

University of Sheffield

Turbulent Flow Induced Vibration



Jah Shamas

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

The University of Sheffield

School of Mechanical, Aerospace and Civil Engineering

June 19, 2026

Acknowledgements

Firstly, I would like to thank my supervisor, Dr. Anton Krynkin. Over the last four years, you have been an incredibly supportive and enthusiastic mentor, always pushing and encouraging me to better myself as a researcher, even at times when I struggled to see ways forward myself. I would also like to thank my secondary supervisor, Prof. Kirill Horoshenkov, your expertise and advice have helped me greatly, and I have valued our shared experiences at the many conferences we have both attended.

These acknowledgements must be extended to Dr. Tim Rogers, Dr. Melika Gul, and Prof. Simon Tait. I greatly appreciate the time you have all given me throughout the research process, the depth of the work in this thesis has benefited immensely from your insight and advice. I must also thank Gavin Sailor for the experimental design and data collection aspects of this work. Without such work, many results and conclusions to be derived in this thesis would not have been possible.

I would like to thank Michael, Joanna, and Yicheng for their support and friendship throughout my time in Sheffield. I extend this to the wider Dynamics Research Group, which has been an incredibly supportive environment in which to grow as a researcher.

To my friends at home in Liverpool, catching up during our excursions at one of the countless pubs in town has always been much-appreciated and much-needed. To the members of 43 Walton Road over the last two years, Jack, Innes, Beth, Mia, Stan, Adam, and Carys, thank you all so much. I feel privileged that I have gotten to know you all so well. The many house meals and evenings of playing the ‘hat game’ will be greatly missed.

A final thanks to my mum, dad, sister, and the extended Wood and Shamas families. I cannot thank you enough, as I definitely would not be where I am today without your unwavering support. Thank you for the many visits over the last four years, and I consider myself lucky to have you all as my family.

Abstract

Storm and wastewater pipes form a significant proportion of the UK's hydraulic infrastructure, the maintenance of which presents key challenges. Ageing pipes face stresses from climate change and urbanisation, demanding robust monitoring solutions. Current monitoring practices utilise spot sensing methods, which fail to provide real-time continuous data on the internal hydraulics of these turbulent pipe flows. While emerging studies utilise free-surface dynamics, characterising turbulence via the fluctuating wall pressure in partially-filled pipes remains under-explored. Fibre-optic sensing non-invasively monitors these conduits from within the pipe wall through its response to the wall pressure. This thesis details the development, implementation, and analysis of a novel flush-mounted fibre-optic sensor designed to infer properties of turbulent partially-filled pipe flows directly from its response to wall pressure fluctuations.

Firstly, the Approximate Bayesian Computation-sequential Monte Carlo scheme is adapted to infer a high-dimensional binary system parameter encoding boundary condition information of a thin plate structure housing a fibre-optic sensor. Applied to vibrational data, it infers boundary condition configurations, with metrics demonstrating how convergence depends on user-design choices.

This is followed by the development and application of a novel flush-mounted fibre-optic sensor in an experimental facility, where its response to two partially-filled pipe flow regimes was analysed. Cross-correlation analyses of sensor data enabled accurate recovery of bulk flow velocity. A score-based Metropolis-Hastings algorithm was developed to identify optimal frequency bands retaining flow dynamics of interest. Azimuthal variation of sensor response to wall pressure fluctuations was recovered, including streamwise correlation structures, frequency-wavenumber spectra, and large-scale motion signatures. Their organisation enabled linkage to secondary flow motions. These findings were compared against Large-Eddy Simulation data of equivalent flow regimes. While showing reasonable qualita-

tive agreement, simulation limitations led to slight quantitative differences between results. Regardless, findings support the sensors ability to recover complex turbulent phenomena in partially-filled pipe flows.

Statement of Originality

I, the author, Jah Shamas, confirm that the work submitted is my own, except where work that has formed part of jointly-authored publications has been included. The contribution of myself and the other authors to this work has been explicitly indicated below. I confirm that appropriate credit has been given within the thesis where reference has been made to the work of others.

Examples:

Chapter 4: I, the candidate, planned and co-wrote the whole manuscript with Tim Rogers. The sections ‘Introduction’, ‘Approximate Bayesian Computation’, ‘ABC-SMC’, and ‘Combinatorial ABC-SMC’ were partially written by Tim Rogers. The co-authors reviewed and provided comments on the manuscript.

Chapter 5: I, the candidate, planned and wrote the manuscript. The co-authors reviewed and provided comments on the manuscript.

Chapter 6: I, the candidate, planned and wrote the manuscript with Yan Liu. The section ‘Numerical Framework and flow conditions’ was written by myself and Yan Liu. The co-authors reviewed and provided comments on the manuscript.

I, the candidate, Jah Shamas confirm that, aside from the collaboration mentioned, this thesis is my own work and that this work has not been previously presented for an award at this, or any other, university

Publications

Journal papers

- **Shamas, J.**, Liu, Y., Krynkin, A., Sailor, G., Gul, M., Tait, S.J., & Horoshenkov, K.V. Azimuthal variation in wall pressure statistics of turbulent partially-filled pipe flows: an experimental and numerical study *In preparation*.
- **Shamas, J.**, Krynkin, A., Sailor, G., Gul, M., Tait, S.J., & Horoshenkov, K.V. Fibre-optic sensor design and strain-response analysis to wall pressure fluctuations in turbulent partially-filled pipe flows. *Flow Measurement & Instrumentation*, 2026.
- **Shamas, J.**, Rogers, T., Krynkin, A., Prisutova, J., Gardner, P., Horoshenkov, K.V., Shelley, S.R., & Dickenson, P. Calibrating the Discrete Boundary Conditions of a Dynamic Simulation: A Combinatorial Approximate Bayesian Computation Sequential Monte Carlo (ABC-SMC) Approach. *Sensors*, 2024.

Conference Proceedings

- **Shamas, J.**, Krynkin, A., Sailor, G., Tait, S.J., Liu, Y., & Horoshenkov, K.V. Fibre-Optic Sensing for Analysis of Turbulent Pressure Fluctuations in Partially Filled Pipes. In *Proceedings of the Institute of Acoustics*, volume 46, 2024.
- **Shamas, J.**, Sailor, G., Tait, S.J., Liu, Y., Horoshenkov, K.V., & Krynkin, A. Distributed Fibre-Optic Sensing for Buried Partially-Filled Pipes. In *Proceedings of the 10th International Symposium on Environmental Hydraulics*, volume 10, 2024.
- Horoshenkov, K.V., Sailor, G., Tait, S.J., **Shamas, J.**, & Krynkin, A. Turbulence-Induced Pressure Pulsations in a Partially-Filled Pipe. In *Proceedings of the 10th International Symposium on Environmental Hydraulics*, volume 10, 2024.
- **Shamas, J.**, & Krynkin, A. Resolving uncertainty in mechanical vibration of thin structures: A novel Bayesian approach. In *Proceedings of the Institute of Acoustics*, volume 45, 2023.

Contents

Abbreviation list	1
Nomenclature list	3
1 Introduction	7
1.1 Motivation	7
1.2 Aims and objectives	9
1.3 Thesis layout	9
2 Literature review	11
2.1 Partially-filled pipe flows	11
2.1.1 Basic description and turbulence	11
2.1.2 Free surface behaviour and secondary flow motions	14
2.1.3 Wall pressure distribution	18
2.2 Flow measurement techniques	20
2.2.1 Free surface measurement	21
2.2.2 Flow visualisation techniques	23
2.2.3 Wall-mounted measurement	24
2.3 Fibre-optic sensors	26
2.3.1 Technology and operating principle	27
2.3.2 Towards deployment in flow monitoring	30
2.4 Inverse problems in engineering	33
2.4.1 A Bayesian perspective on inversion	33
2.4.2 Parameter estimation via Bayesian methodologies	34
2.5 Literature review summary	36

3	Methodology	39
3.1	A statistical description of wall pressure fluctuations	40
3.1.1	Space-time correlations	41
3.1.2	Taylor’s hypothesis	43
3.2	The Bayesian approach	44
3.2.1	Markov Chain Monte Carlo	44
3.2.2	Approximate Bayesian Computation via Sequential Monte Carlo (ABC-SMC)	45
3.3	Experimental methodology considerations	48
4	Calibrating the discrete boundary conditions of a dynamic simulation: A combinatorial ABC-SMC approach	51
4.1	Abstract	51
4.2	Introduction	52
4.3	Approximate Bayesian Computation	54
4.4	ABC-SMC	56
4.5	Combinatorial ABC-SMC	58
4.6	Problem statement	60
4.7	Results and discussion	64
4.8	Conclusions	78
4.9	Acknowledgements	79
5	Fibre-optic sensor design and strain-response analysis to wall pressure fluctuations in turbulent partially-filled pipe flows	81
5.1	Abstract	81
5.2	Introduction	82
5.3	Experimental facility and sensor design	84
5.3.1	Rig description	85
5.3.2	Sensor design and installation	86
5.3.3	Flow conditions	89
5.4	Data analysis	90
5.4.1	Flow velocity recovery	92
5.4.2	Automation of filtering process	95
5.4.3	Effect of time interval and outlier removal	99
5.4.4	Streamwise correlation structure	104

5.4.5	Large-scale motions	109
5.4.6	Out of water comparison	111
5.5	Conclusion and future work	115
6	Azimuthal variation in wall pressure-induced strain statistics of turbulent partially-filled pipe flows: an experimental and numerical study	117
6.1	Abstract	117
6.2	Introduction	118
6.3	Experimental setup and flow conditions	120
6.4	Numerical Framework and flow conditions	124
6.5	Analysis procedure	127
6.6	Data processing	130
6.6.1	Spectral comparison	130
6.6.2	Filter optimisation process	132
6.7	Results & discussion	137
6.7.1	Streamwise correlation	138
6.7.2	Frequency-wavenumber spectra	146
6.7.3	Large-scale motions	152
6.8	Conclusions	162
7	Conclusions	165
7.1	Discussion	165
7.2	Final conclusions	170
7.3	Future work	171
7.3.1	Sensor design choices	171
7.3.2	Further hydraulic testing	172
7.3.3	Further analyses of FOS data	173
	References	176

List of Figures

2.1	Schematic diagram and flow geometry of a partially-filled pipe.	12
2.2	Isovel patterns of the streamwise velocity distribution about the cross-sectional area in a half-filled pipe flow showcasing the velocity dip phenomenon [1]. .	15
2.3	Example of the streamwise velocity pre-multiplied power spectrum at $Re = 66400$, displaying peaks at multiple spatial scales in terms of pipe radius R [2].	17
2.4	Azimuthal variation in normalised mean wall shear stress of a partially-filled pipe flow for different values of friction Reynolds number [3].	19
2.5	Diagram of experimental setup for acoustical measurements in [4].	22
2.6	Array of pressure sensors used for measurement in [5].	25
2.7	Schematic diagram of TIR in a fibre-optic cable [6].	27
2.8	Schematic diagram of FBG strain response [7].	30
2.9	Schematic diagram of FBG sensor developed in [8].	31
3.1	(a): Sluice gate at the pipe outlet controlling the flow depth. (b): Tape measure placed around pipe circumference to infer flow depth from wetted perimeter of flow.	49
4.1	Normalised absolute displacement profile for the sensor subjected to mechanical excitation of 0.041–0.044 N operating at 5 Hz. Impact located at 45 cm along the 90 cm sensor length.	61
4.2	Simplified model of LVV sensor used for predicting response to mechanical excitation. Labelled clamping sections and shaker impact area are highlighted in blue in (a). The mesh used is displayed in (b).	64

4.3	Accepted curves for each wave of ABC-SMC using a difference of area metric. Curves in (a)-(e) correspond to accepted simulated data Y^* from waves 1-5, respectively. Experimental data are shown with circle markers.	66
4.4	Intermediate distributions for each dimension of θ for each wave of ABC-SMC, using a difference of area metric. Distributions in (a)-(e) correspond to the belief of the fixity of each boundary condition after each wave 1-5, respectively.	67
4.5	Accepted curves for each wave of ABC-SMC using a least-squares metric. Curves in (a)-(e) correspond to accepted simulated data Y^* from waves 1-5, respectively. Experimental data are shown with circle markers.	68
4.6	Intermediate distributions for each dimension of θ for each wave of ABC-SMC, using a least-squares metric. Distributions in (a)-(e) correspond to belief of fixity of each boundary condition after each wave 1-5, respectively.	69
4.7	Normalised absolute displacement profile for the sensor subjected to mechanical excitation of 0.035–0.037 N operating at 5 Hz. Impact located at 70 cm along the 90 cm sensor length.	72
4.8	Accepted curves for each wave of ABC-SMC using a least-squares metric applied to data from Figure 4.7. Curves in (a)-(e) correspond to accepted simulated data Y^* from waves 1-5, respectively. Experimental data are shown with circle markers.	73
4.9	Intermediate distributions for each dimension of θ for each wave of ABC-SMC, using a least-squares metric applied to data from Figure 4.7. Distributions in (a)-(e) correspond to belief of fixity of each boundary condition after each wave 1-5, respectively.	74
4.10	Accepted curves for each wave of ABC-SMC using a split least-squares metric. Curves in (a)-(e) correspond to accepted simulated data Y^* from waves 1-5, respectively. Experimental data are shown with circle markers.	76
4.11	Intermediate distributions for each dimension of θ for each wave of ABC-SMC, using a split least-squares metric. Distributions in (a)-(e) correspond to belief of fixity of each boundary condition after each wave 1-5, respectively.	77
5.1	Schematic view of ICAIR pipe facility including inlet and outlet tanks.	85

5.2	(a): External pipe view of the FOS containment structure rotated 90^0 from the working position (0°) of the FOS used for data collection. Schematic diagrams of (b): cross-section and (c): lengthwise view of the FOS structure. FBG element locations and indices highlighted in (c) (not to scale).	87
5.3	One-sided PSD of FOS strain response from the pipe bottom to (a): flow regime 1, (b): flow regime 2, and to (c): empty pipe. Time series of FOS strain response from the pipe bottom to (a): flow regime 1, (b): flow regime 2, and to (c): empty pipe.	91
5.4	Distribution of flow velocity estimates obtained through manipulating cross-correlation between grating readings from the pipe bottom in (a): flow regime 1, (b): flow regime 2, both filtered in the range of 0.1-15 Hz.	93
5.5	Distribution of flow velocity estimates obtained through manipulating cross-correlation between grating readings from the pipe bottom in (a): flow regime 1, (b): flow regime 2. Both filtered in the range of 2-7 Hz.	94
5.6	(a): Heatmap indicating areas of the parameter space Θ searched by the scheme. (b): Density heatmap of accepted samples (normalised to integrate to 1 over Θ). The colour gradient of blue to red represents an increasing number of samples.	98
5.7	Velocity histograms for flow regime 2. Filtered in (a): 2.5-8 Hz, (b): 2.5-11 Hz, (c): 2.5-14 Hz.	99
5.8	Distribution of flow velocity estimates obtained for different values of M in flow regime 1. $M = 30, 60, 300, 600, 900, 1800$ seconds (a-f). Colour legend given for each grating pair.	100
5.9	Distribution of flow velocity estimates obtained for different values of M in flow regime 2. $M = 30, 60, 300, 600, 900, 1800$ seconds (a-f). Colour legend given for each grating pair.	101
5.10	Velocity estimates plotted as a function of streamwise separation between sensor elements. Bold blue and red lines indicate true bulk flow velocity in flow regime 1 and 2, respectively. $M = 1800$	102
5.11	(a-b): Distribution of flow velocity estimates obtained for $M = 120$ in (a-c): flow regime 1, (b-d): flow regime 2, excluding FBG pairs (1,2), (2,3), (1,3), (3,4). (c-d): Normal distribution fit to velocity histograms.	103

5.12	Overlaid filtered time series of FBG 1 (black line) and FBG 2-6 (grey lines) (a-i) in flow regime 1. Normalised cross-correlation between respective gratings (b-j).	105
5.13	Streamwise correlation estimates in flow regime 1 plotted against streamwise separation, $M = 120, 600, 900, 1800$ seconds (a-d). The average of r_{ij} at each streamwise separation was fitted to Equation (5.9) (left axis). The average time lag at each streamwise separation was fitted to a straight line (right axis). Grey dashed line indicates cut-off point for separations used in the fitting. Time lags at which maximum correlations were attained are plotted against streamwise separation (right axis).	107
5.14	Streamwise correlation estimates in flow regime 2 plotted against streamwise separation, $M = 120, 600, 900, 1800$ seconds (a-d). The average of r_{ij} at each streamwise separation was fitted to Equation (5.9) (left axis). The average time lag at each streamwise separation was fitted to a straight line (right axis). Grey dashed line indicates cut-off point for separations used in the fitting. Time lags at which maximum correlations were attained are plotted against streamwise separation (right axis).	108
5.15	Normalised pre-multiplied spectra from (a) flow regime 1, and (b) flow regime 2 as a function of wavelength scale. Parameters U_b and h used in each figure are taken from values in Table 5.2 for each regime. Wavelength axis displayed in logarithmic scale.	110
5.16	Overlaid filtered time series of FBG 1 (black line) and all other gratings (grey lines) ((a)-(i)) in flow regime 1 at the 104° position (above the free surface). Normalised cross-correlation between respective gratings ((b)-(j)).	112
5.17	Distribution of flow velocity estimates obtained through manipulating cross-correlation between grating readings from above the free surface in (a): flow regime 1, (b): flow regime 2, both filtered in the range of 0.1-15 Hz. Histograms computed from 2 hours of fully developed flow, split into 2 minute sections.	113
6.1	Schematic view of the ICAIR pipe facility.	120
6.2	Schematic diagrams of (a), cross-section and (b), lengthwise view of the FOS structure. FBG element locations and indices highlighted in (b) (not to scale).	122
6.3	cross-section schematic of a partially-filled pipe flow showing quantities of central angle γ and flow depth h	122

6.4	cross-sectional views of (a) flow regime 1, (b) flow regime 2. Red markings indicate angular positions at which the FOS collected data in each regime.	123
6.5	(a) Streamwise-averaged PSD obtained from FOS strain time series (converted from FBG wavelength shift via Equation (6.2)) at the pipe bottom position in both flow regimes. (b) Streamwise-averaged PSD obtained from LES wall pressure fluctuation time series at the pipe bottom position in both flow regimes. Frequency cut-off in (a) to match the LES frequency domain. Frequency axis displayed in logarithmic scale.	131
6.6	Recovered streamwise correlation profiles and fittings to Equation (6.22) at the pipe bottom position for (a): flow regime 1, and (b): flow regime 2. Measurements taken over the four 30-minute time series were filtered in the range of 4-9 Hz in each case. Points left of the grey line indicate separations not used in fitting, and an in-depth explanation as to why is given in Chapter 5.	133
6.7	Heatmap of correlation lengths L from fitting the recovered streamwise correlation profiles to Equation (6.22) on Θ . The dataset used is from the pipe bottom position in flow regime 1. The maximum value of L is denoted with a black marker.	136
6.8	Recovered bandpass windows maximising correlation length of FOS data for (a) flow regime 1, (b) flow regime 2.	137
6.9	Recovered best fit correlation lengths at all angular positions for (a) flow regime 1, (b) flow regime 2.	138
6.10	Correlation decay in time plotted as a function of 5 FBG separations at the pipe bottom in: (a) flow regime 1, (b) flow regime 2. Red markers indicate maximum correlation for a given separation ξ fitted to Equation (6.25).	139
6.11	Recovered best fit time correlation lengths at all angular positions for (a) flow regime 1, (b) flow regime 2.	140
6.12	FOS-recovered convection velocity estimates at all angular positions for (a) flow regime 1, (b) flow regime 2. Values normalised against U_b .	141
6.13	Streamwise correlation profile of LES wall pressure data recovered from the pipe bottom position. All measurements taken against the reference point at the $x = 0$ position.	142

6.14	(a): Streamwise correlation profile and fitting to Equation (6.22). (b): Correlation decay in time as a function of streamwise separation fitted to Equation (6.25). Both taken from LES flow regime 1 at the pipe bottom position.	143
6.15	LES-recovered streamwise correlation length L in (a): flow regime 1, (c): flow regime 2, and time correlation length τ_b in (b): flow regime 1, (d): flow regime 2.	144
6.16	Streamwise-averaged root mean square LES wall pressure at all angular positions for (a): flow regime 1, (b): flow regime 2.	145
6.17	Normalised frequency-wavenumber spectra obtained from raw LES data at the pipe bottom in (a): flow regime 1, (b): flow regime 2. Convective ridge highlighted with dashed line.	147
6.18	Normalised frequency-wavenumber spectra obtained from raw FOS data at the pipe bottom in (a): flow regime 1, (b): flow regime 2.	148
6.19	Frequency-wavenumber spectra-recovered convection velocities U_c in (a): LES flow regime 1, (b): LES flow regime 2, (c): FOS flow regime 1, (d): FOS flow regime 2.	151
6.20	Normalised pre-multiplied spectra from the FOS pipe bottom position in (a) flow regime 1, (b) flow regime 2.	154
6.21	Peak locations in FOS pre-multiplied spectra for all angular positions in (a) flow regime 1, and (b) in flow regime 2.	155
6.22	Spatial scales of identified LSMs in (a-b): flow regime 1, (c-d): flow regime 2.	156
6.23	Ratio of normalised pre-multiplied spectra values between peaks 1 and 2 in (a): flow regime 1, and (b): flow regime 2.	157
6.24	Normalised pre-multiplied spectra from the LES pipe bottom position in (a): flow regime 1, (b): flow regime 2.	158
6.25	Peak locations in LES pre-multiplied spectra for all angular positions in (a) flow regime 1, and (b) in flow regime 2.	159
6.26	Peak locations in LES pre-multiplied spectra for all angular positions in (a) flow regime 1, and (b) in flow regime 2.	160
6.27	Overlaid LES pre-multiplied spectra from (a): flow regime 1 at angles of $\gamma = 63, 70, 77^\circ$, and (b): flow regime 2 at angles of $\gamma = 28, 35, 42^\circ$	161

7.1	Plots of Equation 7.2 for varied FBG combinations at the pipe bottom position in flow regime 1. Darker lines denote increased streamwise separation. Vertical dashed lines denote approximate frequency limits where Taylor's hypothesis is satisfied.	174
-----	--	-----

List of Tables

5.1	Centre position of each FBG element from the starting position of the FOS plate (in the streamwise direction), and Bragg wavelengths of FBGs lying at the plate/gel boundary.	89
5.2	Properties of flow conditions in this study.	89
6.1	Properties of flow conditions in this study.	122
6.2	Hydraulic properties of the simulations	126
6.3	Streamwise separation (mm) of all FBG combinations inside the FOS listed in ascending order.	130

Abbreviation list

ABC Approximate Bayesian Computation

ABC-SMC Approximate Bayesian Computation–Sequential Monte Carlo

CFD Computational Fluid Dynamics

CSD Cross-Spectral Density

DFOS Distributed Fibre-Optic Sensor

DIC Digital Image Correlation

DNS Direct Numerical Simulation

FBG Fibre-Bragg Grating

FEM Finite-Element Model

FFT Fast Fourier Transform

FOS Fibre-Optic Sensor

IB Immersed Boundary

LDV Laser Doppler Vibrometer

LES Large-Eddy Simulation

LSM Large Scale Motion

LTI Linear Time-Invariant

MCMC Markov Chain Monte Carlo

MH Metropolis–Hastings

OCF Open Channel Flow
ODE Ordinary Differential Equation
OFDR Optical Frequency-Domain Reflectometry
OTDR Optical Time-Domain Reflectometry
PIV Particle Image Velocimetry
POD Proper Orthogonal Decomposition
PTV Particle Tracking Velocimetry
PZT Piezo-Electric Transducer
RANS Reynolds-Averaged Navier–Stokes
RI Refractive Index
RF Random Forest
RMS Root-Mean Square
SGS Subgrid-scale Stress
S-PIV Stereoscopic Particle Image Velocimetry
TDE Time Delay Estimation
TIR Total Internal Reflection
VLSM Very Large Scale Motion
WALE Wall-Adapting Local Eddy Viscosity

Nomenclature list

Roman symbols

A	Cross-sectional area of flow	m^2
$A_c(\cdot)$	Area underneath observed data curve	m^2
a	Bandpass window lower bound	Hz
B	Channel bed width	m
b	Bandpass window upper bound	Hz
$D(\cdot, \cdot)$	Metric function	[-]
d	Pipe diameter	m
f	Temporal frequency	Hz
f_θ	Forward model	Variable
$f.$	LES volume force	N
$g.$	LES gravitational acceleration	m s^{-2}
H	Heaviside step function	[-]
h	Flow depth	m
$I.$	Indicator function	[-]
J	Total number of ABC-SMC distributions	[-]
$K_j(\cdot \cdot)$	Markov Kernel on j 'th wave of ABC-SMC	[-]
k	Streamwise wavenumber	m^{-1}
L	Streamwise correlation length	m
L_q	Turbulent length scale	m
M	FOS signal duration	s
$\mathcal{N}(\cdot, \cdot)$	Normal distribution	Variable
N	Total number of samples accepted in each ABC-SMC wave	[-]

n	Refractive index	[-]
n_{eff}	Effective refractive index	[-]
n_m	Manning's roughness coefficient	$\text{m}^{-1/3}\text{s}$
n_p	Histogram mode counter	[-]
P	Wetted perimeter of flow	m
$P(\cdot)$	Probability distribution	[-]
$p(\cdot, *)$	Fluctuating wall pressure	Pa
Q	Flow rate	$\text{m}^3 \text{s}^{-1}$
q	Turbulent velocity	m s^{-1}
$q(\cdot \cdot)$	Proposal distribution	[-]
R_h	Hydraulic radius	m
$R_{pp}(\cdot, *)$	Space-time cross-correlation of p	[-]
Re	Reynolds number	[-]
r	Pipe radius	m
S	Flow slope	[-]
$S_{\cdot\cdot}^r$	LES filtered strain-rate tensor	s^{-1}
T	Time series length	s
t	Time variable	s
U_b	Bulk flow velocity	m s^{-1}
U_c	Convection velocity	m s^{-1}
$U(\cdot, \cdot)$	Uniform distribution	Variable
$u_{\cdot}$	LES velocity components	m s^{-1}
u_*	Friction fluid velocity	m s^{-1}
$\mathbf{v}$	Velocity field	m s^{-1}
$\bar{\mathbf{v}}$	Mean velocity field component	m s^{-1}
$\tilde{\mathbf{v}}$	Turbulent velocity field component	m s^{-1}
w_j^i	Weight of i 'th sample in j 'th wave of ABC-SMC	[-]
X	System input variables	Variable
$\mathcal{X}$	Input space	Variable
x	Streamwise variable	m
Y	Observed data	Variable
Y^*	Simulated data	Variable
$\mathcal{Y}$	Output space	Variable

y_m^+	Spanwise wall unit	[-]
y	Spanwise variable	m
Z	Normalising constant in Bayes' theorem	Variable
z_m^+	Vertical wall unit	[-]
z	Vertical variable	m

Greek symbols

α_e	Thermal expansion coefficient	K^{-1}
$\beta(\cdot, \cdot, \cdot)$	Histogram scoring function	[-]
$\Gamma_{pp}(\cdot, *)$	Cross-spectral density of the wall pressure p	$Pa^2 Hz^{-1}$
γ	Central pipe angle	rad (or $^\circ$)
γ_i	Angle of incidence	rad (or $^\circ$)
γ_r	Angle of refraction	rad (or $^\circ$)
ϵ	ABC threshold	[-]
ε	Axial strain	[-]
$\eta(\cdot, \cdot)$	Signed distance function	m
Θ	Space of ABC-SMC system parameters	[-]
θ	Binary boundary condition vector	[-]
θ	System parameters	Variable
Λ	Grating period	m
λ	Wavelength	nm
λ_B	Bragg wavelength	nm
μ	Mean value	Variable
$\hat{\mu}$	Dynamic viscosity of the fluid	Pa s
ν	Kinematic viscosity	$m^2 s^{-1}$
ξ	Streamwise separation	m
ξ_o	Thermo-optic coefficient	K^{-1}
$\pi.$	Histogram scoring function input parameter	[-]
ρ	Fluid mass density	$kg m^{-3}$
ρ_e	Photo-elastic coefficient	Pa^{-1}
σ	Standard deviation	Variable
τ	Time lag	s
τ_b	Time correlation length	s

τ_{eddy}	Eddy turnover time	s
$\Phi_{pp}(\cdot, \cdot)$	Frequency-wavenumber spectrum of p	$\text{Pa}^2 \text{s m}^{-1}$
$\phi_{pp}(\cdot)$	Power-spectral density of p	$\text{Pa}^2 \text{Hz}^{-1}$
χ	Half thickness of LES air-water interface	m
ω	Angular frequency	rad s^{-1}

Chapter 1

Introduction

1.1 Motivation

The UK’s underground pipe infrastructure is a vast, ageing asset, with an estimated length of 1.5 million km. This incorporates 400,000 km of sewage network, and a similar length of clean water network [9, 10]. Effectiveness of these networks degrades with age, resulting in leakages and structural degradation. Existing condition monitoring methods prove slow, invasive, and spatially limited, incapable of providing continuous, high resolution data needed for robust monitoring. This leaves much of the network’s condition unknown. The consequence of this uncertainty is frequent failures, often requiring utility trenching to fix the problem, resulting in an estimated 4.25 million holes dug up on UK streets every year at a cost of over £1.5 billion [9].

Future projections indicate a significant escalation in UK water demand, with the population expected to reach 75 million by 2050. This demographic pressure, combined with warmer summers from climate change and the water requirements for new infrastructure, will place a far greater strain on available water resources [11]. Compounding this demand challenge, the system continues to suffer substantial losses. Despite water companies reducing leakage by 7% since the 1990s, approximately 20% of the public supply, equating to some 3000 megalitres per day in 2020-2021, is still lost through the distribution network [12, 13, 14].

Accurate monitoring of the pipe networks conveying this water is essential for the prevention of the pressing economic and environmental issues mentioned above. This necessitates

the production of enhanced sensing platforms capable of extracting live hydraulic data directly from inside the pipe infrastructure. A key hurdle is that many of these operational conduits, particularly drainage pipes, flow partially-filled. Flows of this type typically experience high levels of turbulence. As such, the sensors to monitor this infrastructure are required to interpret complex turbulent flow dynamics, and the potential influence these dynamics may have on structural loading and degradation, to aid in predictive monitoring solutions. Currently, most conventional methods are intrusive, distorting the flow dynamics in order to obtain a measurement. Therefore, there is a demand for robust non-invasive sensing systems capable of uncovering turbulent flow dynamics.

Unfortunately, only a sparse body of literature exists characterising the internal flow dynamics of turbulent partially-filled pipe flows [15, 16, 17, 18]. In many of these studies, coherent turbulent motions contributing to spatial patterns of fluid loading are identified; these are phenomena not observed in pressure-driven pipe flows. So far, our ability to monitor such processes is dependent on either high-fidelity numerical simulations [19], or laboratory techniques capable of resolving the whole fluid domain [1]. These turbulent flow structures play a direct role in the state of the dynamic pressure field at the wall that encodes information about the flow state. Consequently, it is of interest to recover properties of this fluctuating wall pressure as a non-intrusive means to access the dynamics of the overarching turbulent flow.

Fibre-optic sensing technologies present themselves as a suitable candidate for this measurement challenge in such an environment. With immunity to electromagnetic interference, small diameter, arbitrary shaping, and capabilities for distributed monitoring, these cables have become increasingly popular for sensing purposes in varied environments [20]. For example, they have been utilised for mixture concentration monitoring, [21], seismic activity monitoring [22, 23, 24], and structural health monitoring [25, 26]. Recently, researchers have shown their utility for flow measurement purposes [27, 28, 29], showcasing the versatility of these sensors. In order to recover undisturbed flow properties, it is postulated that these sensors can be installed directly at the pipe wall to respond to the pressure fluctuations, containing information on the flow turbulence and conditions. As such, this technology has the potential to perform as a robust non-invasive sensor designed to relate fluctuating wall pressure with flow properties, a technology much-needed to aid in advanced sensing of live water infrastructure. Here, the term non-invasive is used to enunciate that the sensor does not protrude into the flow in any capacity, leaving the flow

field undisturbed.

1.2 Aims and objectives

The aim of this thesis is to produce a non-invasive fibre-optic sensing solution that measures strain within a flush-mounted sensor system, in response to wall pressure fluctuations in turbulent partially-filled pipe flows. The goal is to unambiguously recover bulk flow conditions and spatial scales about the flow cross-section, from the sensor signal response to flow-induced vibration. With this in mind, the specific objectives are as follows:

- Propose an alternate Bayesian methodology to infer discrete boundary conditions of a complex engineering structure housing a fibre-optic sensor.
- Design and establish a novel fibre-optic sensor to respond to pressure fluctuations at the interior wall of a turbulent partially-filled pipe.
- Apply cross-correlation based analyses in time and frequency domain to strain time series data to recover spatial and temporal properties of wall pressure fluctuations about the pipe azimuth for two distinct flow regimes.
- Recover multi-scale turbulent behaviour of partially-filled pipe flows, and azimuthal changes thereof, from discrete strain measurements within the fibre.
- Validate against numerical Large-Eddy Simulation (LES) data.

1.3 Thesis layout

This thesis is organised in a thesis by publication format. Each individual research chapter is either a published piece of work or a piece of work in the submission process.

- **Chapter 2:** Provides a review of the literature relevant to fibre-optic sensing for partially-filled pipe flows, including current understanding of partially-filled pipe flow dynamics, non-invasive measurement techniques, and current applications of fibre-optic sensing for monitoring of hydraulic infrastructure. A short section is included on Bayesian inference methodologies as applied to inverse problems in engineering contexts.
- **Chapter 3:** Introduces the relevant methodologies utilised in this thesis. The

Navier-Stokes and Poisson equations are introduced, followed by an introduction to statistical quantities derived from time series measurements of wall pressure fluctuations, as well as Taylor’s hypothesis. This is followed by an introduction to Bayesian procedures used to sample from the posterior distribution of model parameters via Markov Chain Monte Carlo (MCMC) and Approximate Bayesian Computation-Sequential Monte Carlo (ABC-SMC) methods.

- **Chapter 4:** Presents a novel adaptation of the well-known ABC-SMC parameter estimation scheme utilised to recover a distribution of complex boundary conditions imposed on a fibre-optic sensor structure subjected to harmonic loading. This chapter focuses on statistical methodology development as applied to a calibration problem of the initial sensor structure. This sensor was only subjected to laboratory-controlled dynamic testing before an improved, hydraulically-tested sensor iteration was developed for the proceeding chapters.
- **Chapter 5:** Introduces the fibre-optic sensor utilised throughout the majority of this work, designed to respond to wall pressure fluctuations in turbulent partially-filled pipe flows. It covers the sensor design, experimental facilities and flow conditions. It also outlines the data processing steps undertaken to recover bulk flow properties from the wall pressure-induced strain measurements, as well as initial identification of large-scale motions through application of Taylor’s hypothesis.
- **Chapter 6:** Investigates the azimuthal variation of the fluctuating wall pressure-induced strain statistics through a combined experimental and numerical approach. The framework used to produce the complementary LES wall pressure fluctuation data will be given. The combined datasets are used to infer azimuthal change in flow dynamics derivable from the fluctuating wall pressure properties. These include convection velocity, correlation structures, and frequency-wavenumber spectra. This is extended to recover azimuthal variation of large scale motions across both datasets.
- **Chapter 7:** Contains conclusions and recommendations for future avenues to build on the research presented in this thesis.

