to destabilise the flow; the waviness in the case of the test HX1A is at frequencies well under 1 Hz, whereas the receptive frequency of the boundary layer would be expected to be approximately those of the Orr-Sommerfeld modes. As a quick order-of-magnitude determination of these frequencies, the stability results for the steady entrance flow of Abbot & Moss (1994) can be used as a guide¹. Here, at the minimum critical point the wavenumber and celerity of the most unstable mode are $\alpha = 2.206$, $c_r = 0.458$, while the critical Reynolds number is $Re_{crit} = 11667$. The temporal frequency (in Hertz) of the most unstable disturbance is given by

$$f = \frac{\alpha c_r R^2}{2\pi\nu}, \quad (8.1)$$

where R is the pipe radius, ν and ν is the kinematic viscosity, which for water at 25° C is $\nu = 0.89267 \times 10^{-6} \text{ m}^2/\text{s}$. For the current values given above, the critical frequency is $f_{crit} \approx 100 \text{ Hz}$. Thus the waviness differs in frequency from the receptive range of the flow by two orders of magnitude.

¹For comparison, see also the discussion in §8.3.1.

²The later analysis will show that the steady stability behaviour is close to that of the accelerating flow.

or extracted LDV mean-flow data, but was suspected to be present² in the flow. Whatever the actual cause of the unevenness, its presence in the entrance flow, is marked. In contrast the parallel development region ($2 < t < 20$ seconds) is still smooth and regular. This indicates that the unevenness is not a gross mass-flow effect, since it will also have been evident in the parallel flow had it been so. This wavy behaviour of the flow is as before of a random nature: the assembly of traces in the figure do not represent one single global 'snapshot' of the flow. All traces in the thinning region remain properly laminar. The waviness is in this case in the free-stream flow and outer boundary layer: flow near the wall remaining smooth. This is a strong indication that the waviness is not a boundary layer phenomenon. The border between the parallel and thinning portions of the flow is marked by a distinct ridge, indicated by (1) in the figure. Profiles immediately after this ridge show a distinctly lower centreline velocity, although the boundary layer development remains undisturbed. The ridge is the so-called entrance hump seen in the other test: its strong presence and the intensity of the subsequent variability seem to go hand in hand.

In the light of the discussion concerning the other flows (sections 8.2.1 and 8.2.2 above), the hallmarks of the exponential acceleration are still evident despite the strong fluctuations in the flow. For instance, at about $t = 70$ seconds in the current figure, boundary layer breakdown ensues as before, leading to a fully turbulent profile by the end of the test. The general structure of this breakdown phenomenon is identical to that seen in figures 8.3 and 8.4³. Thus despite the large uncertainties in the data, it can be successfully interpreted using the same approach as before. Despite the randomness, the flow behaves predictably *in the mean*; furthermore the transition phenomena seems completely unrelated to, and unaffected by the waviness. The waviness in the individual traces seemed to magnify as it progressed downstream. Observation of data from a further upstream station - IIAJ at position $x/R = 32.99$, given in figure 8.10 - attests to this. In this figure the waviness was restricted temporally to the region about the transition point between parallel and thinning flow, and spatially to the outer edge of the boundary layer. There was however still a ridge present between these flows - due quite probably to the exact nature of the start-up impulse. Most of the later thinning region in this case was almost completely unperturbed, and the transition phenomenon was thus clearly visible.

For the current test (IXPA), profile measurements at five different axial stations were conducted in total. The appendix B contains the other three data - IIAI, IIAJ and IIAK, given as figures B.33, B.34 and B.36 respectively. These data clearly show the progressive increase in magnitude of the waviness downstream, as well as the transition to turbulent flow towards the end of the test.

²This was inferred by attempting lower accelerations than IXPA, the result being visible saw-tooth waves in the flowmeter data. Furthermore direct observation of the control valve during test IXPA showed it to be sticking slightly towards the start of the test.

³This breakdown is not merely wash down transition. As will be shown in the next chapter, the flow breakdown for the test IXPA occurred almost simultaneously everywhere within the entrance.

using compressed air to drive a large cylinder attached to the valve needle-and-seat arrangement. This air cylinder had a range of 'velocities' - time rates of change of valve position - over which the output movement was smooth. At low rates of valve opening, the sealing arrangement of the pneumatic piston caused stickiness, resulting in uneven movement of the piston, and as a result jerky valve opening. This problem was thought to be the primary cause of poor flow behaviour at low accelerations.

The exponential shape of the required flowrate curve worsened the effect of the sticky valve. An exponential flow requires low initial acceleration, followed by increasing rates of acceleration. Thus any valve 'jerkiness' was most likely to occur at the beginning of the test, when valve accelerations were at their lowest. If the flow start-up could be ensured to be sufficiently smooth, then the remainder of the test invariably remained so. This valve problem resulted in a stepped flowrate variation, as the valve alternately stuck and then slipped. The presence of this saw-tooth wave had a profound effect on the flow, generating variable and wavy velocity profiles. Waviness seemed to be aggravated by the low acceleration itself, because the stabilising effects of acceleration were simultaneously minimised. A good demonstration of the effect of the low acceleration is the test HXPA. This test used the lowest acceleration value (see figure 7.16 on page 126), and the velocity profile results from the downstream station are shown in figure 8.9.




Figure 8.9: Normalised velocity profiles for the test HXG ($x/R = 113.57$), represented as a surface.

The profiles shown in figure 8.9 were measured at axial station $x/R = 113.57$. The unevenness of the flow variation was not discernible in either the flowmeter

To illustrate the convection of the bump, figure 8.8 gives a series of velocity-time traces, each from different axial stations in the test EXPD, and showing its evolution. Although the data in this figure is not from a single test, raising the same objections concerning classification of an erratic phenomenon as mentioned in the previous chapter, the general development trend is still evident.




Figure 8.8: The development of the entrance 'bump' with axial position, for the test EXPD. The dotted line shows the approximate locus of the inlet bump.

This composite was constructed from traces showing the largest magnitude 'bumps'. The evolution of this structure was very similar to one seen in a series of experiments conducted by Wood (1991), in which axisymmetric disturbances were introduced into an impulsively started pipe flow using an iris-diaphragm mechanism. In that study, huge spatially localised laminar structures persisted far downstream of the disturbance generator, and developed independently of the 'carrier' base flow. It can be clearly seen that the later flow breakdown (at about 60 seconds) is unrelated to the inlet bump.

8.3.2 Flows at low accelerations

The twelve tests carried out ranged over a number of exponential accelerations; these were given in the previous chapter. Their range was limited by the physical characteristics of the experimental apparatus. At the high acceleration end, the response of the flow control apparatus set a limit above which an exponential flow variation could no longer be accurately maintained. At the low acceleration end, other effects limited the range of acceleration, these being primarily 'jerkiness' of the mechanical valve, combined with the innate sensitivity of such low acceleration flows to changes in acceleration. The attainment of an exponential flow variation in time required, as explained earlier, the electro-pneumatic control valve to be moved continuously under closed loop control from a position of high flow restriction to a more open configuration. This was achieved by

8.3.1 The inlet bump

Apart from the variability described above, there was another phenomenon ascribed to the inlet: an inlet 'bump' that convected downstream and was associated with the interface between the parallel flow and the following entrance flow. This was intermittent but seemed also to be related to the details of the initial impulsive start. Unfortunately it was not possible to select suitable start-up conditions freely, thus some tests were more susceptible to this effect than others. A good example of a test showing this effect strongly is FXPD. Profiles from the station FD3 ($x/R = 36.31$) are shown in figure 8.7.




Figure 8.7: Velocity profiles for the test FD3 ($x/R = 36.31$), represented as a surface. (1), the inlet bump.

Figure 8.7 presents data from a station close to the inlet. In this figure a large temporally localised wave can be seen at time 10 to 15 seconds. Its temporal location is at the parallel-entrance interface, and its radial location is approximately in the outer part of the boundary layer and adjacent free-stream flow. As with the variability described in the previous section, the appearance and magnitude of this wave was erratic. Thus, as before, the wave shown in the figure is *not* a global representation of the phenomenon; rather it is an assembly of separate traces in which the disturbance occurred erratically. While the structure of this phenomenon could not be inferred from the profile data, observation of the individual traces showed that it was laminar and possibly vortical in nature. The structure persisted and maintained its identity far downstream, always appearing as a simple 'bump' on the velocity trace. Although quite large, the bump seemed to convect through the flow, leaving behind a velocity field unaffected by its passing. This is evident in figure 8.7, where the traces after about 15 seconds simply revert to their normal evolution.

measurement device, was capable of measuring the flow variation at only a single point in the pipe for any single run. The process of measuring a velocity profile resulted in moving the measuring volume through the boundary layer, positioning it at different radial stations, and repeating the run. Thus with the *global* variability associated with the flow, the assembly of separate traces (from different runs) would result in a velocity profile whose boundary layer showed waviness in the radial direction, *even though the separate runs had smooth profiles differing only slightly from one another*. Thus the waviness of the boundary layer in the thinning region of the flow, as for example shown in figure 8.4, is in all probability simply due to assembly of traces obtained from flows which differed slightly, and is not innate to the system. If it were possible to globally capture a velocity profile from the current flow, it would show smooth regular behaviour with no waviness.

The three-dimensional velocity profiles shown earlier were constructed directly from the smoothed traces, whereas derived data such as the mean flow rate or centreline velocity were extracted by fitting an interpolated line through all the boundary layer points at each time. Thus to calculate derived boundary layer quantities, the *mean* velocity variation of the entire boundary layer was effectively used. Figure 8.6 shows the profile from the station FB5 ($x/R = 90.54$) at $t = 22$ s, showing an interpolated curve to the raw data points.




Figure 8.6: The velocity profile from the test FB5 ($x/R = 90.54$) at time 22 seconds; $\circ$, measured data; —, interpolated curve.

This profile is from the most variable region of the station FB5; nevertheless the scatter is insignificant when compared to the interpolated curve. Thus in terms of the extraction of velocity profiles, the deviations in the individual traces were wholly insignificant.

thickness increases. The variation of this parameter for all three stations is given in figure 8.5.




Figure 8.5: The temporal variation of dimensionless centerline velocity in time, for the test FXPB: $\circ$, FB3 ($x/R = 13.69$); $\bullet$, FB4 ($x/R = 44.61$); $\circ$, FB5 ($x/R = 90.51$).

It is clearly evident that all three curves depart from a single curve - the development curve for flow infinitely far from the pipe inlet - and that the various times of departure are in relative proportion to the station's distance from the inlet. The variation of the boundary layer in time for all the tests will be presented in a more complete form later in the chapter, where the boundary layer displacement thickness and dimensionless wall shear stress will be used as characterising parameters. Firstly the issue of inlet variability and inlet effects will be addressed.

8.3 Inlet variability and inlet effects

Whereas the parallel development region was very smooth, uniform and repeatable, the boundary layer thinning flow showed far greater variability. This phenomenon was mentioned in the previous chapter and earlier in the current one, and was ascribed to a peculiarity of the inlet contraction in conjunction with the nature of the initial impulsive start.

The most probable explanation for the observed variability involves a small random inlet effect, affecting slightly the global shape of the velocity profiles, in conjunction with the point-measurement nature of the LDV method. Simply put, slightly different inlet boundary conditions to the flow at time zero occurred every time a run was conducted, which resulted in the boundary layer shape during the thinning region (the region in which inlet effects had been convected downstream) differing slightly between runs. The LDV, being a point

upstream to convert into this region; it is for this reason that this data shows no advanced transitional behaviour.

For the test FXPB, a third set of profiles was measured further downstream, at $x/R = 90.54$, and this data is presented in figure 8.4. Again the sequence of events is familiar: an initial increase in boundary layer thickness commensurate with the parallel flow development; next a region of boundary layer thinning, followed by the development of boundary layer perturbations leading finally to a turbulent velocity profile. As before, the relative extent of the regions differs according to the axial station: the current data show a far larger parallel development region than either of the two previous data sets.




Figure 8.4: Normalised velocity profiles for the test FB5 ($x/R = 90.54$), represented as a surface.

8.2.3 Similarity between the parallel development regions

Of interest is the parallel (early time) developing region of each profile set. It would be expected that the profile development at each of the stations would be identical here. In this region, the flow should be developing as if it were an indefinite distance from the inlet, with no radial velocity component, and as a result the conditions at all stations should be identical. The start of the boundary layer thinning region, as heralded by the arrival of entrance flow, marks the time at which the flow starts to be influenced by the inlet boundary conditions, and thus the point at which differences between the measuring stations should become apparent. A crude yet clear indicator of the state of the velocity profile within the pipe is the dimensionless centreline velocity. The value of this parameter is unity at start-up, increasing monotonically as the boundary layer




Figure 8.3: Normalised velocity profiles for the test FB1 ($x/R \approx 41.61$), represented as a surface.

the end of this phase is larger than at the inlet station. As before the parallel region ends when the pipe inlet region converts past the measuring station.

The region defined by $12 < t < 28$ seconds is the boundary layer thinning region of the flow. Again the thinning effect is gradual, with the boundary layer showing a slow yet consistent decrease in dimension with time. Noticeable again is the slight variability in this region - the traces at duplicate stations within the boundary layer at this point do not lie concurrently, but deviate slightly from one another. As mentioned earlier, this variability is discussed in detail in section 8.3.

At $t = 28$ seconds, boundary layer perturbations are noticeable, and the boundary layer becomes more and more disturbed, until by 35 seconds almost the entire velocity profile is turbulent. This flow breakdown progresses from a distinctly laminar waviness - limited to a small region of the boundary layer - to a more widespread transitional process finally resulting in a turbulent boundary layer. The fully turbulent profile present after $t = 28$ s quite likely originated further upstream, from perturbations in the boundary layer that had grown while converting downstream. This conclusion can be supported by results of the inception data, presented later. These results show that a substantial portion of the boundary layer within the entrance region seems to show a laminar breakdown more or less simultaneously. The results of such a global breakdown would be exactly those seen in figure 8.3 - initial laminar waviness (local), followed by transition to turbulence, triggered by the presence of converted disturbances from further upstream. This is corroborated by considering again figure 8.1, where because of the proximity to the inlet plane, there was no instability far




Figure 8.2: Superimposed normalised velocity traces for the test FB3 ($x/R = 13.69$), from selected radial stations: (1), $r/R = 0.975$; (2), $r/R = 0.932$; (3), $r/R = 0.866$; (4), $r/R = 0.838$; (5), $r/R = 0$ (centreline).

The velocity profiles shown in figure 8.1 have been assembled from individual LDV traces; thus each velocity-time reading corresponds to a separate experimental run. As discussed in the previous chapter, the events during and after transition were not repeatable in structure to any great extent - the timing of their inception seemed predictable with some statistical variation, but subsequent deviation was not unique in nature. For this reason it was not possible to draw any strong conclusions about global structure of the flow breakdown from this data. The radial extent of the perturbation was, however, localised to the boundary layer alone. At this current - far upstream - station, the scale of the flow breakdown was still minor; progressing downstream, these fluctuations, while similar in nature, become far more pronounced and are discussed below.

8.2.2 Profile development further downstream

While the flow at the inlet maintained a very thin boundary layer, that further downstream showed far more marked development during the test. Figure 8.3 shows this development for the station FB4, located at axial position $x/R = 44.61$.

This data shows analogous trends to the inlet flow; as before, the boundary layer initially develops away from the 'top-hat' condition at start-up, this regime being followed by a region of boundary layer thinning, and finally flow breakdown. The initial parallel development however, occurs over a far longer time, up until $t = 12$ seconds. This prolonged development time is due to the measuring station's larger distance from the pipe inlet, resulting in a greater time before the entrance flow reaches the station. Again, because of the extra duration of the parallel development of the flow, the maximum boundary layer thickness at




Figure 8.2: Superimposed normalised velocity traces for the test FB3 ($x/R = 13.69$), from selected radial stations: (1), $r/R = 0.975$; (2), $r/R = 0.932$; (3), $r/R = 0.866$; (4), $r/R = 0.838$; (5), $r/R = 0$ (centreline).

The velocity profiles shown in figure 8.1 have been assembled from individual LDV traces; thus each velocity-time reading corresponds to a separate experimental run. As discussed in the previous chapter, the events during and after transition were not repeatable in structure to any great extent - the timing of their inception seemed predictable with some statistical variation, but subsequent deviation was not unique in nature. For this reason it was not possible to draw any strong conclusions about global structure of the flow breakdown from this data. The radial extent of the perturbation was, however, localised to the boundary layer alone. At this current - far upstream - station, the scale of the flow breakdown was still minor; progressing downstream, these fluctuations, while similar in nature, become far more pronounced and are discussed below.

8.2.2 Profile development further downstream

While the flow at the inlet maintained a very thin boundary layer, that further downstream showed far more marked development during the test. Figure 8.3 shows this development for the station FB4, located at axial position $x/R = 44.61$.

This data shows analogous trends to the inlet flow; as before, the boundary layer initially develops away from the 'top-hat' condition at start-up, this regime being followed by a region of boundary layer thinning, and finally flow breakdown. The initial parallel development however, occurs over a far longer time, up until $t = 12$ seconds. This prolonged development time is due to the measuring station's larger distance from the pipe inlet, resulting in a greater time before the entrance flow reaches the station. Again, because of the extra duration of the parallel development of the flow, the maximum boundary layer thickness at




Figure 8.18: The variation in experimental shape parameter and displacement thickness, as a function of time, for the test EXPE ($k = 0.03$, $Re_0 = 3018$): o, EF1 ($x/R = 12.24$); +, EF2 ($x/R = 20.25$); *, EF3 ($x/R = 61.52$); ■, EF5 ($x/R = 320.50$).




Figure 8.19: The variation in experimental shape parameter and displacement thickness, as a function of time, for the test EXPF ($k = 0.05$, $Re_0 = 3018$): o, EF1 ($x/R = 63.69$); +, EF2 ($x/R = 40.01$); *, EF3 ($x/R = 13.49$).

are accurate. The continuing increase of curve FA3 during and after transition was puzzling. It was seen in the shape parameter data that the inlet front arrived here at about 30 seconds, thus a reversion to decreasing ν^* would be expected. The continued increase in this parameter was possibly due to inlet effects (the 'blip') or the initial transitional process inducing profiles with larger than usual centreline velocities. The result was that the displacement thickness downstream seemed to continue developing as normal.

8.7 Comprehensive variations of the derived parameters

The previous section specifically used the results from the test FXPA, in order to introduce both the velocity profile shape parameter, and the displacement thickness. This test presented a variation in these parameters that was typical of the general behaviour of all the exponentially accelerating flows measured. This test was primarily chosen because it showed large variations in the shape parameter. As well as being used for a general introduction, it was also used to demonstrate the correction procedure for the shape parameter data. What follows below is a presentation of all the data - both displacement thickness and shape parameter - for each of the remaining eleven tests. The data is presented in sequential order, and illuminating commentary is presented where necessary. Data from the last test - HXPA - was however grouped with the 'E' series of tests, because of the similarity in initial impulsive Reynolds numbers of the former test with the latter ones.

8.7.1 The tests EXPE, EXPF, EXPG and HXPA

The figures 8.18, 8.19, 8.20 and 8.21 below represent the variation of both the measured displacement thickness and velocity profile shape parameter in time, for the 'E' series of tests. The approximate exponential acceleration constants k and R_{α} are indicated on each figure.

8.6.2 Displacement thickness

The dimensionless centreline velocity and hence the displacement thickness were not subject to the same systematic variability as was the shape parameter. It is shown in section A.3 that the error in the wall position has a negligible effect on the mean flow rate and hence on the displacement thickness. The error analysis in section A.1 of the same appendix shows the uncertainty in the centreline determination to be about 1%. This was borne out by the extracted data - curves of displacement thickness at different stations from a single test showed concurrence in the displacement thickness development in their parallel developing regions. Figure 8.17 shows the variation of the displacement thickness in time for the same test FXPA as used in the previous section.




Figure 8.17: The variation in experimental displacement thickness as a function of time, for the test FXPA: o, FA0 ($x/R = 25.52$); •, FA1 ($x/R = 13.07$); ×, FA3 ($x/R = 67.81$); —, mean fitted lines.

The trends in the above data were similar to those shown in the variation of the centreline velocity in figure 8.5; this was expected because the displacement thickness was derived directly from the centreline velocity. The value of δ^*/R for all three curves increased rapidly from zero at time $t = 2$ s - the time of flow inception, thereafter the most upstream curve reverted to boundary layer thinning behaviour, with the displacement thickness decreasing gradually from the maximum point. The station FA0 - the next most downstream station - showed further boundary layer development before thinning behaviour commenced. The most downstream station FA3 showed no decrease in the displacement thickness.

The transition in the flow was seen to occur at stations FA0 and FA1 at about 22 seconds, and at FA3 after some 30 seconds. The behaviour of FA0 and FA1 became more variable at transition, but in both cases the displacement thickness continued reducing. Unlike in the case of the shape parameter, and because of its integral nature, the extraction of δ^* remained relatively accurate throughout transition and within the turbulent regime, thus the full extent of these curves




Figure 8.16: The variation of the corrected velocity profile shape parameter with time, for the test FXPA: $\circ$, FA0 ($x/R = 25.52$); $\bullet$, FA1 ($x/R = 13.07$); $\times$, FA3 ($x/R = 67.81$); —, mean fitted lines.

shape parameter reached a lower value before reversion to the thinning region occurred. The increase in S in the thinning flow was more gradual than at station FA1, and slight waviness was evident, possibly due to inlet effects. The transition between the parallel and thinning regimes was also more gradual than at the upstream station. The station FA3 - curve FA3 - was positioned far downstream, and the parallel development of the flow thus persisted for the predomance of the test. As would be expected, the rate of decrease in S reduced until at 30 seconds the curve was becoming asymptotic to some final value of S . At 30 seconds the inlet front arrived, again reversing the development trend. Shortly afterwards transitional flow from upstream was convected through the station. This resulted in a rapid increase in S , not due to the thinning flow behaviour but caused by transition to turbulence.

Transition and the shape parameter

In all three curves flow breakdown occurred - simultaneously at about 22 seconds in curves FA0 and FA1, and after 30 seconds in curve FA3. The process of extracting the shape parameter was not however suited to measuring its value in highly variable velocity profiles⁹, neither was the positioning of measuring stations adequate to extract the wall shear stress in transitional or turbulent boundary layers. The curves after the point of transition shown above were thus inaccurate. Realistically, the turbulent part of the flow should have shown a very high shape parameter here.

⁹The flow breakdown was a random process, the events occurring at slightly different times for different tests. Thus the composite velocity profiles from this region were not correct and showed a large scatter in the boundary layer.




Figure 8.15: The corrected experimental shape parameter as a function of time, for the test FXPA: (1), FA0 ($x/R = 25.52$); (2), FA1 ($x/R = 13.07$); (3), FA3 ($x/R = 67.84$).

8.6 Typical behaviour of the derived parameters

The presence of a systematic experimental error was accounted for and eliminated in the previous section. In this section, representative variations of displacement thickness and shape parameter from the test FXPA are given, to illustrate the generic experimental behaviour of these parameters.

8.6.1 Velocity profile shape parameter

The general behaviour of the accelerating pipe flow is immediately apparent in the corrected data of figure 8.15, which is repeated below as figure 8.16. Data points plus mean fitted lines have been given in the figure below, in order to emphasise the general experimental trends⁸.

The identical parallel flow development at all measuring stations is evidenced by the decrease in S with time – decreasing S is consistent with a thickening boundary layer. The appearance of the entrance flow at each of the three stations is immediately evident. The first station to undergo reversion to a boundary layer thinning regime was FA1 – the station FA1 – and this occurred at time 6 seconds. At this point, the decrease in S with t was halted, the curve showing a sudden change to an increasing trend. This increase in S was far more gradual than the decrease – consistent with the slow decrease in boundary layer thickness in the thinning region seen in both the profile and centreline data. At the intermediate station FA0, parallel development persisted for longer; hence the

⁸The crude representation of figure 8.15 was specifically intended to illustrate the process of curve alignment.