Chapter 2

Literature review

In the first part of this chapter, I aim to introduce the necessary background of partially-filled pipe flows and the underlying dynamics associated with their motion. To begin, basic geometric descriptions of these flows will be given. I will then discuss the implications of turbulence, focusing on the underlying dynamics that emerge due to the presence of the free surface, arriving at a description of the wall pressure fluctuations. A review of the current literature will be given to identify the present understanding of these phenomena. Next, sensing methods applicable to flow monitoring purposes will be discussed, followed by use cases of fibre-optic sensors in varied measurement scenarios, including flow and pipeline monitoring. Finally, Bayesian approaches to solving inverse problems will be introduced, outlining the most notable schemes used, and where their utility lies in parameter estimation tasks for engineering applications.

2.1 Partially-filled pipe flows

2.1.1 Basic description and turbulence

Open channel flows (OCF) are characterised as flows with a free surface, a boundary between the flowing fluid and the atmosphere above the channel. OCFs occur in various channel geometries, such as rectangular, circular, or irregular (e.g. rivers) conduits. The literature review in this section, which aims to stay close to the topic of this thesis, focuses on OCF in circular conduits. OCFs in circular conduits, or pipes, are called partially-filled pipes. In this case, the flow is driven by gravity due to the presence of the free surface.

In the case of fully-filled pipe flow, where the conduit is filled entirely with fluid, the flow is inherently pressure-driven [30]. I begin by defining some basic geometric parameters of partially-filled pipe flows.




Figure 2.1: *Schematic diagram and flow geometry of a partially-filled pipe.*

For a pipe of circular cross-section with inner diameter d , the flow depth h is the distance from the pipe bottom to the free surface at the centreline of the cross-section. The central angle γ associated with parameters h and d , is given by the angle between the radius r connecting the pipe bottom to the pipe centre, and the pipe centre to the free surface boundary at the wall. The calculation of γ is given by

$$\cos(\gamma) = \frac{r - h}{r}. \quad (2.1)$$

Another important parameter is the hydraulic radius R_h , which can be interpreted as a measure of channel efficiency based on the ratio of cross-sectional flow area A , to the length of the pipe wall boundary in contact with the fluid, also called the wetted perimeter P . Then, $R_h = A/P$, where

$$A = \frac{r^2}{2} (2\gamma - \sin(2\gamma)), \quad (2.2)$$

$$P = 2r\gamma, \quad (2.3)$$

$$R_h = \frac{r(2\gamma - \sin(2\gamma))}{4\gamma}. \quad (2.4)$$

Such expressions are central to hydraulic analyses, forming parts of the Manning equation [31] necessary for calculating hydraulic quantities in open channel flows. Due to the gravity-driven nature of these flows, greater flow depth increases flow rate, as compared to fully filled pipes, whereby greater flow rate is achieved through greater pressure head. As such, bulk flow velocity U_b can be defined as the ratio between flow rate Q and cross-sectional area of flow to give

$$U_b = \frac{Q}{A}, \text{ where } Q = \frac{A}{n_m} R_h^{2/3} S^{1/2}, \quad (2.5)$$

from the Manning equation, which is applicable to OCFs. Here, S is the flow slope, and n_m is Manning's roughness coefficient, a dimensionless quantity which is dependent on the pipe wall material.

Perhaps the most important hydraulic parameter to consider is the Reynolds Number, given by

$$Re = \frac{4U_b R_h}{\nu}. \quad (2.6)$$

Here, ν is the kinematic viscosity of the fluid. Named after Osborne Reynolds for his pipe flow experiments in the 19th century, this dimensionless quantity can be interpreted as the ratio of inertial and viscous forces within the flow [32]. Through non-dimensionalising the Navier-Stokes equations, the ratio between the terms contributing to inertial and frictional forces is equal to the Reynolds number. Indeed, the denominator of Equation (2.6) ν , is a measure of a fluid's resistance to flow under gravity, originating from frictional shear forces within the fluid. The numerator is a combination of inertial forces that arise due to the momentum of the fluid, with flows possessing greater bulk velocity and ratio of area to wetted perimeter produce larger inertial forces relative to viscous forces. The balance between these forces determines whether the flow is labelled as *laminar* or *turbulent*, corresponding to lower and higher values of Re , respectively.

Laminar flows are characterised by layers of adjacent fluid flow throughout the whole domain, with no mixing between layers. In this case, frictional forces dominate, and the shear between layers governs flow behaviour. In the literature, laminar flows are very well understood. Closed form solutions of the Navier-Stokes equations governing incompressible fluid flow are given, under the appropriate conditions. For example, in a fully-filled pipe exhibiting laminar flow, the velocity profile obeys a parabolic profile about the pipe radius, the well known Hagen-Poiseuille flow [33]. However, the same cannot be said for partially-filled pipes. Only at very small Reynolds numbers ($Re < 2400$) is laminar flow observed [34]. Through application of a bipolar coordinate transformation used to account for varied

free surface levels, theoretical solutions for laminar flow in partially-filled pipes have been derived [35, 36]. Although to this day, most literature concerning partially-filled pipe flows has focused on turbulent regimes, as the laminar regime is rarely observed. The transition to turbulence from the laminar regime begins when inertial forces outweigh viscous forces, small perturbations occur between these adjacently flowing fluid layers, which the frictional forces are not strong enough to damp. This leads to instabilities within the flow, which are subsequently amplified. This generates chaotic fluctuations in all dimensions in the flow, resulting in drastically different velocity and pressure fields in the spanwise and depthwise directions as compared to laminar flow. For partially-filled pipe flows, the nature of this turbulence is further complicated by the existence of the free surface. The presence of which, and its consequences, will be investigated in the next section.

2.1.2 Free surface behaviour and secondary flow motions

While the transition to turbulence is a defining feature of partially-filled pipe flows, what distinguishes them most from fully filled-pipe systems is the presence of the free surface. The deformable interface between liquid and air modifies turbulence production and induces secondary flow motions across the cross-section. Properties of the free surface, and the relation to underlying conditions of the flow have only gained significant attention in recent years.

The work of Brocchini and Peregrine [37] paid attention to free surface dynamics, including what features may arise at the surface due to turbulence in different flow regimes. In this work, a framework was developed to classify free surface effects based on the turbulent velocity q , and length scale L_q of turbulence in the flow. The interplay between these two quantities was visualised in a $q - L_q$ diagram, whereby flows existing in different regions of this $q - L_q$ plane exhibit different free surface characteristics. The difference between regions and their corresponding values of q and L_q arise from the relative dominance of different forces. The framework dictates that at low values of q and L_q , there exists little surface disturbance due to sufficient surface tension relative to gravity forces. Whereas for larger values of q and L_q , flows experience extremely strong turbulence, where the turbulent kinetic energy is great enough to overcome both surface tension and gravity. As such, droplets and bubbles form in this case, creating two-phase flow. In the case of high L_q and low q , gravity dominates, leading to small surface deformations. Naturally occurring flows such as rivers and streams exist in this region [38].

Building on insights gained from the interaction between free surface and bulk flow, researchers have also sought to characterise quantities contained beneath the surface. Ng et al. [15] utilised high-speed stereoscopic particle image velocimetry (S-PIV) to construct the velocity distribution in partially-filled pipe flows. In the turbulent regime, the maximum velocity occurred along the pipe centreline below the free surface, otherwise known as the *velocity dip* phenomenon. Using a Preston tube, Knight & Sterling [1] quantified the shear stress distribution about the wetted perimeter for partially-filled pipe flows of varied cross-sectional shape. Information on the velocity distribution throughout the fluid domain was also realised for varied cross-sections, concluding the same velocity dip for filling ratios close to 50%. Similarly, Yoon et al. [39] employed S-PIV and arrived at comparable conclusions. This velocity dip is widely accepted as evidence of secondary flow motions in partially-filled pipe flows.




Figure 2.2: *Isovel patterns of the streamwise velocity distribution about the cross-sectional area in a half-filled pipe flow showcasing the velocity dip phenomenon [1].*

These motions, describing organised motion of fluid in the plane perpendicular to mean flow direction, were first observed in 1926 by Prandtl [40]. As explained by Nezu [41], secondary flow motions originate from the streamwise component of vorticity, which can be derived from the Reynolds-averaged Navier–Stokes (RANS) equations. These flow features are classified as counter-rotating vortices symmetric about the flow centreline, and have been studied extensively in open channel flows [41, 42, 34]. In OCFs, it is understood that their relative strength in relation to mean flow is highly dependent on the aspect ratio B/h ,

where B is the bed width, and h is the flow depth. For small aspect ratios $B/h < 5$, the side walls of the OCF exert greater influence, inducing the velocity dip as well as secondary flow motions [43]. In contrast, for larger B/h ratios, the secondary flow motions are negligible in comparison to the strength of the mean flow. Thus, the ratio B/h , and consequently the geometry of the flow, is clearly a contributing factor in the presence and strength of secondary flow motions.

The available literature on not only secondary flow motions, but also flow behaviour in partially-filled pipe flows, is sparse in general. However, in recent years, there have been more studies appearing that focus on secondary flow motions. Liu et al. [44] performed LES of partially-filled pipe flows, confirming that the strength of the secondary flow is positively correlated with the filling ratio of the pipe, and the number of counter-rotating vortices is higher for smaller filling ratios. Results obtained from a RANS model in [45] also confirm these observations.

As mentioned earlier, the velocity dip phenomenon is closely linked to the presence of secondary flow motions. The counter-rotating vortices influence the velocity distribution about the pipe cross-section, which in turn influences the local shear stress at the wall. Knight & Sterling [1] concluded that the distribution of wall shear stress becomes more uniform for greater filling ratios. Results of this study along with those in [39, 46] agree that maximum shear is experienced at the pipe bottom position. However, the authors of [19] conducted a Direct Numerical Simulation (DNS) study on semi-filled pipes for varied Reynolds' numbers. Through a highly resolved dataset, they conclude that the wall shear distribution is heavily influenced by the existence of inner secondary vortices present at the mixed corner, the flow region where the pipe wall and free surface intersect. This phenomenon is not yet observed experimentally. Indeed, experimental works on secondary flow motions in partially-filled pipes report only two large counter-rotating vortices spanning the whole cross-section [15, 46], with no reporting on additional smaller vortices at the mixed corner. Local wall shear stress maxima occur where secondary flow motions direct high momentum fluid towards the wall [47, 48]. Thus, the high-fidelity simulations capable of resolving smaller vortices conclude that the wall shear stress distribution varies based on inner secondary vortices size, a conclusion earlier experimental results were not able to capture. Specifically, it is found that wall shear stress increases towards the mixed corner as these vortices direct momentum towards the wall, allowing for greater variations in the wall shear stress about the wetted perimeter than once initially understood.

The authors of [49] utilised S-PIV measurements in conjunction with Proper Orthogonal Decomposition (POD) to decompose the PIV velocity field into separate scales within a partially-filled pipe flow. Through this they were able to conclude the existence of large, and very large scale motions (LSM/VLSM) in partially-filled pipe flows. These are coherent, elongated turbulent structures that scale with the flow depth. Having received a lot of attention in recent years, they were first noticed by Kim & Adrian [2] in turbulent pipe flow through recovering the pre-multiplied spectra of the streamwise velocity component. Peaks were observed at wavenumbers corresponding to two different energy-containing spatial scales λ within the time series data, this can be seen in Figure 2.3.




Figure 2.3: *Example of the streamwise velocity pre-multiplied power spectrum at $Re = 66400$, displaying peaks at multiple spatial scales in terms of pipe radius R [2].*

The consistent appearance of these scales among many experimental and numerical studies confirms that LSMs typically scale as $r \sim 4r$, and VLSMs as $10r \sim 30r$. Gibeau & Ghaemi [50] showed that VLSMs dominate the wall pressure at low frequencies. Ng et al. [49] further concluded that there exists an interplay between secondary flow motions and VLSM, with the strength of the secondary currents capable of moving VLSMs about the pipe cross-section. As such, even though the number of studies focusing specifically on flow structures in partially-filled pipe flows is small, it is clear that there exists an interplay between large-scale motions and secondary currents. Resolving near-wall quantities directly, and inferring

what aspects of these data are fingerprints of these aforementioned large flow structures present in these flows, remains a key open challenge in experimental work concerning flows of this type.

2.1.3 Wall pressure distribution

Classical models describing the fluctuating wall pressure beneath turbulent boundary layers have existed in the literature for over sixty years. Corcos [51] was one of the first to develop such models in 1962. Since then, many researchers have developed on these ideas to produce empirical models of the fluctuating wall pressure [52, 53, 54]. Assumptions such as a flat boundary and uniformity in all directions are given in these models. Such models are not appropriate for modelling wall pressure fluctuations in partially-filled pipe flows, as there exist significant variations in momentum transfer about the cross-section due to secondary flow motions. These variations strongly affect shear stress and pressure distributions about the wall, violating uniformity. Furthermore, in the flows considered here, the boundary is inherently curved, adding another layer of complexity beyond the scope of classical models.

Despite the vast amount of work on fluctuating wall pressure patterns in turbulent boundary layer flows and pressurised pipe flows, pipes running partially-filled remain relatively unexplored. As outlined in the previous section, the partially-filled nature of the flows considered in this work induces knock-on effects that drastically alter the flow dynamics inside the pipe, spanning from the free surface to the distribution of near-wall quantities about the wetted perimeter, with varied flow regimes resulting in different interplay between each component in the flow. Recent high-fidelity DNS simulations [19] provide reason to believe the wall shear distribution is more complex than initially realised, as seen in the earlier experimental studies of [1, 39]. This is illustrated in Figure 2.4.




Figure 2.4: *Azimuthal variation in normalised mean wall shear stress of a partially-filled pipe flow for different values of friction Reynolds number [3].*

To the best of the author’s knowledge, the distribution of fluctuating wall pressure about the cross-section in partially-filled pipe flows has not been extensively researched, unlike its tangentially-acting counterpart, wall shear stress. It is hypothesised that the azimuthal variation of the fluctuating wall pressure is similarly, if not more affected by the existence of secondary currents and signatures of LSMs and VLSMs. As such, the aim of this work is to recover properties of wall pressure fluctuations in turbulent partially-filled pipe flows using a wall-mounted sensor designed to dynamically respond to this fluctuating wall pressure, and relate recovered properties to key hydraulic characteristics and spatial scales.

While it is important to understand the motions and structures that govern partially-filled pipe flow dynamics, such as LSM/VLSM and secondary flow motions, a flush-mounted sensor at the wall can only respond to the effect they imprint on the wall in the form of pressure fluctuations. As such, the analysis presented in this thesis primarily focuses on statistical descriptors of the data collected at the wall, which will be introduced in later sections. By combining this statistical framework with the azimuthal sensor measurements, this work aims to quantify fluctuating wall pressure behaviour, and uncover the hydraulic information encoded within. Ultimately, this will link the wall pressure fluctuations to turbulent structures without the need for their direct resolution.

2.2 Flow measurement techniques

This section will introduce some common methods used for flow measurement purposes in varying contexts related to partially-filled pipe flows in general. Over the years, a variety of techniques have been developed to characterise different aspects of flow, including velocity, free surface, and the wall pressure. The choice of measurement technique is usually attributed to the phenomena of interest. For example, to isolate velocity distributions arising from secondary currents, optical methods visualising the fluid motion in the spanwise direction would be preferable.

In industrial and field applications, flow measurement in partially-filled pipe systems is typically achieved through a combination of point-based and bulk measurement techniques. Here, it is typical to obtain simple measurements that can be used to inform network operators of the internal state of the infrastructure, rather than detailed hydrodynamic or turbulent information. Common instrumentation includes ultrasonic and electromagnetic flow meters. Particularly in OCFs, but also for sewer systems, the velocity-area method is common [55]. Here, flow rate is approximated through multiplication of the flow cross-sectional area and estimated average flow velocity obtained in the measurement.

While useful for obtaining quick measurements needed for operators, they are often limited in their ability to resolve finer details of flow dynamics. Many techniques are intrusive, require calibration for specific flow conditions, or rely on assumptions regarding velocity distributions that may not hold in turbulent flows. These limitations motivate the need for more robust, non-invasive measurement techniques capable of recovering more than a single output, particularly in complex turbulent flow conditions where intrusive instrumentation may itself influence the flow behaviour. To gain a better understanding of the complexities involved with turbulent flows, it is advantageous to utilise approaches that can provide greater spatial and temporal information. In this subsection, we proceed by providing a review of measurement methods used for uncovering flow dynamics that have found utility in research and industrial settings. The measurement techniques to be presented in this section are not an exhaustive list, I merely aim to highlight the current state of the art, and how they are utilised in practice to gain understanding of hydraulic properties of pipe flows.

2.2.1 Free surface measurement

Measurement of the free surface behaviour is crucial for understanding OCFs and the underlying turbulent processes taking place between the surface and the channel bed. A short summary of relevant examples will be given here, and reviews of additional methods can be found in [56]. Wave probes are point-based sensors commonly utilised to obtain surface elevation measurements. By measuring the current flowing between two wires inside the probe, the recorded current changes in proportion to the level of submergence of the wires. Dolcetti et al. [57] utilised a non-equidistant array of wave probes spanning two orthogonal directions to gather spatial information on the dynamic free surface of a shallow flow with a rough bed. Across multiple Reynolds numbers, data from the wave probes were used to study dispersion relations of waves originating in shallow flows. Similarly, the authors of [58, 59] constructed a non-equidistant array of wave probes arranged into a streamwise array for measurement of water surface elevation in a laboratory flume with a rough bed for a range of flow conditions. From surface elevation data, they were able to construct a correlation model of the space-time pattern of the surface waves in the flume. The model parameters were found to change between flow conditions. While wave probes have clear utility in quantifying properties of dynamic free surface behaviour, they must be individually calibrated and, under extensive testing, cleaned in between measurements to ensure reliable surface elevation measurements. Data from Nichols et al. [60] were used to formulate a model attempting to relate free surface properties to the underlying flow field of OCFs. It was found the spatial oscillation frequency of the recovered spatial correlation patterns of the free surface were linked to the frequency of near-bed bursting events, suggesting interplay between the free surface and underlying coherent turbulent structures in OCFs.

Non-invasive methods for quantifying free surface behaviour can be used in an effort to remove drawbacks associated with invasive measurement techniques such as wave probes. A further extension of the previously mentioned studies introduced an airborne acoustical measurement setup for free-surface measurement [61]. By reflecting an ultrasonic wave from the free-surface of an OCF flume, analysis of the reflected acoustic field, specifically the phase of the scattered field, enabled for free surface quantities to be recovered. Time-varying estimates of the free surface fluctuations (for gravity waves and turbulent flow) recovered from this setup agreed closely with analogous measurements recovered from wave probes placed directly within the flow. Krynkin et al. [4] utilised an acoustic scattering

phenomenon and a speaker-receiver setup in conjunction with the Kirchhoff approximation to predict the scattered sound field reflected from the free surface of a shallow flow across multiple flow regimes, a schematic diagram of the experimental setup in this study is given in Figure 2.5. In the context of a spatially varying surface, the Kirchhoff approximation models the reflection of an acoustic wave by assuming each point on the surface reflects the incoming wave as if it were a tangent plane [62]. Through use of the Kirchhoff integral as a forward model of the acoustic field above the free surface, it was realised that the mean scattered acoustic pressure from the surface was dependent on the mean surface elevation.




Figure 2.5: *Diagram of experimental setup for acoustical measurements in [4].*

Through a minimisation procedure between measured receiver-obtained data and model predictions, they were able to recover the mean surface elevation height within a $\pm 5\%$ error when compared against wave probe measurements used for validation. Since its inception, this method has been expanded upon to build more robust iterations. These are inclusive of incorporating a Bayesian inversion approach to recover parameters of a sinusoidal surface [63], and a Digital Image Correlation (DIC) method to recover frequency-wavenumber spectra of free surface elevation to compare against theoretical dispersion relations [64]. The DIC method was also utilised by Wu et al. [65] to recover free surface properties of a partially-filled pipe flow fed with white titanium oxide colourant, details of which can be found in [66]. Comparison against wave probe data displayed good agreement between the two methods as a potential alternate method for recovery of free surface properties, including surface fluctuations and surface velocity.

Wang et al. [67] utilised the relative conductance measured between pairs of adjacent electrodes embedded in a novel sensor to produce a conductance profile along the length of the sensor. This sensor was placed in a flow containing several interfaces between sediment, water, and air. Through knowledge of the differing lengths of each electrode in

the array, and assumptions of the number of interfaces the electrodes cross, an expression for conductance was derived that can be related back to the relative distance between each interface. Thus, estimations of the water column height were inferred.

2.2.2 Flow visualisation techniques

In the field of turbulence, flow visualisation has long been a key tool for identifying complex flow structures. Many of the studies mentioned in the previous subsection that produce results advancing the current state of knowledge in partially-filled pipe flows utilise these flow visualisation techniques. One such method is Particle Image Velocimetry (PIV). One of the most widely utilised flow visualisation techniques, the flow is seeded with tracer particles that are buoyant in the flow medium, allowing these tracer particles to be carried along with the flow. Through exposure to a laser illuminating the plane of interest of the flow, a high speed camera records successive images of the illuminated tracer particles in this plane. The light scattered by the tracer particles, and successive snapshots of the state of the flow, allow positions of individual particles to be tracked in time. Quantities related to flow motion such as velocity are calculated through post-processing using techniques such as auto or cross-correlation [68, 69, 70]. However, if the flow is three-dimensional, the PIV-obtained flow velocity can be incorrect due to the out-of-plane motion and its projection onto the PIV plane used [71, 72]. S-PIV extends the conventional PIV approach by facing two cameras towards the PIV plane at different angles, allowing for reconstruction of all three velocity components of the flow within the PIV plane used by combining images from both cameras [73, 74]. A further extension of the aforementioned visualisation tools is that of *tomographic* PIV [75]. Similarly to S-PIV, illuminated tracer particles occupying some part of the flow are recorded by multiple cameras, however in this case it is a volume and not a plane that is illuminated. Utilising a tomographic reconstruction algorithm, the recorded images are combined to estimate the 3D light intensity distribution within the measurement volume. Cross-correlation of the 3D position data is applied at successive time steps to obtain the 3D velocity distribution on the volume. One can see the utility of S-PIV/tomographic PIV when considering consequences of turbulence, such as secondary flow motions. As such, many of the studies presented in section 2.1.2 utilise these measurement techniques.

While the methods mentioned here, and adjacent methods such as Particle Tracking Velocimetry (PTV) can resolve two, or three-dimensional behaviour that enhances under-

standing of flow dynamics that would otherwise be difficult to quantify, they are not without their drawbacks. Post-processing can be laborious, as the PIV plane needs to be split into smaller interrogation windows. The size of these windows relative to the size of the PIV plane determines the overall resolution of the resolved field, introducing a balance between resolution and post-processing time. Careful attention needs to be paid to spatial properties of the inserted tracers. The average spacing between particles in relation to the mean distance moved by particles between two successive images, denotes the *particle spacing displacement ratio*, this parameter heavily determines the traceability of individual tracer particles [76]. Furthermore, the optimal particle distribution is specific to the application. For example, to measure surface velocity using PIV, Weitbrecht et al. [77] required specific tracer particle density and dimensions to ensure they remained close to the surface.

2.2.3 Wall-mounted measurement

The optical measurement techniques introduced in section 2.2.2, while being highly effective for identifying coherent structures and cross-sectional velocity distributions if executed correctly, are limited in measuring quantities at the wall. In particular, pressure, and shear stress are difficult to capture with these techniques. As such, there exist wall-mounted sensors designed to capture information at the wall.

Devices such as accelerometers have been installed on the external wall of turbulent pipe flows to quantify relations between pipe vibration magnitude and fluid velocity inside the pipe [78]. While easy to install, installation onto the pipe exterior adds additional layers of transfer between the vibration sources in the flow, and the recorded acceleration of the pipe structure itself. Knight & Sterling [1] used a Preston tube to measure the local wall shear stress in a partially-filled pipe flow. In [79], the authors again used a Preston tube at the fluid-wall interface around the wetted perimeter in conjunction with pattern recognition techniques to predict continuous profiles of the boundary shear stress about the wetted perimeter for various flow conditions. Blake [80] recorded wall pressure measurements beneath a boundary layer in a wind tunnel with a pinhole-mounted condenser microphone to identify the interaction between the wall pressure and either rough and smooth boundaries. The authors of [81] assessed pressure fluctuations of developing turbulent pipe flows through piezo-electric transducers (PZT). They were installed behind pinholes inside in the wall, pressure taps were installed in these pinholes to operate as a piston, transferring

wall pressure fluctuations to the PZT elements. Clinch [82, 83] flush-mounted pressure transducers to the interior wall of a smooth-walled turbulent pipe flow to measure the wall pressure in streamwise and spanwise directions. From pressure transducer time series measurements of the wall pressure, streamwise and spanwise correlations were recovered and used to estimate bulk velocity of the flow. Abraham & Keith [84] flush-mounted a streamwise linear array of PZT sensors to estimate the frequency-wavenumber spectra of the wall pressure fluctuations in a water tunnel. From this they identified the main convective ridge and provided comparisons against well known spectra models including the Chase and Corcos models. Included in the analyses was an investigation into the effect of the sensor size in attenuating signatures from small scale structures. A similar study was conducted by Hu [5].




Figure 2.6: *Array of pressure sensors used for measurement in [5].*

This study assessed a range of sensors designed to respond to wall pressure fluctuations in turbulent boundary layers. These included wall-mounted and pinhole-mounted combinations of microphones and Kulite pressure transducers of various radii to assess the sensor-size-related attenuation of wall pressure reading, images of the sensors used are given in Figure 2.6. It is known that spatial averaging of signals occurs when recorded over a sensor area [85, 86, 87]. As such, recorded wall pressure spectra typically represent a convolution between the true wall pressure and a transfer function dependent on the sensor sensitivity [88]. This justifies the need for a correction factor unique to the sensor surface area and installation method. Indeed, pin-hole installations minimise exposure of the sensor surface area to wall pressure, but the cavity present in these configurations induces Helmholtz resonances [5].

In the context of partially-filled pipe flows, and where they may occur outside of laboratory or research environments, many of the sensing techniques introduced here, would not be applicable. Pipe flows exist in closed conduits by definition, and so it would be impractical to attempt installations of speaker-receiver systems above the free surface to recover free

surface properties as given in section 2.2.1 in operational pipes. For flow visualisation techniques, it is paramount that the flowing medium be clear, otherwise tracer particles would not be illuminated. Again, many operational pipes outside of laboratory environments are opaque, and so application of any camera-based techniques would not be possible. In contrast, wall-mounted sensors offer a practical solution for recovering hydraulic information in operational conditions, as they can be installed directly in the conduit. With appropriate design, such in-field sensors would be exposed to similar conditions as their laboratory-controlled counterparts. This capability makes this class of sensors an attractive approach for not only research purposes, but also for practical application of monitoring hydraulic infrastructure through information contained at the wall. To the best of the authors knowledge, most studies characterising wall pressure through the use of wall-mounted sensors take place in turbulent boundary layer flows. As mentioned previously, numerical models of the fluctuating wall pressure assume a flat boundary and spatial uniformity. The availability of literature inferring hydraulic conditions from data collected at the wall in pipe flows, and especially partially-filled pipe flows, is scarce.

2.3 Fibre-optic sensors

The scarcity of wall-mounted sensor systems for hydraulic infrastructure monitoring beyond those based on pressure transducers is driving the search for alternative technologies that are robust, minimally intrusive, and capable of surviving deployment in operational conduits. Fibre-optic cables compete as a potential candidate for monitoring purposes such as this. Since emerging as a technology in the 1960's, fibre-optic cables have traditionally been used as a communications device. Consisting of a glass core several μm in diameter, surrounded by plastic cladding layers, these cables enable communication at the speed of light, acting as waveguides through total internal reflection of light pulses propagated through the core. Their sensitivity to mechanical strain and temperature variations has been leveraged in recent years to convert them into a new family of sensor systems. In this subsection, the operating principles behind fibre-optic cables will be introduced, as well as the optical phenomena that can be leveraged to convert physical perturbations into measurable fluctuations of this unique optical waveguide. This will be followed by a review of several use cases to illustrate where these robust family of sensors have found utility in monitoring hydraulic infrastructure, establishing how the work in this thesis advances on current use cases and operating principles currently employed by fibre-optic sensors.

2.3.1 Technology and operating principle

Snell's law describes the change of propagation direction when a ray of light is transmitted between two media of different refractive indices, stated as

$$n_1 \sin \gamma_i = n_2 \sin \gamma_r. \quad (2.7)$$

Here, n_1 , n_2 are the refractive indices (RI) of the first and second media, respectively. The RI of a given medium is the ratio of the speed of light in a vacuum and the speed of light in the medium. The angle γ_i is the angle subtended between the direction of incoming light in the first medium and the line perpendicular to the boundary between the two media, also known as the *angle of incidence*. The angle γ_r is the angle subtended between the direction of outgoing light in the second medium and the line perpendicular to the boundary between the two media, also known as the *angle of refraction*. This angular change in propagation is merely a consequence of Fermat's principle [89], whereby path taken by a ray of light is the one minimising the time taken. For particular combinations of the parameters in Equation (2.7), the ray of light may not pass into the second medium. This occurs when $n_1 > n_2$ and $\gamma_i > \arcsin\left(\frac{n_2}{n_1} \sin \gamma_r\right)$. When the ray intercepts the boundary for these conditions, the ray is reflected back into the first medium from the boundary. This effect, known as *total internal reflection* (TIR), is the fundamental mechanism confining light within the glass core of fibre-optic cables, this is demonstrated in Figure 2.7 between the core and cladding boundaries of the fibre-optic cable.




Figure 2.7: Schematic diagram of TIR in a fibre-optic cable [6].

Total internal reflection occurs when light inside the core intersects with the boundary between the core and the outer cladding layer. The low refractive index of the cladding relative to that of the core causes TIR, constraining the light to propagate only down the core. While TIR enables communication at the speed of light, the cables ability to act as a sensor is due perturbations in the cables physical environment directly altering properties of this light. Mechanical strain, pressure, temperature, humidity, among others, are all parameters capable of altering core geometry or refractive index, which in turn manipulate the optical signal in response to some physical measurand. The optical or geometrical properties that are changed in response to external parameters define the type of fibre-optic sensor (FOS), of which there exist many. Comprehensive reviews of the available types of fibre-optic sensing can be found in [90, 91, 92, 93]. To keep the content of this thesis focused within the realm of monitoring hydraulic processes and infrastructure, I will introduce the two main FOS classes used most appreciably in this field: scattering-based and wavelength-modulated.

Rayleigh scattering is a phenomenon whereby an electromagnetic wave encountering small density fluctuations in a medium produces a scattered component of that wave of identical wavelength to the incident wave [94]. In a FOS, this occurs when incoming light interacts with density fluctuations in the core material that are of a much smaller size than the wavelength. Measurement of this type of scattering can be separated into two classes. The first is optical time-domain reflectometry (OTDR) [95]. This sensing method sends a pulse of light down the core, light scatters back from each section of the fibre. By analysing the time of flight of the light pulse, the location of scattering can be identified. In phase-sensitive OTDR (φ -OTDR), the phase of the backscattered signal from a specific fibre location is tracked over time. A dynamic strain event alters the optical path length, inducing a phase shift, which is proportional to the applied strain. While this method is capable of operating on multiple kilometres of fibre at once, it generally suffers from poor spatial resolution [94]. The frequency domain counterpart, optical frequency-domain reflectometry (OFDR) sweeps the frequency of the input signal linearly in time [96]. The distance travelled before being scattered backwards towards the source is recovered from a Fast Fourier Transform (FFT) of the signal [6]. The combined outward and scattered signal is composed of reflections from different sections of the fibre, creating a beat frequency spectrum where each frequency component corresponds to a specific distance. The frequency content of this signal can be related back to the scattering points based on the time of flight and sweep rate of the input signal [94]. While scattering-based methods pro-

vide long-range capabilities, they suffer from poor localisation. As such, sensors operating on these larger scales are typically referred to as distributed fibre-optic sensors (DFOS), making them applicable for perimeter monitoring and other large-scale systems [97, 98].

The aforementioned wavelength-modulated sensing overcomes poor localisation issues through alteration of the core's optical properties at specific positions along its length. This creates individual elements at specified positions sensitive to external perturbations. Most notably in this category are fibre Bragg gratings (FBG), first introduced by Hill et al. [99]. An individual grating is written into a fibre through periodic modulation of the refractive index of the fibre's core over a specified length. This is achieved by exposing the core to an intense optical interference pattern at these points of modulation. When a broadband light source intercepts the grating, reflecting components from each of the boundaries between the points of modulated refractive index occur in phase, reflecting a singular wavelength backwards along the core, this dichroic effect transmits all remaining wavelengths. The particular value of this reflected wavelength, called the *Bragg wavelength*, is dependent on the properties of the grating itself [100], this is given by

$$\lambda_B = 2\Lambda n_{\text{eff}}. \quad (2.8)$$

In the above expression, Λ is the grating period, and n_{eff} is the effective refractive index of the fibre core, both of which can be sensitive to changes in temperature T or axial strain ε . Strain alters Λ through elongation of the cable, and alters n_{eff} through the photo-elastic effect. Temperature alters Λ through thermal expansion of the core, and alters n_{eff} through the thermo-optic effect [100]. Any changes to either n_{eff} or Λ , induces a Bragg wavelength shift $\Delta\lambda_B$, proportional to the change in the corresponding temperature ΔT , or strain ε . This relation can be expressed as

$$\frac{\Delta\lambda_B}{\lambda_B} = (1 - \rho_e)\varepsilon + (\alpha_e + \xi_o)\Delta T. \quad (2.9)$$

This expression details the full contribution from both strain and temperature in the Bragg wavelength shift. Here, ρ_e is the photo-elastic coefficient of the core material, α_e is the thermal expansion, and ξ_o is the thermo-optic coefficient [101]. This cross-sensitivity to both strain and temperature may require compensation strategies to isolate the effect of each variable in the evolution of $\Delta\lambda_B$, depending on the specific application. Aside from this additional consideration, multiple FBG elements can be written into a single fibre core, a technique known as multiplexing.




Figure 2.8: Schematic diagram of FBG strain response [7].

Provided the Bragg wavelengths of each grating are spaced sufficiently far from each other, the transmitted spectrum from each FBG does not interfere with neighbouring elements. This allows for FBGs to be transformed from a single-point measurement sensor into what is called a *quasi-distributed* fibre-optic sensor, with spatial resolution dependent on the FBG spacing.

2.3.2 Towards deployment in flow monitoring

In the previous subsection, the most common types of fibre-optic sensing, and their measurement principles, were introduced. Now, our attention is turned to their deployment and capabilities in monitoring live infrastructure, with a focus on the utility of FBG technology in hydraulics and flow monitoring use cases. All sensors monitoring the health or condition of some structure typically recover the response of the structure to some physical process, whether it be dynamic loading, temperature variations, or chemical changes. In the context of monitoring hydraulic infrastructure, the fields inducing loading on some structure are naturally spatially and temporally evolving fields. Embedding of FOS onto these structures enables spatial and temporal variations to be resolved, allowing for an abundance of data to be collected from a single sensor, making it a promising tool for advancing flow monitoring and condition assessment. To the best of the author's knowledge, work leveraging fibre-optic sensing to monitor hydraulic processes in partially-filled pipes is sparse, even more so when attempting to utilise this sensing method to collect data at the interior pipe wall. As such, although fibre-optic sensing is common among structural

health monitoring [102, 103], use cases relevant to the general area of monitoring hydraulic infrastructure shall be introduced as a means to provide insight into the current capabilities of fibre-optic sensing in this area.

Multiple studies have utilised FBG technology to respond to changes in hoop strain of a pipe structure. This is a consequence of the internal pressure inside the pipe and changes in wall thickness. The authors of [104] demonstrated that externally-installed FBGs were capable of identifying leaks in a pressurised pipe due to the relation with hoop strain. An ‘optical fibre accelerometer’ system utilising FBGs was attached to the exterior of a buried steel pipeline, and was shown to track changes in natural frequency of the pipe structure when subjected to localised corrosion over time [105]. Studies such as this showcase how monitoring of external parameters leads to sound inference of internal conditions.