Figure 8.14: The uncorrected experimental shape parameter as a function of time, for the test FXPA, showing the systematic deviation between results from the different measuring stations: (1), FA0 ($x/R = 25.52$); (2), FA1 ($x/R = 13.07$); (3), FA3 ($x/R = 67.81$).

8.5.2 Correction of systematic errors in S

Application of the correction factor (8.4) was made, and corrections were applied to all three curves. The three factors were chosen such that their position *in the mean* remained unaltered. As there was no fixed reference value by which the correctness of shape could be ascertained, this mean positioning of the curves was the most natural method to use. The factors were adjusted by visually observing the data until the parallel portion of the curves became concurrent.

For the curve (1), the correction factor was $\delta p_{10} = 0.005$, curve (2) had no correction applied, while for (3) the factor was $\delta p_{10} = -0.005$. The result of the application of these correction values was the data shown in figure 8.15.

What is immediately apparent in the corrected data is the almost exact correlation between the curves during the parallel developing part of the flow. Apart from the first value - which was subject to large error - introduced by the smoothing of the impulsive start to the velocity traces - the three curves are virtually concurrent.

The correction procedure for the shape parameter thus showed itself to be accurate at correcting discrepancies in the measured shape parameter. It was, however, pointed out that no independent verification of the correct shape parameter was available. An ideal 'calibration' tool for the correction process would have been a wall mounted shear stress probe or similar device; unfortunately none was available for the rig.

8.5.1 Errors in the velocity derivative at the wall

The process of extracting the wall shear stress was prone to large systematic errors, due to the unknown position of the wall. In the previous chapter, the wall position was inferred by extrapolation, while here the derivative at the wall was calculated by a similar curve-fitting procedure. In the current determination of slope, the inferred wall position was included among the points used for the curve fit. While the two procedures - finding the wall and finding the slope at the wall - may seem identical, the former was used on a selection of profiles, the wall value from each being combined and the mean wall value taken as the inferred wall position. The velocity points used to determine the slope at the wall thus contained a wall point determined from the average extrapolation for all the velocity profiles. Using this mean wall position was both physically correct (the wall did not move during a run) and acted to eliminate any idiosyncrasies at particular wall stations. For instance, if a velocity point near the wall showed a deviation (due to the flow variability), a curve fit through the data would be drastically different to one through data showing no such variability. The addition of the wall point supplied 'extra' global information about the most likely positioning of the wall, reducing fluctuations in the derivative at the wall caused by variability of the velocity traces.

While including the wall point in the data reduced variability, any error in the position of this point would introduce highly systematic errors into the measured slope at the wall. For instance if the wall had been determined to be displaced further outwards than it actually was, all the wall derivative values would be lower than they should be. For small errors in wall positioning, the error in the slope would be approximated by a constant multiplied by the square of the slope, that is, of the form

$$\left. \frac{du}{dr} \right|_{r=R} = \left[\frac{du}{dr} \right]_{r=R}^t + \left[\frac{du}{dr} \right]_{r=R}^2 \delta_{\text{err}} \quad (8.4)$$

Here the primed derivative represents the measured slope and the derivative on the left-hand side of the equation the actual slope. Section A.3 of the appendix derives this error estimate.

The analysis in this appendix showed that systematic deviations of approximately $\Delta S \approx 1$ should be expected in the data. Precisely this level of deviation was found, fortunately at only some measuring stations. Figure 8.14 below shows shape parameter data for the test FXPA, which showed discrepancies between all three curves.

For the early, parallel portion of the flow, it would be expected that the variation of shape with time would be identical for all three curves. As can be seen, the three curves are certainly *not* concurrent at this end; in fact curve (3) - from the furthestmost downstream station - cuts across the other two.




Figure 8.13: The power spectrum for the centreline flow at the inlet, for the test FXPA. Data re-sampled at 100 Hz.

bances in the range of flow receptivity is thus low. The record length used to obtain the spectrum above comprised the entire experimental data. Although this data was normalised to make it a stationary random process, some doubt must be expressed about the validity of extracting the spectrum from the entire data, which physically was *not* stationary. Unfortunately the use of shorter record lengths introduced variability into the spectra, due to the low number of points. A comparison of sub-spectra from a single record did, however, indicate that not much change in the spectral distribution occurred during the test. Thus as a crude measure of the frequency distribution of the flow, the data shown in figure 8.13 is approximately valid.

8.5 Further data processing

The means of extracting the displacement thickness and velocity profile shape parameter from the experimental data were described in section 7.6. The extraction of the displacement thickness yielded accurate results, but the determination of the shape parameter was subject to unforeseen errors. This section describes the further data processing required to correct the measured shape parameter values. The presentation of these extracted parameters for all the tests are presented in subsequent sections of this chapter.

8.4.3 Power spectral density

The RMS disturbance level calculated above represented a gross *integral* estimation of the disturbance magnitude. Of interest too was the frequency distribution of these disturbances. Obtaining a frequency spectrum from LDV data is problematic; because of the random time base of this data, a direct Fourier transform method can not be applied.

Many approaches have been devised to overcome this difficulty of the LDV data. The early correlation-based slotting technique of Gaster & Roberts (1975) yielded a non-aliased, exact result for the spectrum of irregularly spaced data. This technique required a very large population of data to obtain accurate results, and involved lengthy computational times. An improvement by Gaster & Roberts (1977) reduced the computational time by use of a direct transform, but this method required an even longer record length. Improvements on this technique by Roberts & Gaster (1980) and by Ajmani & Roberts (1990) involved digital filtering (band-pass and band-stop filtering). These methods reduced substantially the high variability associated with the use of short record lengths of the former techniques. Buchave *et al* (1990) contains a brief survey of these and other techniques. McDonnell & Fitzpatrick (1992) presented a study in which the above, exact, techniques were compared to an approximate one using data re-sampling and a FFT method. These authors concluded that, for high data rates the spectra deduced by both techniques were comparable. Biasing errors in the FFT data were predictable and depended on the data rate. Furthermore, the FFT spectrum showed less variability than the auto-correlation techniques for similar record lengths.

The high data rate and short record length of the present data, coupled with the need to extract only an approximate spectrum, indicated the use of the FFT technique. The irregular time base of the LDV data was made regular by employing a simple sample-and-hold procedure. Here the last value of velocity preceding the chosen sampling point was taken as the value of velocity at that point. The data, initially at an average sampling rate of between 200 and 300 Hz, was sub-sampled to a regular time interval at 100 Hz. The power spectral density of the data was obtained by multiplication of the Fourier transform obtained from this re-sampled data by its complex conjugate. The spectrum was obtained for the entire record length for each test. Spectra from the various tests were very similar in structure, and a typical power spectrum - that for test FXPA - is given below in figure 8.13.

This spectrum shows some large low-frequency waves - below 3 Hz. The position of the individual peaks was not repeatable between tests, although strong low frequency waves were present in all tests. Another region of prominent spectral activity was the range $5 < f < 20$ Hz: the data in all the tests showed a slight increase in amplitude here over the higher frequency components - those above 20 - 30 Hz. The high frequency range showed a small and uniform magnitude. Considering the discussion in section 8.3.3 above, the magnitude of the distur-




Figure 8.12: The variation in disturbance intensity with rate of exponential acceleration, for the flow at centreline near the pipe inlet.

test FXPA. Considering the remaining data, rectification of the photodetector problem resulted in a signal with very low instrument induced noise. As a result the value of the mean RMS component for the remaining tests was taken as an accurate indicator of the actual perturbations present. The resulting variation of the RMS value with acceleration given by figure 8.11 shows a clear increase in the magnitude of the mean RMS value with increasing values of exponential acceleration, from a value just above 1% at the lowest acceleration test, up to a maximum value of 1.7% for the largest acceleration.

Except for the two points in the mid-range of the figure, the data correlates approximately to a straight line. This trend occurs despite factors such as the irregular control and subsequent large scale waviness of the test HXPA (the left-most point on the figure). The observed trend indicates that the level of background disturbance at the inlet is unaffected by valve-induced oscillations in the flow, and is merely dependent on the rate of exponential acceleration. The three right-most points in figure 8.11 represent the data for EXPG, FXPC and GXPA. Despite the difference in the initial impulsive start-up, these three data show an extremely close correlation. The basic level of disturbance in the tests was below 1.7%, and thus within the region of secondary instability [Herbert (1988)]. This region is characterised by the non-linear interaction of the linear disturbance modes. Herbert named the range of disturbance amplitudes $0.5\% \leq A \leq 2\%$ the range of *medium* amplitude. For this range, disturbances are three-dimensional, but have a wavelength associated with the linear mode. The range of amplitudes, although high, is low enough not to be associated with by-pass transition. As a result, one would expect the instability mechanism to be non-linear and thus sub-critical. A detailed discussion of the flow breakdown, with specific reference to the level of background disturbances, is deferred until part IV, at which point numerical stability results are available.




Figure 8.11: The variation in disturbance intensity with time at centreline near the pipe inlet, for the test FXPJ.

The data shown in this figure is similar to the variation in time of the RMS value for all the tests. The data in figure 8.11 shows no distinct trend with time, but rather maintains an approximately constant value, in this case about 0.012. A scatter of about 0.002 is apparent - approximately 12% of the range. The Reynolds number in this test was of course increasing exponentially; thus the constant value of the normalised RMS term meant that the absolute magnitude of the disturbance was increasing in direct proportion to the Reynolds number. From the point of view of the normalised (mean flow equal to unity) velocity profiles, the disturbance environment was of course constant. The lack of any significant trend with time was seen in all the data, but each data showed an average value of the RMS velocity which depended on the rate of exponential acceleration for that particular test. For this reason it was decided to compare the variation of the mean RMS velocity with increasing exponential acceleration for all the tests. This variation is given in figure 8.12 below.

The value taken for the rate of exponential acceleration was the constant k set in the experimental run-file. While the actual rate of acceleration differed slightly from this factor, this variation was small⁶. The first two experiments of the 'E' series of tests (E1 and E2) utilised an inferior optical set-up, in particular there was a problem with the photodetector. As a result, broadening due to this inadequacy resulted in a high apparent magnitude of the fluctuating component in these tests relative to the others. Due to the fact that this signal broadening was spurious, the results for the mean RMS component for these tests was disregarded⁷. There was also no suitable centreline inlet data for

⁶A more accurate estimation of the rate of acceleration is used in this regard later in part IV.

⁷In spite of this, these early data yielded satisfactory results for the other aspects of the tests (laminar development and flow breakdown), that is, their mean and low-frequency variation was unaffected by the broadening.

8.4.1 Normalisation of the traces

Due to the highly exponential nature of the flows, the centreline variation of velocity with time showed a strong increase. Of primary interest was the level of background disturbances *relative to the mean velocity*. In order to eliminate the strong scaling effect, the trace data was divided through by the mean flowrate variation. In order to further obtain only the fluctuating component of the velocity, a polynomial function was fitted to the resultant data and the mean value of this polynomial was then subtracted from the data. The result of this process was a distribution of data points about a mean of zero. The processed data was simply the normalised fluctuating component of the velocity, denoted as u' . The processing outlined above simply reduced a non-stationary random process to a stationary random one.

8.4.2 Disturbance intensity

The disturbance intensity is defined [Schlichting (1979)] as the ratio of the RMS fluctuating component to the mean flow rate, or

$$u'_{\text{rms}} = \frac{\sqrt{u'^2}}{U_m}. \quad (8.2)$$

Here u' is the fluctuating component of velocity, and U_m is the mean cross-sectional velocity in the pipe. In the present case the data was already normalised with respect to the mean cross-sectional velocity, and division by U_m was thus unnecessary. For the discrete LDV data (t_i, u'_i) , $i \in [0, N-1]$, equation (8.2) reduced to a summation form, given by

$$u'_{\text{rms}} = \sqrt{\sum_{i=0}^{N-2} \frac{1}{2} [(u'_i)^2 + (u'_{i+1})^2] (t_{i+1} - t_i)}. \quad (8.3)$$

The summation term in equation (8.3) is simply the trapezoidal rule integration of the square of the LDV data points. The use of the more general form of this integration - incorporating the time increment $(t_{i+1} - t_i)$ - was necessary because of the random nature of the time base in LDV data. The RMS component of the velocity was calculated for the normalised data for all the traces. This was done for a 5 second window, which was moved through the data in increments of 2.5 seconds. The result of this process was a series of RMS values at increments of 2.5 seconds. With an average sampling rate of between 200 and 300 Hz, each window contained approximately 100 to 150 data points. These data for the test FXPID⁵ is shown in figure 8.11 below.

⁵The naming convention for the various tests is given in table 7.1.




Figure 8.30: The approximate variation of the final parallel flow shape parameter with acceleration factor k : $\circ$, 'E' series of tests; $\square$, 'F' series of tests; $\triangle$, 'G' series of tests; —, linear regression line; shaded region, region of uncertainty.

was somewhat arbitrary, and some results were bound to be inaccurate. This, combined with the crude estimation of the limiting value S_∞ , possibly accounted for the spread of data.

8.8 Summary

In this chapter, the final data pertaining to the variation of the velocity field in time for all twelve tests were presented. Velocity profiles were given, showing a qualitative view of the entire velocity field development within the pipe inlet: the impulsive start, the subsequent parallel developing flow, showing the thinning behaviour, and finally the transition to turbulence in some cases. A successful quantification of the level of background disturbances was also carried out. The displacement thickness and shape parameter showed themselves to be good parametric representations of the velocity variation in the pipe. The lack of certainty in the wall position resulted in a large error in the shape parameter. However, this was highly systematic, and data from each station could be formally corrected relative to the data from other measuring stations from the same test. However, a large uncertainty remained in the absolute value of S . The displacement thickness was not subject to the same errors, and the values of δ^* were thus accurate reflections of the true variation of S parameter. The section was not concerned with the timing and description of the transition itself; rather the variation of the undisturbed base-flow prior to growth of perturbations. Detailed description of this base-flow environment sets the scene for the following chapter, in which the growth of disturbances within the pipe entrance is considered.

Table 8.5: The axial stations, approximate acceleration constants and impulsive start-up Reynolds number for the 'G' series of tests.

Test name	Dimensionless axial station (x/R)	Mean acceleration constant k	Impulsive Reynolds number Re_0
GAG	11.51	0.05	6510
GAA	41.40		
GBH	11.74		
GBI	41.91		
GBJ	80.79		
GCE	11.49	0.04	6510
GCF	24.69		
GCG	41.28		
GDF	67.01	0.03	6510
GDG	30.91		

finally for GXPB its value is about $S \approx 21$. In the same manner as was done earlier, the velocity profile shape parameter was corrected for the uncertainty in the wall position, and the adjusting factors are listed in table 8.6 below.

Table 8.6: The adjustments for the shape parameter for the 'G' series of tests.

Test name	δ_{corr}	Test name	δ_{corr}
GAG	-0.0035	GCE	0
GAA	0.0035	GCF	0
GBH	-0.004	GCG	0
GBI	0.007	GDF	0
GBJ	0	GDG	0

Correlation of all three sets of tests so far indicated that the initial Reynolds number was not a strong parameter affecting the final parallel shape parameter. Although highly approximate, these final values of S indicated a trend that seemed to be dependent on the acceleration constant k alone. Figure 8.30 presents the approximate final values for all the tests (as extracted visually and by inference throughout the above section) plotted as a function of the acceleration factor k . The data from the 'E', 'F' and 'G' series of tests are denoted separately, and a best-fit line is shown superimposed.

It should be recalled that the extraction of the final values of S shown on the figure was by extremely crude means; nevertheless, a clear trend toward higher final parallel shape parameters at higher accelerations was evident. While the lower end of the graph showed fairly good correlation between S_k for the different runs, the data from the maximum acceleration case showed more scatter. Although this indicated some dependence of the final shape parameter on the initial Reynolds number, this was not so; the position of the three highest points did not correlate with Reynolds number at all but seemed to be randomly scattered. The cause of the discrepancy at the high end was quite probably misalignment of the shape parameter data for some of the tests; the adjustment process




Figure 8.28: The variation in experimental shape parameter and displacement thickness, as a function of time, for the test (GXPC) ($k = 0.01$, $Re_\theta = 6510$); $\circ$, GCE ($x/R = 11.49$); $\triangle$, GCF ($x/R = 24.69$); $\times$, GCG ($x/R = 41.28$).




Figure 8.29: The variation in experimental shape parameter and displacement thickness, as a function of time, for the test (GXPD) ($k = 0.03$, $Re_\theta = 6510$); $\circ$, GDF ($x/R = 67.01$); $\times$, HDG ($x/R = 30.91$).

The 'G' series of tests were those with the highest initial impulsive Reynolds number, of 8800. Each of these four tests had respectively the same four values of the acceleration factor as had the 'F' test series. The tests in order of increasing acceleration were: GXPD, GXPC, GXPA, and then GXPB. The axial positions and various other parameters concerning these tests are given in table 8.5.

As before the emergent pattern is analogous to the previous two sets of tests. The inferred final value of the shape parameter for the parallel portion of the flow again increases with increasing acceleration factor. For test GXPD the value is $S \approx 10$ (the lower values after about $t = 15$ s are due to transitional behaviour), for GXPC it is roughly $S \approx 13$, for GXPA S was somewhere below 19, and




Figure 8.26: The variation in experimental shape parameter and displacement thickness, as a function of time, for the test GXPA ($k = 0.05$, $Re_0 = 6510$): $\circ$, GAG ($x/R = 11.51$), $+$, GAA ($x/R = 44.40$).




Figure 8.27: The variation in experimental shape parameter and displacement thickness, as a function of time, for the test GXPB ($k = 0.07$, $Re_0 = 6510$): $\circ$, GBI ($x/R = 11.74$), $+$, GBI ($x/R = 41.91$), $\times$, GBI ($x/R = 80.79$).

eration tests reached steady state in their parallel developing phase far quicker than the lower acceleration ones. Again the lack of sufficiently far downstream stations limited extraction of quantitative data for the development time of the flow.

The first two tests - FXPA and FXPB - showed typical behaviour in both shape parameter and displacement thickness variation in time. The final two tests - FXPC¹² and FXPD - unfortunately each consisted of only two sets of measurements, and both were missing their most upstream data¹². Because of the relative axial proximity of the two stations to each other in both cases, discrimination between the data was poor. In the test FXPC in particular, the measured shape parameter demonstrated peculiar behaviour. While the displacement thickness data in figure 8.24 for station PC3 departed from the parallel development curve at about 7.5 seconds, the shape parameter showed atypical behaviour in that it remained virtually concurrent with, but displaced slightly from, data from station PC4 until after 10 seconds.

Closer examination of the velocity profiles at station PC3 showed that velocity traces near the wall were strongly affected by the inlet variability described earlier, and that the shape parameter for the boundary layer thinning region was thus subject to great error. For this reason the shape parameter results were thus considered inaccurate for this station. The displacement thickness measurements, largely unaffected by wall variability, remained an accurate description of the flow variation in all cases. Table 8.4 lists the correction factors applied to the shape parameter data, in accordance with the procedure of section 8.5.1.

Table 8.4: The adjustments for the shape parameter for the ‘F’ series of tests.

Test name	δ_{err}	Test name	δ_{err}
FA0	0.005	FB5	0
FA1	0	PC3	0.0024
FA3	-0.005	PC4	-0.0024
FB3	0	FD2	-0.001
FB4	0	FD3	0.001

8.7.3 The tests GXPA, GXPB, GXPC and GXPD

The figures 8.26, 8.27, 8.28 and 8.29 below represent the variation of both the measured displacement thickness and velocity profile shape parameter in time, for the ‘G’ series of tests.

¹²Tests were conducted far upstream for both tests, but both sets of acquired data were corrupted in the file storage process, as were their backups.




Figure 8.25: The variation in experimental shape parameter and displacement thickness, as a function of time, for the test FXPD ($k = 0.03$, $Re_D = 4143$); $\circ$, FD2 ($x/R = 79.51$); $+$, FD3 ($x/R = 36.31$).

The ‘F’ series of tests, encompassing the tests FXPA, FXPB, FXPC and FXPD, concentrated more on readings closer to the pipe inlet. The impulsive start was to a higher Reynolds number than were the ‘E’ tests before. Table 8.3 gives the axial position, acceleration value and impulsive Reynolds number of these tests.

Table 8.3: The axial stations, approximate acceleration constants and impulsive start-up Reynolds number for the ‘F’ series of tests.

Test name	Dimensionless axial station (x/R)	Mean acceleration constant k	Impulsive Reynolds number Re_D
FA0	25.52	0.01	4143
FA1	13.07		
FA3	67.81		
FB3	13.69	0.05	4143
FB4	44.61		
FB5	90.54		
FC3	38.38	0.07	4143
FC4	79.96		
FD2	79.51	0.03	4143
FD3	36.31		

The tests ranked in order of acceleration - lowest to highest - as FXPD, FXPA, FXPB and then FXPC. Due to the fact that most of the readings were close to the pipe inlet plane, limiting asymptotic behaviour for large times of the parallel portion of the flow was difficult to infer. The limiting shapes in order of increasing acceleration were approximately $S \approx 10$ for FXPD, $S \approx 12$ for FXPA, $S \approx 15$ for FXPB, and then $S \approx 17$ for FXPC. The displacement thickness did not develop sufficiently in any of the tests to infer limiting behaviour. The trend in development time was similar to that seen before, namely the higher accel-




Figure 8.23: The variation in experimental shape parameter and displacement thickness, as a function of time, for the test FXIPB ($k = 0.05$, $Re_\eta = 41.43$): $\circ$, FB3 ($x/R = 13.69$); $+$, FB4 ($x/R = 44.61$); $\times$, FB5 ($x/R = 90.54$).




Figure 8.24: The variation in experimental shape parameter and displacement thickness, as a function of time, for the test FXIPC ($k = 0.07$, $Re_\eta = 41.43$): $\circ$, FC4 ($x/R = 38.38$); $+$, FC4 ($x/R = 79.96$).

maintaining approximately the same mean. Because of this, the *absolute* uncertainty of the data was still the same as that for the individual measurements, namely $\delta S \approx 4$. Such large uncertainty was unavoidable in the current context, because of the exclusive use of the LDV as a flow measurement tool.

8.7.2 The tests FXPA, FXPB, FXPC and FXPD

The figures 8.22, 8.23, 8.24 and 8.25 below represent the variation of both the measured displacement thickness and velocity profile shape parameter in time, for the 'P' series of tests. As before, the approximate exponential acceleration constants k and Re_0 are indicated on each figure.




Figure 8.22: The variation in experimental shape parameter and displacement thickness, as a function of time, for the test FXPA ($k = 0.04$, $Re_0 = 41.43$): $\circ$, FA0 ($x/R = 25.52$); $+$, FA1 ($x/R = 13.07$); $\times$, FA3 ($x/R = 67.84$).

this parameter. However it was likely that the final value for δ^* would have been above 0.18, had a sufficiently far downstream measuring station existed¹¹.

For the test IIXPA the development in the displacement thickness within the thinning region was more distinct than in the equivalent shape parameter data. The two most upstream stations - IIAJ and IIAK - were clearly following separate trends, in agreement with data from the other tests. The data for the other three stations were, however, superimposed to a great degree, and no clear trends were evident. Clearly the variability in the flow was greater than the difference between the stations, thus swamping any variation between data. The presence of the inlet 'blip' was also marked in the data IIAJ: a distinct decrease in the shape parameter at about 15 seconds marked the passing of this event. In both the shape parameter and displacement thickness data, the thinning region showed a common type of behaviour. The further downstream stations showed a steeper rate of decrease in displacement thickness in time than did those near the inlet. This trend was also evident in the shape parameter data, but these were more inclined to be disturbed in this trend by the transition phenomenon than was the case with the displacement thickness.

The correction of the shape parameter curves

Some of the shape parameter data presented above and below had been corrected as discussed in detail in section 8.5.1. The correction factors were applied to some of the S data, in order to make the parallel portion of the data concurrent. Table 8.2 lists the correction factors δ_{err} applied to each data set, according to the correction equation (8.4).