More specific to direct flow monitoring, a unique FOS was developed in [8] in an effort to measure flow rate in pipe flow. From the mathematical details given in the study it is assumed the flow is pressurised. A hollow cylindrical cantilever with a circular target at the free end was attached to the interior pipe wall, oriented perpendicularly to the flow direction, a schematic of this setup is given in Figure 2.9. Three FBGs were attached, two responding to strain induced by cantilever bending in response to flow, with the third uncoupled to calibrate against thermal changes. Through a simple mathematical formulation, recovered strain was converted into corresponding flow velocity at the pipe centre, agreeing well with simulation results.




Figure 2.9: Schematic diagram of FBG sensor developed in [8].

The main shortcoming of this unique sensor however is the lack of spatial resolution, and intrusion into the flow. In the context of monitoring live infrastructure, a dense coverage in the sensor network is advantageous. By way of the construction of this specific sensor, it cannot capture spatial variations along the pipe wall, limiting its utility in the field. Takashima et al. [106] presented a water flow meter that uses FBG sensors and a cross-correlation technique to measure flow velocity based on the time delay of vortex signals. The mentioned signals were produced by fluid moving past a bluff body manually installed in the flow domain. While a more enhanced method of analysis is utilised to recover flow properties, the method requires manual disruption of flow, making it unsuitable for live and buried pipelines.

A multi-parameter sensor utilising FBG technology was produced in [107], capable of responding to pressure, acceleration, and temperature changes. Installed on the exterior of a hydraulic pipe system, the results from the FBG sensors were in good agreement with measurements obtained from reference electrical sensors. FBG-recovered modal properties such as natural frequencies and damping ratios displayed excellent agreement with accelerometer-equivalent results. Looseness of the clamping system was also identified by tracking changes to the natural frequencies recovered by the FBG sensors. An externally-mounted FBG sensing system was constructed to work on toroidal sections of pipe [108]. Centripetal forces induced from the curved path of flow exerts forcing onto the pipe wall, resulting in displacement at the FBG location. Wavelength shifts in response to changes in flow velocity were recorded to empirically relate the two, showing errors of only 4.6%.

Collectively, the studies presented in this subsection present the utility of FBG technology in hydraulic infrastructure monitoring, while simultaneously highlighting an all-important trade-off between ease of operation and sensor capability. While non-invasive monitoring through external attachment is relatively easy to install on operational infrastructure, the steps taken to achieve inference of internal flow properties are laborious, owing to the indirect measurement of flow properties. Conversely, direct access to flow inside the infrastructure necessitates intrusion of the sensor into the flow domain, disrupting the flow dynamics. Even measurement principles that recover flow properties directly have only been developed to operate on pressurised pipelines, a condition less common in many operational drainage and sewer conduits, which are often partially-filled. Dynamics of these flows is fundamentally more complex than fully-filled counterparts, often being turbulent in nature. As such, modern sensing techniques must be produced to directly monitor flows

departing from the idealised pressurised flow dynamics. It is precisely this gap, the need for a direct, non-intrusive flow measurement technique capable of recovering hydraulic conditions in partially-filled pipes, that this thesis aims to address.

2.4 Inverse problems in engineering

In some cases, the direct measurand provided by sensors may not be sufficient to provide the information necessary to achieve the required monitoring task. It is most common for sensors to monitor the response of the system to some external process or loading. In other words, sensors record the output of a system, which is dependent on some underlying inputs that are currently unknown to the user. These inputs can be a loading profile induced by wind, or material properties of some structure, such as natural frequencies and damping ratios. In either case, the task of inferring the unobservable inputs, given the output data available to the user, is known as an *inverse problem*. This is in contrast to the *forward problem*, where observable outcomes are computed given the inputs and, usually, a set of equations governing the dynamics of the system in question. Inverse problems are, by nature, susceptible to *ill-posedness*. In these types of problems, there may be many potential inputs θ explaining the observed data Y , or none at all. Traditional optimisation methods, such as least-squares optimisation, aim to seek the input θ minimising the error between Y and model-predicted outputs. However, approaches such as this are often vulnerable to noise in the observed data Y , and do not address the issue of multiple minima existing.

2.4.1 A Bayesian perspective on inversion

First, consider an experiment producing a set of measured data Y that are dependent on some unknown parameters θ ; these data may additionally depend on some known input variables, which will be termed X . One may have access to a mathematical model $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ that perfectly describes the dynamics of the system, allowing for an optimal θ to be chosen that minimises the error between measured data and model-produced data. Here, $\mathcal{X}$, $\mathcal{Y}$ denote the spaces of the inputs X and the output data Y , respectively. Probabilistically, however, one could instead construct, using the same model f_θ , a probability density function $P(Y|\theta; X)$ which expresses how likely the measured Y is given the known X and choice of θ ; this expression will be termed the *likelihood* of the model.

In Bayesian statistics, the *posterior distribution* is related to the likelihood via the relation

$$P(\theta|Y;X) = \frac{P(Y|\theta;X)P(\theta)}{P(Y)}. \quad (2.10)$$

In Equation (2.10), one’s prior beliefs about the distribution of the parameters θ are contained within the *prior* $P(\theta)$ [109]. The posterior $P(\theta|Y;X)$ is the probability distribution which expresses belief over the unknown θ once data have been observed. It balances the influence of the user’s *a priori* belief and the information contained in the observed data through the likelihood, and is the central object providing inference on parameter values given measured data Y . The denominator is known as the *marginal*. This term serves to normalise, ensuring the posterior distribution integrates to one. Under the Bayesian paradigm, attention shifts from inferring a ‘best’ estimate of the unknown quantities θ to understanding the chance of θ taking every possible value, as expressed through the posterior distribution.

Armed with the essential components of a Bayesian analysis, consider how to compute the posterior distribution. The likelihood and prior are both determined by the user. In the context of an engineering problem, the likelihood may consist of the model of the physical process one seeks to understand f_θ and some probabilistic specification of the noise on the observed data. For a certain subset of models with specific choices for $P(Y|\theta;X)$ and $P(\theta)$ it is then possible to compute $P(\theta|Y;X)$ directly and in closed form. This can be a very useful and powerful route for analysis. However, in the majority of cases of interest, the nature of the models (f_θ) which are being used necessitates a form of $P(Y|\theta;X)$ which is not amenable to this exact solution. Such cases of *intractable models* require a numerical approach to solving for $P(\theta|Y;X)$. Introductory texts on the general frameworks of Bayesian inference for inverse problems can be found in [110, 111], however a few use cases of frameworks relevant to the content of this thesis shall be introduced next. The types of inverse problems addressed in this thesis are those of parameter estimation. As such, tasks of this type shall be the focus of the following subsection.

2.4.2 Parameter estimation via Bayesian methodologies

One of the most common Bayesian methods for parameter estimation is MCMC methods [112]. The authors of [113] utilised the Metropolis-Hastings (MH) algorithm to estimate the coefficients of polynomial terms in a linear time-invariant (LTI) system. Variations of adaptive and parallel adaptive MCMC methods were also utilised in conjunction with

priors imposed on the variance of the to-be-estimated parameters to estimate kinetic rate parameters of Ordinary Differential Equation (ODE) models in systems biology [114]. By posing the bending stiffness as a damage-sensitive parameter discretised along the length of an Euler-Bernoulli beam, the authors of [115] developed a modified MH scheme that outperformed conventional MH on the given task. In a problem of acoustic scattering, Johnson et al. [63] utilised an adaptive MH scheme to recover the parameters of a sinusoidal rough surface from the surface’s scattered acoustic field. A requirement of MCMC methods is the ability to compute the likelihood function, how likely observed data Y are given inputs θ . Especially in cases of high-dimensionality or unknown likelihoods, this computation may be expensive or impossible to complete. For the most part, this limits the applicability of standard MCMC schemes for low-dimensional problems, or where there exists a relatively simple likelihood. This limitation motivates the need for other inference methods that can provide sound inference of model parameters (through computation of the posterior distribution) in the presence of unworkable likelihoods.

Among the most popular likelihood-free methods is ABC [116]. Methods such as this circumvent the need for a likelihood through simulation of artificial data Y^* from inputs sampled from the prior. Direct comparison between simulated and observed data under some definition of distance characterised by a metric D , enables assessment of inputs in their plausibility of being responsible for the observed data. If the simulated data Y^* agree with each other up to some threshold ϵ under D ($D(Y, Y^*) \leq \epsilon$), the input θ^* is accepted into an approximate posterior distribution. The threshold ϵ controls the acceptance rate and accuracy of the approximated posterior distribution. On its own, ABC experiences very low acceptance rates, and improvements are based on enhancing the sampling procedure to target more appropriate subsets of the parameter space [117].

Building on this, one key advancement is that of ABC-SMC [118]. This scheme operates by projecting samples θ through a series of intermediate distributions. The distribution from which inputs are sampled in a given wave is the posterior of the accepted samples from the previous wave. The gradual change of the distribution from which to sample potential inputs is the main driving force behind enhanced efficiency. By concentrating efforts on increasingly appropriate areas of the parameter space, less effort is wasted testing suitability of samples occupying areas of parameter space sampled from less-so in the previous wave. The technical details have been spared here, but can be found in greater depth in Chapter 4.5. This enhanced method has found utility in estimating parameters of the repressilator

model for gene regulatory systems [119]. Used to estimate a four-dimensional quantity parametrising a set of six ordinary differential equations, the posterior distribution obtained over 19 waves agrees well with analogous results from previous studies. An improved ABC-SMC (IABC-SMC) scheme was used for model selection and parameter estimation on models of infectious diseases, showcasing improved accuracy of parameter estimates and model selection speed [120]. The authors of [121] combined ABC-SMC with the random forest (RF) algorithm to infer parameters of the Lokta-Volterra model. It was found that their combined architecture typically results in better posterior approximations than previous methods solely relying on RF methods alone, simultaneously maintaining performance comparable to standard ABC schemes. These applications demonstrate the versatility of ABC-SMC as a tool for probabilistic inference tasks, the major drawback of these schemes is that the parameters to be inferred, are continuous. In practice, it is ubiquitous that the goal of inverse problems is to gain an understanding of discrete, categorical, or mixed parameters. Part of the contribution of this thesis is to address the inapplicability of ABC-SMC to varied parameter types by introducing a novel adaptation *combinatorial* ABC-SMC designed to perform inference on high-dimensional binary, rather than continuous variables. The details of which are presented in more detail in Chapters 3 and 5.

2.5 Literature review summary

The literature reviewed in this chapter serves to highlight not only the complex and multi-scale properties of turbulent partially-filled pipe flows, but also the complexities involved in measurement challenges of such flows in the laboratory and out in the field. The key feature complicating the dynamics of partially-filled flows as compared to their fully-filled counterparts is the free surface. The presence of which introduces secondary flow motions, non-uniform wall shear stress distributions, and consequently non-uniform wall pressure fluctuation distributions. Experimental and numerical studies on turbulence in these flows utilising methods ranging from S-PIV to LES have significantly advanced understanding of these flows, yet most approaches are restricted to simulation or laboratory environments only, and do not address the wall pressure distribution. As of yet, the number of studies further advancing understanding of partially-filled pipe flows with data collected from the wall, while also providing new measurement techniques applicable to monitoring of these flows in the field, remains limited.

Ranging from free surface characterisation, velocity field reconstruction, and wall-based measurements, many measurement technologies exist that are utilised in the laboratory and in the field to this day. However, where many of these methods introduced in this chapter fall short is their intrusion into the flow, or sparse coverage of pipe lengths. Fibre-optic sensing technologies offer a promising alternative due to their non-invasive nature, durability, and potential for distributed measurement. Having been adapted for various sensing scenarios such as structural health monitoring and monitoring of hydraulic infrastructure, they have so far not been implemented in a non-invasive fashion with the goal of uncovering turbulent phenomena from partially-filled pipe flows through their response to fluctuating pressure at the wall. In particular, there is a lack of established methodologies for interpreting FOS response in terms of the underlying turbulence in these flows.

More broadly, the problem of inferring unknown system behaviour from raw sensor data, which is inherently a response to some unknown source, can be viewed as an inverse problem. Such problems may be ill-posed and sensitive to uncertainty in sensor measurements, modelling choices, or unknown parameters governing system behaviour. The Bayesian approaches reviewed here encompass some of the more common methods leveraged to perform data-driven inference on real world engineering scenarios. A focus has been given on parameter estimation schemes to outline how the extension of the well-known ABC-SMC scheme presented in this thesis was a necessity in resolving the niche form of uncertainty observed when calibrating the primitive FOS response to controlled excitation.

From the literature reviewed, clear gaps can be identified that this thesis aims to address. Firstly, there is limited understanding of the fluctuating wall pressure spanning the wetter perimeter in turbulent partially-filled pipe flows, especially when considering how the organisation of these pressure fluctuations may be fingerprints of secondary flow structures encompassing the whole flow field. There is also a lack of non-invasive sensing methodologies capable of capturing the targeted dynamics of these flows. Finally, there is a need for robust inference frameworks to interpret sensor response in the presence of uncertainty, particularly when measurements involve unique and complex calibration procedures. The initial focus is placed on resolving uncertainty in a primitive version of this FOS through the application (and extension) of Bayesian inference methods, before progressing to the development of a novel fibre-optic sensor embedded into the wall of a partially-filled pipe for hydraulic testing. This enables for the FOS response to wall pressure to be captured for the first time in flows of this type, resulting in the recovery of key flow properties, including

bulk velocity, spatial correlation structures, and signatures of turbulent large-scale motions and secondary flow motion, directly from the wall. In doing so, the work contributes to both the advancement of non-invasive monitoring technologies for hydraulic infrastructure and the broader understanding of turbulent partially-filled pipe flows.

Chapter 3

Methodology

This chapter introduces the overall methodological frameworks utilised throughout this thesis. This includes experimental, mathematical, and statistical methodologies that have been leveraged to obtain the results to be presented in the proceeding chapters. As the current study adopts methodologies spanning the above mentioned three broad components, it is the aim of this chapter to provide an overall narrative to connect the chapters together to form a single study.

Chapter 4 centres on developing Bayesian methodologies applied to datasets obtained from experiments on a preliminary fibre-optic sensor developed by nuron Ltd, a company who were tied to the initial stages of this project through their research on fibre-optic sensing technology for live pipeline monitoring. Sensor design considerations had already been completed before I began my research. As such, the contribution of this chapter is purely statistical, and was a relevant step needed to resolve uncertainty that was identified in the dynamic response of this preliminary sensor during testing. While the sensor configuration investigated in Chapter 4 differs from the final design adopted in subsequent chapters, the work provides an important methodological foundation for addressing uncertainty in fibre-optic sensor response. Building on insights gained here, a revised FOS was later developed in the subsequent chapters that was deemed appropriate for hydraulic testing, as seen in Chapters 5 and 6. As such, Chapter 4 represents an initial stage in the overall research, where challenges associated with sensor response are addressed in a controlled setting, before being extended to more complex hydraulic environments in the remainder of the thesis.

Chapters 5 and 6 centre around experimental design, data collection, and analysis of the FOS response to fluctuating wall pressure of turbulent flows in a partially-filled pipe rig. As such, one of the methodologies relevant to these chapters is the signal processing used to manipulate raw sensor data for physics-based inference. It is also important to consider the measurement techniques used to setup the flow conditions, and give details of design decisions that needed to be implemented before the data collection process could take place for these works. Inclusion of these details of the experimental methodology will also be given, to guide the reader through the research process from experimental setup, to raw data collection, right up to the construction of the presented results through data analysis pipelines.

To provide an overall methodological framework of the work in this thesis, the current chapter separates the methodology into three distinct sections: mathematical and statistical methods utilised for analysis of wall pressure fluctuations in chapters 5 and 6, Bayesian methodologies used throughout chapters 4, 5, and 6, and experimental design decisions that were central to the data collection processes carried out in chapters 5 and 6.

3.1 A statistical description of wall pressure fluctuations

Content contained within the current sub-section is necessary to build up the mathematical and statistical formalism needed to understand the nature of the wall pressure fluctuations, and the standard signal processing techniques applied to raw sensor data to arrive at meaningful physical interpretations of the collected data. Describing the balance of momentum in an incompressible fluid, the Navier-Stokes equations provide an analytic formulation of the relation between pressure p and velocity field $\mathbf{v}$, typically written as

$$\rho \frac{\partial \mathbf{v}}{\partial t} + \rho (\mathbf{v} \cdot \nabla) \mathbf{v} = -\nabla p + \hat{\mu} \nabla^2 \mathbf{v}. \quad (3.1)$$

Here, $\hat{\mu}$ is the dynamic viscosity of the fluid, and ρ is the fluid density. The incompressibility condition demands that $\nabla \cdot \mathbf{v} = 0$. After taking the divergence of these equations and some rearranging, one arrives at the Poisson equation for the pressure

$$\nabla^2 p = -\rho \nabla \cdot (\mathbf{v} \cdot \nabla) \mathbf{v} - \rho \nabla \cdot (\mathbf{v} \cdot \nabla \mathbf{v}) \quad (3.2)$$

Equation 3.2 tells us that the pressure, even at the wall, isn't purely local. It is determined from contributions of the velocity field across the whole fluid domain, which is especially complex in turbulent flows, where multiple spatial scales exist within the flow. This makes exact prediction of the pressure field impractical in practice, as it would require knowledge of the whole velocity field. This limitation motivates the statistical approach taken throughout this work, the wall pressure is characterised through its measurable statistical properties (introduced in the following subsection), which prove useful in recovering more global flow dynamics.

Near the wall in fully developed turbulent flows, the velocity field $\mathbf{v}$ is assumed to be composed of a constant component, overlaid with a fluctuating component arising from turbulent fluctuations in the velocity field. The consequential variations in pressure at the wall also fluctuate about a mean value, with these fluctuations being stationary in time. This stationarity implies that statistical moments such as mean and variance remain constant over time. Spatial data is usually sparse in comparison to temporal data. As such, the stationary process is usually assumed to be ergodic to allow for time averages to be equated to ensemble averages.

3.1.1 Space-time correlations

The stationarity assumption enables us to look at differences in wall pressure between two points separated in the direction of mean flow as a function of streamwise separation ξ and temporal separation (or time lag) τ only, and not their absolute positions in space or time.

For the fluctuating wall pressure $p(x, t)$, stationary in space and time, the space-time cross-correlation between two points x_1, x_2 such that $\xi = |x_1 - x_2|$ is given by

$$R_{pp}(\xi, \tau) = \langle p(x_1, t) p(x_2, t + \tau) \rangle. \quad (3.3)$$

Here, the subscript denotes correlation of p with itself, and $\langle \cdot \rangle$ denotes the ensemble average, which can be replaced by a time average under the ergodicity assumption, such that

$$R_{pp}(\xi, \tau) = \frac{1}{T} \int_0^T p(x_1, t) p(x_2, t + \tau) dt, \quad (3.4)$$

for a given (x_1, x_2) pair, where T is the length of the function in time. The function R_{pp} characterises similarity between p at varied spatial and temporal shifted versions of itself.

the value of τ at which R_{pp} is maximised corresponds to the time lag at which $p(x_1, t)$ is most similar to $p(x_2, t)$. In the context of fluctuating wall pressure signals separated in the direction of mean flow, this value τ^* is the time delay that denotes the convection of wall pressure fluctuations between the two positions. This, along with the known distance ξ between the two pressure signals, can be used to estimate the velocity of the wall pressure, also known as the convection velocity $U_c = \xi/\tau^*$. Removing dependencies of separation and time lag, the auto and spatial correlations are defined, respectively. These are given by

$$R_{pp}(0, \tau) = \frac{1}{T} \int_0^T p(x_1, t) p(x_1, t + \tau) dt, \quad (3.5)$$

and

$$R_{pp}(\xi, 0) = \frac{1}{T} \int_0^T p(x_1, t) p(x_2, t) dt, \quad (3.6)$$

respectively.

Cross-correlation is, by definition, a time domain method. Transforming its various forms into the frequency domain via the Fourier transform allows for a more detailed view of the temporal and spatial properties, as frequency domain representations can be viewed as a distribution over available frequencies, isolating different scales present in the data.

Taking the Fourier transform of $R_{pp}(\xi, \tau)$, the cross-spectral density (CSD) of wall pressure for a given (x_1, x_2) pair such that $\xi = |x_1 - x_2|$, is given as

$$\Gamma_{pp}(\xi, \omega) = \frac{1}{2\pi} \int_{-\infty}^{\infty} R_{pp}(\xi, \tau) e^{-i\omega\tau} d\tau. \quad (3.7)$$

The CSD provides the degree of correlation between $p(x_1, t)$ and $p(x_2, t)$ across all angular frequencies ω , where $\omega = 2\pi f$ and f is the temporal frequency in Hz. Performing the above transform on the autocorrelation provides the power-spectral density (PSD), given by

$$\phi_{pp}(\omega) = \frac{1}{2\pi} \int_{-\infty}^{\infty} R_{pp}(0, \tau) e^{-i\omega\tau} d\tau. \quad (3.8)$$

The PSD represents the distribution of variance with frequency ω , and is equivalent to the CSD evaluated at $\xi = 0$.

Taking the spatial Fourier transform of the CSD along ξ ,

$$\Phi_{pp}(k, \omega) = \frac{1}{(2\pi)^2} \int_{-\infty}^{\infty} \Gamma_{pp}(\xi, \omega) e^{-ik\xi} d\xi \quad (3.9)$$

gives the frequency-wavenumber spectrum of p . In the context of fluctuating wall pressure signals separated in the streamwise direction, k is the streamwise wavenumber. This function provides a complete spectral representation of the temporal and spatial scales contained in $p(x, t)$. The magnitude of Φ_{pp} at a given (k, ω) indicates the energy of structures present in the fluctuating wall pressure with that specific streamwise wavenumber and oscillation frequency. The main contributor of energy in Φ_{pp} is the *convective ridge*. This indicates the convection of flow of pressure past the wall at the convection velocity U_c , which manifests as a ridge of high energy, and can be expressed as $U_c = \omega/k$ in $\omega - k$ space.

3.1.2 Taylor's hypothesis

The presence of a convective ridge of precise structure suggests a relation between temporal and spatial scales in the wall pressure fluctuations; this relation is formalised through *Taylor's hypothesis* [122]. This hypothesis reasons that, if fluctuations in a turbulent quantity F are small in relation to the mean flow velocity, then time histories of this quantity at a fixed point can be viewed as a spatial pattern frozen in time, translating past the fixed point. For example, a turbulent velocity field $\mathbf{v} = \bar{\mathbf{v}} + \tilde{\mathbf{v}}$ with mean component $\bar{\mathbf{v}} = (V, 0, 0)$, and turbulent component $\tilde{\mathbf{v}}$ such that $|\tilde{\mathbf{v}}| \ll |\bar{\mathbf{v}}|$, Taylor's hypothesis states that

$$\frac{dF}{dt} = \frac{\partial F}{\partial t} + V \frac{\partial F}{\partial x} = 0, \quad (3.10)$$

establishing that

$$\frac{\partial}{\partial t} = -V \frac{\partial}{\partial x} \quad (3.11)$$

so that the time evolution of F at a fixed point is related in full to the convection of the spatial field by the mean flow [123, 124].

Much like white light can be thought of as a superposition of all wavelengths, turbulent flow can be considered similarly as a superposition of structures of varied spatial scales λ . The inherent drawback of Taylor's hypothesis is the assumption that all components travel with velocity V , independent of frequency f , or wavelength λ . Particularly close to the wall where there exists high shear, larger-scale structures tend to propagate at

velocities faster than the local mean V . In the inverted domain, the conversion is given as $k = 2\pi f/V$, this mapping imposed by Taylor’s hypothesis then maps these larger structures to higher than realistic wavenumbers, owing to a misrepresented value of their true velocity. This skews the converted wavenumber spectrum, shifting energy away from low wavenumber (longer wavelength), obscuring the true, flatter energy distribution at large scales [125, 126, 127]. Even with this limitation, Taylor’s hypothesis is still a very powerful assumption that enables recovery of spatial properties of turbulence from single-point time series measurements alone.

3.2 The Bayesian approach

Owing to the difficulty in computation of the posterior distribution in Equation (2.10) for all but simple cases, practical Bayesian inference techniques typically rely on numerical methods to approximate $P(\theta | Y; X)$. More specifically, samples are generated through a stochastic sampling approach. After sufficient time, generated samples are distributed in line with the posterior. In this section, the sampling procedures of the two main Bayesian methodologies used in this thesis shall be introduced.

3.2.1 Markov Chain Monte Carlo

The dominant methodology for sampling-based approaches when dealing with intractable distributions is that of an MCMC approach. When attempting to sample from a target distribution, MCMC methods produce a Markov chain, a sequence of random samples, such that each successive sample depends only on its predecessor. The chain possesses a stationary distribution that converges to the target distribution. This means, after a sufficiently long time, samples drawn from the Markov chain are distributed according to the target distribution.

Perhaps the most well-known MCMC scheme is the Metropolis-Hastings algorithm. Removing dependency on X in Equation (2.10), Bayes’ theorem becomes

$$P(\theta | Y) = \frac{P(Y | \theta)P(\theta)}{P(Y)}. \quad (3.12)$$

Considering the application of Metropolis-Hastings in a Bayesian setting, the goal is to sample from the LHS of Equation (3.12). Computation of the denominator of the RHS is

laborious, as it is typically a multidimensional integral that is expensive to compute. This can be reframed as attempting to sample from a distribution

$$\pi(\theta) = g(\theta)/Z, \quad (3.13)$$

where g represents the numerator of Equation (3.12), which can be sampled from, and $Z = P(Y)$ is a normalising constant. The Metropolis-Hastings algorithm then proceeds as follows: from a current sample θ , a new state θ' is proposed, which is sampled from a proposal distribution $q(\theta'|\theta)$, centred on the current sample. The transition $\theta \rightarrow \theta'$ is accepted with acceptance probability

$$\alpha = \min\left(1, \frac{g(\theta')q(\theta|\theta')}{g(\theta)q(\theta'|\theta)}\right) = \min\left(1, \frac{P(Y|\theta')P(\theta')q(\theta'|\theta)}{P(Y|\theta)P(\theta)q(\theta|\theta')}\right). \quad (3.14)$$

The crucial aspect of this computation is that the normalising factors Z , or corresponding marginals in the posterior probabilities, cancel out. By assessing the relative values of each sample under the likelihood and prior, the scheme can explore the parameter space and generate samples from the target posterior using only the numerator of Bayes' theorem in Equation (3.12). If a proposed state is accepted, the Markov chain moves to θ' , otherwise it remains at θ , and a new state is proposed from $q(\cdot|\theta)$. This simple acceptance condition ensures that, over time, the chain explores the parameter space in proportion to the posterior probability.

3.2.2 Approximate Bayesian Computation via Sequential Monte Carlo (ABC-SMC)

Even though the Metropolis-Hastings blueprint circumvents the need for evaluation of complex terms, knowledge of the likelihood may also be infeasible. This, unfortunately, is the reality for many complex systems, where the likelihood itself is intractable or expensive to evaluate. When the model f_θ used to map inputs to outputs cannot form part of a probabilistic model capable of providing likelihood values, attention is turned to likelihood-free methods that utilise all possible information contained in f_θ .

The core idea of ABC is to infer posterior probabilities directly from artificial data Y^* procured as the output of some forward model $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ governing the system dynamics. The need for a likelihood is bypassed through implementation of a metric D and a threshold ϵ . A sampled θ' and the corresponding simulated data Y^* , are assessed under D . An outline of the scheme is given below:

1. Sample θ' from $P(\theta)$
2. Compute $Y^* = f_{\theta'}(X)$
3. Compute $D(Y, Y^*)$.
4. If $D(Y, Y^*) \leq \epsilon$, store θ' . Return to step 1 and repeat.

The recovered population of samples stored from this scheme originate from the distribution [128]

$$P_\epsilon(\theta | Y) \propto \int P(\theta) P(Y | \theta) \mathbf{I}_{A_{\epsilon, Y}}(Y^*) dY^*, \quad (3.15)$$

where $\mathbf{I}_{A_{\epsilon, Y}}$ denotes the indicator function on the set $A_{\epsilon, Y}$ such that

$$A_{\epsilon, Y} = \{Y^* \in \mathcal{Y} | D(Y, Y^*) \leq \epsilon\}. \quad (3.16)$$

One can see, $\epsilon \rightarrow 0$ forces $Y \rightarrow Y^*$, and consequently $P_\epsilon(\theta | Y) \rightarrow P(\theta | Y)$. In the other direction, taking $\epsilon \rightarrow \infty$ provides no restriction on the acceptance criteria. In this case, $P_\epsilon(\theta | Y) \rightarrow P(\theta)$. Thus, there is a balance in how tight to make the threshold to balance accuracy with acceptance rate.

To enhance the acceptance rate, ABC-SMC is employed. As mentioned in Chapter 2.5, samples are propagated through a series of J intermediate distributions, accompanied by a sequence of decreasing tolerances $\epsilon_1 > \epsilon_2 > \dots > \epsilon_J$. The core idea is that the distribution of accepted samples evolves through a sequence of distributions, bridging the gap between prior and posterior as the scheme progresses. The first distribution ($j = 1$) is recovered via the standard ABC procedure, where samples are drawn from $P(\theta)$, and a tolerance of ϵ_1 is used. After all the required number of samples N are accepted, all are assigned an equal weight of $1/N$.

For all subsequent populations ($j \geq 2$), a sample θ^* is drawn from the previous waves' pool of accepted samples with a probability proportional to its weight on the previous wave. This allows the scheme to focus on areas of parameter space that are deemed more probable. The selected θ^* is then perturbed using a *Markov kernel* $K_j(\theta' | \theta^*)$ to propose a new candidate sample θ' . This operates as a means for the scheme to explore areas surrounding promising samples. Simulated data Y^* are produced from θ' , which is accepted if $D(Y, Y^*) \leq \epsilon_j$. This proposal step is repeated until N samples are accepted. The weight for the i 'th accepted sample in wave $j \geq 2$ is computed as

$$w_j^{(i)} = \frac{P(\theta_j^{(i)})}{\sum_{l=1}^N w_{j-1}^{(l)} K_j(\theta_j^{(i)} | \theta_{j-1}^{(l)})}. \quad (3.17)$$

This balances the sample's prior probability against how likely it was to be sampled, based on previous weights and K_j . Consequently, a sample receives a larger weight if it has a high prior probability, and correspondingly low probability of being proposed from the previous wave's distribution. An outline of the ABC-SMC algorithm for a sequence of J intermediate distributions is given below:

Algorithm 1: ABC-SMC algorithm

```

1 Initialise first population,  $j = 1$ ;
2 for  $i = 1, \dots, N$  do
3   Sample  $\theta'$  from  $P(\theta)$ ;
4   Generate  $Y^*$  from  $f_{\theta'}$ ;
5   if  $D(Y, Y^*) \leq \epsilon_1$  then
6     Set  $\theta_1^{(i)} \leftarrow \theta'$ ;
7     Set  $w_1^{(i)} \leftarrow 1/N$ ;
8   end
9 end
10 Sequential populations;
11 for  $j = 2, \dots, J$  do
12   for  $i = 1, \dots, N$  do
13     Resample: Select  $\theta^*$  from  $\{\theta_{j-1}^{(i)}\}$  with probabilities  $\propto w_{j-1}^{(i)}$ ;
14     Perturb: Sample  $\theta' \sim K_j(\cdot | \theta^*)$ ;
15     Generate  $Y^*$  from  $f_{\theta'}$ ;
16     if  $D(Y, Y^*) \leq \epsilon_j$  then
17       Set  $\theta_j^{(i)} \leftarrow \theta'$ ;
18       Set  $w_j^{(i)} \leftarrow \frac{P(\theta_j^{(i)})}{\sum_{l=1}^N w_{j-1}^{(l)} K_j(\theta_j^{(i)} | \theta_{j-1}^{(l)})}$ ;
19     end
20   end
21   Normalize weights:  $\sum_{i=1}^N w_j^i = 1$ ;
22 end

```

3.3 Experimental methodology considerations

The experimental approach adopted in this thesis balances two requirements: to obtain statistically robust datasets of the target turbulent phenomena, balanced with practical constraints of the laboratory facility and data collection process. Where appropriate, key experimental decisions are justified below, alongside a discussion of associated uncertainties and limitations.

The experimental datasets presented in this thesis encompass two distinct flow conditions, representing different flow depths, Reynolds numbers, and flow geometries within the same partially-filled pipe rig. Rather than attempting a wider study containing a greater number of Reynolds numbers and flow depths, this smaller number of regimes was selected in order to focus on obtaining high quality datasets for both flow cases. Here, the goal was to obtain very long time series measurements of each flow regime. Most datasets in studies of this type span seconds or minutes, whereas here it was desired to obtain hours. This allowed for more reliable statistical analysis and independent samples of the sensor response, ensuring an abundance of derived quantities that would be used to draw conclusions from the raw data. Particularly in turbulent flows, where low frequency or large-scale motions are present, longer sampling durations are also advantageous to capture the full range of dynamics, as these features are slow-varying, and so greater time series lengths enable for recovery of an appropriate number of signatures of these low frequency features. Aside from this, shorter measurement periods were found to introduce greater variability in correlation and spectral estimates, as can be seen in the proceeding chapters.

A key assumption throughout this thesis is that of uniform flow conditions and fully-developed flow throughout the duration of the experimental datasets. Uniformity was manually targeted through control of the pipe slope and downstream boundary conditions of the sluice gate at the pipe outlet. The pipe was surveyed manually with rulers to ensure a consistent gradient, and adjustments to the outlet weir were made, as seen in Figure 3.1(a), such that the water surface slope aligned as closely as possible with the pipe slope over the measurement region. Initial flow-setting calibration was also conducted using wave probes distributed across the pipe length for a range of flow rates and outlet gate conditions, allowing target flow depths corresponding to neutral water slope conditions to be established. In practice, perfectly uniform flow is difficult to confirm, and so throughout this work this assumption should be interpreted in a statistical manner, whereby flow properties recovered from the raw sensor data are measured as approximately constant over the

sensor region. Experimental determination of flow depth and bulk flow velocity were also conducted manually. The former was obtained using manual ruler measurements during the slope setting procedure. The latter was calculated indirectly from measured flow rate, recorded by calibrated Siemens electromagnetic flow meters at the main inlet tank outlet. Manufacturer-stated accuracies are on the order of $\pm 0.2\%$ of the measured value under ideal operating conditions. As the downstream gate condition directly controls the local flow depth at the sensor position, it is acknowledged that the recovered bulk velocity should be interpreted as an approximate bulk quantity rather than a detailed representation of the local velocity field. However, experimental results shown in Section 5 confirm that FOS-obtained velocity measurements align exceptionally well with the experimentally-set bulk flow velocity values. The uncertainty in depth measurement is expected to be on the order of millimetres, which propagates into the calculated wetted area and velocity. Once measurements were made and deemed to fluctuate minimally, additional manual measurements were taken by extending tape around the pipe circumference up to the water surface. The measured arc length was then used to confirm the recorded flow depth, this is pictured in Figure 3.1(b). During the main experimental runs, intrusive wave probes were removed from the flow to minimise additional disturbances and artificial turbulent structures within the measurement region.




Figure 3.1: (a): Sluice gate at the pipe outlet controlling the flow depth. (b): Tape measure placed around pipe circumference to infer flow depth from wetted perimeter of flow.

Once data were collected, the main consideration during preprocessing was filtering of the

time series in an appropriate frequency range. Upon assessment of power spectra, a manual cutoff of 15 Hz was decided, this will be explored further in Section 5.4. The selected frequency range was chosen to retain the dominant energy-containing scales of the turbulence while excluding higher-frequency noise and components of the structural response not directly related to flow-induced pressure fluctuations as seen in the power spectra. It is acknowledged that the choice of cut-off frequency introduces some subjectivity. However, during preliminary data analysis steps, testing of greater cutoff frequencies past this 15 Hz limit revealed no significant differences in derived quantities, affirming this limit was appropriate for the available experimental datasets.

Overall, the experimental methodology reflects a balance between ideal measurement conditions and the practical limits involved with the laboratory facility and study length. The choices made are considered appropriate for capturing the key flow features of interest, while acknowledging the inherent uncertainties associated with turbulent flow measurements and indirect interpretation of sensor response.

Chapter 4

Calibrating the discrete boundary conditions of a dynamic simulation: A combinatorial ABC-SMC approach

4.1 Abstract

This paper presents a novel adaptation of the conventional approximate Bayesian computation sequential Monte Carlo (ABC-SMC) sampling algorithm for parameter estimation in the presence of uncertainties, coined *combinatorial* ABC-SMC. Inference of this type is used in situations where there does not exist a closed form of the associated likelihood function, which is replaced by a simulating model capable of producing artificial data. In the literature, conventional ABC-SMC is utilised to perform inference on continuous parameters. The novel scheme presented here has been developed to perform inference on parameters that are high-dimensional binary, rather than continuous. By altering the form of the proposal distribution from which to sample candidates in subsequent iterations (referred to as waves), high-dimensional binary variables may be targeted and inferred by the scheme. The efficacy of the proposed scheme is demonstrated through application to vibration data obtained in a structural dynamics experiment on a fibre-optic sensor simulated as a finite plate with uncertain boundary conditions at its edges. Results indicate that the method provides sound inference on the plate boundary conditions, which is validated through subsequent application of the method to multiple vibration datasets.

Comparisons between appropriate forms of the metric function used in the scheme are also developed to highlight the effect of this element in the scheme’s convergence.

4.2 Introduction

The calibration of complex computer models remains a challenging problem within engineering contexts. This issue is exacerbated by two factors: high dimensionality, and lack of differentiability in the parameter space of interest. A methodology for targeting the combination of these two problematic scenarios is the key contribution of this work. Additionally, in almost all scenarios, aligning the outputs of a computer model (simulator) with measured data is complicated by the presence of uncertainties throughout the process, some of which may be reducible (e.g., model parameters such as material properties) or irreducible (e.g., pollution of measured data by noise). The understanding and quantification of this uncertainty is often a key step in the calibration process, and one where problems of high dimensionality and lack of differentiability are exacerbated.

In this work, the authors adopt a Bayesian standpoint on the calibration problem, where the user seeks to find a posterior distribution over the parameters of interest given the combination of evidence from observed data and expert judgement (expressed as prior belief). This Bayesian viewpoint on uncertainty is by no means the only approach for the quantification of uncertainty; however, it has previously been shown to be effective in problems the readers will be familiar with. In the context of structural dynamics, there exist several known methods appealing to Bayes’ theorem to infer structural parameters [129, 130, 131]. These have proved to be popular avenues for engineers to take when attempting to perform inference on parameters of dynamical systems when working with measured data. It is also widely known that Bayesian methodologies can be adopted in the context of structural model updating. For example, the works of [132, 131, 133] all show effective methodologies for the recovery of Bayesian posteriors. Just as the successes of the Bayesian approach are well known, so are some of the shortcomings, particularly the difficulty of inference in high dimensions; see [134].

In the examples above, the parameters of interest, such as natural frequency, exist as a set of continuous variables. Less attention has been given to problems where the variables take discrete values or where the parameter vector θ contains a mixture of continuous and discrete values. This problem of inference over discrete or mixed spaces has received

comparatively less attention in general, notable exceptions being in mixture modelling or switching cases, e.g., [135, 136, 137]. However, in the context of engineering modelling, cases such as these may be encountered in a number of scenarios as seen in [138], where the authors used Bayesian optimisation to inform structural design of a welded beam, encoding welding types as a binary variable, whilst geometric considerations were encoded as continuous variables. The contribution of this paper is to propose a methodology for (approximate) Bayesian inference over computer models (simulators), where the parameters of interest are discrete (or mixed) and the parameter space is high dimensional. The approach is one of a sequential approximate Bayesian computation scheme with proposal kernels developed in a similar manner to that seen in Population Monte Carlo [139]. This allows for inference to be made over a thirty-six dimensional binary parameter space in the case study modelling a fibre-optic sensor housing articulation response in an efficient manner.

Commonly, Bayesian inference will be intractable, owing to lack of access to the marginal likelihood which normalises the posterior; hence, most modern Bayesian approaches will resort to an approximation of the posterior either via sampling methods such as MCMC, importance sampling, or sequential Monte Carlo, or via parametric approximation as in variational inference [140]. In practice, however, even the evaluation of the likelihood function, which models the observed data Y as a function of known inputs X , parametrised by θ , may prove to be impractical. A closed form of a function $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ may be available, but the form of the noise model may not be known, or f_θ may be known to only be an approximation of the true physical mechanisms at play due to the lack of sufficient knowledge of the system or the computational expense [128]. A collection of approaches coined *likelihood-free* methods have been produced that allow for computation of the posterior in situations where the corresponding likelihood function is unavailable. One such likelihood-free method that has proved popular within the last two decades is ABC [141]. This algorithm relies on the user being able to produce simulated data Y^* through a simulating model with inputs θ sampled from some known distribution, e.g., the prior. The simulating model mimics the unknown likelihood function by considering the closeness of the outputs of this model to the observed data. Acceptance of θ into the approximate posterior is conditional on simulated and observed data Y^* , Y being similar ‘enough’ under some definition of distance [128]. The application of this algorithm and its augmented versions such as ABC-MCMC [142] and ABC-SMC [143] can be seen throughout many areas, including genetics [144, 116, 145] and applied dynamical systems. For example, in [146], the authors demonstrated the value of ABC-SMC for parameter estimation in nonlinear system identification.

In [147], the authors used both ABC-SMC and ABC-MCMC to estimate the parameters in the deterministic Lotka-Volterra model [148] to demonstrate the decreased computational cost incurred with these adapted schemes compared to traditional ABC. To the best of the author’s knowledge, all documented uses of the ABC-SMC scheme have been produced to estimate continuous parameters.

Particularly in practical engineering cases, parameters encoding the design of some structure to be assessed may manifest as binary variables as highlighted above. In such cases where likelihood-free methods are the best form of attack, the procedure in question must be assessed in order to best ascertain how the sampling of these parameters can be altered to achieve efficient model calibration from the prior to the posterior distribution through the use of Bayes’ theorem. The presented work provides an adapted version of the well-known ABC-SMC scheme in the case where the parameters to be inferred are high-dimensional sets of binary variables rather than continuous variables. In this case, it is the resampling procedure contained in the sequential component of the algorithm that has been modified to accommodate this parameter change, allowing other components of the algorithm to operate as intended. The adapted scheme is not just applicable to the application shown here but in any scenario when an altered approach of well-known likelihood-free methods is needed to suit the nature of the problem to be addressed.

In Sections 4.3 and 4.5, the ABC, ABC-SMC, and adapted combinatorial ABC-SMC algorithms are presented. Section 4.6 presents a niche form of uncertainty that arose in a structural dynamics experiment considered by the authors, as well as details of numerical modelling of the experiment. Section 4.7 presents results from utilising the developed methodology to quantify the experimental uncertainty, followed by the interpretation of these results and notes on potential future work in Section 4.8.

4.3 Approximate Bayesian Computation

The already discussed Bayesian strategy for learning about unknown parameters in engineering models has been suggested as a valuable approach, owing to its formal incorporation of expert knowledge (through the prior distribution) and added insight from recovery of the full distribution over θ . It stated that it is generally possible, either exactly or numerically, to recover the posterior distribution $P(\theta|Y;X)$, which is the main object of interest. However, it is simple to imagine a case where the modeller has insufficient knowledge to

construct the likelihood of the model $P(Y|\theta; X)$. This situation may arise from (broadly speaking) two areas of ignorance. First, it may not be possible to write down the correct form of f_θ , for example, using a computer model such as a finite element simulation of a physical object that will result in some mismatch between the output of the model and the ‘true’ physical behaviour of the system. This mismatch will be termed *model discrepancy*, which will lead to an incorrect likelihood distribution and result in a biased estimation of the posterior. The second source of ignorance is regarding the form of the uncertainty on the observed data Y . It is commonly difficult to express as a probability distribution the uncertainty associated with the measurements, and often, assumptions will be made about the form of this uncertainty for convenience. Most regularly, models of the form $P(Y|\theta; X) = \mathcal{N}(f_\theta|\sigma)$, i.e., additive Gaussian (white) noise with standard deviation σ determined by the probabilistic specification of the problem, may be assumed. Incorrect assumption regarding the uncertainty or noise process will, similarly to model discrepancy, lead to a bias in the estimation of $P(\theta|Y; X)$.

The challenge in accurately specifying the likelihood $P(Y|\theta; X)$ and the known bias that misspecification can introduce has motivated the development of ABC techniques, which can return reliable results even when it is not possible to specify the likelihood exactly [149]. ABC attempts to retain the benefits of a *fully Bayesian* approach, i.e., recovering a posterior distribution which incorporates prior knowledge, by using an approximate likelihood where the computation of $P(Y|\theta; X)$ is replaced by some other function which describes how well the current set of θ allows the data to be modelled given f_θ . The function which replaces the likelihood should measure in some way the distance $D(Y, Y^*)$ between the data generated by the model for some choice of θ , Y^* , and the observed data Y . There are some restrictions on the choice of this distance function; however, in general, it is far easier to specify $D(Y, Y^*)$ than it would be to specify a correct form for $P(Y|\theta; X)$.

The aim of ABC is then to gather a number of samples N from an approximate form of $P(\theta|Y; X)$ which contains within its probability mass the true posterior distribution. Recovering such a set of samples can be achieved by employing a *rejection sampling* approach [149]. In possession of the simulating model f_θ , one can produce outputs Y^* for a trial set of θ values, θ' , which are the predictions from the model $f_{\theta'}(X)$. The quality of these predictions may be assessed by the previously defined distance function $D(Y, Y^*)$, and then, if this is below some specified threshold ϵ , that sample could be designated as plausible and is stored. If the distance is greater than the minimum acceptable as set by

ϵ , then θ' is discarded (rejected), and another possible trial θ is drawn. This rejection sampling procedure is outlined in Algorithm 2 [150].

Algorithm 2: ABC rejection sampler.

Input: $f_\theta, P(\theta), D(Y, Y^*), N, \epsilon$

Output: n_t Samples from approximate $P(\theta | Y; X)$

```

1  $n_a \leftarrow 0$ ;
2 while  $n_a < N$  do
3   Sample  $\theta'$  from  $P(\theta)$ ;
4   Compute  $Y^* = f_{\theta'}(X)$ ;
5   if  $D(Y, Y^*) \leq \epsilon$  then
6     Store  $\theta'$ ;
7      $n_a \leftarrow n_a + 1$ ;
8   end
9 end
```

Assuming that $\theta \in \Theta$, where Θ is the space of all possible θ , one can imagine that if only a small part of Θ contains valid samples, the rejection sampling approach may be very inefficient, i.e., many guesses of θ are discarded. This issue can often be encountered in practice and is exacerbated when the computation of $Y^* = f_{\theta'}(X)$ is computationally demanding. One may then seek ways in which the efficiency of Algorithm 2 could be improved through more targeted sampling of possible θ' to be assessed. However, it is useful to retain the properties and guarantees of this most basic ABC inference algorithm.

4.4 ABC-SMC

The desire for a more efficient ABC inference approach led to the development of ABC-SMC [143], where SMC refers to *sequential Monte Carlo*. This, as the name suggests, utilises a sequential sampling scheme to arrive at the set of valid target samples. The main challenge of effectively using the aforementioned ABC rejection sampling approach is that very few samples may be accepted at the desired threshold; most trial sets of parameters for which the simulation is run are rejected. When a potential θ is simulated and then rejected, the computational effort associated with that simulation is, in effect, wasted. Therefore, it is favourable to design algorithms which will improve the acceptance

rate and therefore reduce the amount of wasted computational effort. An introduction, more generally, to SMC for Bayesian inference can be found in [151]; however, a brief introduction to ABC-SMC is given here.

ABC-SMC attempts to increase the acceptance rate of samples through generating J intermediate distributions which bridge the gap between the prior and the posterior, each of these may be referred to as a population. These intermediate distributions are formed by sequentially reducing the threshold below which samples are accepted, such that in earlier populations of the algorithm, sets of θ whose simulated outputs lie further from the observed data Y are still accepted. The setup for performing ABC-SMC inference is very similar to that seen for the rejection sampling approach presented in Algorithm 2; the output of the algorithm will be a set of n_t samples of θ that produce outputs from the simulator Y^* which lie within the threshold ϵ in terms of their distance from the observed data Y . Additionally, the user specifies a set of n thresholds given by

$$\epsilon = \{\epsilon_0, \dots, \epsilon_{n-1}, \epsilon\}, \quad (4.1)$$

which are decreasing and whose final value ϵ is the target threshold. Notationally, for population j , a set of N samples will be produced $\left\{\theta_i^{(j)}\right\}_{i=1}^N$.

On its own, this would not improve the efficiency of the overall sampling scheme; instead, the acceptance rate would simply fall with each iteration. To facilitate the improvement in acceptance rate, the SMC component of the algorithm uses the accepted samples from the previous population (after the first) to propose θ' from a more restricted space, which, hopefully, will contain more good samples. The full procedure is described in [143], including, in particular, comments on the choice of the *backward kernel*, which have been omitted here for brevity, although a brief description of the full algorithm is given below.

As the samples are accepted, they are also assigned a weight according to some weighting function $w_j^i(\theta')$ which considers the likelihood of θ' under the prior and the $D(Y, Y^*)$ given the simulated data under θ' . In populations after the first wave, the weights are also a function of the previous weight and likelihood under the move kernel (and backward kernel). The move kernel $M_j\left(\left\{\theta_i^{(j)}\right\}|\cdot\right)$ allows samples of θ to be generated for population $t+1$. In other words, it facilitates the sequential part of the SMC procedure, informing the samples of the next population from those currently selected. The move kernel encompasses the following steps. First, a resampling procedure selects to carry forward those samples with the highest weights (i.e., those proportional to the smallest $D(Y, Y^*)$). After sampling,

each sample weight is reset to $1/N$. The resampling step may be omitted if the diversity in the samples (commonly measured by the effective sample size, ESS) is high enough. It should be noted that it is often beneficial to reduce sample degeneracy. After resampling, the remaining samples are used to propose possible θ' in the next population j by means of a Markov kernel $K_j(\cdot | \{\theta_i^{(j-1)}\})$. The most common choice for this Markov kernel (and the backward kernel) is a Gaussian with zero mean and some variance which corresponds to a random walk centred on a previous sample with some perturbation, the extent of which is governed by the variance. This selection has the advantage of simplifying the calculation of w_j^i , ensuring that the whole parameter space can be reached.

4.5 Combinatorial ABC-SMC

In the previous section, ABC and ABC-SMC schemes have been introduced. ABC-SMC aims to reduce the problem of slow convergence when the parameter space only contains a small subset of appropriate samples, this is the typical setting in an engineering model where the posterior is much tighter than the prior distribution. Typically, ABC-SMC methods have been used in situations where the parameters of interest θ are continuous variables. However, as discussed, it is often the case that parameters in engineering models will be discrete or a mixture of discrete and continuous domains; this poses a problem since the typical random walk proposals seen in ABC-SMC are no longer appropriate. Since the parameter spaces are often high dimensional, it would be beneficial to leverage the same advantages as ABC-SMC usually does, increasing acceptance rates through a series of sampling procedures with shrinking tolerances that is applicable to these binary, or mixed spaces. By considering the contributions from all elements of the variable, their individual states, and the outputs through f_θ , it should be assessed which combination of all variable states is more likely to produce appropriate simulated data. To accommodate this challenge, a novel scheme, coined *combinatorial* ABC-SMC is presented. The root of this terminology is the fact that the algorithm is searching for the posterior over the combination of binary variables, as is necessary to resolve the experimental uncertainty to be presented in Section 4.6. In the proceeding text, the parameters to be inferred shall be denoted by $\boldsymbol{\theta}$ to reflect the fact they are multidimensional. It will later become apparent why this change in variable approach was needed.

Combinatorial ABC-SMC again uses a sequential sampling scheme, projecting samples through a sequence of J intermediate distributions with thresholds introduced in Equa-

tion (4.1). The proposed solution involves developing a proposal that works similarly to that seen in [152, 139]. The idea is to build a proposal distribution which depends on all the accepted samples in the previous wave, and generate new samples to be assessed independently from that distribution on the current wave, for example, computing the sample mean and variance, then proposing from that Gaussian distribution as in [153]. Here, however, this is not the desired method since the variables are binary, and it is not clear how a Gaussian proposal could be leveraged. Instead, the solution explored is to develop a proposal distribution that makes the following assumption; that each dimension can be treated as independent in the proposal, i.e., there is no correlation between elements in the variable. The proposal is then given by the estimated posterior probability that a given element θ_v in $\boldsymbol{\theta}$ is equal to 1 in the previous wave, which is computed in a Monte Carlo manner based on the full set of accepted samples,

$$\theta'_v = \begin{cases} 1 & \text{with probability } \frac{1}{N} \sum_{i=1}^N \mathbf{1}_{\{i\}}(\theta_v^{(j)}), \\ 0 & \text{otherwise} \end{cases}, \quad (4.2)$$

where $\mathbf{1}_{\{i\}}(\theta_v^{(j)})$ is an indicator if $\theta_v^{(j)}$ is equal to one in the i 'th accepted sample from the previous wave. Then, θ'_v is the proposed value of the v 'th dimension of a potential $\boldsymbol{\theta}'$ in wave $j + 1$, and this is proposed with the probability equal to the proportion of instances in the previous wave where that element is equal to one (i.e., $\theta_v^{(j)} = 1$). This can be repeated for each dimension of $\boldsymbol{\theta}'$ as independence is assumed. While this may seem complicated, an intuitive explanation is that the probability that a particular element is equal to one is simply equal to the proportion which is equal to one in the previous wave. Rather than assigning weights to accepted samples, the intermediate distribution formed by the N samples described by the above procedure in each wave is simply sampled from the subsequent wave. By recording which samples are more likely to produce appropriate data, which is contained in each wave's intermediate distribution, sampling from this on the next wave ensures a more appropriate subset of the parameter space is explored.

There is a clear case where this may be a bad choice for the proposal. If, for a particular dimension of $\boldsymbol{\theta}$, there are no accepted instances where θ_v is not zero, then it will be impossible to ever propose that θ_v is not equal to zero. This will cause problems. One potential solution, if this is to occur, is to also perform tempering on the proposal distribution. Denoting the density from which a new $\boldsymbol{\theta}'$ is sampled using Equation (4.2) as $\tilde{q}(\boldsymbol{\theta}'|\Theta_j)$, then

the tempered proposal density can be given by,

$$q(\boldsymbol{\theta}'|\Theta_j) = \tilde{q}(\boldsymbol{\theta}'|\Theta_j)^{1-\phi(j)}p(\boldsymbol{\theta})^{\phi(j)}, \quad (4.3)$$

where $\phi(j)$ is a monotonically reducing sequence from 1 to 0 as a function of the wave number j , e.g., a linearly reducing scheme. This tempering approach ensures there is some influence of the prior to maintain diversity across the waves [154]. The contribution of this paper then is to present the new proposal in Equation (4.2) to be used in the move step of the ABC-SMC scheme with the intention of targeting high-dimensional sets of binary variables. An application in which this methodology has been utilised shall be introduced in the following sections.

4.6 Problem statement

nuron Ltd., formed in 2015, is a company that aims to utilise fibre-optic sensing for leak detection and flow measurement in the UK's wastewater pipes. Their researchers and engineers provided a FOS to be used in an experiment based at the Laboratory for Verification and Validation (LVV) at the University of Sheffield. This experiment aimed to analyse the dynamic response of a fibre-optic sensor laid at the bottom of a dry half-pipe when subjected to localised mechanical excitation. For clarity, a fibre-optic sensor is a uniquely designed structure encasing a fibre-optic cable that can be laid at the bottom of the pipe interior, spanning a given section of the pipe length. The sensor is designed so any external pressure perturbations exerted on the structure are magnified by several orders of magnitude as they propagate through, ensuring the perturbations are picked up by the encased cable. In the field, it is pressure impacts due to the fluid flow above the sensor that induce signal changes in the fibre-optic cable within. These signal fluctuations change when leakages occur, allowing leaks to be detected [155]. It should be noted, however, that for these experiments, it is the dynamic response of the sensor structure that is of interest, rather than the interior cable signal response.

The experimental setup in the LVV facility consists of an 8 m long, 300 mm diameter clay half-pipe laid horizontally, clamped in place to minimise pipe vibration. The sensor has a length of 90 cm and consists of three separate layers. An articulation system spanning each side of the sensor is installed to ensure the sensing structure itself did not become dislodged from the pipe interior during testing. This system is divided into 18 clamp sections on either side, each 5 cm in length. The coupling between the articulation system and the sensor

varies between sections, and it is the application of the previously introduced methodology that resolves this uncertainty in coupling variation. The experiments at LVV involved exerting a harmonic load onto the sensor via a electrodynamic shaker at given points along the sensor length. For a given impact location, time series of the sensor transverse displacement at regular length intervals were recorded using a Laser Doppler Vibrometer (LDV). This in turn is used to build a profile of the maximum transverse displacement of the sensor for a load of given frequency, amplitude, and location. Figure 4.1 contains a maximum displacement profile of the sensor that is recovered through the experimental procedure described above. Note, LDV data were only collected in 1 cm intervals from 48 to 80 cm along the sensor length. This acted as a first step in understanding how nuron’s sensor would respond to simple excitation in controlled conditions.




Figure 4.1: *Normalised absolute displacement profile for the sensor subjected to mechanical excitation of 0.041–0.044 N operating at 5 Hz. Impact located at 45 cm along the 90 cm sensor length.*

Through the large difference between the obtained LDV data and previously obtained numerical data during initial simulation stages, it became clear that the extent of the coupling between the clamps and the sensing structure is unknown, implying any given edge section could move freely if the clamp at that position is not locked in place. Whatever the cause for the discrepancy between the numerical and experimental data, there is uncertainty in how to numerically model the sensor response. If each clamp is modelled as on or off, there are a potential 2^{36} possible clamping configurations that may be responsible for the data in Figure 4.1. If it took one second to numerically evaluate the maximum sensor displacement under the same experimental conditions as outlined in Figure 4.1 for

a given clamping configuration (based on the simplified numerical model outlined in the following subsection), it would take over two millennia to evaluate the displacement for all possible configurations.

The presented problem can be stated as follows: obtained experimentally are data Y that are dependent on some unknown parameters $\boldsymbol{\theta} \in \Theta$, of which there are many. Which members of Θ may be responsible for Y ? Or, in the context of the presented work, given the observed sensor response in Figure 4.1, what is the possible configuration of clamping that is physically responsible for these data? It is the methodology developed in the previous sections that the authors have used to resolve this boundary condition uncertainty. The approach taken is to model the clamping configuration for the sensor as a high-dimensional binary variable $\boldsymbol{\theta} \in \Theta$, where $\Theta = \{0, 1\}^{36}$. Each of the 36 elements corresponds to a particular clamping section such that each element of $\boldsymbol{\theta}$ is either 0 (unclamped) or 1 (clamped), resulting in the corresponding edge section being fixed or free, respectively. The authors believe that 5 cm is a long enough length for the sensor displacement to decrease (as seen in Figure 4.1) so that a specific clamp section may not be affected by the condition of its neighbours, and so independence between clamping sections is assumed. There is currently no known method of recovering the transverse displacement of the sensor when subjected to harmonic excitation with the complex boundary conditions that are presented here, making a likelihood function unobtainable. To circumvent this problem, simulated data Y^* can be produced through the use of the model displayed in Figure 4.2 that replicates the observed data Y shown in Figure 4.1 for the same type of excitation, with a differing configuration of boundary conditions for each simulation. Under a suitable user-defined metric D that not only defines the distance between the two data curves but characterises their similarity, the suitability of simulated curves and the corresponding boundary conditions can be assessed. By repeating this and projecting samples through intermediate distributions, an approximate form of the posterior distribution $P(\boldsymbol{\theta} | Y; X)$ can be produced. This gives the probability of each boundary condition being fixed or free, given observed data Y .

Starting with a uniform probability of all boundary conditions being fixed or free, a configuration of conditions is sampled from Θ . Each of these boundary conditions are inserted into the COMSOL model, which analyses the sensor displacement in the same fashion as in the experiment. Under the user-defined metric, the experimental and simulated curves are compared; if they agree within a certain margin of error, the boundary conditions

from the simulation are accepted. Otherwise, that configuration of conditions is rejected, and a different configuration of conditions is sampled. The process is repeated multiple times, sampling from the previous waves distribution, using a smaller margin of error each time, until 500 appropriate samples are accepted. During the testing of this methodology applied to the experimental data, it is found that five waves of combinatorial ABC-SMC is sufficient to achieve the target threshold ϵ from Equation (4.1) and converge to realistic posterior distributions on the parameters of interest, and so the scheme is implemented over five waves for this work. Inference of the true configuration of boundary conditions would not be achievable by use of classical ABC-SMC due to how the boundary conditions manifest as random variables. The authors hope this application demonstrates how the adapted combinatorial ABC-SMC scheme can be utilised in practice to perform inference on binary rather than continuous variables in an engineering context.

Physically Modelling the Articulation Response

A simplified finite element model of the LVV fibre-optic sensor was built in the finite element modelling (FEM) package COMSOL [156]. The geometry of the model is shown in Figure 4.2. While the actual fibre-optic sensor had a more complex geometry, through FEM simulations, it is established that simplifying the geometry has a negligible effect on sensor response. hence, it is decided to proceed with the simplified geometry that resulted in a considerable reduction in computational cost. The simplifications included not explicitly modelling the half pipe to which the sensor is attached, as well as the clamping system holding the sensor in place. The effect of these geometry elements is replicated via appropriate boundary conditions.




Figure 4.2: *Simplified model of LVV sensor used for predicting response to mechanical excitation. Labelled clamping sections and shaker impact area are highlighted in blue in (a). The mesh used is displayed in (b).*

Note, the sections highlighted in blue that line the left and right sides of the sensors' top layer in Figure 4.2a are where the clamps line the sensor at the pipe bottom. The finer mesh section in Figure 4.2b on the top of the sensor highlights where the harmonic load is exerted to replicate the data displayed in Figure 4.1. The top centreline of the sensor, parallel to the edges in Figure 4.2a, is where the maximum displacement of the sensor is plotted from, in line with the LDV measurements obtained in LVV. The presented model is used to generate simulated displacement curves Y^* for different configurations of clamping and played a pivotal role in resolving the uncertainty of the clamping system. The following section displays the results obtained from utilising the displayed finite element model as a simulating function in the combinatorial ABC-SMC scheme.

4.7 Results and discussion

Combinatorial ABC-SMC is applied twice to the data in Figure 4.1. In each instance, a different metric is used to characterise the similarity between normalised observed data Y and normalised simulated data Y^* . It is noted that both datasets are normalised by the maximum value, restricting values to the interval $[0, 1]$. For a sampled set of boundary conditions in the first instance, the areas beneath the experimental curves Y in Figure 4.1 and Y^* , denoted as $A_c(Y)$ and $A_c(Y^*)$ respectively, were calculated on the interval 48–80

cm. The corresponding metric used for this run takes the form

$$D(Y, Y^*) = \frac{|A_c(Y) - A_c(Y^*)|}{A_c(Y)}. \quad (4.4)$$

Using this form of D allows the areas of the respective curves to be compared, and if $D(Y, Y^*) \leq \epsilon$, the 36 boundary conditions producing Y^* are accepted. Figures 4.3 and 4.4 display results from this run. The 500 normalised simulated data Y^* accepted on each wave are displayed first in Figure 4.3. The corresponding probability of each of the 36 boundary conditions being fixed after all 500 samples were accepted for each wave is displayed in Figure 4.4.

It can be observed in Figure 4.3 that the sequence of accepted curves more closely approximates the shape of the experimental data for successive waves. The corresponding boundary conditions capable of producing these more acceptable data show a trend of favouring free conditions at the shaker impact area. By wave 5 (Figures 4.3e and 4.4e), it appears that only configurations of conditions that are free at the shaker impact area are capable of producing curves within the margin of error that has been set for wave 5.




Figure 4.3: Accepted curves for each wave of ABC-SMC using a difference of area metric. Curves in (a)-(e) correspond to accepted simulated data Y^* from waves 1-5, respectively. Experimental data are shown with circle markers.




Figure 4.4: *Intermediate distributions for each dimension of θ for each wave of ABC-SMC, using a difference of area metric. Distributions in (a)-(e) correspond to the belief of the fixity of each boundary condition after each wave 1-5, respectively.*

The results in Figures 4.5 and 4.6 were produced by using a least-squares metric on the same experimental data. The metric used in this instance takes the form

$$D(Y, Y^*) = \sum_{i=0}^{33} |y_{48+i}^* - y_{48+i}|^2. \quad (4.5)$$

The terms y_l, y_l^* are the values of the experimental and simulated curves Y, Y^* at l cm along the sensor length, respectively. Note, as data were only collected along a subsection of the

sensor length, it is only over this interval that areas of simulated data and least-squares distance are calculated for each run.




Figure 4.5: Accepted curves for each wave of ABC-SMC using a least-squares metric. Curves in (a)-(e) correspond to accepted simulated data Y^* from waves 1-5, respectively. Experimental data are shown with circle markers.




Figure 4.6: *Intermediate distributions for each dimension of θ for each wave of ABC-SMC, using a least-squares metric. Distributions in (a)-(e) correspond to belief of fixity of each boundary condition after each wave 1-5, respectively.*

For any ABC scheme, the metric utilised needs to be carefully chosen to fully incorporate the notion of ‘similarity’ between the objects Y and Y^* being considered. If the outputs of the system in question were real numbers, then a simple function such as the absolute distance $D(Y, Y^*) = |Y - Y^*|$ between the two data would be a perfect mechanism to classify similarity. In this scenario, $D = 0$ would imply $Y = Y^*$, and so it is understood that the parameters θ producing Y^* must also produce Y . Unfortunately, for most cases where ABC schemes are applicable, the parameters to be inferred do not produce objects

as simple as this. Understanding the objects being compared is important, i.e., curves in $\mathbb{R}^2$ in the current scenario. However, the metric needs to fully reflect what properties of the experimental data the user wishes to emulate. For the presented case, the user wishes to recover simulated absolute maximum displacement profiles of the plate that match, as closely as possible, the trajectory of the data in Figure 4.1. Whilst taking the form of D in Equation (4.4) provided the expected inference, it is clear from Figure 4.3e that boundary conditions are still being accepted that produce curves with secondary peaks, even though the target data do not display this feature. The difference between the areas of these curves and the experimental curve are still within the margin of error that has been set, and so the scheme sees the samples generating these curves as appropriate, even though these curves do not follow a trajectory similar to Y . Thus, D has, in a sense, failed to only accept samples producing simulated data that conform to the conditions being targeted within five waves, and more iterations may be needed to produce simulated data more closely approximating experimental data.

It is believed that using the form of D in Equation (4.5) would provide more sound inference when attempting to recover the exact trajectory of Y without the need to increase the number of waves, as this metric aims to minimise the sum of offsets between the two sets of data being considered, ensuring that simulated curves with similar trajectories to Y produce lower values of D .

Utilising a least-squares metric as seen in Figures 4.5 and 4.6 accounted more for the shape of the simulated curves. It is believed that this form of D would penalise curves with secondary peaks more heavily when compared against the experimental data, which consists of only a single thin peak. In later waves, this is hypothesised to help in refining the space of appropriate samples, leading to a more accurate approximate posterior distribution. It can be seen by wave 5 that no curves are accepted with secondary peaks along the points covering the experimental data curve. Intense overlapping of the accepted curves by the final wave in Figure 4.5 displays that for an appropriate choice of metric, more realistic samples can be recovered, ensuring a more accurate form of the final approximate posterior distribution.

Regardless of the slight improvement in results for the second run, both instances arrive at the conclusion that the clamps surrounding the midsection of the sensor must be loose. It can also be seen that clamps 8, 12, 13, 26, 30, and 31 have a much higher probability of being fixed. So it must be the case that these clamps were deemed fixed in a greater

proportion of elements sampled from Θ producing accepted simulated curves, implying it is highly likely that these clamps are in fact fixing the corresponding boundary conditions in the experimental setup, constraining the width of this peak in Figure 4.1. As stated in the introduction, coupling between the clamping system and the sensor varies between sections, and it can be seen that combinatorial ABC-SMC has provided valuable insight into the true nature of this coupling.

For the final waves in Figures 4.3 and 4.5, it can be seen that away from the impact area, the clamps' probability of being fixed or free is confined within the $[40,60]\%$ interval, and so the algorithm provides less certain inference in these areas. This phenomenon occurs due to the fact that the displacement would diminish over a relatively short length scale (for any boundary condition configuration) along the sensor when considering the structure (three layers) and material properties of the sensor (as can be seen in Figure 4.1). Consequently, the clamping configuration in these further away areas would have minimal effect on the displacement profile.

While these results convey the applicability of combinatorial ABC-SMC for providing inference of the coupling variation between the clamping system and the sensor, it is necessary to verify the model used for this work. Additional data were obtained from a variation of the original experimental setup at LVV. These data were collected from a similar experiment, where the impact is located at 70 cm along the length of the sensor, rather than 45 cm. Combinatorial ABC-SMC is applied to these data, which are displayed in Figure 4.7. It is the aim to provide insight into how the more complex-looking data manifests in the approximate posterior at the end of wave 5 in this additional run.




Figure 4.7: *Normalised absolute displacement profile for the sensor subjected to mechanical excitation of 0.035–0.037 N operating at 5 Hz. Impact located at 70 cm along the 90 cm sensor length.*

One can see how the maximum displacement profiles presented in Figures 4.1 and 4.7 are nothing alike. These data were chosen specifically to see if the performance of combinatorial ABC-SMC would yield a similar posterior distribution of boundary conditions for the data in Figure 4.1. Additionally, it is believed that the existence of secondary and third peaks in Figure 4.7 may be due to more variation in coupling between neighbouring clamps in this section of the sensor, as well as the fact that there is greater displacement of the sensor further from the shaker location in Figure 4.7 compared to Figure 4.1.

Data from Figure 4.7 were only collected in 1 cm intervals starting from 4 cm away from the shaker position, unlike previous data, which started 3 cm from the shaker position. The MATLAB optimisation toolbox is used to fit the data at 74–76 cm to a Gaussian function centred at the shaker impact location. This is performed using the `lsqcurvefit` function to replicate the shape of the displacement curve localised to the impact area. Using the combined LDV data and artificial data, a longer displacement profile on which to apply combinatorial ABC-SMC is produced. Equipped with a metric of the form shown in Equation (4.5), combinatorial ABC-SMC is applied to the data in Figure 4.7. This is implemented to provide a comparison of how well the aforementioned metric would perform on two different datasets. Due to the increased complexity of experimental data in this case, it is hypothesised that there would be differing levels of convergence between the two runs. It is the authors' belief that the recovery of samples closely approximating the data

in Figure 4.7 may require additional metric structure to produce samples with the more complex trajectory seen in the data, which a standard least-squares metric similar to that in Equation (4.5) would ignore. Results from this run are displayed in Figures 4.8 and 4.9.