Table 8.2: The adjustments for the shape parameter for the 'E' series of tests.

Test name	δ_{err}	Test name	δ_{err}
EE1	-0.004	EG1	-0.003
EE2	-0.007	EG2	0.001
EE3	0	EG3	-0.002
EE5	0.005	IIAG	0
EF1	0.001	IIAH	0
EF2	0	IIAI	0
EF3	-0.002	IIAJ	-0.004
EG0	0.001	IIAK	0

Due to the lack of means of independent measurement of the wall shear stress, and thus 'calibration' of the shape variation curves, these curves were simply adjusted to be self consistent within a test. For instance the four sets of data EE1, 2, 3 and 5 had adjustments applied to their shape parameter data such that the parallel portion of each curve became mutually concurrent, while still

¹¹And, of course, had the rig had enough capacity in the header tank to maintain such a flow for long enough.

parallel flow has as its final limit the parabolic velocity profile. This final behaviour of the shape parameter during the parallel development showed distinct trends with increasing exponential acceleration. In test EXPE - the second lowest acceleration case - the asymptotic value was around $S \approx 11$, for test EXPF somewhere below¹⁰ $S \approx 20$, while for the highest acceleration test - EXPG - the limiting value was about $S \approx 26$. For the test HXPA the parallel development was cut short - at about 17 seconds, by the arrival of the entrance flow at station HAG, but the trend in the shape parameter seemed to be towards a value below 10.

The shape parameter of the Hagen-Poiseuille parabolic velocity profile is 8, thus all the final limits seemed to be profiles with higher gradient boundary layers than the parabolic flow. The rate at which the particular asymptotic value was reached also seemed to be a function of the acceleration parameter. For EXPE the value was attained somewhere after 20 seconds, whereas in EXPG development had virtually ceased by 15 seconds.

Displacement thickness

The figures showing the displacement thickness revealed identical trends to those shown by S : all data fell initially on a single curve, whose value increased from zero displacement thickness at the impulsive start towards some asymptotic value for large times. Results for successive stations left this common curve at successive times, and reverted to a behaviour where δ^*/R reduced with time. Test EXPE had the largest range of axial stations, and the displacement thickness revealed itself to be a far more sensitive means of differentiating between axial stations than the shape parameter. For instance, where the differences between EE3 and EE5 in the S versus t framework were slight, the equivalent δ^*/R figure showed large discrimination between these data. Furthermore, whereas the S vs. t curve for EE5 had appeared to have levelled off to a fixed value, the displacement thickness equivalent showed that this station had not yet achieved equilibrium.

The data revealed the same trends concerning asymptotic boundary layer shape and rate of development as did the shape parameter. In figure 8.20 - the highest acceleration test - the displacement thickness was virtually asymptotic to $\delta^*/R \approx 0.1$ by time 20 seconds, whereas in EXPE the parallel displacement thickness curve was still increasing at time 30 seconds (the time when entrance effects were manifest at the most downstream station) to a value well above $\delta^*/R \approx 0.18$. For HXPA the displacement thickness was developing steeply with a value of about 0.15 when the entrance flow reversed the development trend. Due to this premature curtailment of the parallel development in this particular case, it was impossible to infer even an approximate final value for

¹⁰The station EF3 was not sufficiently far downstream to reveal proper asymptotic behaviour, but the downward trend in S seemed to be levelling off at about $S \approx 18$ before the arrival of the entrance flow disrupted the development.

Table 8.1: The axial stations, approximate acceleration constants and impulsive start-up Reynolds number for the 'E' series of tests. Included is the data IIXPA.

Test name	Dimensionless axial station (x/R)	Mean acceleration constant k	Impulsive Reynolds number Re_0
EE1	12.24	0.03	3018
EE2	20.25		
EE3	64.52		
EE5	320.50		
EF1	62.5	0.05	3018
EF2	40.04		
EF3	13.49		
EG0	13.49	0.07	3018
EG1	64.52		
EG2	105.03		
EG3	84.81		
IIXG	113.57	0.02	2959
IIXH	81.62		
IIXI	54.56		
IIXJ	32.99		
IIXK	11.41		

Shape parameter

First considering the variation of the shape parameter with time, the first three tests showed typical behaviour. The station furthestmost upstream quickly departed from the parallel flow trend of decreasing S with time, reverting to a steady increase in value. At progressively further downstream stations, this departure occurred at a later time, and thus the shape parameter departed from the parallel curve at a lower value further downstream. At the furthestmost downstream station (EE5 and EG2 in particular - the EF1 station was not that far downstream) the shape parameter seemed to be asymptotic to a fixed value as time progressed, before it too increased due to the arrival of the inlet.

The test with the lowest acceleration - IIXPA - was subject to large inlet effects for the predominance of the entrance flow. The parallel development of the shape parameter in figure 8.21 proceeded as with the others, data from successively downstream stations 'peeling off' from this parallel curve. However the subsequent data was subject to large variability, with the shape parameter and displacement thickness curves crossing one another repeatedly. The S curves especially showed no clear pattern, although on average curves from further upstream maintained higher values of S than those from downstream. This variability as mentioned was innate to the flow system and was not due to mis-measurement of the shape parameter. The behaviour of the most downstream curves in each test seemed to indicate that the parallel flow development was proceeding towards a final, asymptotic, velocity shape, much like steady




Figure 8.20: The variation in experimental shape parameter and displacement thickness, as a function of time, for the test EXPG ($k = 0.07$, $Re_0 = 3018$): $\circ$, EG0 ($x/R = 13.49$); $+$, EG1 ($x/R = 64.52$); $\times$, EG2 ($x/R = 105.06$); $*$, EG3 ($x/R = 84.81$).




Figure 8.21: The variation in experimental shape parameter and displacement thickness, as a function of time, for the test IIXPA ($k = 0.02$, $Re_0 = 2059$): $\circ$, IIA0 ($x/R = 113.57$); $+$, IIA1 ($x/R = 81.62$); $\times$, IIA2 ($x/R = 54.56$); $*$, IIA3 ($x/R = 32.99$); $\cdot$, IIA4 ($x/R = 11.41$).

The four tests all possessed similar start-up Reynolds number, and the value of the exponential acceleration constant increased with successive tests in the order IIXPA, EXPE, EXPF and EXPG. The axial positions, approximate acceleration factors and impulsive start-up Reynolds numbers of the various stations are shown in table 8.1. Here k and Re_0 are the constants defining the acceleration, as given in equation 7.1.

Part III

ANALYSIS

the lower test IXPA, all data were above this critical value. Because of the lack of theoretical understanding of the accelerating flow stability behaviour, detailed discussion of this important and interesting topic must be deferred until the Synthesis, at which point theoretical results will be available.

9.5 Summary

The experimental inception data was presented in this chapter. The lack of a suitable framework for comparisons between inception data from different tests, combined with the lack of theoretical stability data for the exponentially accelerating flows, did not allow any profound comparisons to be made or conclusions to be drawn.

The minimum critical Reynolds number was the only universal means of comparison between the results, and its variation with increasing acceleration factor showed clearly that the higher acceleration tests possessed a larger minimum critical Reynolds number. The downstream portion of the inception curves were inferred to be predominantly caused by 'wash-down' effects; by-pass type instabilities developing here due to the convection into this region of large disturbances developing further upstream. The timing results of section 9.3 supported this conclusion. Some evidence was seen of the leeward side of a stability curve, but the results were too few to draw firm conclusions. Due to of this, the definition of the curve of *inception* was used as the only reliable criterion for quantifying the flow breakdown: it was given by the locus of points within the entrance region at which the first occurrence of flow breakdown (of whatever nature) was observed.

Description of the accelerating flows studied here can proceed no further without a sound theoretical base from which to compare them; no relevant study exists in the literature concerning exponential flows. The following part of this work - part III - is dedicated to the theoretical and numerical analysis of these flows, in order to provide a proper grounding for comparing and verifying the data. This task of further analysis and comparison is thus deferred until the synthesis of part IV.

in the $Re - x$ plane. Its value in the data given previously showed an increasing trend with increasing rates of exponential acceleration. Figure 9.15 below shows the variation of this parameter with acceleration factor k .




Figure 9.15: The variation of the experimentally measured minimum critical Reynolds number with acceleration factor k : $\bullet$, measured data; $-$, regression curve; $\circ$, corrected value for EXPG; $\square$, theoretical linear minimum critical Reynolds number for steady pipe entrance flow [Abbot & Moss (1994)].

The values for the minimum critical Reynolds number were extracted simply by taking the lowest data point from the figures 9.1 to 9.12 above. The value of k used was the prescribed value, which was not quite the actual value realised in the flow⁷. A strong correlation is evident between the acceleration factor and minimum critical Reynolds number. All data were virtually collinear, with a low uncertainty, as shown by the shaded region in figure 9.15. The clear relationship between acceleration and critical Reynolds number were evident here; higher accelerations resulted in increased values of the critical parameter. Thus the data justified the commonly held belief [see Shen (1961), Drazin & Reid (1981)] that acceleration stabilises a flow.

The only poor correlation to the trend was shown by the point for the test EXPG - one of the highest acceleration tests, with $k = 0.07$. The most likely reason for this datum being out was that the lowest data point in Reynolds number was spurious. Observation of figure 9.3 showed that the third station from the entrance had an unusually low critical Reynolds number. Removal of this point made the data in that figure far more regular. Incorporation of the next lowest point $Re_c = 15630$ in figure 9.15 reduced the variability at this end substantially. The corrected datum is shown as an open circle $\circ$ in this figure. An interesting observation is the close alignment between the regression data extrapolated to zero acceleration, and the theoretical minimum critical Reynolds number for zero acceleration flow of 11 667 [Abbot & Moss (1994)]. Excepting,

⁷The variation of the flowrate was not quite exponential, because of peculiarities in the valve, control algorithm and flowmeter. In a later part of the work - the synthesis - a more exact parameter will be used.




Figure 9.14: The inferred shape of the stability curve and wash-down region in the pipe entrance: (a), region of linear or small amplitude instability; (e), convection and growth of disturbances; (b), 'front' at which turbulence occurs; (c), region destabilised by large disturbances from the inlet; (d), relatively undisturbed region of flow seen experimentally.

9.3.2 Stability versus inception

The above interpretation of the data, by its simplicity, ignores some confounding factors, namely that the downstream side of the hypothetical stability curve exists within a region subjected to variable levels of finite-amplitude disturbance. Thus in reality the hope of identifying the leeward portion of a curve corresponding to the stability limit for a *fixed* level of disturbance is remote. In contrast, the upstream portion of the stability curve is quite probably the stability limit of a fixed finite-amplitude level of disturbance - the level of the free-stream environment as determined in section 8.4.

Resulting from this, a pragmatic solution is to simply denote the first occurrence on *any* phenomenon at a particular station the point of inception. This avoids the difficulties of identifying stability curves within a variable environment and leads to an unambiguous definition of the region of flow breakdown. It is not that the individual aspects of transition are unimportant, rather that their identification is difficult in the current context.

9.4 The minimum critical Reynolds number

While the shape and position of the stability curves do not allow for mutual comparisons⁶, the variation of their minimum critical Reynolds number with acceleration was directly comparable. The minimum critical Reynolds number is defined as the Reynolds number of the minimum point of the inception curve

⁶In part III of the study that follows - the analysis - a suitable framework will be introduced in which all these data can be compared.

that this flow breakdown was primarily due to convection downstream at the mean velocity, of large disturbances originating further upstream, and their subsequent destabilising of the flow. This conclusion was corroborated by the nature of the flow breakdown further downstream which showed that progressing downstream, the breakdown phenomenon became more like a hard (convected) turbulent front.

9.3.1 The leeward side of the stability curve

Unfortunately, the presence of the large convected disturbances from the upstream instability reduced any possibility of observing the true shape of the stability curve on the leeward or downstream side of the minimum critical point. Any disturbances originating at this minimum point were likely to convect and grow downstream, triggering rapid by-pass type transition, and thus smothering any incipient growth of small disturbances occurring there⁵. It was not possible to tell whether the downstream disability was due to local breakdown or caused by destabilisation due to convection of large perturbations from upstream. Quite possibly it was a combination of both.

There was a hint of the presence of a steep leeward side of the curve in some tests. This was especially so in figures 9.5 and 9.6. Here shortly after the minimum critical point the data increased in Reynolds number rapidly, while further downstream the values of Re decreased again to come to rest approximately on the convection line. In terms of the single standard deviation error bars shown in these data, these trends were significant. The supposed observation of the leeward part of the stability curve was due possibly to both the innate variability of the flow breakdown and the nature of the convection process. Firstly, concerning the variability, the data was assembled from many distinct runs. If a measurement downstream of the minimum critical point was made when flow breakdown at the critical point was ‘late’, then local transition would be seen at that point. Secondly, and possibly more importantly, slow growth of the initially small disturbances would have left a band of relatively unperturbed flow a short distance downstream before growing large enough to trigger the by-pass mechanism. This hypothesis is indicated in figure 9.14 below.

Here, disturbances that convect downstream from region (a) might be small enough and grow relatively slowly (or remain stable) within region (d), thus allowing the appearance of a small portion of the leeward stability curve. By the time these disturbances reach line (b), they are large enough to undergo rapid transition to turbulence.

⁵In a sense both ‘local’ and ‘wash-down’ transition are really local in nature. Both are initiated by the presence of a disturbed environment: in the case of the ‘traditional’ instability, the disturbances are small whereas in the by-pass transition downstream, these destabilising influences are of a large scale. The *growth by destabilisation* of turbulent slugs was discussed by da Silva (1990).




Figure 9.13: The variation of experimental critical Reynolds number with axial position, for the test GXP1A. Lines of mean flow 'convection' are also shown: (a) parallel development region; (b), entrance region; (c), minimum critical point.

where x_i is the axial position at which the fluid originated at time $t = 0$. Equation (9.1) thus establishes the displacement variation of a particle of fluid travelling at the mean velocity with time. As the Reynolds number was a function of time, the variation of the displacement with Reynolds number could then be directly ascertained. The thicker line in the figure represents fluid that was at the pipe inlet at flow start-up. The region above this, denoted (a), contains fluid that, if travelling at the mean velocity, entered the pipe subsequent to start-up; that is, fluid comprising the entrance flow. The region (b) below this line contains flow that originated within the pipe at the start of the test; i.e. possessing values of x_i from equation (9.1) greater than zero. Fluid in this latter region comprises flow from the parallel development part of the flow.

Clearly the point of flow breakdown is well within the entrance region of the flow. The inception points to the left of the figure, on the upstream side of the curve, are clearly representative of local flow break-down. Flow arriving at these points has 'travelled' approximately along the curves shown in the figure¹, thus it has come from undisturbed regions further upstream. The minimum critical point (c) of the experimental data represents the first time within the entire pipe at which instability occurs (time being proportional to Reynolds number), thus its occurrence could not be influenced by disturbances from any other region of the flow. The leading edge of the inception curve described above could conceivably have been destabilised by the influence of the minimum critical point, but this would have required diffusion of the disturbance influence upstream within the boundary layer. In contrast, the downstream part of the inception curve was virtually coincident with a curve of mean velocity convection. This indicated

¹Downstream convection at a greater velocity than the mean cross-sectional value will be represented by a shallower line than the given curves, and slower convection by a steeper line.

9.2 General

The exponential variation of the flow with time made the two forms of data representation - x vs. t and x vs. Re - almost equivalent. This was because the Reynolds number increased monotonically with time. Starting at the pipe inlet and moving downstream, all data showed a decrease in the transition Reynolds number with axial position. A minimum value of Reynolds number was reached; thereafter the Re increased with the axial position. The axial extent of this variation and the position of the minimum point varied considerably between tests and no clear pattern was shown. The points on the graph are associated with the identification of the first time within the flow that a distinct deviation from the laminar state of development occurred. The process of identification was described in detail in section 7.4. The error bars are shown on the $x - t$ plot. They indicate one standard deviation in the variation of the twelve readings taken at each measuring station. In most tests variability at all stations was similar; however in a few cases the variability at particular points was much higher than for others. This was thought to be due to poor positioning of the LDV at those stations. In general these points show a later mean transition value than the surrounding data, which is indicative of improper alignment of the measuring volume with the point of maximum sensitivity to instability³. In most cases an interpolated curve would fit smoothly through all the data within these bars, for instance in figure 9.9; however in a few tests, the data seemed to be defining two curves: one showing a rapid increase in Re with x away from the minimum point, while the other, lower curve, possessed a more gradual slope, and extended further downstream. Figure 9.5 is a good example of such a test. This issue is addressed below.

9.3 Local and convected influences

The general nature of the transition within the pipe entrance is best illustrated by considering a specific test with a well distributed set of points. A good example is the test GXDA and the $Re - x$ data of figure 9.9 is shown again below.

Superimposed on this figure is a set of lines, representing trajectories of particles moving at the mean cross-sectional velocity. The equation for these lines was simply derived: the variation of mean velocity in time for the flow was known experimentally [see section 7.5], thus by integrating this quantity with respect to time, the displacement of fluid at the mean velocity could be calculated as

$$s_1(t) = \int_0^t U(t) dt + x_0 \quad (9.1)$$

³Misalignment would result in later (and more erratic) measurement of the phenomena, because some growth would have to occur before the perturbations could become visible




Figure 9.10: The variation of inception time and Reynolds number with axial position for the test GXPIB ($k = 0.07$, $Re_0 = 6510$).




Figure 9.11: The variation of inception time and Reynolds number with axial position for the test GXPIA ($k = 0.04$, $Re_0 = 6510$).




Figure 9.12: The variation of inception time and Reynolds number with axial position for the test GXPII ($k = 0.03$, $Re_0 = 6510$).




Figure 9.7: The variation of inception time and Reynolds number with axial position for the test FXPC ($k = 0.07$, $Re_0 = 4143$).




Figure 9.8: The variation of inception time and Reynolds number with axial position for the test FXPD ($k = 0.03$, $Re_0 = 4143$).




Figure 9.9: The variation of inception time and Reynolds number with axial position for the test GXPA ($k = 0.05$, $Re_0 = 5510$).




Figure 9.4: The variation of inception time and Reynolds number with axial position for the test IIXPA ($k = 0.02$, $Re_0 \approx 2959$).




Figure 9.5: The variation of inception time and Reynolds number with axial position for the test FXPA ($k = 0.04$, $Re_0 \approx 4143$).




Figure 9.6: The variation of inception time and Reynolds number with axial position for the test FXPB ($k = 0.05$, $Re_0 \approx 4143$).




Figure 9.1: The variation of inception time and Reynolds number with axial position for the test EXPF ($k = 0.03$, $Re_0 = 3018$).




Figure 9.2: The variation of inception time and Reynolds number with axial position for the test EXPF ($k = 0.05$, $Re_0 = 3018$).




Figure 9.3: The variation of inception time and Reynolds number with axial position for the test EXPG ($k = 0.07$, $Re_0 = 3018$).

Chapter 9

Experimental results - inception

In section 7.4 it was described how inception data were gathered from more or less evenly spaced measuring stations throughout the pipe entrance region. On average about ten different axial stations were investigated per test, and twelve identical runs were conducted from within the boundary layer for each of these stations. The radial position of the point of measurement for the twelve traces was at each station the point that showed the greatest amplitude of disturbance. This chapter presents the extracted times of deviation for all the experimental tests. The detailed nature of the flow breakdown and its categorisation were discussed in detail in section 7.4 and will not be elaborated upon further. The data here presented are simply the first points of deviation¹ from the laminar flow at each measuring station.

9.1 Inception times and Reynolds numbers

The inception data are presented here separately for each test. The following figures show the variation of the raw timing data and the critical Reynolds number with axial position. The data from each test are presented separately because the dimensionality of the scales made superimposition meaningless². However comparison is made later between the approximate minimum critical Reynolds numbers (the minimum point of the data in the $x - Re$ plane) to good effect. The data are given in sequential order below as figures 9.1 to 9.12. The approximate acceleration constants k , Re_0 for each test are indicated on the figures, and the error-bars denote one standard deviation.

¹The definition of this term was given in section 7.4. It was defined as the time of first departure of the flow from the simple laminar state, regardless of the nature of the departure.

²The curves intersect in a random fashion in the $x - Re$ plane when superimposed, indicating that this framework is inappropriate for their mutual comparison.




Figure 10.3: The variation of a general function $w(x, y)$ in the range $x \in [0, k]$, $y \in [0, 1]$.

Thus, regardless of the variation⁵ of w in y , the unbounded behaviour is entirely contained in the term w_0 . Figure 10.4 shows clearly the bounded behaviour of the renormalised function w .

Care needs to be taken if w is acted on by some function: for instance if w is differentiated with respect to x , then its expansion in terms of w and w_0 needs to be dealt with using the product law of differentiation. As long as such care is taken, no generality will be lost in the renormalised representation. If w is differentiated with respect to y , then the parameter w_0 is simply a constant, with respect to this variable, and can be treated in the usual way.

The analogous situation exists in generally accelerating pipe flows: here the flow-rate, and thus the velocity, varies in an unbounded manner. If the term $U \equiv U(t)$ is taken to be the instantaneous mean cross-sectional velocity, then the dimensional coefficients u_0 , x_0 and p_0 in (10.20) are not constant with respect to t . Using the principles outlined above, the velocity u , pressure p and axial co-ordinate x can be renormalised, such that, regardless of the variation in magnitude of u , the dimensionless variables u and p will remain of order unity. The terms x_0 and p_0 are not differentiated with respect to t , thus they behave as constants. The statement for the renormalisation of u is

$$u(x, r, t) = U(t) \bar{u}(x, r, t), \quad (10.21)$$

⁵assuming, of course, you are not using the Laplace transform.

⁶Note that w is not written as the more common product $w = X(x)Y(y)$ in which case generality is lost.

A straightforward substitution by inspection for the dimensional coefficients leads to the following form for these terms:

$$\begin{aligned}
 x_0 &= \frac{UR^2}{\nu} \\
 r_0 &= R \\
 t_0 &= \frac{R^2}{\nu} \\
 u_0 &= U \\
 v_0 &= \frac{U}{R} \\
 p_0 &= \rho U^2.
 \end{aligned} \tag{10.20}$$

For an impulsively started pipe flow, substitution of (10.19), (10.20) into the governing equations will recover da Silva's (1990) relations, but in their primitive variable form. However, for the more general case of accelerating flow, further consideration needs to be given to the use of these groups.

10.2.4 Strongly varying flows

In the dimensionless substitutions (10.20), the tacit assumption is that the dimensional terms U , R , ν and ρ are *constant*. The dimensional variables are then simply scaled by these factors to render them of unity magnitude. However a careful look at the terms (10.19) reveals that there is no mathematical reason why the 'constant' terms (x_0 , r_0 , t_0 , u_0 , v_0 , p_0) need be constant at all, as long as their subsequent substitution is treated correctly.

To illustrate that this is so, consider some general function $w(x, y)$ of two variables x and y in the range $x \in [0, k]$, $y \in [0, 1]$. Assume that the function varies in x in such a way that some integral quantity of the variable [for instance $\int_0^1 w(x, y)dy$] varies without bound. Figure 10.3 gives an example of such a function. The direct consideration of w is not helpful, in that the sheer change in magnitude of the integral quantity swamps any other variation that might be interesting. Not. Using the function with respect to some constant (say the integral quantity at some fixed $x = x_0$) will simply scale the entire function by a constant, yielding no improvement in the depiction of w . Of more interest would be to scale w at each point in x by the integral quantity *at that point*. This renormalisation of w will result in some variable $\bar{w}$ which is of order unity everywhere in the range $x \in [0, k]$ regardless of the variation of the original function w . The variable $\bar{w}$ can be written in general as $\bar{w}(x, y) = w_0(x)w(x, y)$, where $w_0(x) \equiv \int_0^1 w(x, y)dy$ is the integral quantity. Important to note is that $\bar{w}(x, y)$ is still a general function of x and y , but its y dependence is bounded,

everywhere. These initial specifications are only specific to this study, but any initial velocity field could be specified and its temporal evolution studied.

10.2.3 Scaling

The primitive Navier-Stokes equations (10.8), (10.9), and (10.10) describe a flow system in terms of dimensional variables. The use of dimensionless groups allows the independent and dependent variables to be scaled in terms of dimensional factors (such as pipe radius, viscosity, etc.), such that the resulting equations are independent of the scale of the system or type of fluid. These dimensional factors can usually be combined into one or more parametric groups, such that geometrically conformal systems with identical values for these parameters are described by a single solution of the resulting equations. The Reynolds number is perhaps the most commonly known such group, and can be understood as the ratio between the inertia and viscous forces acting in a fluid. Two flows at identical Reynolds numbers in geometrically conformal systems will behave identically.

Furthermore, if in this process the dimensionless variables are made of order one, then the magnitude of any term can be restricted to the dimensional coefficient. This allows ready identification of negligible terms in a system of equations, and thus an understanding of the asymptotic behaviour of a system. For instance, for the case of impulsively started pipe entrance flows, the Reynolds number suffices as the only necessary descriptive parameter [i.e. da Silva (1990)]. By using the principal of dimensionless groups, for high values of Re , it is possible to show that the system becomes independent of the Reynolds number. Thus da Silva was able to solve a single initial-value problem to describe all cases of high Re impulsively started pipe inlet flow.

For the present case, both the dependent and independent variables of equations (10.8), (10.9), and (10.10) are subjected to a simple non-dimensional treatment. All terms are represented as a unitary variable of order one multiplied by a dimensional coefficient. In this way, the magnitude of a term is confined to its coefficient. If the terms u_0 , v_0 , p_0 , x_0 , r_0 and t_0 are taken to be the factors of dimensionality, then the variables can be represented as

$$\begin{aligned} x &= x_0 \bar{x} \\ r &= r_0 \bar{r} \\ t &= t_0 \bar{t} \\ u &= u_0 \bar{u} \\ v &= v_0 \bar{v} \\ p &= p_0 \bar{p} \end{aligned} \tag{10.19}$$

The most important condition is the inflow condition: the pipe is modelled as drawing fluid from a still reservoir through a bell-mouth contraction. It is assumed that this contraction is shaped in such a way that the fluid enters the pipe with a uniform velocity across the pipe radius, and with a very thin boundary layer. The inflow velocity profile shape is assumed fixed in time (except perhaps for mass-flow changes). The simplest approximation to the inflow profile is taken as

$$u(0, r, t) = U(t) \quad (10.15)$$

where $U(t)$ has some given variation in t alone, and is the mean flow rate in the pipe. Strictly speaking, (10.15) violates the no-slip condition at $(0, 0, t)$. Its simple form is only intended to be an approximation to the physical inflow condition - where there will persist a very thin boundary layer from the inlet contraction. Thus, strictly speaking, the more general inlet condition is

$$u(0, r, t) = U_0(r, t) = U_0(r)U'(t). \quad (10.16)$$

Here the inflow condition is taken to be a profile U_0 , self-similar in time, scaled by the mean flow rate $U'(t)$. The profile shape U_0 given as some fixed function of r . The mass flow rate $U'(t)$ then effectively becomes a forcing function for the pipe flow. Unsteady pipe flows will be represented by varying $U'(t)$ in a prescribed way in order to model a flow whose flow-rate is changing with t .

The form of the radial inflow velocity has not yet been determined. In a manner similar to (10.16), the r component at $(0, r, t)$ can be specified as $v(0, r, t) = V'(r, t)$. There are two options available here: firstly, the radial component can be fixed at zero. This is equivalent to the flow entering the pipe through a bank of flow straighteners. Secondly, some arbitrary condition commensurate with the inlet contraction can be specified.

The only condition left unspecified above are the initial-conditions. These conditions describe the state of the fluid at the boundary $V_0 = (x, r, 0)$, or in other words the state of the fluid everywhere at $t = 0$. For the present study, the fluid in the pipe is taken as being instantaneously accelerated from rest impulsively to a high Reynolds number, from where its mass flow might or might not change further. Strictly, the temporal domain is the open set $t \in (0, \infty]$. At $t = 0$, the flow is at rest everywhere; i.e. $u(x, r, 0) = v(x, r, 0) = 0$. The flow is then immediately raised to some finite velocity $U'(0)$ - the mean flow rate in the pipe at $t = 0^+$. This is equivalent to

$$u(x, r, 0^+) = U'(0) \quad (10.17)$$

with the radial component fixed at

$$v(x, r, 0^+) = 0 \quad (10.18)$$

$$\frac{\partial u}{\partial t} + v \frac{\partial u}{\partial r} + u \frac{\partial v}{\partial x} = -\frac{1}{\rho} \frac{\partial p}{\partial r} + \nu \left[\frac{\partial}{\partial r} \left(\frac{1}{r} \frac{\partial}{\partial r} (rv) \right) + \frac{\partial^2 v}{\partial x^2} \right] \quad (10.9)$$

$$\frac{1}{r} \frac{\partial}{\partial r} (rv) + \frac{\partial u}{\partial x} = 0 \quad (10.10)$$

The above relations are simply a restatement of equations (10.1) and (10.2) in scalar form, and transformed into a radial co-ordinate system. Thus they are completely general, except in their exclusion of the azimuthal component of velocity.

10.2.2 Boundary conditions

In order for equations (10.8), (10.9) and (10.10) to represent properly a particular flow of interest, relevant boundary and initial conditions need to be applied. These conditions for non-swirling pipe inlet flow are straightforward: firstly, the velocity at the pipe wall $\mathcal{V}_W \equiv (x, R, t)$ must be zero by virtue of the no-slip condition. This results in the two conditions

$$u(x, R, t) = 0, \quad (10.11)$$

$$v(x, R, t) = 0, \quad (10.12)$$

being satisfied at this boundary.

The collapse of the three-dimensional domain to two necessitates the inclusion of a symmetry condition at the centreline $\mathcal{V}_C \equiv (x, 0, t)$, in order that the solution be single-valued and bounded here. In scalar form, this requires the satisfying of the conditions

$$\frac{\partial u}{\partial r}(x, 0, t) = 0, \quad (10.13)$$

$$v(x, 0, t) = 0, \quad (10.14)$$

The above four conditions effectively *constrain* the fluid; the fluid is *driven* by the inflow and outflow conditions in the pipe. Physically, these conditions are most usually imposed as a pressure difference across a pipe (water flowing from a high-pressure reservoir to one at lower pressure). However, in order to study directly the phenomena of acceleration in such flows, the current choice is to impose these conditions as velocity conditions at the inlet of the pipe.




Figure 10.1: The radial co-ordinate system for pipe flow.

A further natural subset of pipe flows are those in which there is no swirl component to the flow. In this case the system exhibits no θ dependence and thus the radial co-ordinate system of figure 10.1 reduces to an axisymmetric system. The solution domain in this case is $0 \leq r \leq R$, $0 \leq x \leq L$, and $0 \leq t < \infty$, where L is the pipe length, and R its internal radius. This co-ordinate system was used in this study, and is indicated in figure 10.2.




Figure 10.2: The simplified co-ordinate system for axisymmetric pipe flow.

The vector notation of the previous section is now dropped. The two governing equations resulting from the elimination of the azimuthal component of velocity, plus the continuity requirement will be given as three scalar equations. The dependent variables are given generally as the axial velocity $u = u(x, r, t)$, the radial velocity $v = v(x, r, t)$ and the pressure $p = p(x, r, t)$, with the resulting simplified equations given by

$$\frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} + v \frac{\partial u}{\partial r} = -\frac{1}{\rho} \frac{\partial p}{\partial x} + \nu \left[\frac{1}{r} \frac{\partial}{\partial r} \left(r \frac{\partial u}{\partial r} \right) + \frac{\partial^2 u}{\partial x^2} \right] \quad (10.8)$$

Direct implicit implementation of the integral vorticity conditions will lead to a causality problem: the cyclical nature in which the equation (10.4) must be solved (evolving vorticity to the next time step and then determining velocity) implies that the vorticity must be evaluated first. There is however no way of implicitly evaluating the vorticity boundary condition without first calculating the new velocity, but the new velocity depends directly on the vorticity wall condition! The effect of this is that any numerical method applying the vorticity condition strictly will be *precluding* the diffusion of any information in from the boundary with time. Thus the only successful method of solving equation (10.4) is by the mathematically incorrect explicit wall vorticity approach.

The boundary values for the streamfunction approach (10.7) are simple and directly imposed conditions. This is because ψ relates to the velocities in a derivative sense [see the relations (10.5)] and thus any velocity end conditions just become derivative streamfunction conditions. The disadvantage of the streamfunction method lies both in the high order of the equation, and in the structure of the temporal term: the system is of fourth order, and the temporal evolution is in terms of $\nabla^2 \psi$, and not ψ directly. These factors usually make this method non-optimum for a finite-difference approach, because of the inaccuracy of representing such high order derivatives.

The above discussion was concerned with general formulations of the full Navier-Stokes relations. The remainder of this chapter will consider approximations to these full equations that are valid for high Reynolds number pipe flows. The discussion will proceed from the primitive-variable Navier-Stokes relations (10.1), and the elimination of the pressure term will be considered later.

10.2 Unsteady pipe flow

The current study is concerned with the flow of an incompressible, Newtonian liquid in the entrance region of a pipe, in the absence of thermal gradients or body forces. Flows of this class possess a symmetry about the axis of the pipe, thus it is natural to represent their spatial domain using a radial co-ordinate system.

10.2.1 Co-ordinate transformation and equations of motion

In a pipe system, the three most natural orthogonal axes are: the axial co-ordinate x oriented along the centre line of the pipe, the radial co-ordinate r projecting outward from the centre line, and the azimuthal co-ordinate θ - an angular displacement around the pipe. Figure 10.1 shows this co-ordinate system.

- (ii) If the domain is restricted to be two-dimensional, then the streamfunction can be used. Consider the planar velocity field $\mathbf{u} \equiv (u, v, 0)$ in the plane $\mathbf{x} \equiv (x, y, 0)$. In this plane, the vorticity is defined as $\boldsymbol{\omega} \equiv (0, 0, \omega)$. The streamfunction ψ is related to the velocity $\mathbf{u}$ as

$$u \equiv \frac{\partial \psi}{\partial y}, \quad v \equiv -\frac{\partial \psi}{\partial x}, \quad (10.5)$$

and to the vorticity ω as

$$\nabla^2 \psi \equiv -\omega. \quad (10.6)$$

By reformulating the equation (10.1) in terms of the streamfunction (10.5), a single equation in the scalar potential function ψ alone can be derived. This streamfunction Navier-stokes equation is

$$\frac{\partial}{\partial t} \nabla^2 \psi + \frac{\partial \psi}{\partial y} \frac{\partial}{\partial x} \nabla^2 \psi - \frac{\partial \psi}{\partial x} \frac{\partial}{\partial y} \nabla^2 \psi = \nu \nabla^4 \psi. \quad (10.7)$$

This equation is sufficient in itself to describe a divergence-free velocity field, thus no continuity constraint is needed. Thus this equation leads to a very simple means of describing a two-dimensional flow system. Extension of the approach to a three-dimensional system is difficult, and a second streamfunction must be used. The relevant equations for a three-dimensional formulation can be found in Murdock (1986).

The above vorticity-transport formulation (10.1) is an elegant way of removing pressure. The price of this convenience is the complication of the boundary conditions: the application of boundary conditions in the remaining velocity terms is straightforward, but proper application of vorticity boundary values results in integral properties having to be satisfied. To illustrate this, consider the parallel flow between two flat plates. If the integration domain is taken from the centreline to the wall, then the centreline vorticity constraint is $\omega_{CL} = 0$, commensurate with the symmetry of the flow. The wall constraint becomes $\int_{CL}^{W} \omega dy = 0$, because in this parallel flow, vorticity relates to the axial velocity u as $\omega \equiv -\frac{du}{dy}$.

The way to avoid the use of such integrals is to evaluate the wall vorticity explicitly and only employ the velocity conditions implicitly. That is, in a temporal evolution, the internal vorticity field, and then the full velocity field, is calculated at a new time step. Thereafter the wall vorticity is determined from the previous time step. This method of evolving the system is not physically valid or mathematically rigorous, and in effect constitutes a work-around¹.

¹It is suspected that this was one of the problems that caused a noticeable error in the base flow calculations of da Silva (1990), leading to his incorrect determination of the neutral stability curve for the steady pipe entrance system.

encountered in solving the Navier-Stokes equations: the momentum equation (10.1) is a temporal evolution equation for $\mathbf{u}$ alone - p is not evolved by this equation. Thus, although equations (10.1) and (10.2) fully determine the four unknowns $(\mathbf{u}, p)$, the structure of equation (10.1) is such that p is not given in the same sense as $\mathbf{u}$ is. The pressure term is in effect an implicit compensator which adjusts its value instantaneously to ensure that the velocity field is divergence free². The problem with pressure arises from the incompressibility assumption itself: in a compressible system, the continuity equation effectively evolves the pressure in the manner of the other variables. It is the simplification of this relation which elevates the pressure to its unusual status in the incompressible system.

This structure of (10.1) detrimentally affects the ease of numerically formulating the full Navier-Stokes equations. It is for this reason that most common methods of numerical solution for the primitive equations (10.1) evolve the velocity components differently to the pressure field³ and may also lead to associated problems with stability and accuracy unless carefully planned. Details of the intricacies needed to make these formulations work correctly concern substantial parts of such works as Roache (1972) and the reader is referred there for further information.

The simplest way around the pressure problem is to eliminate it from the governing equations altogether. Appropriate methods of doing this are outlined below:

- (i) The equation (10.1) can be reformulated in terms of the vorticity. The vorticity of a flow $\boldsymbol{\omega} = \boldsymbol{\omega}(\omega_1, \omega_2, \omega_3)$ is introduced as

$$\boldsymbol{\omega} = \nabla \times \mathbf{u}. \quad (10.3)$$

If the curl of equation (10.1) is taken, then the result is, after the substitution of (10.3), an equation in $\mathbf{u}$ and $\boldsymbol{\omega}$ alone:

$$\frac{\partial \boldsymbol{\omega}}{\partial t} + \mathbf{u} \cdot \nabla \boldsymbol{\omega} = \boldsymbol{\omega} \cdot \nabla \mathbf{u} + \nu \nabla^2 \boldsymbol{\omega}. \quad (10.4)$$

This equation is usually termed the vorticity-transport formulation. Equation (10.4) is for a general three-dimensional fluid. For the restrictive case of a two-dimensional domain, the term $\boldsymbol{\omega} \cdot \nabla \mathbf{u}$ is identically zero, and the equation simplifies somewhat. In a three-dimensional system, this term represents the effect of vortex stretching.

²Mathematically, p acts as a Lagrange multiplier that ensures the incompressibility constraint.

³In some types of finite-difference method, the pressure is even determined at different nodes to those where the velocity is calculated.

10.1 The incompressible Navier-Stokes equations

The incompressible Navier-Stokes equations describe the relationship between the velocity of the fluid, its pressure, and the action of forces on the fluid boundary. The velocity in a general unsteady fluid of three-dimensional extent within a bounded region $\mathcal{V}$ is represented by a three dimensional vector field $\mathbf{u} \equiv (u_i, u_j, u_k)$. The velocity can vary in space and time; $\mathbf{u} = \mathbf{u}(\mathbf{x}, t)$, with $\mathbf{x} \equiv (x_i, x_j, x_k)$ the orthogonal spatial co-ordinate system, and t the time. This velocity is determined by the general Navier-Stokes equation

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p + \nu \nabla^2 \mathbf{u} + \mathbf{F}(\mathbf{x}, t). \quad (10.1)$$

Here ν is the kinematic viscosity of the fluid. The pressure $p = p(\mathbf{x}, t)$ is a scalar field; that is, at any one point in space and time the pressure p is represented by a single value. The velocity field is further subject to the constraint that it be divergence-free. This restriction is the continuity requirement

$$\text{div}(\mathbf{u}) = 0, \quad (10.2)$$

which must be satisfied everywhere in the fluid. The equations (10.1) and (10.2) constitute a complete set of four independent equations in the four unknowns $(\mathbf{u}, p)$. The derivation of these relations can be found in any standard work on fluid mechanics [Batchelor (1967), Schlichting (1979)].

Any particular solution for $\mathbf{u}$ and p is determined by actions on the fluid body through the boundary $\mathcal{V}$ and via the body force $\mathbf{F}$. The fluid is assumed to adhere to the boundary everywhere: portions of the boundary may be moving, 'dragging' the fluid particles adjacent to it along, or portions may be stationary, restraining the fluid here to zero velocity. These boundary conditions are the driving or forcing functions on the fluid. The body force $\mathbf{F}(\mathbf{x}, t)$ also acts as a forcing function, but here its effect is directly on each element in the fluid. Thus the two external influences, the boundary conditions and the body forces, define the particular internal state of the fluid¹. If the fluid has constant density or no free surface, a simplification can be made to equation (10.1). In this case the weight of each element of fluid is exactly opposed by its hydrostatic buoyancy force, and $\mathbf{F}$ can be cancelled out if the pressure p is taken with respect to a suitable datum. Thus for fully enclosed homogenous incompressible fluids, the body force $\mathbf{F}$ can be ignored.

The pressure p in equation (10.1) only occurs once on the right-hand side. It is the appearance of p in this way that is responsible for most of the problems

¹The difficulty in fluid mechanics, and thus the major reason for its continuing interest, is that a particular set of boundary conditions do *not* define necessarily a unique internal state of the fluid. As fluid velocities increase, the internal motions bifurcate from a single-valued nature, and become chaotic. This is the so-called transition to turbulence.

Chapter 10

The governing equations

Perhaps the central tenet in the study of the motion of fluids is the acceptance of the universality of the Navier-Stokes equations in describing such systems. Disagreements in the field of fluid mechanics stem chiefly from disparities in results attained through their application to model problems, and due to their insolubility in all but the simplest cases. Thus the major problem is not a mismatch between the governing equations and reality, but rather the difficulty of solution of the equations themselves.

Derivation of these generally applicable equations has been ascribed firstly to M. Navier in 1827 and S.D. Poisson in 1831. These two workers derived the relations from an argument extended from basic intermolecular forces. Later B. de Saint Venant and G. G. Stokes separately arrived at the same relations by using the Newtonian assumption: that the normal and shearing stresses acting on the fluid are proportional to the rate of deformation. This assumption of the linearity of the applied force to the rate of strain is basically Newton's law of friction. The veracity of the Navier-Stokes equations is indisputable and has been confirmed in many experiments such as laminar boundary layer flows amongst others.

The motion of a liquid at low Mach numbers (far away from that fluid's speed of sound) allows a substantial simplification to the governing relations. Under such conditions, the fluid can be deemed incompressible. This subset of the Navier-Stokes relations has been the subject of much study, and will be the sole subject of the current investigation. The equations are not generally soluble, but with the availability of powerful numerical machines, highly accurate simulation of the behaviour of incompressible fluids has become possible.

Preamble

In this part of the thesis an analytical framework is established for the general mathematical description of accelerating flows. A general numerical method of solution is developed, and as a result, accurate numerical solutions to the resultant governing equations are obtained, both for the laminar flow variation, and to predict the onset of instability, in exponentially accelerating pipe entrance flows.

Chapter 10 that follows derives the basic governing equations of relevance to the study. The numerical method of solution is introduced in chapters 11 and 12, while the resultant numerical techniques are verified on simple fluid mechanics problems of relevance to the study in chapter 13. Results for the laminar base flow in the pipe entrance are then given in chapters 14 and 15, while the stability results are presented in chapters 16 and 17.

Equation (10.48) is a linear partial differential equation in the disturbance terms.

The mathematical form of the disturbance has been left unspecified up until now, but is now given form as a perturbation in streamfunction,

$$\psi' = \psi(r) \exp[i\alpha(x - ct)], \quad (10.49)$$

where α is a (complex) spatial wavenumber embodying the spatial behaviour of the perturbation, c (complex) is termed the celerity, and embodies the temporal frequency and velocity of the disturbance, and $\psi(r)$ is the shape of the disturbance radially. Both α and c are constant coefficients with respect to x , r and t .

The disturbance streamfunction ψ' is related to the velocities u' , v' as

$$\begin{aligned} u' &= \frac{1}{r} \psi_r' \\ v' &= -\frac{1}{r} \psi_r' \end{aligned} \quad (10.50)$$

Equation (10.49) can be differentiated any number of times with respect to x , r or t . Because of this, the u' , v' terms in equation (10.48) (and their derivatives) can be replaced with the equivalent streamfunction relations using the equations (10.50). This leads to a differential equation in ψ' alone. Substitution of the disturbance form (10.49) into the resulting equation, and the subsequent cancellation of the exponential term from the equation by division, results in

$$\begin{aligned} \frac{1}{\alpha} \left[\psi_{rrrr} - \frac{2}{r} \psi_{rrr} + \frac{3}{r^2} \psi_{rr} - \frac{3}{r^3} \psi_r + \frac{\alpha^2}{R^2} \left(\psi_{rr} - \frac{1}{r} \psi_r \right) + \frac{\alpha^4}{R^4} \psi \right] \\ = (\dot{u} - c) \left[\psi_{rr} - \frac{1}{r} \psi_r - \frac{\alpha^2}{R^2} \psi \right] + \left(\dot{u}_{rr} - \frac{1}{r} \dot{u}_r \right) \psi \\ + \frac{1}{\alpha} \Omega \left[\psi_{rr} - \frac{1}{r} \psi_r \right] \end{aligned} \quad (10.51)$$

or, if the operator

$$L \equiv \frac{\partial^2}{\partial r^2} + \frac{1}{r} \frac{\partial}{\partial r} \quad (10.52)$$

is introduced, (10.51) can be simplified to

$$\frac{-i}{\alpha} \left[\left(L - \frac{\alpha^2}{R^2} \right)^2 + \Omega L \right] \psi = (\dot{u} - c) \left[L - \frac{\alpha^2}{R^2} \right] \psi + [L\dot{u}] \psi \quad (10.53)$$

equations. The second standard assumption of the linear analysis is now invoked, namely that the mean flow variation by itself - equations (10.39) - must also satisfy these relations. Substitution of the mean flow variation (10.39) into the Navier-Stokes relations yields the simplified equation

$$\Omega \bar{u} = -\bar{p}_x + \frac{1}{r} (r \bar{u}_r)_r \quad (10.43)$$

Certain of the base flow variables or their derivatives (i.e. u_x) are zero by virtue of equations (10.39). Deletion of these terms from the expanded equation leads to:

$$\begin{aligned} \Omega(u' + \bar{u}) + u'_t + uu'_x + \bar{u}_r v' &= -p'_x - p_x + \frac{1}{r} (ru'_r)_r + \\ &+ \frac{1}{r} (r\bar{u}_r)_r + \frac{1}{Re^2} u'_{xx} \end{aligned} \quad (10.44)$$

$$v'_t + \bar{u} v'_x = -Re^2 p'_r + \left[\frac{1}{r} (rv')_r \right]_r + \frac{1}{Re^2} v'_{xx} \quad (10.45)$$

The equation (10.43) can now be used to further simplify equation (10.44). The result of this is⁷:

$$\Omega u' + u'_t + \bar{u} u'_x + \bar{u}_r v' = -p'_x + \frac{1}{r} (ru'_r)_r + \frac{1}{Re^2} u'_{xx} \quad (10.46)$$

$$v'_t + \bar{u} v'_x = -Re^2 p'_r + \left[\frac{1}{r} (rv')_r \right]_r + \frac{1}{Re^2} v'_{xx} \quad (10.47)$$

It is possible to simplify equation (10.47) further by invoking continuity for the disturbance terms (The disturbance must obviously satisfy the continuity requirement (10.27)). This allows the replacement of the term $\frac{1}{r} (rv')_r$ with $-v'_{xx}$.