Figure 4.8: Accepted curves for each wave of ABC-SMC using a least-squares metric applied to data from Figure 4.7. Curves in (a)-(e) correspond to accepted simulated data Y^* from waves 1-5, respectively. Experimental data are shown with circle markers.




Figure 4.9: *Intermediate distributions for each dimension of θ for each wave of ABC-SMC, using a least-squares metric applied to data from Figure 4.7. Distributions in (a)-(e) correspond to belief of fixity of each boundary condition after each wave 1-5, respectively.*

From the results in Figures 4.8 and 4.9, it can be observed that the least-squares metric is unable to provide any meaningful inference from five successive waves of sampling on the sensor boundary conditions in locations where experimental data were collected. As outlined previously, correct inference of the target parameters relies heavily on an appropriate metric structure capable of recovering the features of importance contained in the experimental data. A metric globally minimising the average of vertical offsets for the data in

Figure 4.7 clearly does not account for the multiple peaks seen in the data, and so falls short in mimicking the experimental data. This is comparable to the failings of the metric in Equation (4.4). Additional structure is needed to fully capture the local variations seen in the secondary and third peaks, whereas a metric of the form in Equation (4.5) is only capable of minimising the global vertical offsets between two datasets.

To combat poor convergence, the metric used for acceptance is modified significantly to account for the additional structure seen in Figure 4.7. The authors hypothesised that combinatorial ABC-SMC would then detail what boundary conditions would be necessary to produce additional peaks in displacement as seen in Figure 4.7, provided the metric is tailored appropriately to the specific data. When considering how to assess samples, the input data were split into 3 sections: 71–75 cm, 76–80 cm, and 81–89 cm. Again, a least-squares metric is used, but each of these sections is considered separately. This ensured the squared vertical distance between simulated and experimental data is averaged for subdomains of the plate length rather than the whole length itself. Additional terms were added to the metric, ensuring only curves displaying a gradient increase and decrease in the 76–80 cm region were accepted, to encourage more samples to be accepted displaying secondary peaks. The metric D used for this run takes the form

$$D(Y, Y^*) = \left(\sum_{i=0}^4 |y_{71+i}^* - y_{71+i}|^2, \sum_{i=0}^4 |y_{76+i}^* - y_{76+i}|^2, \sum_{i=0}^8 |y_{81+i}^* - y_{81+i}|^2, \mathbf{H}_{\{1\}}, \mathbf{H}_{\{2\}} \right), \quad (4.6)$$

where

$$\mathbf{H}_{\{1\}} = \begin{cases} 1 & \text{if } y_{78}^* - y_{76}^* > 0, \\ 0 & \text{if } y_{78}^* - y_{76}^* < 0 \end{cases}, \quad \mathbf{H}_{\{2\}} = \begin{cases} 1 & \text{if } y_{80}^* - y_{78}^* < 0, \\ 0 & \text{if } y_{80}^* - y_{78}^* > 0 \end{cases}. \quad (4.7)$$

The terms y_j, y_j^* are the values of the experimental and simulated curves Y, Y^* at j cm along the sensor length, respectively. The terms in Equation (4.7) characterise whether there is a secondary peak in the 76-80 cm region of Y^* . Both $\mathbf{H}_{\{1\}}, \mathbf{H}_{\{2\}}$ must be equal to one for this condition to be satisfied. The results for this run are displayed in Figures 4.10 and 4.11.




Figure 4.10: Accepted curves for each wave of ABC-SMC using a split least-squares metric. Curves in (a)-(e) correspond to accepted simulated data Y^* from waves 1-5, respectively. Experimental data are shown with circle markers.




Figure 4.11: *Intermediate distributions for each dimension of θ for each wave of ABC-SMC, using a split least-squares metric. Distributions in (a)-(e) correspond to belief of fixity of each boundary condition after each wave 1-5, respectively.*

As with the previous converging runs, progressing through the waves, accepted simulated data mimic more closely the trajectory of the experimental data as desired, with a large reduction in variance in the final two waves. In the first four waves, results appear to show that clamps 14, 15, 32, and 33, which are positioned between 65 and 75 cm along the sensor length, have a high probability of being fixed. This is in contrast to all of the other clamps, which appear to have little effect on the simulated data, as all of these elements have an equal probability of being fixed or free, even by wave 4. However, the stricter margin of error present in wave 5 results in only much more realistic curves being accepted. This improvement can be seen by the lower variance in the values of the accepted curves in

the 75–90 cm interval. Consequently, the posterior changes quite significantly compared to the previous waves as well. Clamps 14 and 32 at the 65–70 cm position are deemed completely fixed in this final wave, which continues the trend being shown throughout the previous waves. Clamps 15 and 33 at the 70–75 cm position appear to have little effect on the sensor response, despite the shaker position covering this area. Instead, clamps 16 and 34 are now deemed almost certainly fixed. This covers the positions of 75–80 cm along the sensor, which is the location of the secondary peak. The authors conclude that clamps 16 and 34 being fixed is most likely a limiting factor in the height of this secondary peak. Additionally, this also played an important role in constraining the width of the main peak.

Due to limitations and time constraints in the experimental work, LDV data were not collected over the entire length of the sensor. The algorithm has freedom in how to choose the clamps conditions in the ‘blind’ areas where experimental data are not recorded, as the algorithm would not have anything to compare the simulated displacement to at these points. This may be a cause for concern when evaluating the accuracy of the formed posterior distribution in each run, as the sensor response in these unrecorded areas may provide a whole picture of the boundary condition configuration for the entire sensor. The counterpoint to this assertion is that it seems reasonable to have very little change from the prior belief in regions far from the measurements. However, in all cases, the absolute displacement of the sensor diminishes over a very short length as can be seen in Figures 4.1 and 4.7, making it difficult to infer the state of clamping in areas further away from the impact location.

4.8 Conclusions

Through the several sets of presented results, it has been shown that combinatorial ABC-SMC can be successfully applied to resolve uncertainty in calibrating the challenging boundary conditions of this fibre-optic sensor. The procedure is also generally applicable to cases where a calibration of a set of binary variables is needed. The algorithm determines the most realistic configurations of clamping of the fibre-optic sensor given the observed displacement profile when subjected to localised mechanical excitation. The obtained results match reasonable expectations of how the sensor would respond dynamically in different scenarios, subsequently providing inference on the underlying boundary condition parameter θ responsible for the sensor’s dynamic response. Future work could include the repetition of experiments with another impact location, or recording LDV data along

the entire sensor length to observe if additional information could lead to a tighter posterior distribution. The formed posteriors at the end of wave 5 for each run result in roughly symmetrical boundary conditions along each side of the sensor. The authors believe that the size of the parameter space may be considerably reduced if only symmetrical boundary conditions are considered, further reducing the computational cost required for additional runs, even if additional experimental data are collected along the entire plate length.

In its current form, combinatorial ABC-SMC assumes independence between the elements of θ . However, since the elements of θ are likely correlated, some refinement of the proposal could increase the efficiency of the sampler further. Despite the assumptions seen in this modelling approach, it has been seen how in a challenging inference setting, the presented approach has allowed insight into a complex modelling procedure, highlighting areas both of confidence and uncertainty.

The methodology presented in this chapter was developed using an initial FOS under controlled laboratory conditions when faced with uncertainty in the FOS dynamic response. Even though a newer FOS iteration was developed for later chapters, the insights gained here highlight the importance of additional methodologies that may be required when attempting to interpret sensor responses. This uncertainty encouraged the development of an improved FOS structure suitable for hydraulic testing. The following chapters build on this progression by applying sensing and analysis techniques to turbulent partially-filled pipe flows, where the challenges of interpreting flow-induced excitation are extended to more complex and realistic conditions.

4.9 Acknowledgements

The authors would like to acknowledge the support of the UK Engineering and Physical Sciences Research Council (EPSRC) through grant EP/S017283/1 on Distributed Fibre-optic Cable Sensing for Buried Pipe Infrastructure. The authors would like to thank Dr Robin Mills and the technical team at the Laboratory for Verification and Validation for aiding in the experimental design and setup used for this work. The authors would like to acknowledge the work conducted by the nuron technical team in developing and installing the fibre-optic sensor used and to nuron management for support-in-kind over the life of this project. This research made use of The Laboratory for Verification and Validation (LVV) which was funded by the EPSRC (grant numbers EP/R006768/1 and

EP/N010884/1), the European Regional Development Fund (ERDF) and the University of Sheffield. For the purpose of open access, the authors have applied a ‘Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising.

Chapter 5

Fibre-optic sensor design and strain-response analysis to wall pressure fluctuations in turbulent partially-filled pipe flows

5.1 Abstract

There exist many studies analysing flow dynamics of pressure-driven pipelines, yet the number of studies quantifying hydrodynamic relations and conditions in partially-filled gravity-driven pipe flows remains limited. This work presents time-series strain data for two uniform, steady flow regimes from a novel streamwise, flush-mounted fibre-optic sensor embedded in the interior wall of a partially-filled pipe. The sensor responds to turbulence-induced pressure fluctuations in the near-wall boundary layer. Cross-correlation analysis on post-processed data recovered bulk flow velocity estimates across the sensor length within 0.7% of the true velocity. A Metropolis–Hastings-based automated filtering process identified the frequency range retaining bulk flow properties while removing noisy measurements, further refining velocity estimates. The identified streamwise correlation pattern agreed with experimentally observed fluctuating wall pressure correlations in boundary layers. Results suggest streamwise correlations decay at a rate proportional to the hydraulic radius in each regime, with only a 7% discrepancy between regimes. Spectral analysis indicated

sensor sensitivity to coherent structures at multiple spatial scales, suggesting potential for monitoring flow structures beyond bulk conditions. When the sensor section was rotated above the free surface, time-series analysis produced uninformative results, confirming that in-flow placement captures genuine flow dynamics despite measurement noise. These findings demonstrate a new approach to monitoring partially-filled pipe flows, providing both bulk property estimation and potential insights into multi-scale flow structures, making way for a new form of sensor to be deployed for monitoring these flows in varying industrial contexts.

5.2 Introduction

While the flow dynamics in fully-filled pressurised pipelines are well understood, considerably less attention has been given to pipes running partially-filled and driven by gravity. The primary difference between fully-filled and partially-filled pipe flows is the presence of a free surface, which results in significant secondary flow motions [44]. Understanding the hydraulic conditions in partially-filled pipes is of great practical importance for many engineering applications, such as leak detection for sewer pipeline networks and hydraulic flow estimates in stormwater management. Current methods for monitoring partially-filled pipe flows utilise spot sensing, where hydrodynamic measurements are collected at discrete locations, such as flow metering [157]. These methods provide little information on the interior hydraulic conditions in real time.

Recent advancements have utilised properties of the free surface to develop non-invasive monitoring techniques to characterise the dynamics of these flows [57]. These methods rely on acoustic scattering of the free surface recorded via a microphone array, which only provides spatially averaged hydrodynamic information. However, the energetic turbulent structures in a partially-filled pipe, including counter-rotating secondary flow motions, result in significant pressure fluctuations on the interior pipe wall. Akin to the free surface, the pressure fluctuations on the wall are also fingerprints of hydrodynamic conditions. Recovering properties of the instantaneous fluctuating pressure on the wall can provide us with significant information about the potential energy losses along the partially submerged pipe section. Reflecting upon the past few decades of research, many studies have managed to characterise the wall pressure fluctuations in turbulent boundary layers via the formulation of numerical models of the frequency-wavenumber spectra of these fluctuations [158]. Most notable are the Corcos and Chase models used to predict the response of

structures to turbulence-induced excitation originating from pressure fluctuations within boundary layers [51, 159]. While the literature in this area is extensive, and many empirical and analytical relations have been established that accurately describe the spectra of boundary layer pressure, there has been much less attention given to the experimental study of pressure fluctuations in boundary layers, especially in partially-filled pipes, to the best of the authors' knowledge. Even less so, studies attempting to relate the fluctuating wall pressure to bulk hydrodynamic conditions within the pipe in this case.

This motivates the current study, which aims to infer important flow parameters, such as bulk flow velocity, from time-resolved wall pressure-induced strain measurements, with excellent spatial resolution. A novel flush-mounted FOS embedded on the wall was developed and is used as the strain measurement device. Fibre-optic cables provide unique advantages over conventional sensing methods due to their sensitivity to multiple parameters of their external environment, immunity to electromagnetic interference, and their ability to sense over long distances [160]. Multiple studies have been conducted applying fibre-optic sensing to pipeline monitoring, but their application to measure flow in a partially-filled pipe has been largely limited. A comprehensive review of fibre-optic cables for pipe flow sensing is given in [161]. The authors of [162] developed a transducer to measure mean and instantaneous flow rates in a turbulent flow by forming a cantilever beam consisting of a fibre-optic cable and a sensor support anchored onto the interior wall. Standing perpendicular to the flow direction inside the pipe, the drag force experienced by the fibre could be related to the square of the flow velocity, while also being able to detect instantaneous fluctuations in flow direction at the point of sensing. A majority of applications of FOS have been to detect leaks in pressurised pipes. For example, the authors of [163] were able to localise leak events through anomalous pressure values recorded by a FOS installed in a lengthwise helical pattern on the exterior of a pressurised pipeline. By leveraging the relation between hoop strain and internal pressure, anomalous pressure values recorded by the FOS along its operational length signalled localised leak events, providing a means of leak detection for pressurised pipelines by use of a FOS, with varying spatial resolution based on the spatial separation of the hooping structure of the cable. By constructing a sensing enclosure in which to house a fibre-optic cable, the authors of [164] constructed a free-floating submersible FOS to respond to hydrostatic pressure levels in water pipes. When pressure changes characteristic of leaks were simulated in the pipe section, strain readings from the FOS were in good agreement with equivalent changes in pressure values recorded from a pressure transducer installed in the pipe.

The current state of the art of fibre-optic sensing for pipeline monitoring is mostly applicable to pressurised piping systems, whereby the operating principle is mostly reliant on the effect leakages have on the internal pressure and its fluctuations. In scenarios where FOS finds use in non-pressurised pipelines, it is necessary for the FOS to protrude into the flow in some manner to collect the data needed to provide meaningful insight into the hydraulic conditions of the flow. There is a clear gap in the current literature whereby the operating principle of a FOS is yet to respond to the fluctuating wall pressure in partially-filled pipelines, without intrusion into the flow. Designing a FOS in such a way would enable the fluctuating wall pressure and bulk properties of flow to be unequivocally established. The novelty of this manuscript is therefore to report on a new means of non-invasive monitoring technology for these flows capable of utilising the response of a FOS to this wall behaviour to infer the hydrodynamic conditions such as bulk flow velocity. This is achieved by means of cross-correlation analyses of single location strain measurements.

This paper is organised as follows: Section 2 describes the design and installation of the FOS as well as its installation in the experimental facility used. Section 3 discusses the data analysis procedure used to process the raw FOS data. Section 4 details conclusions and discusses avenues for further work.

5.3 Experimental facility and sensor design

The main objective of the present study was to unambiguously relate properties of wall pressure fluctuations in a partially-filled pipe to the interior hydraulic conditions through the design and installation of a novel flush-mounted FOS that responds to the turbulence-induced fluctuating pressure on the pipe wall. By designing a FOS that operates on the principles of Bragg reflection, axial strain of the cable can be recovered at multiple points along its length, allowing it to effectively operate as a streamwise array of strain gauges that respond to the vibration caused by the pressure fluctuations at the wall in the turbulent flow. As the proposed sensor responds to turbulent flow-induced pressure on the pipe wall, the experimental facility was designed in such a way that any unwanted structural vibration in the mechanical components making up the pipe rig was minimised and did not contaminate the FOS response. A new flow facility was designed accordingly (including the pipe rig), which is detailed in the following sections.

5.3.1 Rig description

Experiments were conducted in a gravity-driven pipe flow facility located in the Integrated Civil and Infrastructure Research Centre (ICAIR) of the University of Sheffield. The test pipe was comprised of eight Acrylic 2500 mm sections of 5 mm thick, 300 mm OD sections linked with SC335 310-335 mm flexible couplings (external clamped rubber) to form a 20 m long continuous pipe. Care was taken to avoid misalignment at the joint interfaces. Different fully-developed flow conditions (Reynolds Number and water depth) can be achieved up to a flow rate of 20 L/s. A schematic of the pipe rig with its important components is given in Figure 5.1.




Figure 5.1: Schematic view of ICAIR pipe facility including inlet and outlet tanks.

The water is driven into the straight pipe section via a 940x800x1500 mm inlet tank. The entire 20 meter section was installed onto a sand bed deposited on top of the C-section of welded steel to mitigate both external and flow induced vibration in the pipe walls, while an adjustable hangar system supports each joint (2.5 m) to enable fine tuning and maintenance of a consistent 1:500 slope across all sections. Each section was lightly restrained with cargo webbing to ensure good frictional contact and minimise potential movement. The frames are made of steel, and similarly, they sit on steel supports roughly every 5 m along the 20 m pipe length. The overall weight of the steel sections with sand was estimated at 7 tonnes. The downstream end of the pipe is connected to an outlet reservoir via an adjustable flow control gate that helps achieve uniform flow conditions (i.e. uniform water depth). As detailed below, the FOS spans between 13.3-14.2 m from the pipe inlet.

5.3.2 Sensor design and installation

With the nature of the research aims, it was pertinent to design a sensor that could non-invasively provide measurements of wall pressure-induced strain in turbulent partially-filled pipe flows. As such, the sensor design choices presented in this chapter reflect several key considerations central to the overall thesis. These include the implementation of a simple flush-mounted geometry, inclusive of a compliant base layer for pressure transfer, and the spatial arrangement of fibre Bragg gratings in a streamwise orientation to allow for convection of flow to be the main driving force behind the sensor response. As such, the contribution of this work lies not only in the application of fibre-optic sensing to this class of flows, but also in the design of a sensing system capable of extracting meaningful hydraulic information from wall-based measurements.

It was decided to design a two-layer sensor consisting of a 900x2x20 mm 40% talc-loaded polypropylene plate (spanning the streamwise, wall-normal, and spanwise directions, respectively). With the top face of the plate exposed towards the pipe interior, encased in a bed of a 900x4.5x24 mm magic gel substrate to facilitate transfer of the fluctuating pressure from the pipe wall to the FBGs. This component is a low viscosity gel with strong sealing properties and low stiffness. Once cured into the mold, it forms a soft and flexible medium that deforms under pressure while maintaining structural stability. The low stiffness and compliant nature of the gel are the desired properties needed to ensure effective strain transfer to the sensing fibre, allowing the plate to respond to fluid loading. At the same time, the bed of gel provides a degree of damping to the structure to reduce the influence of structural vibrations of the pipe structure itself, aiding in data quality.

To allow for installation, a 900x24 mm slot was cut out of the pipe wall test section. This layered structure was housed inside a containment system that could be externally attached to the pipe wall. The dimensions of all components of this combined sensing structure ensured that the 900 mm long top layer of the polypropylene plate lay flush with the pipe wall in the streamwise direction. In this study, the angular positioning of the sensor about the pipe wall is defined such that 0° corresponds to the bottom-most point of the pipe interior. When other angular positions of the FOS are mentioned, angles increase in the anticlockwise direction when viewed in the streamwise direction of flow. For example, an angle of 180° lies at the highest point on the pipe well. A photo of this sensor plate installed in the Acrylic pipe and rotated 90° for illustrative purposes is displayed in Figure 5.2(a).




Figure 5.2: (a): External pipe view of the FOS containment structure rotated 90° from the working position (0°) of the FOS used for data collection. Schematic diagrams of (b): cross-section and (c): lengthwise view of the FOS structure. FBG element locations and indices highlighted in (c) (not to scale).

Figure 5.2(b-c) presents schematic diagrams of the structure that supported the FOS used in the reported experiment. At the centreline of the sensor structure and lying at the plate-gel boundary, an engionic Femto Gratings GmbH SM1250BI(9.8/125)P fibre-optic cable

was installed that operates on the principles of Bragg reflection [165]. It was positioned as such to ensure the dynamic response of the cable was coupled with the bending of the plate element. FBGs are a type of wavelength-modulated sensor that are sensitive to temperature and pressure fluctuations. Reviews on their application in various settings can be found in [166, 101]. They are popular in the field of fibre-optic sensing due to the ease of manufacturing and simple operating procedure [167]. the FBGs used in this experiment are 1cm in length and is embedded into the fibre by periodically modulating the refractive index of the fibre core over this length. When an incident ray of light intercepts an FBG, it passes through the sections of varying RI. At a certain wavelength called the Bragg wavelength, contributions from the reflections of this wavelength that occur at each grating section occur in phase and are subsequently amplified, this fully reflects light of this wavelength back in the opposite direction of the incident light, while allowing all other wavelengths to be transmitted through the FBG [168]. The Bragg wavelength is defined as

$$\lambda_B = 2n_{eff}\Lambda, \quad (5.1)$$

where n_{eff} is the effective RI of the core, and Λ is the grating period length. For varying n_{eff} and Λ , the wavelength that is reflected will be unique for that particular FBG. When strain is applied to an FBG along the axial direction, this perturbs the value of Λ , allowing the Bragg wavelength shift in an isothermal environment to be calculated as

$$\frac{\Delta\lambda_B}{\lambda_B} = (1 - \rho_e)\varepsilon. \quad (5.2)$$

Here, $\Delta\lambda_B$ is the Bragg wavelength shift, ρ_e is the photo-elastic coefficient of the fibre core, and ε is the axial strain applied to the FBG. The locations of the FBG elements in the cable used in the reported experiment are known. In principle, the axial strain of the fibre can be directly recovered from the wavelength shift of each grating. However, the raw time series provided by the FBG elements are strain-proportional signals. Since wavelength shift scales linearly with strain as shown in Equation (5.2), the structure of the (normalised) cross-correlations and spectra to be presented in this work directly quantify the spatio-temporal structure of the strain field induced by the wall pressure fluctuations. As such, to reduce processing steps, explicit conversion from wavelength to strain was not carried out, unless stated otherwise. The positions and Bragg wavelengths of each grating

along the sensor centreline are given in Table 1. Note, the cable itself lay 150 mm from the pipe centre due to being beneath the 2 mm polypropylene plate, which was flush with the interior pipe radius of 148 mm.

	FBG index					
	1	2	3	4	5	6
FBG location (mm)	50	66	90	130	186	274
Bragg wavelength (nm)	1510	1518	1526	1534	1542	1550

Table 5.1: *Centre position of each FBG element from the starting position of the FOS plate (in the streamwise direction), and Bragg wavelengths of FBGs lying at the plate/gel boundary.*

5.3.3 Flow conditions

Experiments were conducted for two different flow conditions (having different flow depths and Reynolds Numbers), see Table 5.2. The slope of the rig is 1:500 and remained the same for both flow conditions.

Regime	Slope	Flow depth (m)	Bulk flow velocity (m/s)	Re (-)	Hydraulic radius (m)
1	1:500	0.148	0.698	207200	0.074
2	1:500	0.074	0.520	90200	0.0434

Table 5.2: *Properties of flow conditions in this study.*

When determining the slope, the pipe was surveyed to check for a consistent 1:500 slope both for each section separately and across all eight sections of the pipe together. The slope calibration was achieved with a surveying sight level and rulers inserted into the pipe, held in place by multi-purpose 3D printed wave probe guides attached to the top of the pipe. The aim was to run the experiments in uniform flow conditions, i.e at a water slope coinciding with pipe slope over the FOS length. A honeycomb mesh was installed at the main inlet tank outlet to reduce disturbances originating from the inlet tank, ensuring the direction of flow fully aligns with the pipe axis. After 30 minutes of developing flow,

a ruler was used to spot check each measurement location and confirm the water slope at the expected depths. Fine adjustments were made to the variable horizontal weir at the pipe outlet to get as close as possible to the desired conditions. Another 30 minutes of flow were allowed to pass before the start of the main two-hour data collection process. The flow rate was measured by flow meters at the main inlet tank outlet, controlled by butterfly valves immediately downstream of these meters. There were two and one flow meters and valves for flow regimes 1 and 2, respectively. The bulk flow velocity (U_b) for each flow condition is determined based on the flow rate (Q) and the wetted area (A), i.e. $U_b = Q/A$. The Reynolds Number was calculated via the formula $Re = 4U_b R_h / \nu$. Here, R_h is the hydraulic radius, calculated as $R_h = A/P$, where P is the wetted perimeter of the flow, and ν is the kinematic viscosity of water. The FOS response to the pressure fluctuations was recorded using a Luna si155 ST optical sensing interrogator at a sampling rate of 1 kHz. Data were collected for two hours for each steady, uniform flow condition, whereby the FOS sits at the pipe bottom position.

5.4 Data analysis

Figure 5.3 illustrates the PSD of the strain recorded by all six gratings (Equation (5.2)) obtained over 30 minutes in the two-hour measurement period in both flow regimes (see Table 5.2), as well as a test in the absence of flow. Note, $\phi_{\epsilon\epsilon}$ denotes the FOS PSD. A sharp roll-off in PSD magnitude can be observed at 15 Hz in the response of all gratings to flow, showing that the high-frequency components of the fluctuating wall pressure are attenuated. The length of the individual gratings within the fibre extends over 1 cm, and thus small scale information (fluctuating at higher frequencies), gets averaged out over this length scale, meaning the FOS structure acts as a low-pass filter.




Figure 5.3: *One-sided PSD of FOS strain response from the pipe bottom to (a): flow regime 1, (b): flow regime 2, and to (c): empty pipe. Time series of FOS strain response from the pipe bottom to (a): flow regime 1, (b): flow regime 2, and to (c): empty pipe.*

Due to the nature of turbulence in this type of flow and the low-pass filtering effect of the sensor, the sensor only responds to larger flow structures and ignores small-scale fluctuations within the fluctuating wall pressure. Thus in the subsequent analysis, all work is conducted with a cutoff frequency of 15 Hz. It is important to note the response of the FOS when no flow is present, there is little response across all frequencies, as expected as there is no excitation acting on the sensor. It is noted that the strain signals in Figure 5.3(a-c) exhibit some degree of noise in the high frequency range. This is expected given the nature of the whole experimental setup. Aside from the wall pressure fluctuations inducing a response of the FOS, there are other mechanical elements such as inlet conditions and structural vibrations of the pipe rig itself. It is clear that the consistent behaviour between datasets of active fluid flow (a-b) occupies the low frequency portion of the spectra, far from this range in both cases. This provides further grounding to implement a cut off frequency of 15 Hz to ensure all data analysis steps are conducted on processed data containing within it signatures of the expected low frequency turbulence, while excluding higher frequency components that are more likely to be dominated by measurement noise

or structural response unrelated to the flow.

5.4.1 Flow velocity recovery

For both flow regimes, data from each of the six FBGs within the fibre were collected at a sampling rate of $f = 1$ kHz for two hours. Linear trends in these time series were removed, and then fed through a third-order Butterworth bandpass filter for different intervals in the 0.1-15 Hz range, resulting in a vector of strains $[\varepsilon_i(t)]$ for the i 'th FBG. Note, $\varepsilon_i(t)$ is the instantaneous strain response of grating i to the wall pressure fluctuations at time t . As the flow has uniform flow depth and is assumed to be fully-developed throughout the measurements, turbulence, and hence FOS statistics, remain constant along the streamwise direction. For each FBG, data collected over this two-hour period were divided into smaller non-overlapping windows of M seconds where $M \geq 30$ s. This minimum value of M was chosen sufficiently large such that each individual window could be reasonably treated as an independent sample of the stationary turbulence field. Under this assumption, splitting the time series into sections in this manner enables us to treat each section as a different realisation of the same field on which to perform analysis, consequently resulting in more estimates of the target hydraulic parameters that are the output of this analysis:

Splitting $[x_i(t)]$ up into sections of M seconds allows it to be rewritten in the following matrix form,

$$\hat{\varepsilon}_i(t) = [[\varepsilon_{i,1}] [\varepsilon_{i,2}] \dots [\varepsilon_{i,m}]]. \quad (5.3)$$

Where $m = 7200/M$ is the number of time series segments, and each column $[\varepsilon_{i,k}]$ corresponds to the k 'th segment of the time series of FBG i , each column consisting of fM time steps.

Utilising the built-in 'xcorr' function in MATLAB, the time lag τ^* maximising the cross-correlation between the k 'th M minute section of gratings i and j , $R_{i,j,k}$ was computed as:

$$\tau^* = \underset{\tau \in [0, M]}{\operatorname{argmax}} (R_{i,j,k}(\tau)). \quad (5.4)$$

Where

$$R_{i,j,k}(\tau) = \left(\frac{1}{M} \int_0^M [\varepsilon_{i,k}(t)][\varepsilon_{j,k}(t + \tau)] dt \right). \quad (5.5)$$

Finding the lag τ^* gives a time delay estimation (TDE) between the time series of the two gratings. Computing the ratio

$$U_{b,i,j,k} = \frac{\xi_{ij}}{\tau^*}, \quad (5.6)$$

where ξ_{ij} is the streamwise separation between FBGs i and j , results in an estimate of the velocity of the propagating turbulent pressure fluctuations at the pipe wall obtained from the k 'th section of the time series. I can compute $U_{b,i,j,k}$ for all combinations of pairs (i, j) , $1 \leq i < j \leq 6$ to provide estimates of the bulk flow velocity recorded between all combinations of grating pairs, for all k sections. Below are histograms of velocity estimations obtained from applying this methodology to fibre data collected at the pipe bottom for both flow regimes. A bandpass window from 0.1-15 Hz was used, split into sixty $M = 120$ second sections, resulting in 900 estimates of flow velocity displayed in each subfigure below.




Figure 5.4: *Distribution of flow velocity estimates obtained through manipulating cross-correlation between grating readings from the pipe bottom in (a): flow regime 1, (b): flow regime 2, both filtered in the range of 0.1-15 Hz.*

While there are a significant number of velocity estimates concentrated around the true bulk flow velocity for each regime, there exists a large number of outlying velocity estimates in each case. It was hypothesised by the authors that all information pertaining to bulk flow is contained within a certain frequency band, and any information outside of this range corresponds to either electrical noise or turbulent-induced vibration of the pipe structure itself. For example, filtering the time series in the range 2-7 Hz in each regime yields the following velocity histograms:




Figure 5.5: *Distribution of flow velocity estimates obtained through manipulating cross-correlation between grating readings from the pipe bottom in (a): flow regime 1, (b): flow regime 2. Both filtered in the range of 2-7 Hz.*

The distributions in Figure 5.5 have a significant mass concentrated around a particular estimate of velocity, having modes at 0.7 and 0.48 m/s, respectively. These are reasonable estimates of the flow velocity in both flow regimes recovered from a TDE of the sensor readings from the plate/gel boundary, with an error of 0.3% and 7.7% as compared to the values stated in Table 5.2, respectively. While not all noise has been expelled from these histograms, it is clearly an improvement compared to filtering over the range 0.1-15 Hz. It appears that by applying an appropriate filtering window, outlying velocity estimates were excluded, demonstrating that the sensor is sensitive to flow features within a specific frequency range. This analysis emphasises that the sensor can detect bulk flow velocity, provided the data are properly filtered, and that features travelling at the bulk flow velocity dominate the FOS time series inside this range. The following section outlines a stochastic sampling scheme designed to further refine the search for the frequency range retaining bulk flow information and expelling outlying velocity estimates to the greatest extent.

Finally, it is important to note that, when computing velocity estimates from all possible grating pairs, potential biases that may arise must be considered. There is a slightly greater number of FBG pairs corresponding to smaller streamwise separations than larger ones. When combining velocity estimates across all available pairs, the resulting histograms may then be biased towards velocity estimates recoverable from more closely-spaced gratings. In particular, if there is behaviour recoverable from closely-spaced gratings that largely-

spaced gratings cannot capture, this may impact the overall shape of the recovered velocity distributions, depending on the severity of this bias. To challenge this, results will also be examined as a function of streamwise separation in later sections to observe if any bias due to uneven grating separation produces a notable skew in any derived quantities.

5.4.2 Automation of filtering process

To streamline the analysis process for flow velocity recovery, it was in the interest of the authors to automate the filtering aspect of the procedure outlined in the previous section. Parametric inference was performed on the upper and lower bounds of the bandpass window applied to the experimental data to determine what frequency range retained information on flow to the greatest extent. A Scoring-Based MCMC method, adapted from the conventional MH algorithm [169], was implemented by the authors to achieve this. For a given set of initial bandpass limits a, b , encoded as a random variable $\theta = \{a, b\}$, the scheme iteratively proposes new values for the filter bounds based on a stochastic process. A given parameter value θ was assessed by evaluating the form of the histogram produced from the FOS time series filtered in this window. Owing to the difficulty in finding an analytic expression mapping filter bounds to velocity histograms, and not having a reference dataset I wish to emulate, it was decided to assign a score to velocity histograms. This scoring function $\beta(\theta)$ evaluates the quality of these histograms by examining the number of peaks, and location of the highest mode. It was the idea that more optimal values of θ produce velocity histograms with a concentrated mass around the bulk flow velocity with few outliers, in each case.

This scoring function behaves as a surrogate for the likelihood function as seen in conventional MH, allowing the scheme to explore the parameter space and identify which values of θ maximise the retention of velocity estimates pertaining to bulk flow in the time series. Akin to conventional MH, the transition kernel q along with an acceptance probability A comprised the move step of the scheme. The form of q was chosen to be a two-dimensional normal distribution centred on the current filter bounds, with a standard deviation σ , set by the user. A uniform prior π was imposed on θ . Moving from $\theta \rightarrow \theta^*$ was governed by the acceptance ratio, $\frac{\beta(\theta^*)}{\beta(\theta)}$, where the values of θ produce higher quality histograms (being assigned a higher score through β), were more likely to be accepted, and vice versa. The aim of this scheme was to provide a distribution over θ , allowing us to infer which bandpass windows produce more optimal histograms, under the definition of β that has been set.

In likelihood-free schemes such as ABC and ABC-MCMC [170, 147], it is important that the metric D used for assessing suitability of samples θ is carefully chosen so that the model outputs $Y^*(\theta)$ align closely with the experimental data Y under some definition of distance D between the two sets of data. In our case there is no experimental data that I am attempting to replicate, just a predefined histogram shape I wish to achieve; the scoring function assesses sample suitability by evaluating the features of the histograms based on predefined criteria set by the user. By penalising multimodal histograms (representative of outlying velocity estimates) and peaks of these histograms centred far away from the true bulk flow velocity, β is able to score highly only those samples that produce velocity distributions with characteristics more consistent with our expectations, which guides the parameter estimation process. Owing to these desirable histogram properties, the following form of β was chosen:

$$\beta(\pi_1, \pi_2; \theta) = e^{-\pi_1(|U_b^* - U_b|)} e^{-\pi_2(n_p - 1)}. \quad (5.7)$$

In this expression, U_b denotes the experimentally-determined bulk flow velocity of the flow regime, and U_b^* denotes the mode velocity estimate of the histogram. The first exponential term imposes a penalty on histograms whose mode deviates significantly from the true bulk flow velocity, with the severity of the penalty controlled by the input parameter π_1 . The second term defines a penalty based on n_p , the number of peaks present in the given histogram produced by θ for each regime. The degree of this penalty is tuned by the input parameter π_2 . When considering flow regime 1 for example, a parameter $\theta = \{a, b\}$ producing a histogram with all of its mass centred on $U_b^* = 0.7\text{m/s}$ with standard deviation $\sigma \ll 1$, would score highly under β , given that $n_p = 1$ and $|U_b^* - U_b| \ll 1$ in this case. Whereas θ producing a histogram with all of its mass centred on $x^* = 0\text{m/s}$ (thereby not capable of recovering bulk flow velocity) would score poorly due to an incorrect mode estimation of the bulk flow velocity. As seen in Figure 5.3 and discussed earlier in this section, the sensor attenuates frequencies above 15 Hz. Consequently, the search space for the filter bounds was limited to frequencies below this threshold, defined as $\Theta = \{(a, b) | 0 \leq a < b - 5 \leq 10\}$. The constraint of $b - a > 5$ was imposed to allow the bandpass window to have sufficient width so the resulting filtered time series was not over-constrained in its frequency components. The sampling scheme, outlined in Algorithm 3, was applied to both flow regimes. For brevity, only the results from flow regime 2 shall be displayed in this section.