In order to eliminate the pressure terms, equation (10.46) is differentiated with respect to x , equation (10.47), after division through by Re^2 , is differentiated with respect to x . Thereafter, the results are subtracted one from another. This process results in a single equation without the pressure term.

$$\begin{aligned} \left[\frac{1}{r} (ru'_r)_r \right]_r &= \Omega u'_x + u'_{tr} + \bar{u}_r (u'_x + v'_r) + \bar{u} u'_{xr} + \bar{u}_{rr} v' \\ &+ \frac{1}{Re^2} [2u'_{xrr} + v'_{tr} + \bar{u} v'_{xr}] - \frac{1}{Re^2} v'_{xxr} = 0 \quad (10.48) \end{aligned}$$

⁷Equation (10.45) has been repeated below as equation (10.47) for the sake of clarity.

These equations describe the general behaviour of a fluid in a pipe, where the strongly varying flowrate is encapsulated by the acceleration parameter Ω , while the magnitude of the renormalised velocity u ($\bar{u}$ in the previous section) remains of order one. The flow in the pipe is now taken to consist of an invariant mean flow with a superimposed small fluctuating component, or

$$\begin{aligned}u &= \bar{u} + u', \\v &= \bar{v} + v', \\p &= \bar{p} + p'.\end{aligned}\tag{10.38}$$

The mean flow in the renormalised basis is taken as being steady and parallel; that is

$$\begin{aligned}\bar{u} &= u(r), \\ \bar{v} &= 0, \\ \bar{p} &= p(r).\end{aligned}\tag{10.39}$$

This choice of base flow is analogous to the assumption made in the classic Seshi derivation, where that flow is also assumed steady and parallel. However in the present case, the mean cross-sectional velocity $\bar{U}$ can vary generally in *magnitude* as $\bar{U} = \bar{U}(t)$ with no loss of generality. The superimposed fluctuating component is given a general form, namely:

$$u' = u'(x, r, t),\tag{10.40}$$

$$v' = v'(x, r, t),\tag{10.41}$$

$$p' = p'(x, r, t).\tag{10.42}$$

Thus the actual flow in the pipe consists of the invariant mean flow (10.39) plus a superimposed fluctuating component (10.40).

Equations (10.38) are now substituted into (10.35), (10.36) and (10.37), and expanded. The primary assumption of the linear theory is now invoked; that products of small quantities, such as $u'v'$, are negligible in comparison with terms containing one or no small terms, such as $\bar{u}v'$ or uv' . Products of these small terms are thus eliminated from the resulting equations.

The substitution of the equations (10.38) into the Navier-Stokes relations is tantamount to stating that the full velocity field (u, v, p) has to satisfy these

10.4.2 Assumptions & methods of derivation

The application of the Seshi equation to pipe entrance flows is subject to the assumption that the flow is *nearly-parallel*. That is, that the variation of the velocity profiles in the axial direction is gradual enough to allow the stability analysis from isolated profiles to be used. The effect of the non-parallel velocity field within the pipe inlet on stability has been shown to be negligible by Abbot & Moss (1994). For this reason a parallel analysis is performed below.

Two methods of derivation - yielding the same final results - can be used to generate the Seshi equation: in the first approach, a general mean flow variation (with axial dependence) is used, and the nearly-parallel assumption is used to eliminate the negligible base-flow terms during the course of the derivation, while the second method assumes a parallel base flow from the outset, and thus the non-parallel terms are removed before proceeding with the analysis. The latter approach leads to a much more terse formulation, while, as mentioned, yielding identical results.

The temporal variation of the velocity profile shape (not its magnitude, which is accounted for by the acceleration parameter) has not been mentioned. In fact the assumption of *steadiness* implicit in the stability derivation would be violated if the geometrical shape of the velocity profiles varied in time. However it will be shown later that the exponentially accelerating pipe entrance region is completely *steady* in the renormalised framework; thus this reservation is of no consequence.

10.4.3 Derivation

The stability derivation begins with equations (10.25), (10.26) and (10.27), which fully describe the axisymmetric flow of an incompressible fluid in a pipe, under conditions of arbitrary acceleration. These equations refer to dimensionless variables,

Due to the nature of the derivation, a more compact notation will be adopted for the representation of a derivative, namely $\frac{\partial \eta}{\partial z} \equiv \eta_z$. Furthermore, for clarity, the bars will be dropped from the dimensionless quantities in the above mentioned equations. These equations can be restated for clarity as

$$u_t + \Omega u + uu_x + vu_r = -p_r + \frac{1}{r}(ru_r)_r + \frac{1}{Re^2}u_{xx} \quad (10.35)$$

$$v_t + vv_r + uv_x = -Re^{-2}p_r + \left(\frac{1}{r}(rv_r)\right)_r + \frac{1}{Re^2}v_{xx} \quad (10.36)$$

$$\frac{1}{r}(rv_r)_r + u_x = 0. \quad (10.37)$$

10.4.1 Types of linear instability

In the pipe flow system, the modal shapes of the instability can either be rotationally symmetric about the origin (axisymmetric), or possess an azimuthal dependence (non-axisymmetric). In the latter state a number of different configurations can occur, each with an integer azimuthal frequency. Of all these modes of instability, the two most prominent are the axisymmetric mode and the non-axisymmetric mode with an azimuthal frequency of $n = 1$. In the case of steady flow, the equation for the axisymmetric case is the straightforward Squire relation [Drazin & Reid (1981)]. The non-axisymmetric equation is mathematically more complex and consists of a coupled set of differential eigenvalue equations [see Huang & Chen (1974b)].

Unlike the stability of planar flow - where Squire's (1933) theorem dictates that the two-dimensional disturbances are always the most unstable - there is no guarantee that axisymmetric disturbances in a pipe system are the more dangerous. In fact, work by Huang & Chen (1974a, b) and Garg (1981) indicate that pipe inlet flow is most unstable to the non-axisymmetric disturbance with azimuthal wavenumber $n = 1$. Thus, strictly speaking, the more complex relation should be employed in conjunction with the Squire relation, in order to determine fully the response of the pipe inlet to general small disturbances.

However, Huang & Chen (1974b) showed that there was only a small numerical difference between the axisymmetric and non-axisymmetric minimum critical Reynolds numbers (in their case $Re_{crit} = 19\,780$ for $n = 1$ as opposed to $Re_{crit} = 19\,900$ for $n = 0$). They also showed that the two variations of Re_{crit} with axial distance approach each other towards the inlet, and that towards this limit, the axisymmetrical disturbance has prominence. This is because, as the pipe entrance is approached, the boundary layer thins, and approximates more and more a two-dimensional boundary layer, with the result that the Squire (1933) criterion becomes more valid. This was seen in Huang & Chen's (1974b) results. This boundary layer behaviour in this limit is inherent in the governing equations, and is demonstrated in a later chapter.

The current investigation concerns the development of a stability model for unsteady flows. The lack of any appreciable research in the field fully justifies the consideration of the simpler - axisymmetric - disturbances first, regardless of whether they are the most unstable modes or not. The closeness of results for the two types of disturbance in the steady case will quite probably hold in the more general, unsteady study. Thus in order to satisfy the primary goal of the research - the understanding of the general stability behaviour of unsteady flows - while retaining simplicity, a consideration of the simpler axisymmetric disturbance modes is sufficient. The non-symmetrical study is thus left as a topic for further research. In fact the simpler derivation given here will present the framework upon which the more complex analysis can be built.

$$\frac{1}{r} \frac{\partial}{\partial r}(rv) + \frac{\partial u}{\partial r} = 0. \quad (10.34)$$

The equations (10.33) and (10.34) are sufficient to describe the high Reynolds number unsteady flow in a pipe entrance. The pair form a coupled set of partial differential equations in z and r alone. Because the above equations are derived directly from (10.28), (10.29) and (10.30) without any assumptions, they are *no less* restrictive than the latter. Although the integral and boundary terms in equation (10.33) may seem awkward, it will be shown in chapter 12 that the numerical procedure treats them in a uniform manner, and thus they offer no impediment to the solution. The equations are non-linear; however, the solution technique employed will handle these non-linearities properly, and no linearising will be necessary.

10.4 The unsteady stability equation

This section is concerned with the derivation of a linear stability equation applicable to accelerating pipe flow, from the governing equations given in section 10.2. The derivation follows the classical approach to linear stability [see Drazin & Reid (1981), Lin (1955)]. A general form of disturbance is introduced into the governing equations, the resulting equations are linearised in the disturbance quantities, and a complex differential eigenvalue problem describing the linear stability of the flow results. However, the equation derived here will be seen to be more general in its application than the Squire equation for steady pipe flow, but reducing to the latter in the limit of zero acceleration.

The only known stability study of temporally unsteady flows is that by Shen (1961). He showed in this pioneering work, how inviscid linear stability would apply to a limited set of parallel accelerating flows, in which the velocity changes with time in magnitude only as $u(r, t) = R(r) \cdot T(t)$. He argued that the rate of increase of the energy of the disturbance had to exceed the rate of increase of the energy of the base flow, in order for a disturbance to be deemed unstable. This he termed *momentary stability*. To satisfy this criterion, he showed that for instability in such a system *in the inviscid limit*, $c_i \sim \frac{1}{U} \frac{dU}{dt}$, where U is the mean flow velocity, rather than $c_i > 0$ applicable for quasi-steady flows. Thus, an accelerating flow would have a higher critical Reynolds number than an identical one showing steady-state behaviour, while the converse would apply for a decelerating flow. The $\frac{1}{U} \frac{dU}{dt}$ term is familiar from the analysis of section 10.2 as the quantity Ω - the acceleration parameter. Shen's analysis was limited by its inviscid nature. For viscous (finite Reynolds number) flows, he showed that his analysis was only valid in the limit of very low accelerations.

The above renormalized equations give the impulsively started system of da Silva (1990) for $\Omega = 0$. Furthermore, they are Reynolds number independent, so a single solution for a particular $\Omega(t)$ recovers an entire family of flows with different Reynolds number ranges. As will be shown in Chapter 17, this point is very important for the ease of implementation of the stability analysis, in that a single set of base flow results can be used to determine stability throughout the full range of Re .

10.3.2 Elimination of pressure term

As mentioned above, any numerical solution method for the Navier-Stokes and related equations invariably has difficulties with the pressure term p . The approach adopted here is to use a pressure-integral method which, while eliminating the pressure term, will do so without approximation and will still allow a simple and uniform numerical treatment.

Broadly speaking, the pressure can be eliminated by equating the average pressure at a particular pipe cross-section with the pressure at any one point. This is possible because (10.29) states that the pressure is a function of x and t only. This proceeds as follows:

The statement that the average pressure is equal to the actual pressure is equivalent to the mathematical statement

$$\begin{aligned} \frac{1}{A} \int_A p dA &= \frac{1}{\pi} \int_0^1 p(2\pi r dr) \\ &= \bar{p}. \end{aligned} \quad (10.31)$$

The relation (10.31) is no more than (10.29) in a different guise.

The way forward is to take the integral $\frac{1}{A} \int_A dA$ of both sides of equation (10.28) then by virtue of the equality (10.31), an independent equation for the pressure is obtained. Using the continuity consideration (10.30), certain terms in the resulting equation can be eliminated. The result of this integration is

$$-\frac{\partial \bar{p}}{\partial x} = \Omega + 2 \frac{\partial}{\partial x} \int_0^1 r u^2 dr - 2 \left. \frac{\partial u}{\partial r} \right|_{r=1}. \quad (10.32)$$

Equation (10.32) may now be combined with equation (10.28), in order to eliminate the pressure term. The result is shown below. The continuity requirement (10.30) is also restated below for completeness. The resulting equations are as follows:

$$\frac{\partial \bar{u}}{\partial t} + \Omega(u-1) + u \frac{\partial \bar{u}}{\partial x} + r \frac{\partial u}{\partial r} = 2 \frac{\partial}{\partial x} \int_0^1 r u^2 dr - 2 \left. \frac{\partial u}{\partial r} \right|_{r=1} + \frac{1}{r} \frac{\partial}{\partial r} \left(r \frac{\partial \bar{u}}{\partial r} \right) \quad (10.33)$$

figure 8.17 for a typical such plot]. This is the strongest possible vindication of the power of the renormalised equations, in that a simple mapping from an *invariant* steady-state solution can explain the details of a highly accelerative base flow. This section is a brief overview of the behaviour of the co-ordinate mapping, and detailed numerical results for the accelerating base flow are contained later in chapter 15.

10.3 Laminar base flow equations

The above derivations have been general in nature, with the full axisymmetric incompressible Navier-Stokes equations being considered in the context of unsteady flows. In order to obtain accurate *laminar* results from these general relations it is appropriate to simplify them somewhat to expedite the solution process. This section is concerned with simplifications useful for obtaining accurate laminar solutions for the unsteady pipe entrance flow. In the following section the derivation of the normal mode stability equation for this system is considered.

10.3.1 The high Reynolds number approximation

The equations (10.25), (10.26) and (10.27) are the dimensionless form of the renormalized incompressible Navier-Stokes equations in an axisymmetric system. The relative importance of the various terms for pipe flows is immediately apparent by observation of the terms embodying the Reynolds number in equations (10.25) and (10.26). In the former relation the axial diffusion term $\frac{\partial^2 u}{\partial z^2}$ scales as Re^{-2} and thus becomes negligible at the Reynolds numbers at which transition occurs⁶. Likewise, considering equation (10.26), the high Re approximation states that $\frac{\partial p}{\partial r} = 0$ which implies that $p = p(x, t)$ alone. Thus for high Re , the governing equations reduce to

$$\frac{\partial u}{\partial t} + \Omega u + u \frac{\partial u}{\partial x} + v \frac{\partial u}{\partial r} = -\frac{\partial p}{\partial x} + \frac{1}{r} \frac{\partial}{\partial r} \left(r \frac{\partial u}{\partial r} \right), \quad (10.28)$$

$$\frac{\partial p}{\partial r} = 0, \quad (10.29)$$

$$\frac{1}{r} \frac{\partial}{\partial r} (rv) + \frac{\partial u}{\partial x} = 0. \quad (10.30)$$

⁶Instability occurs at a minimum Reynolds number of approximately 10000, thus the axial diffusion term is smaller by 8 orders of magnitude in comparison with the remainder of the equation.

solution for the equations will mean invariance in this transformed basis. For instance, if the flow is undergoing an exponential acceleration $\Omega = k$, then these equations will produce a steady-state solution (after the transients have died away) in the dimensionless system. However, the mapping $\bar{x} \mapsto x$ is proportional to U ; thus in the *physical* co-ordinate system, the solution domain encompasses more and more of the pipe in an exponential manner. Figure 10.5 shows this mapping at one time, and then at later times as the velocity increases. This figure demonstrates how the strong variation in an accelerating pipe inlet - a substantial and continuing boundary layer thinning - is solely a result of a mapping from an invariant solution.




Figure 10.5: The effect of the variable mapping $x \mapsto \bar{x}$ on the perception of an impulsively started and exponentially accelerating pipe entrance velocity flow (indicated using displacement thickness δ^*) in the dimensionless co-ordinate system, from a position of constant x in the physical co-ordinate system. The dashed lines indicate lines of constant x . The physical variation of δ^* at an x station $x = k$ is indicated in the accompanying plot.

Immediately apparent is that the form of the results produced by this mapping is identical to the experimental results from chapter 8 for exponential flows [See

- (iv) For any functional variation $U = U(f)$, the equivalent $\Omega = \Omega(f)$ can be found from relation (10.24).

The acceleration parameter is a forcing function on the system of equations (10.25), (10.26) and (10.27). It is not determined by the system but is imposed from the outside. As such there is no limitation on its form; functions with arbitrary temporal variation, even discontinuous functions, can be used. Numerically, any value of Ω can be supplied at each time-step. This opens up unlimited possibilities in terms of studying unsteady pipe flows. As an example, an impulsively partially blocked flow (a flow decelerated impulsively from a high Reynolds number to a lower one) can be modelled simply by making $\Omega \ll 0$ for a short time.

This parameter was first introduced by Shen (1961) in his pioneering (and only to date) study on the stability of accelerating flows. This term arose in his inviscid stability analysis of self-similar accelerating velocity profiles, as a ratio between the energy of the base flow and the energy of the disturbance. He did not generate any base flow model, concerning himself mainly with the stability behaviour of arbitrary prescribed inviscid profiles. Moss (1991) developed a laminar model for the temporally developing flow, incorporating the effects of a general acceleration. This was formulated using a momentum-integral technique, and the acceleration parameter arose naturally from this analysis. The simplified nature of the formulation resulted in limited results, although his study showed the qualitative effects of acceleration correctly.

The experimental studies of Lefebvre & White (1989), (1991), also introduced an energy-ratio term analogous to the acceleration parameter. They were concerned with the prediction of transition in a long pipe undergoing a constant acceleration.

The parameter (10.24) thus arises naturally from the formulation of a properly unsteady flow system. It is a parametric quantifier of acceleration in as much as the Reynolds number quantifies the flow velocity.

10.2.6 The behaviour of the axial co-ordinate

While the renormalisation of u is fairly straightforward, it is the scaling of the axial co-ordinate x that needs further elucidation: the dimensionless parameter x_0 is a function of U , and scales as $x_0 \propto U$. Thus for a flow with say, increasing Reynolds number, any constant region $[x_1, x_2]$ in the physical co-ordinate system maps to a *variable* region in the dimensionless system of decreasing size $[x_1/x_0, x_2/x_0]$.

It is perhaps clearer if this mapping is explained in the converse way. Equations (10.25), (10.26) and (10.27) are in terms of the dimensionless variables. As the dimensionless basis is the natural framework for these equations, an invariant

As stated, the other non-constant terms x_0 and p_0 are never differentiated with respect to t and can thus be treated in the more conventional sense as constant coefficients, for the sake of this derivation.

If the parameter

$$\bar{\Omega} \equiv \frac{1}{U} \frac{dU}{dt} \quad (10.24)$$

is defined, then the U terms in equation (10.23) are conveniently encapsulated in the term $\bar{\Omega}$. The incorporation of (10.23), (10.24) along with the other (constant) non-dimensional substitutions (10.19), (10.20), into equations (10.8), (10.9), and (10.10) leads to the dimensionless formulation:

$$\frac{\partial \bar{u}}{\partial \bar{t}} + \bar{\Omega} \bar{u} + \bar{u} \frac{\partial \bar{u}}{\partial \bar{x}} + \bar{v} \frac{\partial \bar{u}}{\partial \bar{r}} = -\frac{\partial \bar{p}}{\partial \bar{x}} + \frac{1}{\bar{r}} \frac{\partial}{\partial \bar{r}} \left(\bar{r} \frac{\partial \bar{u}}{\partial \bar{r}} \right) + \frac{1}{Re^2} \frac{\partial^2 \bar{u}}{\partial \bar{x}^2} \quad (10.25)$$

$$\frac{\partial \bar{v}}{\partial \bar{t}} + \bar{v} \frac{\partial \bar{v}}{\partial \bar{r}} + \bar{u} \frac{\partial \bar{v}}{\partial \bar{x}} = -Re^2 \frac{\partial \bar{p}}{\partial \bar{r}} + \frac{\partial}{\partial \bar{r}} \left(\frac{1}{\bar{r}} \frac{\partial}{\partial \bar{r}} (\bar{r} \bar{v}) \right) + \frac{1}{Re^2} \frac{\partial^2 \bar{v}}{\partial \bar{x}^2} \quad (10.26)$$

$$\frac{1}{\bar{r}} \frac{\partial}{\partial \bar{r}} (\bar{r} \bar{v}) + \frac{\partial \bar{u}}{\partial \bar{x}} = 0. \quad (10.27)$$

10.2.5 The acceleration parameter

The $\bar{\Omega}$ term in equation (10.25) introduced as equation (10.24) conveniently describes the $d(\text{Flowrate})/dt$ behaviour as distinct from the $d(\text{Profile shape change})/dt$ behaviour governed by the rest of (10.25). As such it is termed the *Acceleration Parameter*. The form of $\bar{\Omega}$ for various forms of $\frac{dU}{dt}$ deserves mention.

- (i) For an impulsively accelerated flow, $\bar{\Omega} = 0$, and thus equation (10.25) reduces to the more conventional form given by da Silva (1990) for developing flows with a steady mass flow rate.
- (ii) For systems in which the mean flow varies as $U(t) = at^n$, the acceleration parameter $\bar{\Omega} = \bar{\Omega}(t) = \frac{n}{t}$ defines the effect of a polynomial acceleration. In this case, the effect of the acceleration diminishes in time for any n , and the dimensionless system approaches the same limit as the impulsive case for large t . The constant-acceleration experimental flows of Lefebvre & White (1989), (1991) are described by setting $n = 1$.
- (iii) In systems where U varies exponentially as $U(t) = a \exp(nt)$, the acceleration factor $\bar{\Omega} = n$ is a *constant* parameter in the system. For large t , the system approaches a constant limit different from the impulsive one.




Figure 10.4: The variation of a normalised general function $w(x, y)$ in the range $x \in [0, k]$, $y \in [0, 1]$, showing the encapsulation of the unbounded variation of the original function w in the term w_0 .

with U defined as

$$U = \frac{2\pi \int_0^R r u \, dr}{\pi R^2}. \quad (10.22)$$

The equation (10.21) is equivalent to separating the axial velocity's time dependence into two components: firstly the *strong* magnitude scaling dependence U and secondly the *weak* profile shape temporal variation. This separation of variance allows the governing equations to be formulated analogously to the dimensionless approach of da Silva (1990) for his case of impulsively started pipe entrance flows. The resulting equations will be shown to contain da Silva's relations as a limiting case.

As stated, care must be taken to manipulate relation (10.21) properly. The time-derivative of the product must be evaluated using the product rule of differentiation as

$$\begin{aligned} \frac{\partial}{\partial t} u &= \frac{\partial}{\partial t} (U^* u) \\ &= U^* \frac{\partial u}{\partial t} + u^* \frac{dU^*}{dt} \\ &= U^* \left[\frac{\partial u}{\partial t} + u \left(\frac{1}{U^*} \frac{dU^*}{dt} \right) \right]. \end{aligned} \quad (10.23)$$

The approximate solution for u can then be presented as

$$\hat{u}^{(N)}(x) := \sum_{k=0}^{N-1} a_k \psi_k(x) \quad (11.5)$$

where the ψ_k , a_k are the trial functions and expansion coefficients respectively.

The purpose of the numerical method is to find the values of the expansion coefficients a_k such that the interpolating series - equation (11.5) - is as accurate a representation of the unknown function as possible. To do this, the residual $P := M\hat{u}^{(N)}$ is minimised by enforcing the condition

$$\int_{x_1}^{x_2} [P - M\hat{u}^{(N)}] \psi_k(x) dx = 0, \quad k = 0, \dots, N-1. \quad (11.6)$$

The residual is analytically calculable because of the simple nature of the trial functions (either Fourier or polynomial). Here ψ_k are the test functions, which determine the weights of the residual. By using suitable sets of trial and test functions orthogonal on $[x_1, x_2]$, the integral (11.6) can be made to reduce to a simple form, that is, for terms like

$$\int_a^b (\psi_k \cdot \psi_j) \psi_j dx = \delta_{kj}, \quad (11.7)$$

the term $\delta_{kj} = 0$ for all cases where $k \neq j$. What results is a set of N simultaneous equations in the N unknown coefficients a_k . The auxiliary conditions (11.4) are usually enforced at this stage in a variety of ways: either the trial and test functions are chosen such that these conditions are automatically satisfied, or some of the higher order equations obtained from equation (11.6) are dropped in favour of the auxiliary conditions. The specific application of the boundary conditions particular to each type of spectral method is outlined below.

This is a general description of all so called 'spectral methods'. Within the whole spectral gamut, the choice of trial and test functions is that which separates the various approaches.