Algorithm 3: Score-Based MCMC Sampler

Input: $\beta(\pi_1, \pi_2; \theta)$: Scoring function, σ : Proposal standard deviation, $P(\theta)$: Prior distribution of θ , N : Number of target samples

Output: $\{\theta_1, \theta_2, \dots, \theta_N\}$: Accepted samples

```
1 Initialize  $\theta_{\text{current}} \leftarrow \theta_0$ ;  
2 Initialize  $n_a \leftarrow 0$ , Samples  $\leftarrow \{\}$ ;  
3 while  $n_a < N$  do  
4   Sample  $\theta^* \sim \mathcal{N}(\theta_{\text{current}}, \sigma^2)$  // Propose new sample  
5   Compute  $\beta(\pi_1, \pi_2; \theta^*)$  // Evaluate scoring function for  $\theta^*$   
6   Compute  $\beta(\pi_1, \pi_2; \theta_{\text{current}})$  // Evaluate scoring function for current  
   sample  
7   Compute acceptance probability:  
  
                                     
$$P_{\text{accept}} = \min \left( 1, \frac{\beta(\pi_1, \pi_2; \theta^*)}{\beta(\pi_1, \pi_2; \theta_{\text{current}})} \right)$$
  
8   Sample  $u \sim U(0, 1)$ ;  
9   if  $u < P_{\text{accept}}$  then  
10     $\theta_{\text{current}} \leftarrow \theta^*$ ;  
11    Append  $\theta_{\text{current}}$  to Samples;  
12     $n_a \leftarrow n_a + 1$ ;  
13  end  
14 end  
15 return Samples;
```

The input parameters of Algorithm 3 for inference of the optimal frequency range for flow regime 2 were as follows; $\sigma : 3$ Hz, $\pi_1 : 10$, $\pi_2 : 5$, $\theta_0 : \{4, 9\}$, $P(\theta) = U(\Theta)$, $N = 1000$, where U denotes the uniform distribution. The value of σ was chosen sufficiently large so the scheme was able to explore the whole domain Θ , as can be seen in Figure 5.6(a).




Figure 5.6: (a): Heatmap indicating areas of the parameter space Θ searched by the scheme. (b): Density heatmap of accepted samples (normalised to integrate to 1 over Θ). The colour gradient of blue to red represents an increasing number of samples.

It should be noted that the constraints on the domain in which to search yield the upper left triangle of the domain $[0.1, 15] \times [0.1, 15]$. After the target number of samples n_t was accepted, additional figures signalling the distribution of accepted samples were produced in Figure 5.6(b). One can see the density of accepted samples is concentrated in a thin strip of Θ containing a lower bound of 2.5 Hz and a variable upper bound in the range of 8-15 Hz, indicating a range of possible upper bounds. The fixity of the lower bound indicated that lower bounds straying from 2.5 Hz contribute to velocity histograms inconsistent with the requirements to score highly under the score function, as seen in Equation (5.7). It was believed there is variability in the upper bound due to the roll off seen in the PSD of the data (Figure 5.3b). As the upper bound of the window is increased, the additional present frequency components approach the noise floor in the sensor response, providing little additional information in the sensor response. This decrease in magnitude of the higher frequency components did not hinder the schemes ability to produce sufficiently high scoring histograms. However, it can be seen from additional analysis that the introduction of these higher upper bounds contributed to an increasing number of outlying velocity estimates (Figure 5.7). For this reason, it was decided to implement the lowest of these upper bound ranges to retain the window in which the sensor responds most strongly to

flow.




Figure 5.7: *Velocity histograms for flow regime 2. Filtered in (a): 2.5-8 Hz, (b): 2.5-11 Hz, (c): 2.5-14 Hz.*

In conclusion, the score-based MCMC scheme successfully automated the task of filtering the data in the correct frequency range for bulk flow velocity recovery. By designing a scoring function that operates as a reward system for histograms consisting of accurate bulk flow velocity estimates across the entire time series, the scheme is guided to retain the bandpass windows achieving this characteristic, without use of a reference dataset. The results of the scheme dictate a preference for a particular subset of frequency ranges in the parameter space that yield more consistent bulk flow velocity estimates. The output of the optimised frequency range given in Figure 5.7(a) aligns well with the desired features of the velocity histogram, demonstrating the effectiveness of the approach. For flow regime 1, the same analysis was performed, yielding an optimal bandpass window of 4.5-10 Hz in this case. In all subsequent analyses, it was these bandpass windows that were to be applied to the data for each flow regime.

5.4.3 Effect of time interval and outlier removal

Subsequent analysis was performed to assess the implication of the signal duration M on the resulting velocity histograms. On the previously identified optimal bandpass windows, time series from both regimes were split up into $M = 30, 60, 300, 600, 900, 1800$ second sections, before velocity histograms were produced via the time delay estimate method. Lower values of M provide more data realisations to cross-correlate, resulting in a greater number of flow velocity estimates. Provided these greater number of estimates are concentrated around

a value of the expected velocity, I can confirm that the flow characteristics affecting the sensor response appear over a short time scale, and that the flow information that is recoverable from the data remains present throughout the whole time series. On the other hand, recovering fewer velocity estimates for larger M allowed us to observe if there is any long-term behaviour in the time series causing bias in the velocity estimates. In Figures 5.8 and 5.9 below, the histogram bars have been colour-coded according to the two-point FBG measurement used in each estimate. Darker bars refer to FBGs with a greater streamwise separation, and vice versa.




Figure 5.8: *Distribution of flow velocity estimates obtained for different values of M in flow regime 1. $M = 30, 60, 300, 600, 900, 1800$ seconds (a-f). Colour legend given for each grating pair.*




Figure 5.9: *Distribution of flow velocity estimates obtained for different values of M in flow regime 2. $M = 30, 60, 300, 600, 900, 1800$ seconds (a-f). Colour legend given for each grating pair.*

Even though the mode of the distributions varied slightly for each subfigure in Figure 5.8, the bulk flow velocity estimates remain in good agreement with the experimentally-obtained reading, yielding an averaged error of 4.3% and 7.7% over all tested values of M for each flow regime, respectively. In flow regime 2, I again observed modes of the velocity estimates that were in good agreement with the experimentally-obtained bulk flow velocity reading, staying consistent no matter the time interval used in the calculation in Figure 5.9. The colour coded element of the histograms allowed us to see that for any value of M , the more closely-spaced gratings near the upstream end of the plate (FBGs 1,2,3) consistently overestimated the bulk flow velocity, as seen in the $0.8 - 1$ m/s and $0.6 - 1$ m/s range in Figures 5.8 and 5.9, respectively. Plotting the velocity estimates with respect to the spatial separation used, I can further observe how the velocity estimates change across the sensor length in Figure 5.10. The results plotted in Figure 5.10 suggest that the error in the velocity estimates reduces with increasing separation between sensor elements.




Figure 5.10: *Velocity estimates plotted as a function of streamwise separation between sensor elements. Bold blue and red lines indicate true bulk flow velocity in flow regime 1 and 2, respectively. $M = 1800$.*

Measurements of bulk flow were consistent in each regime, apart from in the smaller separations, as indicated previously. Due to the outlying measurements consistently produced by gratings 0.04 m apart or less in both regimes, it is obvious however that no matter the value for M , closely-spaced gratings may have a shared sensitivity to pressure fluctuations, leading to time lags that are shorter than expected for the flow regime. These gratings are also positioned within a few mm of the end of the plate length (Figure 5.2(b)). It is understood that the two-point correlations of these gratings could also be corrupted due to edge effects of the plate or mounting uncertainties of the sensor structure. Despite these pitfalls, so long as the data were filtered in the identified frequency range for both regimes, the mode of the velocity histogram provided reasonable bulk flow velocity estimates for all values of M , highlighting the sensors ability to recover meaningful bulk flow velocity estimates of flow over a range of timescales. When assessing reliability of these histograms, there is a trade-off between the number of estimates, and any negative effects manifesting in the estimates due to the analysis procedure, as highlighted above. Given the trade-off between the quantity and quality of velocity estimates, $M = 120$ was chosen as the optimal balance, minimising the negative effects from shorter time intervals while preserving sufficient estimates for meaningful analysis.

To further refine the velocity estimates and reduce the impact of outlying measurements, an

additional analysis was implemented to remove estimates from FBG elements spaced less than 0.04 m apart. This specified distance threshold excludes the two-point correlations obtained from FBGs: (1,2), (1,3), (2,3), and (3,4) which were the pairs responsible for overestimation of the bulk flow velocity in Figures 5.8 and 5.9. By excluding these data points, the analysis reduces the influence of shared sensitivity to pressure fluctuations among closely spaced gratings. The resulting histograms in Figure 5.11, free from the distortions caused by closely spaced gratings, provide tighter velocity histograms, allowing for a more precise characterisation of the flow.




Figure 5.11: (a-b): Distribution of flow velocity estimates obtained for $M = 120$ in (a-c): flow regime 1, (b-d): flow regime 2, excluding FBG pairs (1,2), (2,3), (1,3), (3,4). (c-d): Normal distribution fit to velocity histograms.

In both regimes, excluding the overestimated velocities from closely spaced gratings resulted in tighter histograms that better align with a normal distribution. The resolved

values of μ for each regime align closely to the true bulk flow velocity in each case, giving an error of 1.2% and 5.3% for flow regimes 1 and 2, respectively. By quantifying the uncertainty of our bulk flow velocity estimates through the standard deviation σ , I can evaluate the performance of the resulting histograms given the necessary considerations that have been highlighted in this section.

5.4.4 Streamwise correlation structure

In the previous sections, the importance of filtering the FOS data in a correct frequency range to highlight bulk flow dynamics was highlighted. As detailed in Figures 5.6 and 5.7, there exists an optimal filter range that is capable of recovering bulk flow characteristics from the sensor response. Due to the existence of this range, it was decided to analyse the streamwise correlation of the time series data filtered in this specific range only. It has been suggested that correlation lengthscales of wall pressure fluctuations relate to the size of turbulent structures in the flow, and so it was the aim to determine if patterns of correlation obtained from the sensor match realistic expectations of streamwise pressure correlations at the pipe wall. The temporal offset between neighbouring FBGs i, j was presumed to approximate ξ_{ij}/U_b , with the maximum cross-correlation values at these temporal offsets decreasing for an increasing separation ξ_{ij} , showing that the bulk flow manifests in the FOS response. As well as focusing on the time lag between gratings, in this section I focussed on the peak cross-correlation values to assess the reduced coherence of the pressure fluctuations over greater distances. It is understood that the responses of FBGs separated by at most 0.04m, coupled by their joint response to local pressure fluctuations, overestimated the bulk flow velocity by up to 0.3 and 0.4 m/s in flow regimes 1 and 2, respectively (see Figures 5.8 and 5.9). Nonetheless, I will display measurements inclusive of these grating combinations to evaluate their impact on the streamwise correlation patterns recovered by the FOS, and determine if they introduce any unexpected behaviour.




Figure 5.12: *Overlaid filtered time series of FBG 1 (black line) and FBG 2-6 (grey lines) (a-i) in flow regime 1. Normalised cross-correlation between respective gratings (b-j).*

Figure 5.12 presents examples of the time series pairs from the FBG sensor array and corresponding cross-correlation functions. From Figure 5.12 on the right (b-j), I observe the expected behaviour: A decrease in the maximum correlation between two pairwise gratings of increasing streamwise separation, and an increase in the time delay at which this maximum was observed. This reinforces the idea that the degree of coherence of the

FOS to wall pressure fluctuations decreases with increased separation, which will be characterised throughout this section. Due to the assumption of homogeneous turbulence, it is understood that correlation between two points in space only depends on their separation, and not their absolute position in space [171]. Thus, the streamwise array of 6 individual gratings can be utilised to calculate the spatial correlation at $6(6 - 1)/2 = 15$ distinct non-negative separations in the streamwise direction. To achieve this, the time series signals for each non-negative separation pair of gratings (i, j) were cross-correlated, and the maximum value of this normalised cross-correlation

$$r_{ij} = \max \left(\frac{1}{T} \int_0^T [x_i(t)] [x_j(t + \tau)] dt \right) \quad (5.8)$$

was recovered. This corresponds to the highest degree of correlation between two locations ξ_{ij} apart, regardless of time lag between the two signals. By not restricting the lag at which to recover this value of r_{ij} , this form of spatial correlation provides information on how the wall pressure fluctuations correlate at different separations, as the fluctuations convect with the flow.

In order to more formally characterise the behaviour of maximum spatial correlation r_{ij} , an analytical expression $\hat{r}$ was proposed as

$$\hat{r}(\xi) = e^{-\xi/L}. \quad (5.9)$$

In this expression, ξ is the streamwise separation along the length of the sensor, and the additional parameter L is defined as the correlation length. It has been proposed that the spatial correlation of boundary layer pressure fluctuations can be approximated by the analytical expression given above [172, 173, 82, 174]. Profiles of r_{ij} , their fitting to Equation (5.9), and the respective time lags these maximums were attained at, are displayed in Figures 5.13 and 5.14 for values of $M = 120, 600, 900, 1800$ for both flow regimes. These figures plot the maximum correlation coefficient and estimated time lag as a function of the sensor separation. The dashed and solid lines correspond to the fitted dependence for these parameters, respectively. The gradient of the latter provides an estimate of the bulk flow velocity U_b . A range of time intervals is provided to observe if there exists a significant change in the streamwise correlation decay when this is computed over longer instances. It should be noted, gratings spaced less than 0.04 m apart were excluded when fitting the correlation profiles to Equation (5.9).




Figure 5.13: Streamwise correlation estimates in flow regime 1 plotted against streamwise separation, $M = 120, 600, 900, 1800$ seconds (a-d). The average of r_{ij} at each streamwise separation was fitted to Equation (5.9) (left axis). The average time lag at each streamwise separation was fitted to a straight line (right axis). Grey dashed line indicates cut-off point for separations used in the fitting. Time lags at which maximum correlations were attained are plotted against streamwise separation (right axis).




Figure 5.14: *Streamwise correlation estimates in flow regime 2 plotted against streamwise separation, $M = 120, 600, 900, 1800$ seconds (a-d). The average of r_{ij} at each streamwise separation was fitted to Equation (5.9) (left axis). The average time lag at each streamwise separation was fitted to a straight line (right axis). Grey dashed line indicates cut-off point for separations used in the fitting. Time lags at which maximum correlations were attained are plotted against streamwise separation (right axis).*

Focusing on the correlation structure itself, it is obvious from the data presented in Figures 5.13 and 5.14 that the sensor response decorrelates over a shorter length in flow regime 2. Even though the variance is greater in the individual maximum correlation estimates for shorter intervals in both cases, the mean values of these estimates stay roughly constant, as seen through the evolution of the correlation length L in each case. Only minor fluctuations occur in the fitted correlation length for both regimes, hinting that the underlying pressure fluctuations correlate similarly across different time intervals, even with the addition of

larger variance for shorter time intervals. Denoting the correlation length for flow regimes 1 and 2 as L_1 , L_2 , respectively, I have $L_1 \approx 0.1265$ m and $L_2 \approx 0.0795$ m. The hydraulic radii for regimes 1 and 2 are given by $R_{h_1} = 0.074$ m and $R_{h_2} = 0.0434$ m, respectively. I find that the ratio of the correlation lengths and hydraulic radii yield $\frac{L_1}{R_{h_1}} \approx 1.71$, $\frac{L_2}{R_{h_2}} \approx 1.83$ in each case, i.e. it varied only 7% between the two regimes. The consistency of this ratio suggests there exists a scaling relationship between pipe geometry and the flow dynamics of partially-filled pipe flows. Despite the differences in flow regimes and correlation lengths, these normalised values are comparable. This hints that the hydraulic radii are able to quantify the extent to which the wall pressure fluctuations correlate, and that the near-wall pressure fluctuations, and their correlations, are in fact proportional to the flow geometry.

It is also helpful to overlay the time lags at which these maximum correlation values were recovered for each streamwise separation in both flow regimes. It is clear these lags follow a linear trend, with the time lag increasing with streamwise separation, suggesting a consistent flow structure causing peaks in the correlation along the length of the sensor. It is understood that the inverse slope of this trend provides an estimate of the flow velocity in each flow regime. Specifically, this provides another method to visualise how the bulk flow velocity estimates recovered from the sensor are relatively consistent in space and time (bar the overestimation from closely spaced gratings). In fact, in the $M = 1800$ case (Figures 5.13(d) and 5.14(d)), this corresponds to an error of 0.7% and 4.9% in flow regimes 1 and 2, respectively. These lag plots solidify the conclusion that the fluctuating pressure driving the sensor response propagate over the sensor with the bulk flow velocity, affirming that these correlation patterns recovered by the sensor are a consequence of flow dynamics rather than external sources of noise in the sensor response, such as pipe vibration.

5.4.5 Large-scale motions

To further characterise the size of structures present in the flow independently of the streamwise correlation structure, pre-multiplied spectra were computed from the measured FOS signals. In the literature, pre-multiplied spectra are used to visualise and interpret different scales embedded in the data, most notably used to identify LSMs. Since the earlier work of Falco [175], LSMs have been understood as coherent structures present in free-surface flows that extend over multiple flow depths [176, 177, 178], manifesting as peaks in the pre-multiplied spectra at locations corresponding to their spatial scale with

respect to flow geometry. It is understood that LSMs exist at scales between 1 and 4 flow depths, that is $1 < \lambda/h < 4$, where λ is the LSM scale [179, 16].

Typically, LSMs are identified through pre-multiplied spectra and cross-spectra of velocity components within the flow. As I only have access to FOS-obtained data, the power spectral density as seen in Figures 5.3(a-b) shall be our starting point. To obtain approximate information on the spatial scales of coherent structures in the flow, I convert from frequency to wavenumber domain through use of Taylor's hypothesis. Using the relation $k = 2\pi f/U_b$, I have $k\phi_{\varepsilon\varepsilon}(k) = f\phi_{\varepsilon\varepsilon}(f)$. Plotting this as a function of λ/h , where λ is the streamwise wavelength scale derived from $k = 2\pi/\lambda$, I have the following figures:




Figure 5.15: Normalised pre-multiplied spectra from (a) flow regime 1, and (b) flow regime 2 as a function of wavelength scale. Parameters U_b and h used in each figure are taken from values in Table 5.2 for each regime. Wavelength axis displayed in logarithmic scale.

In flow regime 1, the pre-multiplied spectra reveal peaks at $\lambda = 0.87h, 3.92h, 14h$, corresponding to structures of roughly 0.129, 0.58, 2 m in length. Similarly, for flow regime 2, peaks are observed at $\lambda = 1.53h, 4.2h, 39.5h$, corresponding to structures 0.11, 0.31, 2.9 m in length. The latter peaks in these lists would, by definition, correspond to VLSMs, even larger streamwise structures spanning $10 < \lambda/h < 30$ existing in the outer layer of flow [180, 181]. However, the FOS lies at 70-140 flow depths (depending on the regime), and it is reported that VLSMs only begin to fully develop at streamwise positions greater than 130 flow depths from the inlet. Thus, the existence of these peaks shall be ignored, as VLSMs are not expected to be fully developed within 130 flow depths from the pipe inlet [179]. However, the first two peaks in both cases lie in the accepted range of LSMs

in the literature, suggesting that the FOS responds to larger-scale coherent structures associated with LSMs in addition to the bulk flow properties analyses in previous sections. In fact, the smaller scales with $\lambda \approx h$ recovered from both regimes agree well with LSM scales recovered from the near-wall fluid domain, as seen in Kim & Adrian [2, 182, 183]. It has been reported that the scales of LSMs increase with increased distance from the wall, until a maximum is attained, before decreasing again to scales on the order of 2–4 flow depths [181]. The scales $\lambda \approx 4h$ seen for both regimes may therefore correspond to either large-scale features originating near the wall, or to structures in the outer region of the flow imprinting their signature onto the pipe wall.

Even though the FOS is flush-mounted, it cannot be said for certain that these recorded structures are near-wall features. As seen in previous sections, the FOS can recover bulk flow properties far from the near-wall region; the same may apply to LSMs originating in the outer-layer. While the focus of this manuscript has been on demonstrating the utility of the presented novel FOS to infer bulk flow conditions from the pipe wall, the results presented in this section point to the potential of the FOS to resolve multiple scales within turbulent flows. However, a more detailed investigation into its sensitivity across different spatial scales and the precise origin of these structures is beyond the scope of the current study.

5.4.6 Out of water comparison

An important question to raise when looking at the analysis from the previous subsections is, how do I know these results arise in the FOS data as a consequence of bulk flow, and are not to do with some other process present in the data? It is important to verify that the velocity estimates and streamwise correlation patterns obtained pertain to flow, and not to the sensor response to structural vibration of the pipe rig itself.

To address these questions, I aimed to identify any potential processes manifesting in the sensor response from structural vibration by rotating the pipe section in which the FOS lies so that it sat on the pipe wall above the free surface of the flow. The free surface lies at 90° and 60° about the pipe centre for regimes 1 and 2, respectively. The section containing the FOS was rotated to position the FOS at 104° and 88° , lying above the free surface in each regime, respectively. By recording FOS data for two hours of fully-developed flow for both regimes in these respective positions, no water is in contact with the sensor, and the response came solely from the manifestation of pipe vibration in the sensor at the specified

angular position. Following the analyses presented in subsection 5.4.1 to observe that the output velocity histograms and from out of water measurements are uninformative in this case, it was possible to confirm no flow information is recoverable in these cases.




Figure 5.16: Overlaid filtered time series of FBG 1 (black line) and all other gratings (grey lines) ((a)-(i)) in flow regime 1 at the 104° position (above the free surface). Normalised cross-correlation between respective gratings ((b)-(j)).

Where in Figure 5.12 there was a clear decay in the maximum correlation values for in-

creased grating separation in response to flow, this behaviour is not evident in Figure 5.16. There is a slight peak seen in figures (b) and (d), however, I believe this is due to the fact that these gratings lie within 40 mm of each other, and so a very slight degree of correlation is expected in response to the pipe vibration. Even though FOS data at 0 degrees were still contaminated by the pipe vibration, isolating the vibration component and highlighting that the outputs from these data in Figure 5.16 are uninformative for our purposes, indicated that the characteristics that are important to us seen in Figure 5.12 arise from the flow component of the data. The idea that the main response of the sensor when placed underwater is due to flow and not structural vibration will be further reinforced in this subsection.

Even though it is believed there is an optimal bandpass window to apply to the data to retain flow properties, bulk flow velocity still appears in the velocity histograms of the pipe bottom data when these data are filtered in the whole frequency range in which the sensor responds (0.1-15 Hz). Figure 5.4 illustrates that filtering over this wider window still retains strong peaks at the expected bulk flow velocity for both regimes, even with unrealistic velocity estimates produced by the frequency components corresponding to noise. For this reason, the out-of-water measurements were filtered over this whole range to identify the behaviour of the sensor over the whole active frequency range.




Figure 5.17: *Distribution of flow velocity estimates obtained through manipulating cross-correlation between grating readings from above the free surface in (a): flow regime 1, (b): flow regime 2, both filtered in the range of 0.1-15 Hz. Histograms computed from 2 hours of fully developed flow, split into 2 minute sections.*

As expected, in Figure 5.17, there is a large concentration of mass about 0 m/s in both regimes when the FOS was placed above the free surface, with no prominent peaks indicating an element of a particular velocity present in the FOS time series. The absence of distinct peaks in these histograms highlights that the FOS output reflects the response to structural vibration rather than flow in this case. It should also be noted that this observation further validates the distinct peaks observed in Figure 5.5. When the FOS is placed at the pipe bottom, the peaks are indeed due to flow rather than sensor interactions with structural vibration. By comparing these in- and out-of-water measurements, I show that the signal captured in the underwater position aligns with the expected bulk flow velocity estimate, providing different conclusions to those derived from Figure 5.17. These findings establish a reliable basis for interpreting the submerged FOS response as being driven by flow characteristics, further supporting the reliability of the proposed sensor to recover bulk flow conditions whilst only collecting data from the pipe wall.

5.5 Conclusion and future work

This study presented the design, installation, and analysis of a flush-mounted fibre-optic sensor for quantifying hydrodynamic conditions for partially-filled pipe flows. Contrary to typical sensing systems that had disturbed flow dynamics by protruding into the flow, the novel sensor had been equipped to infer these conditions through its response to the hydraulically turbulence-induced pressure fluctuations at the interior pipe wall.

Key insights had included bulk flow velocity recovery through a cross-correlation-based analysis of the sensor time series, which had been proven to operate as desired through application to two flow regimes. In each case, velocity estimate histograms had displayed good agreement with the experimentally-set bulk flow velocities, even in the presence of noise. By optimising the particular bandpass window in which to filter the data, further elements of noise were removed to enhance consistency in the velocity estimates obtained from the sensor. This had resulted in errors as little as 0.7% and 4.9% between recovered and true bulk velocity after post-processing the data in flow regimes 1 and 2, respectively. With differing spatial separation of sensing elements within the cable, and using the assumption of a stationary turbulence field driving the response of the sensor, streamwise correlation patterns had been established. Resulting correlation patterns from both flow regimes had been consistent in their decay at a rate proportional to the hydraulic radius in each case, whereby the ratio of correlation length to hydraulic radius had only differed by 7% between flow regimes, highlighting the sensors ability to capture spatial characteristics of wall pressure fluctuations and providing evidence that the near-wall pressure scaled with flow geometry. The validity of these results had been further aided through out-of-water experiments, where the sensor response had been solely influenced by structural vibrations of the pipe. In this position, no bulk flow velocity had been recovered, reinforcing the conclusion that the sensor output underwater had been primarily driven by flow dynamics rather than structural vibrations of the pipe. It had been understood that closely spaced sensing elements within the cable had possessed a shared sensitivity to the pressure fluctuations from the flow, skewing velocity estimates obtained from these grating pairs. Perhaps for future iterations of this sensor, a spatial separation between gratings greater than 40 mm had been recommended to ensure the readings from individual gratings were independent and free from a joint local response of the FOS structure to the pressure. Indeed, the accuracy of the bulk velocity estimates produced through the cross-correlation analysis had been greatly improved when rejecting any estimates from closely spaced gratings in the

sensor, as demonstrated in Figure 5.11. While the focus of the presented research has been on the applicability of the novel sensor for monitoring bulk flow conditions in partially-filled pipe flows, further spectral analysis provides evidence to suggest that the full response of the sensor is due to multiple spatial scales within the flow. In future work, these spectral features will be investigated in greater detail to quantify the presence and dynamics of multi-scale flow structures in partially-filled pipe flows, compared to recovering only bulk flow information from the FOS, as has been the focus of this manuscript.

In conclusion, the proposed fibre-optic sensor had demonstrated the ability to infer bulk flow conditions in partially-filled pipe flows, while only collecting data from the pipe wall. This non-intrusive approach had minimised flow disruption and had leveraged turbulence-induced wall pressure fluctuations to provide accurate estimates of bulk flow velocity and spatial correlation patterns. The results presented here could have been improved upon through the use of a fibre-optic cable with more fibre Bragg gratings written into the core. Thereby allowing for more two-point measurements to have been taken, and a longer profile of streamwise correlation patterns to have been established. Future work will involve data collection at different angular positions about the pipe cross-section to understand the angular change of pressure fluctuations within these flows. Recovery of bulk flow velocity will also be achieved through computation of the frequency-wavenumber spectra of the unfiltered FOS data at different angular positions for comparison. Finally, further validation of results will be achieved through comparison with analogous results obtained from Large-Eddy Simulation (LES) wall pressure fluctuation data.

Acknowledgements

The authors would like to acknowledge the support of the UK Engineering and Physical Sciences Research Council (EPSRC) Impact Acceleration Account (IAA), and through grant EP/S017283/1 on Distributed Fibre-optic Cable Sensing for Buried Pipe Infrastructure. This research made use of the Integrated Civil and Infrastructure Research Centre (ICAIR) which has been part funded by EPSRC, Grant UKCRIC-National Distributed Water Infrastructure Facility (EP/R010420/1) and the European Regional Development Fund (Project Number 28R15P00608). For the purpose of open access, the authors have applied a ‘Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising.

Chapter 6

Azimuthal variation in wall pressure-induced strain statistics of turbulent partially-filled pipe flows: an experimental and numerical study

6.1 Abstract

This experimental and numerical study leverages fibre-optic strain measurements as a proxy to investigate azimuthal variation of wall pressure statistics for partially-filled pipe flows at two distinct turbulent flow regimes. A novel fibre-optic sensor operating on the principles of Bragg reflection is embedded into the wall surface of an experimental pipe rig in a streamwise orientation. Through correlating measurements obtained between the individual FBG elements within the sensor, statistical characteristics of the strain response of the sensor to wall pressure fluctuations are recovered around the wetted perimeter. Wall pressure data were obtained from a LES study of equivalent flow regimes to provide an analogous numerical dataset for comparison with experimental measurements. Similarly to wall shear, the azimuthal variation in the sensor strain response to wall pressure fluctuations is shown to be influenced by flow geometry. Frequency-wavenumber spectra reveal dominant convective ridges and the corresponding convection velocities, which are in close agreement between experimental and numerical datasets, identifying all temporal

and spatial scales present. Through Taylor’s hypothesis, signatures of large-scale motions are identified, and the energy contribution of different scales is characterised around the wetted perimeter, revealing that larger scales dominate at azimuthal positions closer to the free surface. The results demonstrate that the fibre-optic sensor responds to wall pressure fluctuations, azimuthal variations thereof, and identifies the influence of secondary flow structures, providing new insight into the spatial organisation and multi-scale behaviour of these flows.

6.2 Introduction

The study of turbulence-induced pressure fluctuations inside boundary layers is valued in many engineering contexts, as these pressure fields can yield notable structural consequences. In high-speed aircraft, cabin noise is produced due to a vibratory response of the structure to these pressure fluctuations [184]. Similarly, in fluid transport systems, these fluctuations exert added loadings onto pipe walls, shortening operation lifetimes due to fatigue [185]. A more comprehensive understanding of the structural excitation sources in fluid-structure systems can only be beneficial to operators, as steps can be taken to mitigate the destructive effects of these pressure fluctuations and increase the operational lifetime of the structure in question.

Whether this be through direct measurement inside the boundary layer [84], or through development of models derived from the Poisson equation [186], the study of these pressure fluctuations in varying fluid-structure settings has received considerable attention in recent decades. Most notable is the development of analytical and numerical models of the frequency-wavenumber spectra of wall pressure fluctuations in turbulent boundary layers. These are included but not limited to, the models of Corcos, Chase, Goody, among others [51, 187, 54]. Comprehensive reviews updating the state of the art of these wall pressure fluctuation models are given in [158, 188]. The early works of Clinch, and Willmarth & Wooldridge, instead focus on direct measurement of the wall pressure fluctuations via wall-mounted pressure transducers in a smooth-walled pipe and wind-tunnel, respectively [82, 189]. The analyses in these works focus primarily on the correlations of the wall pressure fluctuations parallel to the direction of bulk flow. In these papers, parameters characterising pressure correlation decay rates were uncovered from the statistical data, yielding empirical space-time correlation estimates with parameters varying with frequency.

While numerous studies have characterised the wall pressure fluctuations in varying fluid-structure systems, there has been little focus on pressure fluctuations at the wall in pipes running partially-filled. The presence of a free surface in these flows introduces secondary flow motions, which redistribute the velocity field around the wetted perimeter. This, in turn is believed to influence the nature of the pressure fluctuations on the pipe wall [190], as are the nature of the secondary flow motion patterns. By analysing the effect of wall pressure fluctuations on the signal response of a FOS mounted in the wall of a partially-filled pipe flow, I aim to recover wall pressure fluctuation properties, and azimuthal variations thereof, in two distinct flow regimes. The FOS is installed in a rotatable pipe section, allowing streamwise measurements at different angular positions around the pipe cross-section. The streamwise correlation can be analysed not only as a function of wall position and flow regime, but in the frequency-wavenumber space as well. Specifically, by defining empirical relations of the streamwise correlations, insight will be gained into the extent to which turbulent structures are propagated along with the flow. Azimuthal patterns of wall pressure-induced FOS strain statistics will be recovered and compared against observations derived from LES wall pressure fluctuation data that match the experimental flow regimes.

Beyond advancing understanding of free surface flows, this work has broader significance. Pipes operating partially-filled are widespread, mainly comprising a large quantity of storm and wastewater conduits. Enhanced understanding of flow dynamics of partially-filled flows, and the corresponding azimuthal variation of pressure around the interior wall, proves to be a valuable first step in linking turbulence-induced fluid loading to structural reliability and the physical processes controlling the hydraulic capacity of individual pipes in such networks. This paper is organised as follows: Section 2 introduces the experimental facility and fibre-optic sensor; Section 3 outlines the LES framework, setup and validation; Sections 4-5 describe statistical analysis and data processing procedures to recover bulk flow properties; Section 6 presents results, including correlation and frequency-wavenumber analysis of bulk flow from both datasets, signatures of LSM recoverable from the time series, and azimuthal variations in FOS-recovered strain response to wall pressure fluctuations and LES wall pressure fluctuations for both flow regimes; and finally Section 7 provides the main conclusions from this work.

6.3 Experimental setup and flow conditions

All experiments were carried out at ICAIR at the University of Sheffield. This was comprised of eight separate, but carefully aligned, 2500 mm long, 5 mm thick acrylic pipe sections. The gravity-driven pipe flow rig measured 20 m in length with an internal diameter of 0.296 m. The pipe sections were installed on top of a sand bed on the C section of the steel framework to isolate the pipe from internal and external sources of vibration. An outlet section contained a flow control gate that was adjusted to achieve the desired flow depth of each regime before the data collection process. A consistent pipe slope of 1:500 was set by adjusting the slope of the stiff steel frame and checking the slope through surveying sights and rulers along all eight sections of pipe. Installed at the inlet tank outlet was a hexagonal honeycomb mesh to encourage symmetrical flow and reduce large-scale disturbances from the inlet tank. The honeycomb mesh was 40 mm thick with a 4 mm cell size.




Figure 6.1: *Schematic view of the ICAIR pipe facility.*

The sensor section was installed 12.5 m downstream of the inlet, roughly 84 pipe radii, to reduce disturbances originating from the inlet and to encourage a well-developed flow to be achieved at the FOS position. The maximum water depth in these partially-filled pipe flows was approximately one pipe radius. A 900 mm x 24 mm slot in streamwise and spanwise directions, respectively, was removed from the sensor pipe section to allow for FOS installation. The two-layer FOS was fitted into the slot such that the face of the plate exposed to the flow lay flush with the pipe wall. Consisting of a polypropylene plate housed among a bed of low stiffness magic gel to ensure plate deformation in response to wall pressure fluctuations, an engionic Femto Gratings GmbH SM1250BI(9.8/125)P fibre-optic cable was installed at the plate-gel boundary along the FOS length. Embedded into


the fibre core were six FBGs that experienced axial strain under deformation of the plate. Each FBG is associated with it a Bragg wavelength,

$$\lambda_B = 2n_{eff}\Lambda. \quad (6.1)$$


Here, n_{eff} is the effective refractive index of the fibre core, and Λ is the spatial grating period. Axial strain ε averaged over the length of an individual FBG (10mm) in an isothermal environment is recovered through the relation

$$\frac{\Delta\lambda_B}{\lambda_B} = (1 - \rho_e)\varepsilon. \quad (6.2)$$

Where ρ_e is the photo-elastic coefficient of the fibre core material, and $\Delta\lambda_B$ is the Bragg wavelength shift recorded directly by the Luna si155 ST optical sensing interrogator connected to the FOS fibre at a sampling rate of 1 kHz. For the analyses presented in this work, I operate on strain-proportional time series when working with FOS data as seen in Equation (6.2). Since the main goal of this work is to characterise the temporal and spatial structure of flow features rather than measure absolute pressure, a direct conversion between strain and pressure was deemed unnecessary.



(a)



(b)

Figure 6.2: Schematic diagrams of (a), cross-section and (b), lengthwise view of the FOS structure. FBG element locations and indices highlighted in (b) (not to scale).

Regime	Slope	h (m)	U_b (m/s)	Re (-)	R_h (m)
1	1:500	0.148	0.698	207200	0.074
2	1:500	0.074	0.520	90200	0.0434

Table 6.1: Properties of flow conditions in this study.




Figure 6.3: cross-section schematic of a partially-filled pipe flow showing quantities of central angle γ and flow depth h .

The FOS collected data for two different flow conditions, outlined in Table 6.1. As mentioned in chapter 3, slope and flow depth were set during the flow calibration procedure, and flow rate was measured by flow meters at the main inlet tank. Knowledge of flow rate Q , flow depth h , and central angle γ of the circular segment formed by the liquid level in

each regime enables calculation of the wetted area and perimeter A and P as

$$A = r^2(\gamma - \sin \gamma \cos \gamma), P = 2r\gamma. \quad (6.3)$$

Thus, an estimate of bulk flow velocity U_b was given as $U_b = Q/A$. The Reynolds' number was calculated by the relation $Re = 4U_b R_h / \nu$, where ν is the kinematic viscosity of water, and R_h is the hydraulic radius in each flow regime, the ratio of flow cross-sectional area and wetted perimeter, given by the relation $R_h = A/P$.

The sensor pipe section was installed with removable couplings to allow for a variable angular position of the FOS around the pipe cross-section, allowing measurements to be taken at different circumferential locations for each dataset. This adjustability enables data collection at multiple positions, providing a greater understanding of azimuthal variations in FOS response to wall pressure fluctuations.




Figure 6.4: *cross-sectional views of (a) flow regime 1, (b) flow regime 2. Red markings indicate angular positions at which the FOS collected data in each regime.*

As outlined in Figure 6.4, the FOS collected data at varying positions characterised by the angle γ of the segment formed between the pipe bottom and cross-section position with respect to the pipe centre point. The maximum angular position of the FOS was 72° and 40° , respectively, on either side of the pipe bottom in each flow regime. This was to ensure that the main process driving the sensor response was from wall pressure fluctuations, not surface effects. Rotating the sensor section so the FOS lay at a particular angle, flow was allowed to develop to an equilibrium state over thirty minutes. Before data collection, flow steadiness was confirmed by verifying a neutral water slope over the sensor section. This was achieved using rulers for depth measurements at two upstream and one downstream

location in an attempt to ensure consistency with the wave probe depth measurements. Following this, all tools were removed from the flow for 10-15 minutes to allow for any disturbances to pass. Once uniform, fully developed flow conditions were attained after an additional 30 minutes of flow development, data from the FOS were recorded for a total of two hours at the given angular position. This process was repeated to obtain an extensive dataset of two-hour strain measurements induced by the plate's dynamic response to the wall pressure fluctuations for all angular positions and flow regimes.

6.4 Numerical Framework and flow conditions

Numerical wall pressure fluctuation time series were obtained from a LES method for flow conditions matching those in Table 6.1. The LES data provides space-time wall pressure fluctuation data spanning the wetted perimeter of each flow regime. The contribution of this thesis with regards to the LES study is to process and analyse the raw LES data to produce results comparable to the experimentally-obtained results. The LES method is employed to simulate turbulent flow in partially filled pipes using the open-source code Hydro3D. As a finite difference solver, Hydro3D has demonstrated its effectiveness in capturing large turbulence structures and water surface deformations in such flow conditions [190, 17]. The code solves the filtered Navier-Stokes equations for incompressible, unsteady and viscous flow, which are described as follows:

$$\frac{\partial u_i}{\partial x_i} = 0 \quad (6.4)$$

$$\frac{\partial u_i}{\partial t} + \frac{\partial u_i u_j}{\partial x_j} = -\frac{1}{\rho} \frac{\partial p}{\partial x_i} + \frac{\partial(2\nu S_{ij}^r)}{\partial x_j} - \frac{\partial \tau_{ij}}{\partial x_j} + f_i + g_i \quad (6.5)$$

Here, u_i and u_j denote the resolved velocity components in the x -, y - and z -axis directions ($i, j = 1, 2, 3$), while x_i, x_j represent the spatial coordinates. p is the spatially resolved pressure and ρ is the density of the fluid; ν is the kinematic viscosity of the fluid; f_i is the volume force introduced in the immersed boundary (IB) method proposed by [191] to impose the no-slip and impermeable constraints on the fluid-solid interface; g_i is the gravitational acceleration. $S_{ij}^r = 1/2(\partial u_i/\partial x_j + \partial u_j/\partial x_i)$ denotes filtered strain-rate tensor. The subgrid-scale (SGS) stresses τ_{ij} , are computed using the wall-adapting local eddy viscosity (WALE) model [192].

The convection and diffusion terms of the Navier-Stokes equations are approximated by 4th-order accurate central differences. An explicit 3-step Runge-Kutta scheme is used to

integrate the equations in time, providing 2^{nd} -order accuracy. A fractional step method is employed, i.e. within the time step convection and diffusion terms are solved explicitly first in a predictor step which is then followed by a corrector step during which the pressure and divergence-free velocity fields are obtained via a Poisson equation. The latter is solved iteratively through a multi-grid procedure [193, 194].

The water-air interface is tracked using the level-set method [195], which employs a signed distance function η defined as:

$$\eta(x, t) \begin{cases} < 0 & \text{if } x \in \Omega_{gas} \\ = 0 & \text{if } x \in \partial\Omega \\ > 0 & \text{if } x \in \Omega_{liquid} \end{cases} \quad (6.6)$$

Here, Ω_{gas} and Ω_{liquid} denote the air and water domain, respectively, while $\partial\Omega$ represents their interface. The interface's motion is predicted by solving a pure convection equation [196]:

$$\frac{\partial \eta}{\partial t} + u \cdot \nabla \eta = 0 \quad (6.7)$$

Discontinuities between density and viscosity at the interface can lead to numerical instabilities. This is avoided by setting a transition zone in which density and dynamic viscosity transition between water (l) and air (g).

$$\begin{aligned} \rho(\eta) &= \rho_g + (\rho_l - \rho_g) \mathbf{H}(\eta) \\ \hat{\mu}(\eta) &= \hat{\mu}_g + (\hat{\mu}_l - \hat{\mu}_g) \mathbf{H}(\eta) \end{aligned} \quad (6.8)$$

The transition zone is defined as $|\eta| \leq \chi$, where χ is half the thickness of the interface. This is implemented through the Heaviside Function $\mathbf{H}(\eta)$ as formulated [197, 198]:

$$\mathbf{H}(\eta) = \begin{cases} 0 & \text{if } \eta < -\chi \\ \frac{1}{2} \left[1 + \frac{\eta}{\chi} + \frac{1}{\pi} \sin \left(\frac{\pi \eta}{\chi} \right) \right] & \text{if } \eta \leq \chi \\ 1 & \text{if } \eta > \chi \end{cases} \quad (6.9)$$

To preserve the level-set function's property $|\nabla \eta| = 1$ over time, a re-initialisation procedure is applied to the transition zone, as proposed by [199]:

$$\frac{\partial d}{\partial t_a} + s(d_0) (|\nabla d| - 1) = 0 \quad (6.10)$$

where $d_0(x, 0) = \eta(x, t)$, t_a is the artificial time and $s(d_0)$ is the smoothed sign function formulated as:

$$s(d_0) = \frac{d_0}{\sqrt{d_0^2 + (\nabla d_0 \cdot \varepsilon_r)^2}} \quad (6.11)$$

This partial differential equation is solved for several iteration steps, $\varepsilon_r/\Delta t_a$, where ε_r is a single grid space. These adjustments to the level set method are applied only in the interface zone [200].

Regime	Q (L/s)	h/D (%)	h (m)	U_b (m/s)	u_* (m/s)	S (-)	Re (-)	Re_τ (-)	Fr (-)
1	23.93	49.7%	0.147	0.70	0.0366	1:541	207200	2617	0.82
2	6.96	25%	0.0735	0.52	0.0283	1:527	89700	1220	0.8

Table 6.2: *Hydraulic properties of the simulations*

The hydraulic properties of the simulation are detailed in Table 6.2. The water depth, h/D , and bulk flow velocity, U_b , are very similar to the experiment, although frictional losses are not. The friction velocity u_* is defined as $u_* = \sqrt{\frac{R_h}{\rho} \frac{dp}{dx}} = 0.0366$, where $\frac{dp}{dx}$ is the time-averaged pressure gradient obtained from the LES. Using the time-averaged pressure gradient, the channel bed slope is calculated as $S = \frac{dp}{dx}/g$.

While the bulk flow velocity and water depth are utilised as direct inputs to match experimental conditions in the LES scheme, the channel slope S is an output of the simulation. Discrepancies in this value compared to the target slope appear from grid resolution constraints at high Reynolds numbers. As such, the grid resolution is relatively sparse, at roughly $\Delta z_m^+ = 14.52$ wall units, where Δz_m^+ represents the mean grid resolution in the vertical direction. To obtain a value of S more closely approximating the bed slope in the experiment, a value of $\Delta z_m^+ \approx 5$ would be needed, which would increase the computation time by 27-fold.

The length of the pipe is $L_x = 10.81r$, this was taken to be sufficiently long to allow full development of secondary currents. The spanwise and vertical dimensions of the computational domain are $L_y = L_z = 2.054r$, i.e. slightly larger than the diameter of the pipe, because extra grid points are required for representing the pipe walls using the IB method. Periodic boundary conditions are applied in the streamwise direction.

The simulations employ a mesh resolution of $640 \times 380 \times 380$ grid points in the x -,

y -, and z - directions, respectively. Uniform grid spacing is used in all directions, with $\Delta x = 0.0025$ m and $\Delta y = \Delta z = 0.0008$ m. Due to the geometry of the pipe, the actual distance from the first grid point to the wall varies with azimuthal angle, but it remains well below the maximum grid spacing at most locations. The grid resolution in wall units is approximately $\Delta x^+ = \frac{u_* \Delta x}{\nu} \approx 87$ and $\Delta y_m^+ = \Delta z_m^+ = 14.52$, where Δy_m^+ represents the mean grid resolution in the spanwise direction. The current grid resolution ensures that at least one grid point lies within the viscous sublayer near the wall, as required by the IB method.

The simulation is initially run for approximately 30 seconds (which corresponds to 13.12 flow-through times) to establish fully developed turbulence, followed by an additional 240 seconds to collect turbulence statistics. The wall pressure at 80 cross-sections with a streamwise spacing of 2 cm is recorded. For each cross-section position, 23 different angular locations ranging from -77° , to 77° , and -42° , to 42° ,s relative to the pipe bottom are recorded, with a uniform interval of 7° , for flow regimes 1 and 2, respectively. The sampling frequency for the wall pressure is approximately 208 Hz. The LES data were processed using the same numerical and analytical framework as was applied to the FOS data, allowing for direct comparison procedure for interpreting the FOS response to wall pressure fluctuations in partially-filled pipe flows.

6.5 Analysis procedure

When understanding statistical properties of the fluctuating wall pressure, a commonly studied object is its correlation between two different points separated only in the direction of the mean flow. It is known that this fluctuating wall pressure field is stationary in space and time in the case of fully developed flow. This assumption enables statistical quantities of wall pressure to be expressed as functions of spatial separation and time lag, rather than absolute position and time instances used in the computation.

For a function $u(x, t)$ satisfying the above assumption, the normalised space-time correlation between two points with spatial and temporal separation ξ and τ , respectively, is typically given by

$$R_{uu}(\xi, \tau) = \frac{1}{T\sigma_u(x, t)\sigma_u(x + \xi, t + \tau)} \int_0^T u(x, t)u(x + \xi, t + \tau)dt. \quad (6.12)$$

Where σ_u represents the standard deviation of u , acting as a normalising factor bounding

R_{uu} between -1 and 1. Similarly, for values of $\xi = 0$ and $\tau = 0$, I denote $R_{uu}(0, \tau)$ and $R_{uu}(\xi, 0)$ as the autocorrelation and spatial correlation, respectively. Due to the assumption of stationarity, R_{uu} is only a function of separation in space and time. As a consequence of the Wiener-Khinchin theorem, the CSD is equal to the Fourier transform in time of the space-time correlation function.

$$\Gamma_{uu}(\xi, \omega) = \frac{1}{2\pi} \int_{-\infty}^{\infty} R_{uu}(\xi, \tau) e^{-i\omega\tau} d\tau, \quad (6.13)$$

where ω denotes angular frequency. Taking the spatial Fourier transform of the CSD yields the frequency-wavenumber spectrum of u .

$$\Phi_{uu}(k, \omega) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \Gamma_{uu}(\xi, \omega) e^{-ik\xi} d\xi = \frac{1}{(2\pi)^2} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} R_{uu}(\xi, \tau) e^{-i(k\xi + \omega\tau)} d\tau d\xi, \quad (6.14)$$

which is also the space-time Fourier transform of $R_{uu}(\xi, \tau)$, where k is the wavenumber in the x direction. Setting $\xi = 0$, I can see that

$$\Gamma_{uu}(0, \omega) = \phi_{uu}(\omega) = \frac{1}{2\pi} \int_{-\infty}^{\infty} R_{uu}(0, \tau) e^{-i\omega\tau} d\tau = \int_{-\infty}^{\infty} \Phi_{uu}(k, \omega) dk. \quad (6.15)$$

Here, $\phi_{uu}(\omega)$ is defined as the PSD of u . This is the time Fourier transform of the autocorrelation function $R_{uu}(0, \tau)$, while also being equivalent to the frequency-wavenumber spectra integrated over all wavenumbers. The PSD ϕ_{uu} can be thought of as a distribution of variance in the function u as a function of frequency ω . Finally, the root-mean square (RMS) of u , σ_u^2 is calculated as

$$\sigma_u^2 = \int_{-\infty}^{\infty} \phi_{uu}(\omega) d\omega = R_{uu}(0, 0). \quad (6.16)$$

The derived quantities presented above play a large part in understanding the properties of wall pressure fluctuations. Due to the available data being discrete in space and time, these relations are approximated through numerical procedures applied to the time series data. From the spatial configuration of the FOS, experimental data are only attainable in the streamwise direction at the six gratings of varying separation; a schematic view of their locations in the FOS is given in Figure 6.2. The spacing of the points at which fluctuating wall pressure was recovered in the LES simulation has been highlighted in Section 6.4. Due to the streamwise configuration of the FOS, the analysis in this study is limited to statistical

quantities obtained from two-point measurements along the streamwise direction. Angular variations in wall pressure fluctuations will be assessed through direct comparison of the streamwise results obtained at different angular positions for both the FOS and LES data.

A discrete formulation of the above statistical quantities will be introduced for a placeholder signal u , denoted as $u_\gamma(x_j, t_k)$ when referring to time series data for a given angular position γ on the pipe wall. Throughout the manuscript, dataset-specific quantities will be denoted by ε for FOS data, and p for LES data. When considering the space-time correlation of the two sets of data, the root-mean square of the time series at position (x_j, γ) is given by

$$\sigma_{u_\gamma}^2(x_j) = \frac{1}{T} \sum_{h=1}^T u_\gamma^2(x_j, t_h). \quad (6.17)$$

Denoting $\xi_{ij} = |x_i - x_j|$, the space-time correlation can be determined as

$$R_{uu}(\xi_{ij}, \tau_k) = \frac{1}{T} \sum_{h=1}^T \frac{u_\gamma(x_i, t_h) u_\gamma(x_j, t_h + \tau_k)}{\sigma_{u_\gamma}(x_i) \sigma_{u_\gamma}(x_j)}. \quad (6.18)$$

Consequently, the CSD is denoted by

$$\Gamma_{uu}(\xi_{ij}, \omega) = \frac{1}{2\pi} \sum_{h=1}^T R_{uu}(\xi_{ij}, \tau_h) e^{-i\omega\tau_h} \Delta\tau \quad (6.19)$$

The frequency-wavenumber spectrum $\Phi_{uu}(k, \omega)$ is obtained by taking the discrete spatial Fourier transform of Γ_{uu} , which is equivalent to the space-time Fourier transform of the correlation R_{uu} . Summing over all distinct streamwise separations (i, j) , I obtain,

$$\Phi_{uu}(k, \omega) = \frac{1}{2\pi} \sum_{i,j} \Gamma_{uu}(\xi_{ij}, \omega) e^{-ik\xi_{ij}} \Delta\xi = \frac{1}{(2\pi)^2} \sum_{i,j} \sum_{h=1}^T R_{uu}(\xi_{ij}, \tau_h) e^{-i(k\xi_{ij} + \omega\tau_h)} \Delta\tau \Delta\xi. \quad (6.20)$$

As the number of available measurement points in the fibre (6 FBGs) is scarce, it was decided to utilise all possible streamwise separations between two different FBGs to attain an array of $6(6-1)/2 = 15$ non-negative separations. Computing the cross-correlation between each pair of gratings enables a denser collection of space-time correlation estimates from the FOS, under the stationary assumption. These spacings and grating pairs for $1 \leq i < j \leq 6$ are given in ascending order in Table 6.3. The LES simulations provided

a dense coverage of the wall pressure fluctuations in the streamwise direction, giving a streamwise spacing of 20 mm for a total length of 1.6 m, and so this pairwise matching procedure was not carried out for LES data.

FBG pair (i,j):	(1,2)	(2,3)	(1,3)	(3,4)	(4,5)	(2,4)	(1,4)	(5,6)	(3,5)	(2,5)	(1,5)	(4,6)	(3,6)	(2,6)	(1,6)
separation (mm):	16	24	40	40	56	64	80	88	96	120	136	144	184	208	224

Table 6.3: *Streamwise separation (mm) of all FBG combinations inside the FOS listed in ascending order.*

Due to the uneven FOS spacing, the space-time correlation measurements were interpolated in space to attain a uniform grid of measurements before transformation into the wavenumber domain. This interpolation attempts to mitigate the effects of the sparse spatial information. It was chosen to interpolate the FOS data onto a uniform grid of 50 locations spanning the operational length of the sensor (224mm), yielding spatial separations of $\Delta\xi = 0.00448\text{m}$, spanning wavenumbers between $k_s = \pm 109.5\text{m}^{-1}$, with a uniform step of $\Delta k = 4.46\text{m}^{-1}$. LES data has a spatial separation of $\Delta\xi = 0.02\text{m}$, spanning wavenumbers between $k_s = \pm 25\text{m}^{-1}$, with a uniform step of $\Delta k = 3.92\text{m}^{-1}$. The measurement durations and sampling periods were $T = 7200\text{s}$ and $1/\Delta t = 1\text{ kHz}$ for the FOS data, and $T = 260\text{s}$ and $1/\Delta t = 208\text{ Hz}$ for the LES data. The Nyquist frequencies for the FOS and LES data were given as $f_s = 500\text{ Hz}$ and $f_s = 104\text{ Hz}$, respectively.

6.6 Data processing

6.6.1 Spectral comparison

This subsection identifies spectral features of both the FOS response to fluctuating wall pressure and the LES wall pressure fluctuation signals. The spectral differences are attributed to differences in spatial resolution between the two data types. LES resolves down to the grid spacing of $\Delta x = 2.5\text{ mm}$ in the streamwise direction, whereas, as with any finite-sized sensor, the FBGs within the FOS inherently average strain over their 10 mm length. PSD of both datasets are presented from the pipe bottom position $\gamma = 0^\circ$ for both regimes to characterise the effective bandwidth of the FOS and identify which frequency ranges may yield reliable information from both data types. Data presented in the following figures are averaged over the whole streamwise domain in each case.




Figure 6.5: (a) Streamwise-averaged PSD obtained from FOS strain time series (converted from FBG wavelength shift via Equation (6.2)) at the pipe bottom position in both flow regimes. (b) Streamwise-averaged PSD obtained from LES wall pressure fluctuation time series at the pipe bottom position in both flow regimes. Frequency cut-off in (a) to match the LES frequency domain. Frequency axis displayed in logarithmic scale.

The data in Figure 6.5(b) can be thought of as excitations of the FOS structure for the corresponding regime, with the data in Figure 6.5(a) interpreted as the output. As expected, it is clear that the presence of the sensor structure induces a low-pass filtering effect when wall pressure fluctuations are transmitted into the FOS. Between regimes in 6.5(a), one can see the difference not only in the overall amplitude of the response, but also in the spectral content. Intuitively, the varied amplitude arises due to the greater magnitude of wall pressure activity driven by increased flow depth and water volume compared to flow regime 2. A distinct roll-off in the FOS response can be observed below 15 Hz for both flow regimes. It is believed this roll-off is governed by the mechanical properties of the sensor structure, the surrounding bed of low stiffness gel, and the spatial averaging effect of the individual FBG measurements. The shift in roll-off frequency between regimes suggests a variable active frequency range of the FOS, dependent on the flow regime, which is also observed between flow regimes in the case of LES data in Figure 6.5(b).

The computational steps involved in the LES study resolve the pressure fluctuations at pointwise locations on a dense grid of $\Delta x = 0.0025\text{m}$, which is then recovered at intervals of 0.02m to reduce computational load. As such, much smaller features down to the grid resolution are resolvable in the LES data compared to the FOS data. These smaller

features tend to oscillate at higher frequencies, explaining the greater bandwidth observable in Figure 6.5(b). In contrast, the inherent length over which individual FBGs operate naturally produces an averaging effect over this length, meaning fluctuations on spatial scales smaller than this cannot be resolved. This spatial averaging acts as a low-pass filter, which directly explains the reduced bandwidth observable in the FOS data in Figure 6.5(a). This phenomenon is widespread when considering measurement of the wall pressure. For well-known sensor types such as circular pressure transducers with either uniform or radially-dependent sensitivities, there exist transfer functions describing the degree of spatial averaging [85, 5]. Indeed, results contained in [85] indicate that application of these transfer functions to measured data to produce PSDs with a sharper roll-off at lower frequencies, as seen in Figure 6.5(a).

Despite this somewhat limited active bandwidth of the FOS, it can be seen in Figure 6.5 that both datasets possess significant energy in the low frequency range (below 15 Hz). It is this frequency range in which most of the energy resides before the roll-off into the energy cascade. As such, as this spatial averaging effect is concerned more with high frequency behaviour, it is deemed that the FOS is still capable of capturing information on low frequency, larger scale structures that dominate the signal in this frequency range, allowing for a meaningful comparison between data types. This, combined with the unique sensing mechanism and design involved with the FOS, it was deemed laborious to proceed with a complete spatial transfer function calibration for this study. Since the main energy-containing scales of interest are well within the FOS bandwidth, such a rigorous correction was not considered essential.

6.6.2 Filter optimisation process

Taking into account the spectral differences between the two data types used in this study, this subsection details a data-driven approach to infer the frequency range that preserves the streamwise correlation structure recovered by the FOS to the greatest extent. By utilising the 15 non-negative separations of all FBG elements within the FOS, a profile of streamwise correlation of the sensor response to wall pressure fluctuations at a given angle γ was constructed as

$$r(\xi_{ij}) = \max(R_{\epsilon\epsilon}(\xi_{ij}, \tau)). \quad (6.21)$$

The above expression corresponds to the maximum correlation observed between all pairs of gratings, regardless of time lag. Through imposing no restrictions on the lag, this

expression provides details on how the wall pressure convects with the flow between all grating pairs. Much work on the correlation of wall pressure fluctuations in the direction of mean flow has indicated that the streamwise correlation function can be expressed as [172, 82, 174]

$$\hat{r}(\xi_{ij}) = e^{-\xi_{ij}/L}. \quad (6.22)$$

Here, the *correlation length* L , governs the rate at which wall pressure correlations decay with increased streamwise separation.




Figure 6.6: Recovered streamwise correlation profiles and fittings to Equation (6.22) at the pipe bottom position for (a): flow regime 1, and (b): flow regime 2. Measurements taken over the four 30-minute time series were filtered in the range of 4-9 Hz in each case. Points left of the grey line indicate separations not used in fitting, and an in-depth explanation as to why is given in Chapter 5.

Examples of typical correlation profiles are given in Figure 6.6 for both flow regimes. A parameter estimation scheme was performed on the bounds of a bandpass filter applied to the FOS data at a given angle γ . The optimal bandpass limits were identified as those limits maximising the degree of streamwise correlation recovered by the FOS. This equates to recovering the bandpass limits that, when applied to the data, produced a correlation profile $r(\xi_{ij})$ with a maximal value of L for the given dataset. The underlying assumption is that a lesser rate of decay in the recovered correlation approximating the form of Equation (6.22) indicates greater retention of flow structure within the signal. To explore the space of possible bandpass bounds, a score-based MCMC method was implemented, mimicking the conventional MH scheme. Here, the bounds of a bandpass filter, written as a random

variable $\theta = \{a, b\}$, were assessed through how well the streamwise correlation profile fit the exponential model given in Equation (6.22). The recovered discrete correlation profile and analytic fitting were numerically integrated over the sensor length. The score $\beta(\theta)$ for given filter bounds θ was computed numerically as

$$\beta(\theta) = 1 - \frac{\int_{\xi_{\min}}^{\xi_{\max}} |\hat{r}(\xi) - r(\xi)| d\xi}{\int_{\xi_{\min}}^{\xi_{\max}} \hat{r}(\xi) d\xi}. \quad (6.23)$$

The fraction effectively characterises the deviation between the recovered and fitted profiles. The numerator gives the total deviation between the model and the measured correlation profile, while the denominator normalises by the area under the model curve. By subtracting this ratio from 1, $\beta(\theta)$ is closer to 1 when the profiles match more closely, and closer to 0 for more poorly fitted cases. The bandpass window 4-9 Hz producing the profiles in Figure 6.6 would score highly under $\beta(\theta)$ due to their strong fitting to the analytical formulation of correlation decay.

Much like the standard MH scheme as introduced in Chapter 3, the sample space is explored through a transition kernel q and an acceptance probability P_{accept} . The kernel q governs how new samples are proposed, given the current sample at each step in the scheme. Noting that the parameter space to be explored is a subset of $\mathbb{R}^2$, q was set as a two-dimensional normal distribution centred on the current bounds in each step. Due to the symmetry of q , the acceptance probability simplifies to

$$P_{\text{accept}} = \min\left(1, \frac{\beta(\theta)}{\beta(\theta')}\right). \quad (6.24)$$

This governs the probability of making the step from the current sample θ' to a proposed sample θ . The ratio ensures that a proposed sample scoring higher than the current sample would have a greater chance of acceptance. The goal of this scheme was to provide a distribution over possible filter bounds, allowing us to identify which bounds retain correlation structure to the greatest extent, for each flow regime and angle. As outlined earlier in this section, it is clear that the FOS attenuates frequencies above 15 Hz for both flow regimes, and so the search space for filter bounds θ was restricted to $\Theta = \{(a, b) \mid 0 \leq a < b - 4 \leq 10\}$. As to not over-constrain the frequency content of the filtered signal, a minimum bandwidth of 4 Hz was imposed on the sample space. The steps of the score-based MCMC scheme are outlined below.

Algorithm 4: Score-Based MCMC scheme

Input: $\beta(\theta)$: Score, σ : Proposal standard deviation, $P(\theta)$: Prior distribution of θ ,

N : Number of target samples

Output: $\{\theta_1, \theta_2, \dots, \theta_N\}$: Accepted samples

1 Initialize $\theta_{\text{current}} \leftarrow \theta_0$;

2 Initialize $n_a \leftarrow 0$, Samples $\leftarrow \{\}$;

3 **while** $n_a < N$ **do**

4 Sample $\theta' \sim \mathcal{N}(\theta_{\text{current}}, \sigma^2 \mathbf{I}_2)$ // Propose new sample

5 Compute $\beta(\theta')$ // Evaluate score for s'

6 Compute $\beta(\theta_{\text{current}})$ // Evaluate score for current sample

7 Compute acceptance probability:

$$P_{\text{accept}} = \min \left(1, \frac{\beta(\theta')}{\beta(\theta_{\text{current}})} \right)$$

8 Sample $u_p \sim \text{U}(0, 1)$;

9 **if** $u_p < P_{\text{accept}}$ **then**

10 $s_{\text{current}} \leftarrow \theta'$;

11 Append θ_{current} to Samples;

12 $n_a \leftarrow n_a + 1$;

13 **end**

14 **end**

15 **return** Samples;

Algorithm 4 was applied to the whole two hours of time series data individually for each angle and flow regime, collecting $N = 1000$ samples for each dataset. The uniform distribution $\text{U}(\Theta)$ was used as the prior for all datasets, and a consistent standard deviation of $\sigma = 3$ was chosen sufficiently large to give the scheme mobility to explore the whole space Θ in each case.




Figure 6.7: Heatmap of correlation lengths L from fitting the recovered stream-wise correlation profiles to Equation (6.22) on Θ . The dataset used is from the pipe bottom position in flow regime 1. The maximum value of L is denoted with a black marker.

One can see that, as a function of filter bounds, the correlation length converges to a maximum, denoted by the black marker. Straying from this peak yields filtered versions of the FOS data that decorrelate over shorter streamwise distances, indicating information is missing on how the transmission of wall pressure fluctuations into the FOS correlates over its operational length. The output of the score-based MCMC scheme yielded maximum correlation length estimates L for every angular position of the FOS in both flow regimes, as well as the filter bounds producing these maximums. The optimal bandpass windows are given in Figure 6.8.




Figure 6.8: *Recovered bandpass windows maximising correlation length of FOS data for (a) flow regime 1, (b) flow regime 2.*

The recovered bandpass windows in Figure 6.8 reveal that the frequency ranges maximising the correlation length vary around the wetted perimeter, with a roughly symmetrical pattern. The trend present in these figures indicates that the FOS is slightly more sensitive to higher frequency wall pressure fluctuations for angular positions closer to the free surface, indicating wall pressure fluctuations at these higher frequencies preserve the correlation structure closer to the surface. Throughout the rest of the paper, all instances of FOS data will be filtered in the identified bandpass windows for the given flow regime and angular position, unless stated otherwise. The LES data to be compared shall also be filtered in the equivalent window to ensure consistency between preprocessing steps applied to both data sources, unless stated otherwise.

6.7 Results & discussion

This section contains results obtained from application of the filter windows as seen in Figure 6.8 to all FOS and LES datasets at the corresponding azimuthal location. The results in this section focus on profiles of streamwise correlation, recovery of the frequency-wavenumber spectra, and evidence of large-scale motions present in the time series from both FOS and LES data. Conclusions are drawn from both dataset types, and where necessary, limitations and discrepancies are acknowledged between the recovered quantities. Throughout this section, emphasis is placed on the azimuthal variation of the quantities of interest, and how these azimuthal variations may relate back to wider flow features such

as secondary flow motions.

6.7.1 Streamwise correlation

Owing to the optimal filter ranges identified in the previous section, the correlation lengths of the FOS response could be recovered for all angular positions and flow regimes. For a given angular position and flow regime, streamwise correlations were recovered over four 30-minute instances of FOS data, whereby the ensemble average was fitted to Equation (6.22).




Figure 6.9: *Recovered best fit correlation lengths at all angular positions for (a) flow regime 1, (b) flow regime 2.*

While the angular change in correlation length is not perfectly symmetric in either case, there exist clear trends in the change of this parameter around the wetted perimeter, suggesting that the wall pressure fluctuations inducing these correlation structures remain coherent over greater distances at varying points around the wetted perimeter. This is approximately centred about the 0° position. In flow regime 1, it is clear that the maximum length over which the wall pressure-induced strain correlates occurs at roughly $\pm 24^\circ$ from the pipe bottom, with additional local minima seen at $\pm 40^\circ$. Straying closer to the free surface yields wall pressure-induced strain correlating over shorter distances, confirming that the local distance to the free surface controls the correlation length of the structures in this part of the flow. Flow regime 2 exhibits similar behaviour of increasing correlation length up to around $\pm 32^\circ$ about the pipe bottom, and then begins to decrease past this position. Looking at the scale of the correlation lengths for each regime, one can see the

values in flow regime 1 are roughly double the values in flow regime 2 for angles up to $\pm 40^\circ$, reflecting the ratio of flow depth between flow regimes. Regardless of the noisy behaviour, it is clear that variability in the wall pressure-induced correlation lengths around the wetted perimeter is recoverable by the FOS, even with the slightly asymmetric characteristics observed.

Instead of focusing on streamwise correlation directly, it is possible to analyse the correlation decay as a function of lag. Under the assumption of stationarity of the fully developed flow, it is understood that the convection velocity U_c remains constant along the FOS length. Here, convection velocity is defined as the mean speed at which eddies of varied sizes travel past a fixed point on the pipe wall. As such, the delay between FBGs of specified streamwise separation ξ can be approximated by $\tau \approx \xi/U_c$. By plotting this correlation as a function of lag for increased separations, the correlation decay in time of the FOS response to wall pressure can be recovered, the envelope of which is given by $R_{\varepsilon\varepsilon}(\tau U_c, 0)$, the auto-correlation of an observer convecting along with the flow [82]. Given the assumption of an exponentially decaying streamwise correlation, it is appropriate to fit the wall pressure correlation to an exponential decay in time, described by:

$$R_{\varepsilon\varepsilon}(\tau U_c, 0) = e^{-\tau/\tau_b}, \quad (6.25)$$

where the time correlation length, τ_b controls the rate of decay in time.




Figure 6.10: Correlation decay in time plotted as a function of 5 FBG separations at the pipe bottom in: (a) flow regime 1, (b) flow regime 2. Red markers indicate maximum correlation for a given separation ξ fitted to Equation (6.25).

The curve centred around $\tau = 0$ in both cases of Figure 6.10 is the autocorrelation of

FBG 1, defining the start point of the streamwise domain and time domain, as this correlation peaks at 0 time lag by definition. Each successive correlation $R_{\varepsilon\varepsilon}(\xi, \tau)$ peaks at progressively greater lags as a consequence of the flow taking longer to convect between the points. Simultaneously, the peak correlation diminishes with increasing separation, showing diminished correlation in space and time of the wall pressure.




Figure 6.11: *Recovered best fit time correlation lengths at all angular positions for (a) flow regime 1, (b) flow regime 2.*

Through fitting of Equation (6.25) to time correlations as a function of streamwise separation for all angular positions, the variation in time correlation length τ_b around the wetted perimeter can be recovered. When compared to Figure 6.9, I can see the angular changes in τ_b and L share a similar azimuthal dependence for both regimes as minima and maxima occur at roughly similar positions. While their absolute scales differ, this overarching similarity suggests a strong relationship between the spatial and temporal scales of the fluctuating wall pressure, meaning pressure fluctuations correlating over greater distances also correlate over greater timescales, and vice versa. Comparison between flow regimes shows smaller values of L and τ_b in flow regime 2, meaning a reduction in spatial and temporal scales of the wall pressure as compared to flow regime 1. This is to be expected when considering the difference in flow depths and, therefore, the overall strength of the fluctuating wall pressure in each case.