11.3.1 Collocation

The collocation method uses shifted Dirac functions, placed suitably at nodes or collocation points, for the test functions. This reduces equation (11.6) to the requirement that the approximating function should satisfy the function exactly at these points, which are specially placed in order to ensure high accuracy in the solution domain. This simplification makes this method seem very much like a finite-difference scheme, because the expansion coefficients are in fact the values

Table 11.1: Convergence results of da Silva (1990) for the Orr-Sommerfeld equation applied to plane-Poiseuille flow, with $\alpha = 1$ and $R = 10\,000$, and compared to the 'exact' result of Orszag (1971a).

Order	Result	Absolute error in imaginary component (%)
da Silva (1990)		
$N = 100$	$0.236991 + 0.000364i$	2.7
$N = 200$	$0.237394 + 0.003718i$	0.5
$N = 300$	$0.237468 + 0.003730i$	0.2
$N \rightarrow \infty$ (extrap.)	$0.237520 + 0.003740i$	0.008
Orszag (1971a)		
$N = 14$	$0.23713751 + 0.00563644i$	50.7
$N = 20$	$0.23752676 + 0.00373127i$	0.14
$N = 23$	$0.23752670 + 0.00373982i$	0.004
$N = 26$	$0.23752648 + 0.00373967i$	'exact'
$N = 29$	$0.23752649 + 0.00373967i$	'exact'

in the grid density from 200 to 300 points only yields an improvement of the error in the imaginary component from 0.5% to 0.2%. Furthermore, both the real and imaginary values for $n = 300$ are only accurate to at most 4 significant figures. This should be contrasted with the results of Orszag (1971a), who achieved convergence to 8 significant figures with a Chebyshev approximating polynomial of order 29, and at order 20 the results were nominally better than that of da Silva.

11.3 The method of weighted residuals

The method of weighted residuals (MWR) is based on representing the exact function $f(x)$ in a compact range $[x_1, x_2]$, by a series expansion in a basis e^{ν} , of mutually orthogonal functions of order ν , with ν chosen such that the approximating series, $f^{(0)}(x) = \sum_{\nu=0}^{\nu} a_{\nu} e_{\nu}$ on $[x_1, x_2]$ is a sufficiently close approximation to $f(x)$. The terms a_{ν} are the ν unknown expansion coefficients, which will be determined by the solution procedure. Solution is achieved in the following way. The argument proceeds for a simple one dimensional example. Assume a general linear operator M operating on an unknown function $u = u(x)$. The solution for u is determined by the differential equation

$$Mu = P \quad (11.3)$$

along with b auxiliary (boundary, symmetry, initial-value, etc.) conditions

$$f_i(u) = k_i, i = 0, \dots, b-1 \quad (11.4)$$

In the discretisation decay as $(\Delta x)^r$ for an arbitrary finite r as the grid spacing Δx approaches zero. Increasing the order of discretisation (for instance to a five-point grid from a three-point one) will not yield any asymptotic improvement in rate of convergence, because the order of discretisation p relates to r in a linear sense. Another way of putting it is to say that finite-difference methods are not *efficient*⁵ – a large increase in the order is necessary to achieve any commensurate increase in accuracy. Having said that, many problems require only approximate solutions, and in these cases the finite-difference method is deemed adequate⁶.

The solution of the Orr-Sommerfeld equation for the temporal stability of plane Poiseuille flow serves as a good point of comparison for the various numerical approaches; because it is a high-order and thus difficult numerical system, and because many researchers have investigated this case. Two *test cases* for this problem have emerged. Firstly, calculation of the minimum critical Reynolds number of $Re = 5772.22$ (obtained accurately by Orszag (1971a) using a Chebyshev-Tau method), and secondly, calculation of the eigenvalue spectrum for the somewhat arbitrary case of $\alpha = 1$ and a Reynolds number of either $Re = 5\,000$ or $Re = 10\,000$ [for instance Davey & Drazin (1969)]. One of the earlier stability calculations to be carried out on plane Poiseuille flow were those by Thomas (1953) for the estimation of the most unstable symmetric disturbance mode at a Reynolds number of $Re = 10\,000$, and $\alpha = 1$. His finite-difference formulation used a five point discretisation, and achieved a truncation error of order $(\Delta x)^4$. da Silva (1990) presented a finite-difference formulation, which for the solution of the stability of plane Poiseuille flow, was able to approach the accuracy of Orszag (1971a) for the test case of $\alpha = 1$, $Re = 10\,000$ on plane Poiseuille flow, but was only achieved by using an extremely fine grid, large computational times, and extrapolation of the results to zero grid spacing. His truncation errors were of $O(\Delta x^4)$. Recently Gustavsson (1991), investigating the energy density of a disturbance in plane Poiseuille flow, used a shooting method. He integrated the Orr-Sommerfeld equation temporally using a fourth order Adams implicit method, and was able to extend the accuracy of his calculation to $O(\Delta x^{10})$. He utilised 432 nodes to solve the stability problem using this exotic method.

Finite-difference schemes, while able to achieve good accuracy when formulated with extreme care, are thus inherently limited in their efficiency, as evidenced in the calculation of fourth-order systems such as the Orr-Sommerfeld equation. As a result of their inefficiency, they require large storage space and use of powerful computing facilities to achieve accurate results. Again da Silva's results highlight the limitation of a finite-difference method. Considering his analysis for the stability of plane Poiseuille flow [da Silva (1990) – these results are shown in table 11.1], it is clearly evident that the rate of convergence of the most unstable eigenvalue to the exact value is of an algebraic nature. An increase

⁵This is probably an historical perspective – the finite-difference paradigm possibly stems from the calculus of Newton and its method of derivation (the limit of a derivative is a difference). If global analytical methods such as the conceptually simple operational Tau approach had been adopted early on by engineers and other applied mathematicians, then the 'ease of use' attributed to the finite-difference method might be seen to be somewhat of an exaggeration.

For instance, consider again the one-dimensional grid shown in figure 11.1. Let the representation of an unknown function $z(x)$ on this grid be given by a linear equation $Gz = 0$ with one of the boundary conditions being of simple Dirichlet form, specifically $z(x_0) = 1$. For the finite-difference formulation, the equation would be replaced by its matrix form analogue $\tau G = c$, where τ is a row vector of the N unknown values z_i at the nodes x_i , c is defined as $c = [0 \ 0 \ 0 \ \cdots \ 0 \ 0]$ and G is an $N \times N$ matrix⁴. The boundary condition would be imposed directly by replacing the first column of the operator G by the column vector $[1 \ 0 \ 0 \ \cdots \ 0 \ 0]^T$ and the first entry of c by 1. More complicated (Neumann, Robin) boundary conditions can be asserted in similar, but possibly more difficult ways.

11.2.2 Solution procedure

Most formal operators on a discretised function can be represented in matrix form, and the principle is analogous for all forms of finite-difference discretisation. The solution to the problem is found by solving the system that results, of the form

$$A_0 U B_0 + A_1 U B_1 + \cdots + A_n U B_n = C \quad (11.2)$$

for the unknown U . The boundary conditions would have been augmented to the operators A_i , B_i prior to this in order to satisfy the given constraints.

In the simpler approaches, the left hand side of (11.2) is constructed so that U is separable, and the system reduces to the form $AU = B$, which can be solved explicitly for U by the inversion of A . For the more general case, dependent upon the complexity of discretised system, diagonalisation or iteration are successful ways of solving for the unknown values.

Most finite-difference approaches are never explicitly formulated in a matrix form but are rather solved on a per-node basis. It is fairly elementary to show that any such approach can be restructured in a matrix operator form. In fact all the numerical methods outlined in this chapter result in the solution of a system identical to equation (11.2)⁵. The solution of such systems is in the realm of matrix algebra, and will not be taken any further here.

11.2.3 Convergence & accuracy

The convergent behaviour of any finite-difference approach is fundamentally linked to the discretisation process itself. Finite-difference methods can be said to yield only 'finite-order accuracy' [Orszag (1971a)] in that the errors involved

⁴If $G \propto \frac{d^2}{dx^2}$, then its finite-difference analogue G would be the matrix (11.1).

⁵The commonality of all the numerical methods is discussed in section 11.4.




Figure 11.1: A typical one-dimensional finite-difference grid.

line. If these values are represented as a vector $U = [u_0 \ u_1 \ \cdots \ u_{N-2} \ u_{N-1}]$, then the derivative $\frac{\partial u}{\partial x}$ on this domain is simply represented by a matrix operator on U as: $\frac{\partial}{\partial x}(U) = U \cdot D$ where D for a second order scheme is the tri-diagonal matrix of order $N \times N$, given as

$$D = \frac{1}{2\Delta x} \begin{bmatrix} -2 & -1 & & & & & & 0 \\ 2 & 0 & -1 & & & & & \\ & 1 & 0 & -1 & & & & \\ & & 1 & 0 & & & & \\ \vdots & & & & \ddots & & & \vdots \\ & & & & & 0 & -1 & \\ & & & & & 1 & 0 & -1 \\ & & & & & & 1 & 0 & -2 \\ 0 & & & & & & & \vdots & 2 \end{bmatrix}. \quad (11.1)$$

If the nodal values represent a variation in a two-dimensional domain, then U is a matrix, and differentiation with respect to the two independent space variables is implemented by a pre- or post-multiplication matrix operator. For instance the mixed first derivative with respect to x and y is $\frac{\partial^2}{\partial x \partial y} U = D^T \cdot U \cdot D$.

11.2.1 Boundary conditions

Boundary conditions of any complexity do not avail themselves to any uniform treatment although these methods allow a great deal of flexibility in their implementation. These conditions are usually embedded in the operator matrices at the boundaries, or again, virtual nodes outside the solution domain are employed to ensure certain relations at the boundary elements are met.

Direct turbulence modelling has been carried out for certain systems such as curved channel flow [Moser & Moin (1987)]. A systematic investigation into turbulence in a simple boundary layer has been conducted by Spalart & Leonard (1985). Simulations for isotropic turbulence have been produced by Érikacher *et al* (1987), where the high accuracy and spatial resolution of the approach is ideal if computer solution times are not to be too prohibitive. The detailed work by Caimto *et al* (1988) provides a good summary of the current state of affairs in this area.

An alternative approach to the same methods as above has been followed by Ortiz and his school [*inter alia* Ortiz (1969), Ortiz & Samara (1981, 1983), Liu & Ortiz (1982, 1986, 1987)] for some time. This approach, which is termed *operational* by Ortiz, results in identical formulations to the Tau method in its more conventional guise, but neatly side-steps the MWR mathematics conventionally used in deriving the matrices. His method could be safely termed a philosophy, in that its formalised systematic approach lends understanding to the finite-difference and collocation methods as well. It is this method which will be the subject of a detailed description in chapter 12. The different classes of numerical method are discussed in detail in the sections that follow.

11.2 The finite-difference approach

As stated above, the finite-difference method simulates a differential equation by representing it at discrete points in the solution space. Derivatives and other functional operators within the equation are represented as relational operators between the discrete values; for instance, the derivative $\frac{d^n f}{dx^n}$ is replaced by its discretised operator, $(f_i^n)^p$, the n th derivative operator of order p at node i .

The different classes of method are classified according to the choice of the function $(f_i^n)^p$. Thus a second-order finite-difference discretisation of the operator $\frac{d^2 f}{dx^2}$ would be $(f_i^{(1)})^{(2)} = \frac{1}{2\Delta x^2}(f_{i+1} - f_{i-1})$. Derivatives at the boundaries are replaced by skewed or one-sided differences, or extrapolation is made to 'virtual' nodes outside the boundary in order to better represent the boundary variation. Figure 11.1 shows a typical one dimensional finite-difference grid, with the continuous function $f(x)$ shown as a solid line, and the node points shown as solid circles.

Finite-difference systems are perhaps the most intuitive of all numerical methods. The effect of, for instance, differentiation is implemented exactly as would be anticipated:¹ the derivative at a point is calculated simply by looking at the local slope, or difference in height between two local nodes. The figure shows the implementation of the double sided difference at node x_i , as an estimate of the derivative there. The effect of differentiation on the nodal values can perhaps be better represented in matrix form. Consider a one-dimensional domain like that shown in figure 11.1, with the function given as values at N points along this

Stability of flows to linear disturbances, and the solution of the eigenvalue problem resulting therefrom, have been the subject of many spectral approaches. Grosch & Salwen (1968) investigated the stability of various classes of flow using expansions in orthogonal Chandrasekhar-Reid functions [Chandrasekhar (1961)]. However, their results did not show spectral accuracy due to these functions' peculiar end behaviour. Davey & Drazin (1969) also incorrectly used a Bessel function type expansion in the belief that orthogonal functions close to the structure of the eigenvalue problem itself would yield optimum results. Their results did not show spectral convergence for similar reasons. Orszag (1971a) was the first to develop a Chebyshev-Tau method for the Orr-Sommerfeld equation, and his solution has stood as the definitive benchmark for plane-Poiseuille flow. He outlined reasons why the previous authors' results were not accurate. He also gave the formulation for pipe-Poiseuille flow although it is possible that this is flawed².

More recent work on the topic of stability includes a study by Schmid & Henningson (1994), who investigated the energy density growth of disturbances in Hagen-Poiseuille flow using a collocation approach. They obtained 'extra' eigenvalues for the axisymmetric case of pipe-Poiseuille flow, which may be attributable to their using the full non-axisymmetric system to solve the simpler axisymmetric relation. Abbot & Moss (1994) obtained accurate stability curves for pipe inlet flow using a modified operational Tau-Chebyshev method. Their results agreed substantially with the accurate study of Garg (1981), revealing a systematic error in the finite-difference formulation of da Silva (1990). In particular, stability results were obtained for profiles close to the inlet, and the most downstream unstable profile was determined to have a critical Reynolds number of over 500 000. The generation of results in both these limits (thin boundary layer and high Re) required the extremely high efficiency of the spectral method for their accuracy.

Spatial stability in a similar boundary-layer has been investigated using Chebyshev methods by Bridges & Morris (1984a, 1984b). Von Kerczek (1982) studied the stability of oscillatory flow using a Chebyshev-Tau method and the Floquet theory.

The direct simulation of finite-amplitude disturbances and other transition phenomena using spectral methods has been well addressed [Orszag & Kells (1980), Patera & Orszag (1981), Orszag & Patera (1983), Herbert (1983a, 1983b, 1984, 1985), Malik & Hussini (1990), Kleiser & Zang (1991), Sandham & Kleiser (1992)]. In particular the last authors were able to successfully model boundary layer transition phenomena from the linear stage and T-S waves, through non-linear A vortex interactions, all the way to the onset of actual turbulence.

²The current research has shown that the Chebyshev-Tau approach is very sensitive to the inclusion of singularities in the domain. The Seshi equation applicable to pipe stability has a regular singular point at the origin. It is shown in chapter 17 that the Seshi equation must be transformed in order to remove this singularity before spectral convergence can be guaranteed. It is not clear from his derivation whether Orszag did remove the singularity analytically.

employs many simple elements to represent the solution, the former uses few (or one) for the entire discrete domain and ensures an accurate representation by dramatically increasing the *order* of the interpolating function within this domain. Bases of orthogonal functions suitable for this approach are trigonometric (Fourier) or classes of orthogonal polynomials such as the Legendre or Chebyshev polynomials¹.

An alternative way of looking at the spectral methods is to consider them as dealing with values at discrete points (the coefficients) much like the finite-difference approach, but in a transformed domain. The operators are, in a similar way to finite-difference operators on the physical domain, of a simple and local form in the transformed domain. The coefficients represent a continuous function in the real domain, and the link between domains is usually via a Fourier or related transform. In fact a class of spectral methods, the collocation methods, is implemented at specially placed node points in the physical domain - much like a finite-difference method - but the differential operators are actually the spectral ones transformed back to the physical domain.

11.1 Historical development

Early work on the solution of differential equations using a spectral approach was initially by Slater (1934), Kantorovic (1934), and then by Frazer *et al* (1937). These authors all used a collocation approach, with the latter devising a general scheme for the solution of ordinary differential equations. Lanczos (1938) was the first to formalise the collocation method, by establishing a regular framework by which to correctly ascertain choice of trial functions and collocation points. The method has become popular because of its ease of application to variable-coefficient and non-linear problems. It is also known as the *pseudo-spectral* method after Orszag (1972). Galerkin formulations for the solution of partial differential equations were initially attempted by Silbermann (1954), but Orszag (1970) was the first to produce a practical method of application of such methods, by deriving methods for the convolution sums that arise in the application of this approach to non-linear terms. Orszag (1971b, 1971c, 1971d) demonstrated the application of these methods within simple flow boundaries. These references contain detailed discussion on the accuracy and mathematical properties of the Galerkin approach. Lanczos (1938) was the first to formulate the Tau method. Fox [see Fox (1962), Fox & Parker (1968)] brought a systematic treatment to the principle of the application of the Tau coefficients. Orszag (1971a) used this method to great effect in solving the Orr-Sommerfeld equation for plane-Poiseuille flow, attaining the first authoritative minimum critical Reynolds number of 5 772.23.

¹The latter two are specific examples of sets of polynomials generated as the eigenfunctions of singular Sturm-Liouville problems [see Canuto *et al* (1988)] - any set of eigenfunctions from such a problem would serve as an acceptable orthogonal basis for a spectral approach.

Chapter 11

Numerical methods

The solution of differential equations is necessary in order to simulate events of interest occurring in the outside world. All methods of solution of the governing equations in fluid dynamics have in common that they try to represent (continuum) variations in properties in a finite, *discretised* domain accurately enough to ensure that the finite representation emulates enough features of the real domain to be useful. The sense of the word *useful* may have various levels of connotation, from the scheme merely depicting the coarse features of the flow under investigation, to accurately representing many levels of the real flow's innate behaviour.

The current study is concerned with accurate solutions to the equations from chapter 10. To this end, this chapter will discuss the merits of various numerical approaches, and then outline an applicable method which is both extremely simple to implement and highly accurate for their solution. All numerical approaches attempt to represent a continuum fluid variation in a discretised domain. It is the choice of domain that distinguishes the various classes of solution. The methods can be divided into broad categories. The first set of techniques can be classed under the heading *finite-difference*. As their name implies, they represent the functional values at discrete points in a domain, and the differential operators by *difference formulae* locally between the nodes in this domain. The second general group can be classified as the *Method of Weighted Residuals* or MWR. Within the MWR are two further categories: the *Finite-element* methods and the *Fourier* or *Spectral* methods. The first of these approaches divides the solution space into discrete sub-domains or elements, and in each element represents the continuous function by low order orthogonal linear functions (for instance polynomial or trigonometric) with unknown coefficients. The values of these coefficients for each element is calculated by asserting that some residual or error condition with respect to a suitable weight is minimised for each element. The continuity between elements is enforced by ensuring that the solutions conform on the boundaries between the elements. The Fourier or spectral approach is similar in spirit to the MWR, but whereas the latter

10.5 Summary

The equations applicable for pipe flow undergoing general accelerations were derived from the axisymmetric Navier-Stokes equations. In particular a new parameter Ω emerged from the derivation which described the acceleration. Accelerating pipe entrance flows were thus shown to be described by two parameters: the Reynolds number, and the acceleration parameter. This parameter was identical to Shen's (1961) criterion for instability, and to the parameter α determined by Moss (1991) for laminar flow development between parallel plates.

Various levels of approximation were made to the derived equations. Consideration of the high Reynolds number limit resulted in the elimination of the Reynolds number, analogous to the manner in which da Silva (1990) disregarded the Re term in his vorticity-transport formulation of the steady relations. The use of a pressure-integral formulation allowed the pressure term to be eliminated formally in this high Re limit.

A linear stability derivation for the generally accelerating axisymmetric system was conducted, yielding a stability equation that contained the steady Sexl relation as a zero-acceleration limit. It was shown to be as applicable to the accelerating pipe entrance flow as is the steady Sexl relation to the steady pipe entrance.

This ends the formal derivation of the applicable flow equations. The following chapters are devoted to the development and verification of a numerical method of solution. The solution of the base flow and stability equations are thus left until chapters 15 and 17, where some of the concepts introduced here will finally become clear.

(10.26) and (10.27) instead of the usual Navier-Stokes relations (10.8), (10.9), (10.10). As was shown in chapter 10, the renormalised relations were derived with no restrictive assumptions as to the nature of the pipe flow further than incompressibility, axisymmetry and constant fluid properties. Thus they are applicable for an axisymmetric stability derivation.

- (ii) The derived stability equation is formally applicable only to parallel flows that satisfy the requirement (10.39), that is flows that vary as $u(r,t) = R(r)T(t)$. It is thus the *viscous* equivalent of Shen's (1961) relation. However, it is approximately valid over a far broader range of flows, as will be explained in the point below.
- (iii) The stability equation can be applied approximately to the non-parallel flows occurring in the entrance region of an accelerating pipe flow, in as much as the steady analysis can be applied approximately to the velocity profiles within the non-parallel region of steady pipe flow. This can be argued by considering the behaviour of the exponentially accelerating flow from the perspective of the renormalised framework. An exponential pipe entrance flow will, after sufficient time, reach a point of steady state⁸ behaviour in the renormalised system. This steady state flow field will be shown in a later chapter to possess the same order of non-parallelism as its steady counterpart. Since the stability derivation was performed in the renormalised framework, the current stability equation is as valid as the Seshi equation is valid for steady entrance flows.
- (iv) As a generalisation of the above point, the current analysis is valid to the same extent as the parallel one is, if the respective base flows to which they apply show the same amount of deviation from the set preconditions, as long as that deviation of the accelerating case is compared from within the renormalised framework.

There is much precedent for application of the steady linear analysis within the steady pipe inlet; for instance, Huang & Shen (1974a, b), Moss & da Silva (1993), da Silva & Moss (1994), Abbot & Moss (1994). In particular the last paper formally showed that non-parallel effects were insignificant. Thus the unsteady analysis is as applicable to the unsteady pipe entrance as the steady one is to the steady entrance, because this unsteady base flow is in reality completely steady and sufficiently parallel from the perspective of the renormalised variables.

The above points become more evident if ones predisposition to relate events to the physical system is avoided. The best way to perceive the accelerating system - both base flow and stability - is to consider it as really occurring within the renormalised framework, and that all the unsteadiness is simply associated with the variable projection between it and the physical system.

⁸This will be shown properly when the numerical base flow results are presented later. For now this must remain an assumption, although an obvious one, because the accelerative base flow possesses constant boundary conditions and constant forcing terms for exponential flow, involving a constant limiting behaviour for large times.

The spatial wavenumber α appears in equation (10.53) closely related to the Reynolds number terms. To simplify the equation even further, this wavenumber can be *normalised* using the Reynolds number as

$$\tilde{\alpha} = \frac{\alpha}{Re}. \quad (10.54)$$

The new definition of the wavenumber is a function of Re and thus has time dependence. This however does not lead to any lack of generality in the stability system, because all differentiation was carried out with only the constant α present. Equation (10.54) is simply a convenient redefinition.

Using (10.54), the equation (10.53) can readily be transformed to a form that is an extension to Sexl's stability equation, but applicable to accelerating flows:

$$\frac{-i}{\tilde{\alpha} Re} \left[(L - \alpha^2)^2 + \Omega L \right] \phi = (\tilde{u} - c) [L - \alpha^2] \phi - [L \tilde{u}] \phi. \quad (10.55)$$

For flows with constant mean flow rate, the acceleration parameter Ω is zero, thus equation (10.55) is identical to the Sexl equation in this limit.

The boundary conditions applicable to the steady Sexl equation are also valid here. These are [see Drazin & Reid (1981)]

$$\begin{aligned} \frac{1}{r} \phi(0) &= 0, \\ \frac{1}{r} \phi_r(0) &= 0, \\ \phi(1) &= 0, \\ \phi_r(1) &= 0. \end{aligned} \quad (10.56)$$

The conditions at the pipe wall ($r = 1$) are straightforward. As with the steady system, the centreline constraints assert the uniqueness of the solution at this singular point.

10.4.4 Applicability and limitations

The above derivation is novel and leads to some interesting insights. A number of issues need to be addressed to clarify the validity and range of applicability of the derived equation:

- (i) The derivation is identical in method to any classical approach to linear stability. It differs in its use of the renormalized basic equations (10.25),

is made, a relation equivalent to equation (12.1) in r is obtained according to

$$\phi(r) = \vec{d} \cdot r, \quad (12.7)$$

where $\vec{d}$ is the coefficient vector in the canonical basis. Equation (12.7) is termed the canonical polynomial representation of the function $\phi(r)$.

There is now equation (12.5) representing a transformation from the orthogonal to the canonical basis, and equation (12.6) denoting the relation between the coefficients in both bases. Thus the representation of $\phi(r)$ can be equivalently associated with both equations (12.1) and (12.7).

The matrices and vectors in the above analysis are infinite, because the series expansion equation (12.1) is infinite. Obviously such matrices cannot be treated numerically, and means of dealing with a truncated system will be discussed later. Basically, truncation at a certain order yield an approximate series with the same order. The infiniteness of the operators is not a problem when performing analytic operations, for instance as performed above in equations (12.1) to (12.7), hence detailed discussion of truncation will be reserved for later.

12.2 Operators

12.2.1 Differentiation & polynomial pre-multiplication

The effect of simple operations on the series representation of $\phi(r)$ will now be examined. Consider a general linear operator $\mathcal{P}$ onto ϕ of the form

$$\mathcal{P}\phi(r) = \left[r^m \frac{d^n}{dr^n} \right] \phi(r), \quad m, n \in \mathbb{Z}. \quad (12.8)$$

For any given m, n in $\mathbb{Z}$, the functional $\mathcal{P}$ is a member of the general set of differential operators $\mathcal{Q}$ with polynomial coefficients. Members of $\mathcal{Q}$ can be represented in the canonical basis as

$$\mathcal{P}\phi(r) = dDr = \vec{d}\eta^a\mu^m r \quad (12.9)$$

where the matrices η and μ are of particularly simple form, namely

$$\eta = \begin{bmatrix} 0 & & & \\ 1 & & & \\ & 2 & & \\ & & 3 & \\ & & & \dots \end{bmatrix}, \quad \mu = \begin{bmatrix} 0 & 1 & 0 & & \\ 0 & & 1 & & \\ & & & 1 & \\ & & & & \dots \end{bmatrix}, \quad (12.10)$$

where

$$\vec{r} \equiv [r_0 \ r_1 \ r_2 \ \cdots]^T \quad (12.2)$$

is the suitable linearly independent polynomial basis orthogonal on the range $[0, 1]$ with respect to a suitable weighting function²; and

$$\vec{b} \equiv [b_0 \ b_1 \ b_2 \ \cdots] \quad (12.3)$$

is the vector of unknown expansion coefficients. The *canonical* basis is defined as the vector

$$\vec{r} \equiv [1 \ r \ r^2 \ \cdots]^T. \quad (12.4)$$

Because the elements of (12.2) are polynomial, the basis r can be defined in terms of the canonical basis (12.4) by a lower triangular matrix operator V such that

$$r \equiv V\vec{r} \quad (12.5)$$

The elements in the i 'th row of V are simply the coefficients of the polynomial representation of r_{i-1} - the basis function of order $i-1$ - taken in order from lowest to highest power; that is, if $r_{i-1} \equiv a_0 + a_1 r + a_2 r^2 + \cdots + a_{i-1} r^{i-1}$ is the polynomial representation of r_{i-1} , then the i 'th row of V is simply $[a_0 \ a_1 \ a_2 \ \cdots \ a_{i-1} \ 0 \ 0 \ \cdots]$. If the basis r is linearly independent, then V is invertible, and V then defines a unique mapping between the canonical and orthogonal basis. Dependent on the basis functions used, the relevant form of V can be constructed, but for now, V is left unspecified.

Thus it is possible to reformulate $\psi(r)$ in the canonical basis $\vec{r}$ as $\psi(r) \equiv (\vec{b}^T V) \vec{r}$. If the substitution

$$\vec{a} \equiv (\vec{b}^T V) \quad (12.6)$$

²There is in fact no need for the basis to be orthogonal on the range; it is sufficient that the basis functions be linearly independent. The operational approach will - for any reasonable polynomial basis - ensure uniform convergence to the exact result, not in the usual least-squares sense, but rather in a point-wise manner. The operational method is thus completely independent of the constraints of the MWR; however, for the Chebyshev polynomials, the operational approach produces an equivalent matrix formulation to the more usual MWR approach.

uses a matrix-operator or *operational* approach to the Tau method, which is intuitive and simple to implement. This method by its simplicity allows a ready understanding of the Tau method, by avoiding the MWR type of explanation altogether. However, the basic form of the operational approach as implemented by Ortiz has limitations. Firstly, it has only been used to address Tau systems, even though its range of application is far broader, and secondly, it is severely limited in terms of the order to which solutions can be obtained, and by inference is thus restricted to simpler problems where accuracy can be obtained at low order. These limitations are fortunately not innate, and this chapter extends the operational philosophy¹ in two ways. Firstly, in the context of the Tau method, the order dependent limitation is removed, allowing application to difficult problems. Secondly, the collocation approach is implemented in an operational sense, and its close relationship with the Tau and other methods is shown.

This chapter contains only a description of the basic principles of the method, along with its various components. It presents a set of 'tools', the use of which are demonstrated in subsequent chapters. In chapter 13 which follows, three problems are solved using the operational Tau approach: a differential eigenvalue problem, the Orr-Sommerfeld equation; a simple initial-value problem, the temporal development of pipe flow far from the inlet; and a non-linear ordinary differential equation, the Blasius equation. The collocation discretisation is used in chapter 14 to discretise the axial co-ordinate for the general pipe entrance model.

12.1 Notational description

For the sake of simplicity, this description centres around the determination of an unknown function $\psi \equiv \psi(r)$ on the interval $[0, 1]$. The reason for this interval, rather than the usual $r \in [-1, 1]$ applicable to the Chebyshev or Legendre polynomials, is that the former is the range of the radial co-ordinate in pipe flows (normalised with respect to radius). Only a simple algebraic mapping is needed to project the latter interval onto the former. In fact the Chebyshev polynomials of the second kind, T_n^* [Fox & Parker (1968)] are simply the normal Chebyshev relations projected onto the smaller range. Thus there already are analytic relations defined for these functions, making the derivation easier. The dependent variable $\psi(r)$ in the compact interval $[0, 1]$ is replaced by an infinite series expansion in a linearly independent polynomial basis. As was explained in section 11.3.4, as long as the function is smooth, and the basis functions are correctly chosen, then the series expansion will converge uniformly to the function. In vector notation, such a series expansion can be represented as

$$\psi(r) \approx \hat{b}^T \mathbf{e}, \quad (12.1)$$

¹The term philosophy is used because the approach is far broader than a 'method', and represents a whole way of thinking.

Chapter 12

The operational philosophy

This chapter discusses the derivation of the operational Tau method, illustrating its advantages over the standard Tau approach. Improvements to the operational method of Ortiz [Ortiz & Samara (1981, 1983)] are shown which remove its order-based restriction, resulting in a numerical approach which is as powerful as the standard Tau methodology, but which is straightforward to use. It also shows how the collocation approach can be treated in a uniform manner using the operational philosophy, and gives the interrelationship between the two methods.

The basic principle of the application of the Tau and other spectral methods to the solution of differential equations is very simple. A high order functional series is used as an approximation to the solution of the equation, and differentiation and other operations are implemented as analytic operations on that series. Unfortunately the MWR notation only serves to obscure this simple fact, and the methods are thus avoided by many because of the perceived difficulty of their implementation. A good illustration of this problem with the standard Tau method is shown in its application to the Orr-Sommerfeld equation by Orszag (1971a). While the results he obtained for this problem were exceedingly accurate, his notation is obscure, and to the uninitiated the simplicity of the underlying method is completely hidden. This has resulted in many researchers using inappropriate techniques for fear of not understanding the 'difficult' spectral method. This was discussed in the previous chapter, and the quote from Maslowe (1981) on page 224 still describes the general attitude towards these methods.

Spectral methods are still largely used by applied mathematicians, and as a result the formal notation employed by them is generally inaccessible to engineers and other scientists. For instance, the comprehensive summary of spectral methods given by Canuto *et al* (1988) while authoritative, does nothing to popularise the method. In contrast to the normal approach, that of Ortiz [Ortiz & Samara (1981, 1983), Liu & Ortiz (1982, 1986, 1987)] is unusual in that it

11.5 Summary

The various types of numerical methods of solution to differential equations in general were presented, and the various merits and disadvantages were discussed. All methods were shown to possess a basic commonality, an issue usually avoided in the literature. Here a distinction was made between the discretisation and the numerical solution of these systems.

The spectral approach, showed itself to be the most accurate of the methods described, making it the obvious choice for accurate solution of the equations developed earlier. In the following chapter, the Tau approach of Ortiz [Ortiz (1969), Ortiz & Samara (1981, 1983), Lin & Ortiz (1982, 1986, 1987)] is modified and improved, and its advantages are demonstrated.

The separation of the solution process into two stages is common to all the categories of numerical method. Thus the categorisation of da Silva (1990) - where he distinguishes spectral, initial value, and finite-difference methods as separate but equal entities - is flawed. In reality spectral and finite-difference methods describe the method of *discretisation* (literally, making discrete), whereas the initial value approach is just one means of solution of the resulting matrix system. Another equally valid means of solving both spectral and finite-difference systems would be a direct approach, such as the Q-R or Q-Z algorithm.

The confusion arises from the methods commonly employed for the construction of the numerical system within the computer. Traditionally, finite-difference schemes deal directly with the interaction between nodal values. For instance, a computer algorithm will usually calculate the derivative of the nodal values on a per-point basis, even though the entire operation is representable using the matrix notation. While the matrix nature of some of the proofs and constructions used in spectral methods might give their users insights into this principal of commonality, the perceived unique and obscure nature of these methods to the common numericist results in an ongoing misunderstanding of their simplicity.

Obviously the accuracy of the different methods is profoundly related to the type of discretisation. While different approaches might be combining and manipulating matrices in the same way, the actual numbers in the respective matrices has a vast effect on the resulting solution. With regard to accurately representing a smooth differential system, the spectral methods are simply far more efficient than the finite-difference approaches. This is achieved by judicious choice of interpolating function or placement of grid points, and use of accurate operator matrices. The analogy is the same here as between the convergence of, for instance, a Taylor series as opposed to a Fourier series. While both may be evaluated to the same number of terms, the latter is far more efficient at describing the variation of a function while using many fewer degrees of freedom than the former.

In some three-dimensional systems, for instance the direct solution of the full Navier-Stokes equations, the unknowns usually form a *cube* of numbers, and simple (i.e. two-dimensional) matrix algebra does not apply directly. However the argument is still completely valid. Any form of discretisation will result in the same form of discrete (although not matrix) system to be solved.

ical values of the approximated function at positions in the solution domain. In the Fourier type systems (finite-element, Tau, etc.) these discrete values are more abstract entities. Here they are the coefficients of an interpolating series, for instance a Fourier or polynomial series. In either case, in the *computer* representation of these unknowns is in the form of an array of numbers or matrix.

Operations in the differential equations, for instance differentiation or multiplication, are represented simply as operations of these unknown values. These operations can be described using the rich language of the matrix algebra. Effectively the commonality between all the methods is the common usage of matrix manipulation to solve the resulting algebraic systems. There are various kinds of analytic equation; broadly speaking, initial value, boundary value and characteristic value (eigenvalue) systems. The nature of each type of system is the only factor that effects the structure of the resulting numerical system, not the numerical method employed. For instance a differential eigenvalue problem would result in the solution of the discrete system

$$Ax = \lambda Bx, \quad (11.10)$$

where λ is the characteristic value, x is the eigenvector, and the matrices A and B vary according to the numerical formulation used. The size of the matrices determines the order of solution.

All numerical methods of formulation on an identical analytic eigenvalue problem will lead to a matrix formulation identical to (11.10) above. In some cases this may be explicit [i.e. the numerical system that is solved is (11.10)], while other methods of solution might implicitly reformulate this equation into another form that allows simpler solution. The shooting method employed by Ng & Reid (1979) or Gersting (1980, 1987) is no more than this equation rearranged as an initial-value problem. The shooting method is thus not specific to the finite difference (or other) method, rather it is simply a particular way of solving the matrix problem. The separation of the numerical formulation and solution is shown in figure 11.2 below.




Figure 11.2: The process of numerical formulation and numerical solution.

more common Sturm-Liouville problems are in fact polynomials. Examples of such polynomials are the Chebyshev and Legendre polynomials. Both allow a simple extension of the behaviour of the Fourier series to non-periodic smooth functions.

The use of any arbitrary set of orthogonal functions will not guarantee spectral convergence for any general problem. Amongst others, Davey & Drazin's (1968) choice of basis functions was inappropriate and did not allow spectral convergence [Orszag (1971a)]. Furthermore, the type of problem to be modelled must match the method of solution; for instance, a Fourier expansion is ineffective when used on a non-periodic problem, and all spectral methods' convergent behaviour reduces to algebraic near a discontinuity - due to the Gibbs phenomenon [Canuto *et al* (1988)]. Thus, unlike a finite-difference method which, although inaccurate, is generally robust, analytic properties of the system can profoundly affect a spectral method, rendering it no better than the former in accuracy. However, the class of solutions sought in incompressible fluid flow are generally smooth. Therefore as long as due care is taken in their formulation, spectral methods for such problems are usually straightforward. The argument that the spectral method's difficulty in implementation should somehow be traded for a finite-difference scheme's ease of implementation [Maslowe (1981), da Silva (1990)] is wholly without truth. Firstly, spectral methods are **not** hard to implement⁸, and secondly, the level of accuracy attainable from a spectral *a priori* reach is so great⁹ that such comparisons are rendered fallacious. However, finite-difference methods are still optimum in certain classes of system such as shock-wave phenomena, where the discontinuities in the solution domain render the spectral approach non-optimal.

11.4 Commonality of solution procedure

Traditionally, the different numerical schools (finite-element, finite-difference, spectral, etc.) are presented as being vastly disparate in formulation *and* numerical solution. While the conceptual formulation of each method might be rather different, the numerical solution of the resulting *numerical system* is in each case much the same. This commonality is usually hidden behind the nomenclature (or in some cases jargon) that each philosophy uses.

All numerical methods require the determination of a set of discrete unknowns. In a finite difference or collocation method, these unknowns will be direct physics

⁸The operational Tau method, explained in section 12, is as intuitive as the finite-difference philosophy, and much easier to implement.

⁹The infinite-order accuracy attainable with spectral methods allows one to solve certain problems to any level of accuracy required, (limited by machine precision). For instance, a spectral solution to a linear ordinary differential equation can be made to approach the *analytic* solution to as many significant figures as is desired. In contradistinction, in a finite-difference formulation, it is not possible to specify some arbitrary precision and then attain it - the order of the system necessary to do this might be so large as to render it impossible within a lifetime!

around the perimeter. A general approach for the two dimensional method of solution is contained in Canuto *et al* (1988). This description shows how to satisfy piece-wise continuous boundary conditions around the perimeter, but is very theoretical and rather difficult to comprehend. Unfortunately they give no example of the more general method.

The standard method of formulation for a Tau implementation is to develop the equations analytically, directly from the MWR condition equation (11.6). This can be tedious for any but the simplest problems [the derivations of Orszag (1971a) are a perfect example of the mathematical complication involved]. Even the current main-stream thinking on this matter [see Canuto *et al* (1988)] is little changed from the original approach of Lauzeos (1938) and the formalisms of Fox & Parker (1968). As will be shown later, there exists an ingenious alternate method of derivation for this approach which largely removes the mathematical obfuscation associated with the standard form above. It is unfortunate that the perceived mathematical complexities of the method have discouraged many researchers from adopting this extremely simple yet highly accurate approach. For instance Maslowe (1981) makes the following statement:

"There are some rather involved manipulations required in the Chebyshev-Tau method, so this method should only be used if great accuracy is required and/or high values of Re are to be considered". (Italics present author)

Statements such as this are highly dissuasive to new researchers in the field.

11.3.4 Accuracy

The accuracy and convergence of a well formulated spectral approach is related to the convergent behaviour with order of an interpolating Fourier series to a smooth periodic function. That is to say, if a function $u(x)$ in the range $[x_1, x_2]$ is periodic and infinitely differentiable with respect to x , then the j -th Fourier coefficient of this function's Fourier series will decay faster than any power of j [see Orszag (1971a), and the corollary in Canuto *et al* (1988), p36]. Thus the Fourier series to $u(x)$, truncated at order N , will converge point-wise to the function $u(x)$ at a rate greater than any power of N . This behaviour is what is termed *infinite-order accuracy* by Orszag (1971a).

The convergence of a Fourier series to a non-periodic smooth function cannot however be guaranteed to be of infinite-order at the end points. This is because the symmetry embedded in the Fourier series does not allow for different end values in the range. However, certain classes of orthogonal polynomials when used as an interpolating series for a non-periodic, but smooth function, will guarantee spectral accuracy *no matter what the end conditions*. A singular Sturm-Liouville operator can serve as a generator of such functions - its eigenvalues will serve as suitable candidates. Moreover, the eigenfunctions of the

11.3.3 Tau

The Tau approach is similar to the Galerkin method, but here the trial and test functions need not be identical. The test functions do not necessarily satisfy the boundary conditions, and the auxiliary conditions must be applied separately.

In order to incorporate the boundary relations, some of the higher order inner equations are dropped in their favour. It can be shown that this is equivalent to perturbing the original equation by a polynomial function with a number of unknown coefficients τ_i . The discarded relations can be used to determine the value of these coefficients. Due to the introduction of these perturbations, the problem is not over-specified by the application of the boundary conditions. As can be seen, the name of the method stems from the denotation used for the perturbing coefficients.

It is necessary to perturb the differential equation under investigation, because a truncated series such as equation (11.5) will not necessarily satisfy this equation exactly and the boundary conditions simultaneously. However, if the series contains enough terms, these perturbing values will be very small indeed, indicating that the series has converged close to the correct result. In fact the perturbation coefficients τ_i can be used as an estimate of the error in approximation. Fox & Parker (1968) contains a comprehensive description of the Tau method.

Since they are implemented functionally, the approach to boundary conditions in this method allows great flexibility. Thus it is an ideal choice for problems in which there are unusual constraints. There is no real limitation on the type of boundary condition that can be applied, save that they can be represented as an analytic function. Thus integral, derivative and other conditions can be represented properly with great ease, and in a completely uniform manner. Furthermore, there is no restriction as there was in the collocation approach, concerning multiple boundary conditions at a single point. All boundary conditions are alike in that they are linear operators onto the unknown coefficients. Thus in general any boundary operator that can be represented as

$$f(u)|_{x_a} = 0, \quad (11.9)$$

can be uniformly incorporated into a Tau approach. Here f represents a general operator onto the unknown function $u = u(x)$, evaluated at a point in the range $x = x_a$.

The boundary implementation, although highly flexible in one-dimensional Tau problems, becomes virtually impossible if a two-dimensional formulation is used. This is primarily due to incompatibility of these conditions when applied in the two co-ordinate directions. Thus two-dimensional Tau discretisations are rare, and are usually only seen in the solution of simple relations like the Poisson equation on the square, and where the boundary conditions are zero everywhere

- (iii) Multiple simple boundary conditions at a single point in the domain pose a problem. For instance, the Orr-Sommerfeld equation for plane-Poiseuille flow requires four boundary conditions - two at each end of the domain - for a proper solution. Thus a collocation method has to assert both $\phi|_{y_{j+1}} = 0$ and $d\phi/dy|_{y_{j+1}} = 0$ at the same node. For such a situation, the algebraic equations at some inner nodes have to be removed to make 'space' in the operator matrix for these extra conditions. This is inelegant and physically incorrect. Canuto *et al* (1988) describes how to avoid problems with this particular formulation.
- (iv) The above limitations are perhaps offset by the advantages of a collocation approach in a two- or three-dimensional domain. Here the flexibility in the application of boundary conditions allows use of non piecewise-continuous conditions around the edges of the domain⁶.

11.3.2 Galerkin

In the Galerkin approach, the test and trial functions are both global smooth functions of similar form, and are chosen so as to satisfy the boundary conditions automatically. Thus there is no need to apply the auxiliary conditions. As a result, such schemes can only be used in systems where boundary conditions are both simple and 'periodic'⁷, and they are thus highly restrictive in terms of their applicability. Typical systems where such methods can be applied are amongst others: convective flow in a closed box, and the transverse (y) direction discretisation of flow between parallel plates.

The most common orthogonal basis for spectral-Galerkin methods in fluid mechanics is the Fourier series, or its sine and cosine counterparts. This allows use of a lot of available - and highly optimised - numerical machinery, such as fast Fourier transform algorithms. It is perhaps the familiarity with the Fourier series that leads many people to utilise this approach, despite the severe restrictions with boundary conditions entailed by its adoption. Both the Collocation and Tau approaches are, despite their respective limitations, far more flexible than the Galerkin method in this regard. Galerkin approaches are however extensively used in finite-element systems [see Zienkiewicz (1986)], it seems mainly for historical reasons.

⁶In the two-dimensional pipe-inlet flow problem of chapter 10, a discontinuity exists at the pipe wall at the inlet. This is because the inlet flow is constrained to be uniform across the pipe inlet, which violates the no-slip condition at the wall here. This dilemma led to the use of a mixed Γ -collocation solution procedure for this system.

⁷By periodic is meant that the conditions on either side of the domain are identical.

of the interpolating function at points (the collocation nodes) in the domain. The distinguishing feature about this approach is the structure of the operators. Here, instead of the first derivative operator, for instance, being derived locally as in a finite-difference scheme, it is constructed as an analytic function under the MWR approach, and then *transformed* into the physical domain. The result of this is that such operators are global (full matrices).

The transformation between the physical and spectral domain is - dependent on the nature of the trial functions - usually a Fourier or similar transform. Thus, unlike a finite-difference approach, the collocation method is closely related to, and indeed equivalent to, the Galerkin and Tau methods (these two methods are explained below). The coefficients of a Galerkin (or Tau) problem can be transformed via a linear matrix-operator (or, even quicker, by a fast Fourier transform) to the space-values for an equivalent collocation set-up, and vice versa. This flexibility is extremely useful when considering the non-linear terms in a Tau method. These involve performing a (numerically expensive) convolution in the orthogonal basis; but it is substantially quicker to transform these terms into the physical space, perform a term-wise multiplication, and then transform them back again.

Alternatively, the collocation values can be thought of as proper coefficients onto the Lagrange polynomials, that is polynomials of the form

$$P_k = \frac{\prod_{i \neq k}^{N-1} (x - x_i)}{(x_k - x_i)} \quad (11.8)$$

which are identically zero at all the collocation points x_i except the node x_k . An alternative way of deriving differential and other operators for the collocation method is to consider the analytic behaviour of the polynomial (11.8) under differentiation and other operations.

The similarity of the collocation approach to a finite-difference method allows the same philosophy to be adopted regarding the application of boundary conditions. Thus the collocation approach shows similar advantages (and limitations) as does a finite-difference one:

- (i) Simple Dirichlet boundary conditions can be enforced directly at the node of preference, usually by replacing the relevant row or column in the operator matrix by the boundary condition (implicit solution), or simply by setting the relevant nodes to the applicable values (explicit).
- (ii) More involved conditions (Neumann, Robin or other) may require complicated manipulation of boundary nodes. Canuto *et al* (1988) contains a detailed discussion of how to deal with such conditions although no single uniform approach seems to exist.

and using (12.44), the inverse is

$$V^{-1} = \begin{bmatrix} 1 & & & & \cdots \\ 0 & 1 & & & \\ \frac{1}{2} & 0 & \frac{1}{2} & & \\ 0 & \frac{3}{4} & 0 & \frac{1}{4} & \\ \frac{3}{8} & 0 & \frac{1}{2} & 0 & \frac{1}{8} \\ \vdots & & & & \ddots \end{bmatrix}. \quad (12.46)$$

Under the normal MWR formulation using Chebyshev polynomials, the test functions are $w_k(x) \equiv 1/(1-x^2)^{1/2}$ to ensure the orthogonality condition (11.7) in the range $[-1, 1]$. The operational-Tan method never deals explicitly with test functions, so no further consideration need be given to them.

12.6.3 Chebyshev polynomials in $[0, 1]$

The Chebyshev polynomials in the restricted range $r \in [0, 1]$, termed the Chebyshev polynomials of the second kind [Fox & Parker (1968)], can be easily generated from the algorithms (12.1) and (12.44) via substitution of the algebraic relation $r \equiv \frac{1}{2}(x+1)$. These are the basis polynomials that will be used in all the numerical solutions in this work. The Lanczos notation of using an asterisk (*) to denote these polynomials will be used here.

They are given recursively by

$$\begin{aligned} T_0^*(r) &= 1, \\ T_1^*(r) &= 2r - 1, \\ T_{i+1}^*(r) &= 2(2r-1)T_i^*(r) - T_{i-1}^*(r), \quad i = 1, 2, \dots \end{aligned} \quad (12.47)$$

and the first few terms are

$$\begin{aligned} T_0^*(r) &= 1 \\ T_1^*(r) &= 2r - 1 \\ T_2^*(r) &= 8r^2 - 8r + 1 \\ T_3^*(r) &= 32r^3 - 48r^2 + 18r - 1 \\ &\text{etc.} \end{aligned} \quad (12.48)$$

The forward mapping V^* is given by the algorithm (12.2) below.

From (12.10) it is possible to formulate a trivial recursive rule to derive the infinite matrix V . If the elements of V are denoted as $v_{i,j}$, then V can be generated by algorithm (12.1) which follows:

Algorithm 12.1: V

```

 $v_{0,0} \leftarrow 1$ 
 $v_{1,1} \leftarrow 1$ 
for  $i = 2 : \infty$  {
     $v_{i,0} \leftarrow -v_{i-2,0}$ 
    for  $j = 1 : i$  {
         $v_{i,j} \leftarrow -v_{i-2,j} + 2v_{i-1,j-1}$ 
    }
}

```

The generation of V^{-1} is possible in two ways. Firstly, it can be inverted to yield V^{-1} . However it is possible to define V^{-1} *non-recursively* [Luke (1968a), p 298, equation² (22)] as

$$V^{-1} : v_{i,j-2k} = \frac{1}{2^k} \delta_{i-2k} \binom{i}{k}, k = 0, \dots, [i/2], \quad (12.11)$$

where

$$\delta_0 = \frac{1}{2}, \quad \delta_i = 1, \quad i \geq 1.$$

Generation of the leading terms of the infinite matrices using the algorithm (12.1) gives

$$V = \begin{bmatrix} 1 & & & \dots \\ 0 & 1 & & \\ -1 & 0 & 2 & \\ 0 & -3 & 0 & 4 \\ 1 & 0 & 8 & 0 & 8 \\ \vdots & & & & \ddots \end{bmatrix} \quad (12.12)$$

²A modification is made for the leading term in Luke's formulation. The coefficient of the δ_0 term must be multiplied by $\frac{1}{2}$. This is due to the fashion of taking the first Chebyshev coefficient with a factor 2.

12.6.2 Chebyshev polynomials in $[-1, 1]$

The set of polynomials $T_n(x)$, $x \in [-1, 1]$ are termed the Chebyshev polynomials. They can be represented explicitly as

$$\begin{aligned} T_n(x) &= \cos(n\theta) \\ \theta &= \cos^{-1}(x). \end{aligned} \quad (12.39)$$

An extensive study has been made on the properties of these functions [see Fox & Parker (1968), Rivlin (1990)], which are simply trigonometric relations with a change of independent variable. Thus they share the same convergent behaviour as a Fourier series, along with the property of infinite differentiability and minimax behaviour. They can be evaluated recursively as follows:

$$\begin{aligned} T_0(x) &= 1, \\ T_1(x) &= x, \\ T_{i+1}(x) &= 2xT_i(x) - T_{i-1}(x), \quad i = 1, 2, \dots \end{aligned} \quad (12.40)$$

Thus the first few terms can be simply evaluated as:

$$\begin{aligned} T_0(x) &= 1 \\ T_1(x) &= x \\ T_2(x) &= 2x^2 - 1 \\ T_3(x) &= 4x^3 - 3x \\ T_4(x) &= 8x^4 - 8x^2 + 1 \\ &\text{etc.} \end{aligned} \quad (12.41)$$

There are many variations of recursion formula applicable to these functions [See Luke (1969a, 1969b), Fox & Parker (1968), Rivlin (1990)], that are useful in the derivation of operators in the operational-Tau approach. A singularly useful set of results which will be used in the development of the numerical framework, is the inverse formulation of (12.41), but firstly, it is useful to consider (12.41) in vector notation

$$e = Ux, \quad x = [1 \ x \ x^2 \ \dots]^T, \quad e = [T_0 \ T_1 \ \dots]^T, \quad (12.42)$$

and the inverse relationship for x as a function of T

$$x = V^{-1}e. \quad (12.43)$$

12.6.1 Legendre polynomials

These polynomials are explicitly defined in $[-1, 1]$ as

$$L_k(x) = \frac{1}{2^k k!} \sum_{l=0}^{\lfloor k/2 \rfloor} (-1)^l \binom{k}{l} \binom{2k-2l}{k} x^{k-2l}. \quad (12.35)$$

With $L_0(x) = 1$ and $L_1(x) = x$, the series can be defined recursively as

$$L_{k+1}(x) = \frac{2k+1}{k+1} x L_k(x) - \frac{k}{k+1} L_{k-1}(x). \quad (12.36)$$

Thus, using either equation (12.35) or (12.36), it is elementary to derive either an explicit or recursive algorithm for assembling an infinite matrix V , such that

$$v = Vx; \quad v = [L_0 \ L_1 \ \cdots]^T, \quad x = [1 \ x \ x^2 \ \cdots]^T. \quad (12.37)$$

For the Legendre system, V defines a unique mapping; thus V is invertible, and an inverse transform is given by:

$$x = V^{-1}v. \quad (12.38)$$

The matrix V^{-1} can be generated either by direct inversion of V , or by making use of the identities contained in Szegő (1939).

The Legendre polynomials have potentially attractive properties with respect to their use in a Tan formulation; namely that they show their maximum amplitude in the centre of the range. This makes their error distribution largest near the centre of the range, with the lowest error density near the boundaries, the areas where better resolution is required in most pipe stability or boundary layer calculations. This can be contrasted to the Chebyshev polynomials described in the next section, where the error is evenly distributed over the range [Wright (1993)]. The accuracy gain is however insignificant in practice, because the Chebyshev solution can just be taken to an extra term or so, to better the accuracy of an equivalent Legendre formulation. In other words, any small accuracy gain due to the use of a Legendre basis over a Chebyshev one is swamped by the *rate* of increase in accuracy which both show when the order of solution is increased. Thus both approaches are effectively equivalent in terms of accuracy.

is important. When seeking a solution, one is not looking for the solution to the truncated series (in which case one would assemble matrices with initial dimension $N \times N$), but rather the projection of the infinite series solution onto the subspace $\mathcal{P}_{N-1}$. The resulting solution is guaranteed to converge spectrally above certain N (once the essential features of the problem are properly represented) as long as the function is smooth and the system is complete in a suitable Hilbert space. These two conditions are met by all reasonable systems in incompressible fluid dynamics, and certainly by all the problems considered in this work.

The augmentation carried out in equation (12.34) is equivalent to ignoring the highest order inner equations in favour of the auxiliary conditions. As mentioned before, this approach is valid (and indeed is the essence of the Tau method) because there generally is no exact solution in the truncated basis ϕ_{N-1} to the original equation (12.26) which also satisfies the auxiliary conditions (12.30). In order to make the series representation an exact solution, the system (12.26) is perturbed by a high order polynomial with (a small number of) coefficients τ_i . That is, a solution is sought to $\mathcal{L}\phi + \mathcal{T} = 0$. The perturbation $\mathcal{T} = \sum_{i=1}^{n+d} \tau_i \phi_i$ is such that the value of d is small. The solution for the values of the unknown τ_i 's can be readily obtained by utilising the disregarded relations from the inner problem (The augmentation carried out in equation (12.34) effectively removed ν relations from the matrix by 'pushing' them to the right). Knowledge of the values of these coefficients can be useful because they give an explicit representation of the error of the approximation. However, this is unnecessary because the magnitude of the highest order coefficients of the series solution will supply as good an error estimate. The above perturbation approach is fully implicit in the operational-Tau approach, and is considered here only for completeness. In fact there is no straightforward way in which the τ coefficients can be extracted from the operational-Tau method.

12.6 Suitable sets of polynomials

The above discussion has proceeded in a general fashion, with no specific reference to the polynomial basis. This section considers applicable candidates for such a basis. Two polynomial bases commonly used in spectral methods are considered - the Chebyshev polynomials and the Legendre polynomials. Both sets of polynomials can be generated from specific classes of singular Sturm-Liouville operator [see Canuto *et al* (1988)]. Both are considered in the range $x \in [-1, 1]$. However, and as will be demonstrated for the Chebyshev polynomials, it is trivial to convert both sets of polynomial to the restricted range $r \in [0, 1]$ by the algebraic transformation $r = 1/2(x + 1)$.

where g_i is a known column vector evaluated as $g_i = \tilde{F}_i^T c|_{r=r_i}$. Equation (12.31) can be rewritten as

$$\tilde{b} \dot{q}_i = \tilde{c}_i, \quad i = 0, 1, \dots, \mu - 1, \quad (12.32)$$

The μ separate conditions can now be combined into a single relation by augmenting the separate vectors g_i into a matrix as

$$\tilde{b} \dot{q} = \tilde{k}, \quad \dot{q} = [\dot{q}_0, \dot{q}_1, \dots, \dot{q}_{\mu-1}]^T, \quad \tilde{k} = [\tilde{c}_0, \tilde{c}_1, \dots, \tilde{c}_{\mu-1}]. \quad (12.33)$$

The equation (12.33) is now in an analogous form to the inner problem, equation (12.29). The boundary conditions can now be augmented to the inner equation, and the system to be solved now becomes

$$\tilde{b} \dot{q} = \tilde{\psi}, \quad \dot{q} = [\dot{q}; \dot{g}], \quad \tilde{\psi} = [\tilde{k}; \tilde{g}]. \quad (12.34)$$

12.5 Solution procedure

The equation (12.34) from the previous section is an algebraic system in the orthogonal basis c analogous to the original differential equation. It is however an infinite system, because the unknown function ϕ is replaced by its infinite series equivalent. As was mentioned before, provided the function ϕ is smooth, and so are all its derivatives, then its series expansion in the basis c will converge everywhere point-wise to the actual function at a rate greater than any power of N , as N - the number of terms in the series - is increased. That is to say, the infinite series representation of the differential equation - equation (12.34) - will be equivalent to the original equation. Furthermore, due to the infinite-order accuracy of the series representation, the series expansion truncated at some finite N will be an exceedingly good approximation to the original function - the values of the higher coefficients ignored by this truncation will be very small. As computer resources are finite, and because the solution of the infinite system holds no real benefits in terms of accuracy, the solution is now sought to the system truncated at some finite value.

A solution is sought to (12.34) in the subspace P_{N-1} of polynomials of order $N - 1$, or in other words by representing the solution in the truncated basis c_{N-1} . This is achieved by truncating both the operator matrices $\tilde{b}$, $\tilde{\psi}$ at $N \times N$ and at N respectively. As mentioned before, the derivations concerned *infinite* matrices, thus to achieve the correct formulation, the final system (12.34) must be truncated *after* assembly from infinite operator matrices⁴. This issue

⁴Practically, one needs only work during construction with matrices of order $N + k \times N + k$ where k equals or exceeds the power of the largest polynomial pre-multiplier onto the function.

easily achieved through cancelling τ from each side of the equation (12.28) by division. This step is equivalent to the equating of coefficients in the standard Tau approach, and consequently the inner problem becomes

$$\tilde{h}\hat{G} = \tilde{g} \quad , \quad (12.29)$$

Equation (12.29) is soluble by inversion of, and post-multiplication by $\hat{G}$. However, no boundary conditions yet have been augmented, and thus in all likelihood, $\hat{G}$ will be singular. Further to this problem, the matrices at this stage are all assumed to be infinite in extent; and for this reason the seeking of a solution is a little premature. However, equation (12.29) is in the correct form to have boundary conditions augmented - the subject of the next section.

12.4 Boundary conditions

The solution for a particular differential equation such as equation (12.26) needs the implementation of boundary conditions in order to yield a unique solution. Such conditions generally take the form of linear functionals imposed at some fixed point in the domain, namely

$$P_i \psi(r)|_{r=r_i} = b_i, \quad i = 0, 1, \dots, \nu - 1 \quad , \quad (12.30)$$

Here P_i is some linear functional; for example such an operator could be $P_0 := r \frac{d^2}{dr^2} + \frac{d}{dr}$. The result of this operator on ψ is evaluated at some point r_i , usually at one end of the domain, but generally anywhere within it. The term b_i is some constant. The ν boundary operators P_i can be composed of members of the set $\mathcal{Q}$ of differential operators.

Using the same approach as before, equation (12.30) can be reformulated in the orthogonal basis as

$$\tilde{h}\tilde{P}_i \left[\psi|_{r=r_i} \right] = b_i, \quad i = 0, 1, \dots, \nu - 1 \quad , \quad (12.31)$$

At this point a slightly different approach is taken. The boundary operators need to be augmented to the system (12.29), which is the inner problem *after* the basis vectors have been removed. Thus equation (12.31) must *not* be reformulated in terms of the basis vectors. This is the simpler option, because this equation is already in this form (this is because the basis vectors have been re-evaluated at a point). The left-hand side of equation (12.31) can be simplified to

$$\tilde{h}\tilde{P}_i \psi|_{r=r_i} = \tilde{b}_i.$$

transformed operator $\tilde{J}$ is simply

$$\tilde{J} = \begin{bmatrix} \int_0^1 r_0 dr & 0 & 0 & \\ \int_0^1 r_1 dr & 0 & 0 & \cdots \\ \int_0^1 r_2 dr & 0 & 0 & \\ \vdots & & & \end{bmatrix}. \quad (12.25)$$

The values for the matrix (12.25) can easily be calculated dependent on the basis used.

12.3 The inner problem

It has been shown above that a generalised linear differential equation can be represented as an equivalent set of infinite matrix operators in an orthogonal polynomial basis.

Any linear ordinary differential equation can be written in operator notation as

$$\mathcal{G}\phi(r) = f(r) \quad (12.26)$$

where $\mathcal{G}$ is any linear combination of operators which are members of $\mathcal{Q}$, or operators representable within the system such as the limit operator (12.15) or the integral operator (12.21). The term $f(r)$ is some known (polynomial) function of r (it may be the solution from the previous time-step in an initial value problem). As such, $f(r)$ can be represented in the canonical basis as $f(r) = \tilde{f}r$ where $\tilde{f}$ is the coefficient vector. The function $f(r)$ can be easily projected into the basis e by the mapping V as

$$\begin{aligned} f(r) &= \tilde{f}[V^{-1}r] \\ &= \tilde{f}e. \end{aligned} \quad (12.27)$$

Thus, replacing ϕ in equation (12.26) with its series representation, and the operators in $\mathcal{G}$ by their operational-Tau analogues, results in

$$\tilde{h}\tilde{\mathcal{G}}e = \tilde{f}e, \quad (12.28)$$

The object of the endeavour is to obtain a solution for the unknown coefficients $\tilde{h}$. In order to do this, the coefficient vector e must be eliminated. This is

12.2.3 The definite integral

The definite integral

$$\int_1^1 \phi(r) dr \quad (12.21)$$

is much like the limit operator (12.15) in that its result is not a function of r . It can however be treated in an identical manner to the limit operator.

In order to find the form of the integral operator matrix, consider the integral of the expanded series, equation (12.7). This can be implemented as

$$\begin{aligned} \int_0^1 \tilde{a} r dr &= \int_0^1 (a_0 + a_1 r + a_2 r^2 + a_3 r^3 + \dots) dr \\ &= \left[.5r + \frac{1}{2}a_1 r^2 + \frac{1}{3}a_2 r^3 + \frac{1}{4}a_3 r^4 + \dots \right]_0^1 \\ &= [a_0 \ a_1 \ a_2 \ a_3 \ \dots] \left[1 \ \frac{1}{2} \ \frac{1}{3} \ \frac{1}{4} \ \dots \right]^T \\ &= \tilde{a} \left[1 \ \frac{1}{2} \ \frac{1}{3} \ \frac{1}{4} \ \dots \right]^T. \end{aligned} \quad (12.22)$$

Thus using the same approach as for the limit operator, equation (12.17) is post-multiplied onto equation (12.22), in order to render the format of the integral operator equivalent to the operators from the set $\mathcal{Q}$.

This results in the definite integral being defined as

$$\int_0^1 \phi(r) dr \Rightarrow \tilde{a} J r \quad (12.23)$$

where J is given by

$$J = \begin{bmatrix} 1 & 0 & 0 & 0 & \dots \\ \frac{1}{2} & 0 & 0 & 0 & \dots \\ \frac{1}{3} & 0 & 0 & 0 & \dots \\ \frac{1}{4} & 0 & 0 & 0 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}. \quad (12.24)$$

By using the identity (12.11), J can be projected into the orthogonal basis. Again, if the orthogonal basis has the zero'th order polynomial $r_0 \equiv 1$, then the

In the canonical basis, the operator $\mathcal{L}$ can be written as

$$\vec{a} \cdot \vec{r}|_{r=r_0}, \quad (12.16)$$

Equation (12.16) has no basis vector (because this vector has been evaluated at a point in the region) and thus this form cannot be treated uniformly along with terms such as the elements of the set $\mathcal{Q}$. This is easily rectified by post-multiplying equation (12.16) by the quantity

$$\begin{aligned} & [1 \ 0 \ 0 \ \dots] r \\ = & [1 \ 0 \ 0 \ \dots] [1 \ r \ r^2 \ \dots]^T \\ = & 1 \end{aligned} \quad (12.17)$$

This results in

$$\begin{aligned} \mathcal{L} \phi & \Rightarrow \vec{a} \left[\vec{r}|_{r=r_0} [1 \ 0 \ 0 \ \dots] \right] r \\ & = \vec{a} \left[[1 \ r|_{r=r_0} \ r^2|_{r=r_0} \ \dots]^T [1 \ 0 \ 0 \ \dots] \right] r \\ & = \vec{a} L r \end{aligned} \quad (12.18)$$

where the matrix L consists of zeroes except for the first column which is $r|_{r=r_0}$. For example, if $r_0 = 1$, then L will evaluate to

$$L_1 = \begin{bmatrix} 1 & 0 & 0 & \dots \\ 1 & 0 & 0 & \dots \\ 1 & 0 & 0 & \dots \\ \vdots & & & \end{bmatrix}. \quad (12.19)$$

By using the identity (12.11), L can be projected into the orthogonal basis. It is however trivial to show that L , for an orthogonal basis where the zero'th order polynomial $r_0 \equiv 1$, is simply³

$$\hat{L}_{r_0} = \begin{bmatrix} r_0|_{r=r_0} & 0 & 0 & \dots \\ r_1|_{r=r_0} & 0 & 0 & \dots \\ r_2|_{r=r_0} & 0 & 0 & \dots \\ \vdots & & & \end{bmatrix}. \quad (12.20)$$

As will be shown later, assembly directly in the orthogonal basis is important numerically.

³All reasonable linearly independent bases - certainly both the Chebyshev and Legendre series - have as their zero'th order function $r_0 \equiv 1$.

$$\begin{aligned}
\hat{D} &= VDV^{-1} \\
&= V\eta^u\mu^mV^{-1} \\
&= \underbrace{(V\eta V^{-1})(V\eta V^{-1})\cdots(V\eta V^{-1})}_{n \text{ terms}} \underbrace{(V\mu V^{-1})(V\mu V^{-1})\cdots(V\mu V^{-1})}_{m \text{ terms}} \\
&= \hat{\eta}^n\hat{\mu}^m.
\end{aligned}$$

Thus by theorem 12.1, any general linear differential equation with operators elements of $\mathcal{Q}$, given by

$$[\mathcal{D}_0 + \mathcal{D}_1 + \cdots + \mathcal{D}_N]\phi(r) = f(r), \quad (12.13)$$

where $f(r)$ is some general (polynomial) function of r , can be represented in the basis $\bar{v}$ as the system

$$\bar{b}[\hat{D}_0 + \hat{D}_1 + \cdots + \hat{D}_N]\bar{v} = \hat{f}\bar{v}, \quad (12.14)$$

where $\hat{f}$ represents the known coefficients of the applied polynomial function f , in the basis $\bar{v}$.

As can be seen, any such differential equation can be readily formulated as an operational Tau method, simply by assembling the matrices $\hat{\eta}$ and $\hat{\mu}$ in accordance to the operator notation of the original differential equation.

The application of the method is by no means limited to members of $\mathcal{Q}$. Unusual operators such as definite and indefinite integrals, and limit-operators are some of many linear operators that can be subjected to a uniform treatment under the method. The two types of operator that will be treated here are the limit operator and the definite integral operator - because of their application to the current problem. It is elementary to extend the principal to other unusual operators.

12.2.2 The limit operator

Consider an operator on ϕ of the form

$$\mathcal{L}\phi(r) = \phi|_{r=r_0} \quad (12.15)$$

Such an operator occurs within the equation (10.33) for the solution of the base flow. As will be shown, the current approach will treat such an operator in a uniform way, *even though the result is not a function of r* .

and represent the effects on the coefficients of a polynomial due to differentiation and pre-multiplication by r , respectively.

The matrices (12.10) can easily be generated by considering the effects of differentiation and multiplication on the series expansion of equation (12.7) $\phi(r) = \sum_{i=0}^{\infty} a_i r^i$. Important to note is the order of the operators η and μ in equation (12.9): these matrices operate on the coefficient vector $\vec{a}$, therefore, because differentiation precedes multiplication in equation (12.8), the derivative matrix η must post-operate on $\vec{a}$ first in this case.

The operators η and μ are now projected into the orthogonal basis. By virtue of the mapping V , equation (12.9) can be represented in the basis $\vec{e}$. A similarity transformation is defined such that for any matrix A ,

$$\hat{A} \equiv VAV^{-1} \quad (12.11)$$

is termed the conjugate of A in $\vec{e}$.

Theorem 12.1 *If $\mathcal{D}$ is a member of the set of differential operators $\mathcal{Q}$ with polynomial coefficients, defined by equation (12.8), then in the orthogonal basis $\vec{e}$, the effect of $\mathcal{D}$ on the series representation of ϕ - equation (12.1) - is*

$$\mathcal{D}\hat{\phi}(r) = \hat{b}\hat{D}\vec{r} = \hat{b}\hat{\eta}^n\hat{\mu}^m\vec{r} \quad (12.12)$$

Proof 12.1 Consider the generalised operator $\mathcal{D} = r^m \frac{d^n}{dr^n}$ operating on $\phi(r)$, $r \in [0, 1]$. If ϕ is replaced by its series expansion in the canonical basis, $\phi = \vec{a}r$, then by equation (12.9), $\mathcal{D}\phi$ can be rewritten as

$$\begin{aligned} \mathcal{D}\phi &= \vec{a}Dr \\ &= \vec{a}\eta^n\mu^m\vec{r}. \end{aligned}$$

Equations (12.6) and (12.5) allow the respective replacement of $\vec{a}$ and r in the above equation, and the subsequent use of the definition (12.11) leads to

$$\begin{aligned} \vec{a}Dr &= (\hat{b}V^*)D(V^{-1}\vec{r}) \\ &= \hat{b}(VDV^{-1})\vec{r} \\ &= \hat{b}\hat{D}\vec{r}. \end{aligned}$$

The proof is completed by considering the expansion of D according to