Figure 6.12: *FOS-recovered convection velocity estimates at all angular positions for (a) flow regime 1, (b) flow regime 2. Values normalised against U_b .*

Now that distributions of both L and τ_b around the wetted perimeter have been estimated, I aim to use this information to recover convection velocity estimates around the wetted perimeter. I draw from the approach seen in [82], who utilised knowledge of convection velocity to recover the streamwise correlation length, through the relation

$$L = U_c \tau_b. \quad (6.26)$$

Assuming the convection velocity changes around the wetted perimeter due to the secondary flow motions, I instead invert this relationship to infer the distribution of U_c around the wetted perimeter. This approach allows us to access a direct comparison against the experimentally-determined U_b whilst relying solely on quantities derived from correlation structures within the FOS data. Again, due to the similarities in the distributions of L and τ_b , but differences in relative scale, I see peaks in the same angular locations, showing that the convection velocity peaks at around the $\pm 24^\circ$ positions in flow regime 1. While the distribution for flow regime 2 is not as pronounced for all quantities shown so far, there is still a noticeable increase in convection velocity further from the pipe bottom position, which is consistent with the distributions of L and τ_b . In both cases, the recovered velocity estimates at the pipe bottom positions are $U_c = 0.705\text{m/s}$ and $U_c = 0.505\text{m/s}$ for flow regimes 1 and 2, respectively. While not a perfect match, these are reasonable estimates of bulk flow velocity in both cases, yielding errors of 1 and 3% at the pipe bottom position, respectively. While the FOS measurements provide reasonable fits to the pressure decay models introduced in this section, there is always going to be some degree of mismatch

due to experimental conditions and noise. Even with the mismatches, overall trends in the displayed parameters around the wetted perimeter can be realised, the values of which scale appropriately between flow regimes as well.

Equivalent analysis was performed on the LES-obtained wall pressure fluctuation data to compare against the observations made with the experimental results in this section. Before displaying results, it is important to consider the implications of the imposed periodic boundary conditions and length in the streamwise direction of the LES domain. While useful for simplifying the domain and aiding in computation time, periodic conditions result in less than realistic decorrelation properties of the wall pressure fluctuations, as the domain is essentially self-repeating. Nonetheless, azimuthal variations of LES-recovered wall pressure statistics will prove as a valuable qualitative comparison against FOS-obtained equivalents.




Figure 6.13: *Streamwise correlation profile of LES wall pressure data recovered from the pipe bottom position. All measurements taken against the reference point at the $x = 0$ position.*

To illustrate the issue presented with periodic boundary conditions, Figure 6.13 displays a typical streamwise correlation profile of the LES wall pressure data from the pipe bottom position. As expected, the correlation decays from a maximum value of 1 at $\xi = 0$ with increased separation. However, rather than eventually decaying to 0 with increased separation, a minimum of roughly 0.1 is reached at the $\xi = 0.4$ m mark, then plateaus before rising again to values approaching 1 towards the end of the streamwise domain. This rise is a clear artefact of the periodic boundary conditions. Points separated by $\xi = 1.6$ m are treated as adjacent due to domain wrapping, resulting in a falsely high correlation despite

significant physical separation. As such, the streamwise domain is not long enough to allow for complete decorrelation of the wall pressure before the periodic boundary conditions artificially increase the correlation values back to 1 as $\xi = 1.6$ m is approached.

To mitigate the influence of this artefact, correlation analysis from LES data is restricted to a streamwise domain spanning no more than $\xi = 0.2$ m, before the plateau occurs. This ensures that only the physically meaningful decaying portion of the correlation profile is used for fitting. It should be noted that the shape of the correlation profile in Figure 6.13 is invariant, regardless of the chosen reference point from which to correlate. By using this restricted range, parameters such as τ_b , L and U_c are therefore estimated from regions free of periodic contamination as best as possible. This approach prioritises fitting to the underlying flow physics and ensures that LES correlation results remain comparable to those obtained from experimental measurements.




Figure 6.14: (a): Streamwise correlation profile and fitting to Equation (6.22). (b): Correlation decay in time as a function of streamwise separation fitted to Equation (6.25). Both taken from LES flow regime 1 at the pipe bottom position.

It is clear that the correlation decay in space and time is not a perfect fit to the exponential decay model that is common in the literature. Experimental measurements of wall pressure are typically made using in-pipe sensors such as wall-mounted and pinhole sensors, which naturally average the pressure field over their finite operational area. As a result, these sensors are not capable of responding to spatial scales smaller than their operational radius/width [85]. In contrast, the presented LES data resolves smaller-scale fluctuations that would be filtered out in experimental sensor responses. Consequently, the spatially unfiltered LES data exhibit faster decorrelation than the experimental measurements even

when the LES data are equivalently bandpass filtered in time as in the experiment. Regardless of this mismatch in scales, the more significant feature of the correlation structure is its variability with angle around the wetted perimeter, which still varies appreciably even with the mismatch to the pure exponential decay models as seen in Figure 6.15.




Figure 6.15: *LES-recovered streamwise correlation length L in (a): flow regime 1, (c): flow regime 2, and time correlation length τ_b in (b): flow regime 1, (d): flow regime 2.*

The space and time correlation profiles in Figure 6.15 are to be compared directly to those seen in Figures 6.9 and 6.11. In flow regime 1, I can see clear symmetric peaks in LES values of L around the wetted perimeter, at values of $\pm 21, 63^\circ$, and local minimums at $0, \pm 40^\circ$. FOS-recovered values of L yield peaks at the $\pm 24^\circ$ position, with slight indication of local minimums at around $0, \pm 40^\circ$. However, the relative amplitude variation between minima and maxima in the FOS data is far less pronounced than in the LES results, obscuring the

visibility of these peaks. The angular change in τ_b is nearly identical to that of L within each dataset, showing consistent peaks and minimums in both, further reinforcing the idea that the decay of the wall pressure in space and time is directly correlated.

Focusing on flow regime 2, the magnitude of the correlation parameters is much more agreeable between FOS and LES-derived quantities. The main inconsistency here is the pronounced peak at the pipe bottom in the FOS-derived correlation parameters, which is not visible in the LES cases. As this behaviour is not observed in any of the other derived correlation parameters, it is reasonable to interpret this as an anomalous measurement, rather than physically meaningful. Disregarding this, both LES and FOS correlation parameters in this regime reach a peak in the range of $\pm 20 - 30^\circ$ from the pipe bottom, although the lack of symmetry in the FOS case makes the location of the peak difficult to quantify. This asymmetry is likely due to measurement noise and slight misalignments of positioning when rotating the pipe section.




Figure 6.16: *Streamwise-averaged root mean square LES wall pressure at all angular positions for (a): flow regime 1, (b): flow regime 2.*

The RMS of the wall pressure fluctuations from both LES flow regimes symbolises the extent to which the local wall pressure fluctuates about the mean value for a given angular position. Clear peaks can be seen at $0, \pm 40^\circ$. Given that the horizontal axes are scaled according to flow depth in Figure 6.16, this shows that the RMS of the pressure fluctuations obey roughly the same variation around the wetted perimeter between regimes up to $\pm 40^\circ$. Beyond this point, the RMS declines up to the angle denoting the free surface, in each case. As mentioned in [190], secondary currents exist in all partially-filled pipe flows, their strength increasing with flow depth. The downward motion of these counter-rotating flow

structures on each side of the pipe centreline separates in opposing directions, approaching the wall at roughly the $0, \pm 40^\circ$ positions. The increased intensity of these flow structures at these positions leads to maxima in the wall shear stress distribution [3, 190]. These peak locations are consistent with resolved values of maximal wall shear stress as seen in [190], as well as maxima in the turbulent kinetic energy [15], with the addition of local minima at the $\pm 20^\circ$ positions. Relating this change in RMS wall pressure back to the recovered correlation parameters, I can see where RMS is low in Figure 6.16, correlation parameters are maximised in Figure 6.15, and vice versa. This inverse relation suggests that greater RMS wall pressure leads to pressure readings that decorrelate over shorter streamwise separations and time lags due to greater fluctuation magnitudes about the mean pressure. Conversely, regions with lower fluctuations around mean pressure (near the $\pm 20^\circ$ positions) exhibit longer correlation lengths, indicating more coherent pressure signals along the pipe wall. This variation in correlation parameters, in turn, is a response to the change in wall-normal strength of secondary currents around the wetted perimeter. The findings shown here align with the results on the interplay between secondary flow motions and wall shear stress distribution from previous studies [1, 47].

While this behaviour is apparent for LES data, it is more obscured for the FOS data due to a slight asymmetry about the pipe bottom, and smaller magnitude changes between neighbouring minimums and maximums. Nonetheless, overall trends can still be observed around the wetted perimeter for the FOS data. Unfortunately, due to the narrow operational bandwidth of the FOS, analysis of the frequency dependence in the correlation structures and parameters could not be undertaken as has been done in previous studies. This limitation restricts insight into how wall pressure correlates at varying frequency scales from the FOS data due to the natural filtering properties of the structure. Even so, the angular change of wall pressure-induced statistics is still recoverable from the sensor, and provides clear changes in the FOS response to wall pressure fluctuations around the wetted perimeter, yielding comparable results to LES-obtained wall pressure data. These agreements suggest that, despite the restricted bandwidth, the FOS still proves a valuable tool for evaluating wall pressure correlation structure in the streamwise direction.

6.7.2 Frequency-wavenumber spectra

The process for recovering the two frequency-wavenumber spectra of both datasets is given in Equations (6.18) and (6.20). As previously mentioned, $R_{\varepsilon\varepsilon}(\xi_{ij}, \tau)$ recovered from the

raw, unfiltered FOS data was interpolated onto an evenly spaced grid of 50 elements over the FOS length to enhance the spatial resolution of $R_{\varepsilon\varepsilon}$. This was necessary due to the non-uniform spacing of streamwise separations within the FOS (Table 6.3) to increase smoothness of the spectra, not in an attempt to enhance scale resolution. A two-dimensional Hanning window was then applied to the interpolated space-time correlation function prior to the 2D Fourier transform, producing $\Phi_{\varepsilon\varepsilon}(f, k)$. In the case of raw, unfiltered LES data, no interpolation was necessary due to the dense, uniform spacing in the streamwise direction. Similarly, a two-dimensional Hanning window was applied prior to recovery of $\Phi_{pp}(f, k)$ in this case. Details on the wavenumber and frequency resolutions and scales are given in Section 6.5. After computing, the spectra are normalised at each frequency by integrating over the wavenumber axis for both datasets. Normalisation is denoted by a hat symbol throughout the rest of this work. This highlights the relative amplitude change of Φ_{uu} over spatial scales by removing the influence of amplitude changes over different frequencies.

$$\hat{\Phi}_{uu}(\omega, k) = \frac{\Phi_{uu}(\omega, k)}{\sum_{i=-k_s}^{i=k_s} \Phi_{uu}(\omega, k_i) \Delta k}. \quad (6.27)$$




Figure 6.17: *Normalised frequency-wavenumber spectra obtained from raw LES data at the pipe bottom in (a): flow regime 1, (b): flow regime 2. Convective ridge highlighted with dashed line.*

Firstly, the common attribute shared by both plots in Figure 6.18 is the convective relationship $f = U_c k$, denoted by the high energy ridge in each plot. The slope of this in $f - k$ space is associated with the convection velocity U_c of flow near the pipe wall. While the bulk of the energy in the $f - k$ domain is derived from structures propagating at

this convection velocity, the spread of energy from this ridge indicates there are a range of velocities that the pressure fluctuations are convected with in the streamwise direction. The identified parameters from flow regime 1 are: $U_c = 0.595\text{m/s} = 0.85U_b$, and for flow regime 2: $U_c = 0.465\text{m/s} = 0.89U_b$. The values of $0.85 - 0.9U_b$ for the convection velocities are slightly higher than those reported elsewhere in the literature. Common ranges for convection velocities resolved from wall pressure fluctuations in turbulent boundary layers differ between studies. In [82], convection velocities were found to range between $0.59 - 0.69U_\infty$, a range of $0.53 - 0.8U_\infty$ was reported in [201], where U_∞ denotes the free stream velocity in these studies. A number of experimental studies have concluded that U_c is in fact frequency dependent, which is a consequence of eddies of various sizes propagating at different velocities [189, 171, 202, 82, 203]. Here, the slope of the convective ridge does not appear to change directly with frequency. This lack of variation in ridge slope suggests that the structures contributing to the wall signal in this lower frequency range are convecting at relatively uniform velocities U_c for all frequencies. Aliasing of convective energy can be observed mapping high wavenumber energy of $k > 25$ back onto negative wavenumbers. This is a reflection of the streamwise spacing of the LES domain being too coarse to properly resolve high-wavenumber content in the pressure fluctuations, where fluctuations with wavelengths shorter than twice the grid spacing are misrepresented. The onset of aliasing at different frequencies in each case reflects the differences in convection velocity, whereby structures of the same spatial scale travel at different speeds between regimes.





Figure 6.18: *Normalised frequency-wavenumber spectra obtained from raw FOS data at the pipe bottom in (a): flow regime 1, (b): flow regime 2.*

Figure 6.18 displays the normalised frequency-wavenumber spectra for both experimental flow regimes. I primarily expect to observe a convective ridge at positions $f = U_c k$, which appears as a ridge of high energy that corresponds to convection of flow over the FOS. This ridge extends from $f \approx 3$ to $f \approx 8$ and $f \approx 2$ to $f \approx 7$ in flow regimes 1 and 2, respectively. The small range of these ridges is a consequence of the limited bandwidth of the sensor response, as given in Figure 5.3. However, the slight change in the active bandwidth between regimes is supported by the filter processing scheme outlined in Section 6.6.2. In both cases, energy is only contained in the region $|k| < 20$, owing to the finite spatial resolution of the FOS, which is unable to resolve smaller spatial scales (larger k) due to the relatively sparse spacing of individual sensor elements within the FOS. The recovered values of the convection velocity U_c are $U_c = 0.593\text{m/s} = 0.85U_b$, and $U_c = 0.501\text{m/s} = 0.96U_b$ in flow regime 1 and 2, respectively. These values and their scale relative to U_b provide excellent agreement with LES-obtained convection velocities. From these spectra, however, there is no clear indication of a range of advection velocities of the wall pressure propagating over the FOS. I associate this with the physical sensing mechanism, whereby the sensor predominantly responds to pressure fluctuations associated with the dominant convective scale. This is a stark contrast to the LES-obtained spectra, inherently containing pressure fluctuations of varied scales, that consequently convect at a range of velocities. The LES-obtained spectra are essentially an isolated pressure field generated only by flow, free from external influence such as measurement noise and structural vibration of the wall itself, explaining why only the convective ridge is present. As such, these spectra act as a valuable tool for interpreting the FOS-obtained equivalent plots, making for easy identification of artefacts in the experimental data arising from external factors other than wall pressure.

Ridges can be seen in both FOS cases concentrating around $k = 0$ extending from $f = 8$ and $f = 7\text{Hz}$ in flow regime 1 and 2, respectively. A ridge of this type at 0 wavenumber indicates a fluctuating component present among all FBG elements. As this source is present among all elements within the FOS, I believe this may originate from the vibrational response of the pipe structure housing the FOS. Such a response could be induced through a response of the pipe structure to either the wall pressure or other sources of vibration, such as pumps upstream pushing water from the inlet tank. Nonetheless, the fact that this broadband feature is shared between both flow regimes in Figure 6.18 enforces the idea that the ridge is a result of experimental conditions, and not flow.

Similarly, there exist signatures of convection distinct from the convective ridge not seen in the LES data for both flow regimes in Figure 6.18. These manifest as sharp spikes in energy concentrated on a singular (f, k) pair, corresponding to waves travelling with frequency f and wavelength $1/k$ [204]. These appear at $(1.6, -8.75)$ and $(0.35, 4.4)$ in both flow regimes, the former denoting an upstream convecting component of the time series. These features were not identified during the earlier analyses, as this range was excluded from the bandpass region used to filter the data in the bulk of our analysis. In terms of energy content, the relative prominence of these features differs between the two regimes. In flow regime 1, the convective ridge dominates the total energy in the spectrum more so than the isolated signatures. This indicates that the overall convection of flow is the main driving force behind the FOS response in this case. In flow regime 2, the energy associated with the isolated signatures matches that of the convective energy, suggesting that the total energy is spread more evenly around these features, and the FOS response in this case cannot be attributed solely to the convection of bulk flow over the sensor. These signatures are visible at the aforementioned locations for all angular locations of the FOS data, although for brevity, these additional spectra will not be shown. The component at $(1.6, -8.75)$ may correspond to reflections originating from the downstream flow control gate due to the direction of propagation, whereas the low frequency feature at $(0.35, 4.4)$ may correspond to a long wavelength, slowly oscillating mode, indicative of large-scale flow structures. Aside from speculation, it proves difficult to recover the exact nature of these features, due to the sparsity of spatial measurements. However, based on the fact that these features are inherently low frequency, it is plausible they are signatures of large-scale motions, this reasoning will be explored in the following subsection.






Figure 6.19: *Frequency-wavenumber spectra-recovered convection velocities U_c in (a): LES flow regime 1, (b): LES flow regime 2, (c): FOS flow regime 1, (d): FOS flow regime 2.*

Focusing on LES-derived convection velocities, Figure 6.19(a-b) displays clear patterns around the wetted perimeter. For both flow regimes, I observe local minima around the $0, \pm 40^\circ$ positions, with an increase seen past $\pm 40^\circ$ towards the mixed corners in the case of flow regime 1. This is somewhat of an inverted pattern to the RMS variation in Figure 6.16, and aligns with the overall variation in correlation parameters displayed in Figure 6.15. From this, I can conclude that regions of higher wall pressure fluctuations correspond to lower convection velocities, as well as shorter correlation lengths of the filtered time series. This is in line with expectations of secondary flow behaviour, where fluid directed towards the wall at the $0, \pm 40^\circ$ positions develops stronger spanwise components, reducing convection velocity and consequently correlation length. The FOS-obtained convection velocities in Figure 6.19(c-d), however, do not highlight the same azimuthal variation as other FOS-derived statistics, such as those in Figures 6.9, 6.11 and 6.12. Owing to the sparsity of measurement locations in the FOS, and the need for interpolations between unevenly spaced gratings, the resolution of the frequency-wavenumber spectra is very limited. As such, fewer points are available to estimate the gradient of the convective ridge within the $f - k$ domain, as compared to the LES case. This makes the computation much more susceptible to irregularities, obscuring the true estimate. Consequently, the FOS-obtained convection velocity variation should serve as an estimate rather than an absolute profile of its azimuthal variation, owing to the lack of resemblance to the profiles seen in Figures 6.9, 6.11 and 6.12. Even so, the recovered magnitudes still display good agreement with

LES-obtained estimates, providing a method to obtain estimates of convection velocity straight from raw, unfiltered time series data.

6.7.3 Large-scale motions

In partially-filled pipe flows, there are multiple phenomena to take into consideration when assessing the pattern of pressure exerted onto the pipe wall. As mentioned previously, secondary flow motions transfer high-momentum fluid from the free surface towards the wall at particular azimuthal locations. Other coherent structures present in open channel flows are LSMs and VLSMs. These coherent structures have characteristic length scales extending up to multiple pipe radii in length, their presence manifesting as ‘hills’ or ‘shoulders’ in the pre-multiplied spectra of either streamwise velocity or pressure. The position of this hill in the pre-multiplied spectra determines the corresponding size of the structure λ , with LSMs typically reported to lie in the $1 < \lambda/h < 4$ range, and VLSMs occurring in the $12 < \lambda/h < 30$ range, where h is the flow depth [180, 49]. The authors of [190] suggest that the existence of secondary flow motions suppresses VLSMs in partially-filled pipe flows, although they do not hinder the production of LSMs. As such, it is the aim here to recover footprints of these coherent structures within the fluctuating wall pressure field by analysing both FOS and LES datasets, and to assess how their presence varies with angular position around the wetted perimeter.

To obtain information on the spatial scale of structures influencing the wall pressure fluctuations from point measurements, I use the assumption of Taylor’s hypothesis to convert between frequency and wavenumber domain. Power spectral densities from individual FBGs were first recovered, this was converted to wavenumber domain through the transformation $k = 2\pi f/U_c$. In practice, it is common to either utilise local convection velocity U_c , or bulk velocity U_b to convert between frequency and wavenumber domains. Utilising U_b allows for a consistent conversion, but disregards the azimuthal variation in convection velocities around the wetted perimeter. This implies the conversion to wavenumber domain should change in accordance with the local convection velocity around the wetted perimeter, if the goal is to recover accurate spatial scales. Both conversion methods have been tested, yielding similar qualitative trends around the wetted perimeter, as the absolute variation in U_c is small relative to U_b . As such, the results shown here will be from local convection velocities U_c calculated in the previous sections to enunciate azimuthal variation in the selected quantities. This follows the approach taken in previous studies

[179, 205, 206, 2]. From this conversion, the pre-multiplied spectra $k\phi_{\varepsilon\varepsilon}(k)$ can be computed directly from the FBG strain time series $\varepsilon(x, t)$. Plotting the pre-multiplied spectra as a function of spatial length scale λ allows us to see what spatial scales contribute the most energy to the FOS response. Recording the values of λ corresponding to peaks in the pre-multiplied allows for a picture of the dominating spatial scales driving the FOS time series, and the change in these scales around the wetted perimeter.

In this study, a segment length of $T_{\text{seg}} = 3.5\text{s}$ was used when computing the pre-multiplied spectra with Welch’s method [207]. This choice increases ensemble averaging as compared to longer segments and yields smoother, more consistent spectra between angular positions. To justify that the individual segments provide independent realisations of the flow, the segment length is here compared to τ_{eddy} . This represents the turnover time of the largest eddies that cover the whole water depth, referring to the time taken for an eddy to complete one cycle. For the largest eddies, $\tau_{\text{eddy}} \approx h/U_b$. Here, $T_{\text{seg}}/\tau_{\text{eddy}} \approx 17$ and 25 for flow regime 1 and 2, respectively. These ratios demonstrate that the chosen segment length is sufficient for multiple lifetimes of the largest eddies to exist within each segmented time series.

The pre-multiplied spectra at each azimuth were normalised such that the integrated area of the pre-multiplied spectra up to the Nyquist limit equals unity. This transforms each spectrum into a distribution, whereby the relative prominence of each of the spatial scales relative to the total energy in the data at each azimuth can be assessed across azimuths. Peaks in these normalised spectra then provide a heuristic means of quantifying the relative contribution of the corresponding spatial scale to the overall energy contained in the data at that azimuth, independent of total energy. Performing this transform to FOS (and LES) data at each azimuth allows us to track variations of the contribution level of the dominating spatial scales about the wetted perimeter. A comparison between normalised and un-normalised pre-multiplied spectra confirmed that the normalisation procedure does not alter the azimuthal variation or relative ordering of the dominant peaks to be extracted, it simply rescales the spectra to allow for a comparison of relative contribution to total energy recovered at each azimuth.




Figure 6.20: *Normalised pre-multiplied spectra from the FOS pipe bottom position in (a) flow regime 1, (b) flow regime 2.*

The pre-multiplied spectra shown here exhibit more isolated peaks when compared to the typical ‘hill’ or ‘shoulder’ features that are attributed to different spatial scales in the literature. However, in both regimes, three main peaks can be identified in Figure 6.20(a). There exist peaks at $\lambda/h = 0.87, 3.92, 14$, and at $\lambda/h = 1.53, 4.2, 39.5$ in Figure 6.20(b). Note, the value of h used corresponds to the water depth for the particular regime. In both cases, I see clear peaks in the typical range of LSMs, suggesting that these structures leave a footprint on the fluctuating wall pressure-induced strain field despite the influence of secondary flows and the presence of a free surface. Even though it has been reported through studies mentioned previously that the interplay between secondary flows and VLSMs results in the suppression of VLSMs, the peaks seen in the pre-multiplied spectra recovered from the FOS may conclude otherwise.




Figure 6.21: Peak locations in FOS pre-multiplied spectra for all angular positions in (a) flow regime 1, and (b) in flow regime 2.

The spatial scales corresponding to peaks in the pre-multiplied spectra were recovered to identify dominant spatial scales driving the FOS response for different positions around the wetted perimeter. The first two peaks are consistent in their scale for each regime, however the location of peak number 3 ($\lambda/h > 10$) in flow regime 1 exhibits variations of up to three orders of magnitude between angular positions. As such, the observed variability is unlikely to be a true feature of the flow in this case, and may originate from low-frequency noise. In flow regime 2, this third peak remains relatively consistent in the range of ($30 < \lambda/h < 40$), obeying a roughly symmetrical distribution around the wetted perimeter. The authors postulate that the differing behaviour of this third peak between regimes is linked to the effect of inlet conditions from the honeycomb mesh in relation to flow geometry in each case. Even with honeycomb meshes designed to break up large-scale disturbances in the flow, there may still exist significant smaller-scale disturbances generated through high shear experienced when fluid propagates through the mesh. These smaller disturbances can propagate downstream, only settling after sufficient streamwise distance before the disturbances level out [208, 209]. In turbulent flows, larger Reynolds numbers cause honeycomb-induced turbulence intensity to remain elevated for greater streamwise distances as compared to lower Reynolds numbers in the turbulent regime [209]. In flow regime 1, disturbances introduced by the honeycomb persist further downstream of the inlet. In flow regime 2, however, the smaller flow geometry (and Re value) encourages settling of these disturbances over a shorter distance. While at the same physical distance from the inlet, this corresponds to a much longer length of flow

development for smaller flow geometries. Specifically, the FOS is located at $168R_h$ and $288R_h$ for flow regimes 1 and 2, respectively. As a result, VLSM signatures are more coherent in regime 2, yielding consistent spectral peaks, whereas in regime 1, there exists greater variability in the apparent third peak, as structures of this size relative to the flow geometry have not been able to fully develop. For this reason, the first two peaks shall be the focus of our attention going forward. The greatest streamwise separation between individual FBG elements within the sensor is only 0.22m. The FOS length itself ($1.5h - 3h$ depending on the regime) is much smaller than that of the length of pipe preceding the FOS ($85h - 170h$ depending on the regime). As such, the flow can be considered homogeneous over the FOS length. These spatial scales are more consistent between angular positions, with there being two scales concentrated around $\sim 0.85h$, $4h$ present in flow regime 1, and $\sim 1.5h$, $4.3h$ in flow regime 2, well within the reported LSM range.




Figure 6.22: *Spatial scales of identified LSMs in (a-b): flow regime 1, (c-d): flow regime 2.*

In the case of flow regime 2, the two peaks remain at relatively stable values around the wetted perimeter, showing only small fluctuations and exhibiting a roughly symmetrical distribution about the pipe centreline. This contrasts the case seen for flow regime 1, where even though the second peak shows a roughly consistent and symmetric variation around the wetted perimeter (Figure 6.22(b)), there exists a more significant deviation when compared to the first peak (Figure 6.22(a)). Since the consistency of the first peak

in regime 1 is comparable to that of both peaks in flow regime 2 (Figures 6.22(c-d)), I can suggest that it is inlet condition sensitivity contributing to the slightly more pronounced variance in the scale of peak 2 in the case of flow regime 1.




Figure 6.23: *Ratio of normalised pre-multiplied spectra values between peaks 1 and 2 in (a): flow regime 1, and (b): flow regime 2.*

To further investigate the interplay between these two dominating spatial scales in the FOS time series, I look back to the relative peak prominence of these spatial scales in the pre-multiplied spectra as seen in Figure 6.20 as an example. Focusing on the first two peaks, which I will now denote as λ_1, λ_2 , I wish to assess which scale dominates, and how this relative importance changes around the wetted perimeter. To achieve this, I denote P_1 and P_2 as the normalised pre-multiplied spectra values at locations λ_1 and λ_2 , respectively. I compute the ratio P_2/P_1 between the pre-multiplied spectra values at these peak locations to isolate the relative importance of each scale in terms of energy contribution; these ratios as a function of azimuth are given in Figure 6.23. I can see a clear pattern in both flow regimes. Towards the mixed corners, the ratio P_2/P_1 increases sharply, indicating the energy contribution of λ_2 grows relative to that contained in λ_1 . This implies that towards the free surface, larger spatial scales become the main driving force behind the FOS time series. In Figure 6.23(a), it can be seen that the ratio achieves local minima close to the $\pm 20^\circ$ positions, which then increases towards the pipe bottom position, which cannot be seen for flow regime 2. This suggests that in a half-filled pipe flow, the relative importance between the two scales actually flips near the $\pm 20^\circ$ positions, and the relatively smaller scale structures may dominate the wall pressure at these locations. By contrast, near the pipe bottom, both scales contribute with roughly equal prominence in flow regime 2. In Figure

6.23(b), $\lambda_2/\lambda_1 > 1$ for the whole wetted perimeter, suggesting that even near the pipe bottom position, it is the larger scale that dominates. This further reinforces the hypothesis that the flow is more developed in the case of flow regime 2, as larger-scale structures relative to the scale of the flow geometry have been established and clearly dominate in terms of energy contribution. Care should be taken when attempting to compare multi-scale behaviour outlined here to FOS correlation behaviour in Section 6.7.1. The correlation behaviour presented earlier was obtained from bandpass-filtered time series. LSM/VLSM typically manifest at very low frequencies outside of this range, as such it is mostly the behaviour of scale λ_1 driving the filtered time series, and consequently the correlation behaviour presented earlier. Larger scales λ_2 dominate towards the free surface, whereas λ_1 holds less influence, reflecting a decrease in correlation parameters of the filtered time series as the mixed corners are approached, as seen in Figures 6.9 and 6.11.




Figure 6.24: *Normalised pre-multiplied spectra from the LES pipe bottom position in (a): flow regime 1, (b): flow regime 2.*

The raw LES wall pressure fluctuation data, however, only detail fingerprints of LSMs, with clear hills in the pre-multiplied spectra indicative of these scales only; no VLSM signatures were present in these data. Their absence can be explained by the length of the streamwise domain used in the simulations. A length of $10.81r$ was used, which is of the order VLSM scale. Previous studies have demonstrated that domains on larger orders than this are required for VLSMs to develop, owing to their size [179, 180]. As such, their absence in Figure 6.24 is an expected outcome of the simulation setup. Due to the large pipe diameter and high Reynolds number, it is not possible to conduct a longer streamwise length without compromising on acceptable temporal and spatial resolution of the simulation. Even with

this limitation, the changes in LSM scale can be tracked around the wetted perimeter for each case.




Figure 6.25: Peak locations in LES pre-multiplied spectra for all angular positions in (a) flow regime 1, and (b) in flow regime 2.

The absolute scales associated with the LES peaks in Figure 6.25 do not directly match those observed experimentally in Figure 6.22. In particular, energy contribution in the pre-multiplied spectra from LES wall pressure fluctuations comes from the range of $0.4h$, $0.7h$ in flow regimes 1 and 2, respectively. These scales are roughly half the size of the smallest scales identified in the FOS-obtained data, and considerably smaller than scales reported in the literature. This underestimation in spatial scales is somewhat expected given the constraints imposed by the setup of the LES study. The limited streamwise domain length prevents the full development of LSMs, preventing comparison against LES and FOS-obtained spatial scales within the flow. Even with this limitation, the authors still believe that the azimuthal variation of the LES-resolved scales is worth examining, while treating the absolute values of these scales with caution. There is still significant variation in the LSM scale azimuthal distribution in both regimes as seen in Figure 6.25. Although this variation is less pronounced than in the streamwise correlation decay parameters discussed in Section 6.7.1, clear similarities can be drawn between Figures 6.25 and 6.15. Dips can be seen in all quantities at the pipe bottom position, which increase to roughly the $\pm 20^\circ$ positions. In the case of flow regime 1, all quantities then reach a local minima around the $\pm 40^\circ$ positions, the additional maxima seen in Figure 6.15(a) and (c) at $\pm 60^\circ$ are not observed in the LSM scale variation. Past $\pm 20^\circ$ in flow regime 2, the correlation parameters show a much stronger decay towards the mixed corners than the spatial scales in Figure

6.25(b). Even with slight differences between recovered correlation parameters and spatial scales present in the LES pre-multiplied spectra, roughly similar azimuthal variations can be inferred, suggesting that larger-scale structures result in wall pressure fluctuations that correlate over greater distances. This is expected, as larger streamwise structures dominate the wall pressure over greater separations, keeping the wall pressure correlated while the structure passes.

As drawn to in Section 6.7.1, the inverse relationship seen between Figures 6.15 and 6.16 suggests that greater RMS wall pressure results in greater streamwise decorrelation of the wall pressure, and vice versa. In a similar fashion, the extracted maxima of the pre-multiplied spectra can be interpreted as a measure of the energy contribution of the corresponding spatial scale, as the spatial scale at the peak value represents the dominant contributor to the overall energy. Since the RMS is the integral of the power spectrum (Equation (6.16)), this approach yields a quantity derived directly from the pre-multiplied spectra that can be compared with angular variations of the RMS.




Figure 6.26: Peak locations in LES pre-multiplied spectra for all angular positions in (a) flow regime 1, and (b) in flow regime 2.

It is the change in peak heights around the wetted perimeter that is of interest in Figure 6.26, rather than their absolute value. Similarities between Figures 6.16(a) and 6.26(a) show that high RMS aligns with regions where energy in the pre-multiplied spectra is more localised around the identified spatial scale (higher peak value). These show three dominant peaks at the $0, \pm 40^\circ$ positions, indicating that, at least in the case of flow regime 1, regions with larger fluctuations about the mean pressure have a greater proportion of their energy coming from a singular spatial scale at that angle (Figure 6.25). By contrast,

at the minima of $\pm 20^\circ$, the dominant scales are slightly larger (Figure 6.25(a)), but the total energy is distributed over a greater range of scales as determined by the smaller peak, and therefore greater spread in the pre-multiplied spectra. In flow regime 2 (Figure 6.26(b)), while not as pronounced, a comparable trend is observed with maxima near $0, \pm 35^\circ$ positions, and minima at $\pm 20^\circ$. The somewhat inverted patterns between RMS and spatial scales around the wetted perimeter lead us to believe that at positions where secondary flow motions are pointed directly towards the wall, the RMS is greater and the wall pressure is concentrated over a slightly tighter distribution of smaller spatial scales as compared to other locations. Upon approaching the mixed corners, the decreasing maxima correspond to a slight flattening of the pre-multiplied spectra, resulting in energy spreading slightly towards larger spatial scales, as illustrated in the region $\lambda/h > 1.1$ in Figure 6.27.




Figure 6.27: Overlaid LES pre-multiplied spectra from (a): flow regime 1 at angles of $\gamma = 63, 70, 77^\circ$, and (b): flow regime 2 at angles of $\gamma = 28, 35, 42^\circ$.

This azimuthal variation of spatial scales towards the extremities of the wetted perimeter manifests differently between LES and FOS data, as there are different numbers of scales present between datasets. Even so, both display evidence that larger-scale structures hold greater influence, more so towards the mixed corners, as shown experimentally in Figure 6.23, supporting evidence given in [16]. Even without full flow development, and a single spatial scale present in the LES case, Figure 6.27 arrives at this same conclusion, albeit more subtly and relying on shifts in the shape of the pre-multiplied spectra rather than the relative height of the peaks corresponding to spatial scales in the data.

Taken together, the results suggest that the dominant scales of the fluctuating wall pressure, and the corresponding recovered FOS strain, vary around the wetted perimeter, in

conjunction with other statistical measures such as correlation parameters and RMS. It is believed this variation is primarily driven by the organisation of secondary flow motions around the wetted perimeter. The LES further suggests that where fluid is directed normally towards the wall, energy becomes more tightly concentrated into smaller scales. However, this narrowing effect cannot be resolved experimentally from the FOS spectra, limiting direct comparison of the two datasets. Even so, both datasets agree that larger-scale structures become more prominent when the mixed corners are approached. The exact form and shape of the secondary flow motions in both datasets is unknown, however a difference in their arrangement between LES and FOS datasets is plausible, given the differences in flow development and azimuthal variation of the quantities of interest. Even with these differences, the agreement between datasets is stronger in terms of overall trends, and I conclude the results in this section demonstrate the variability of spatial scales around the wetted perimeter and the consequence this variability has on the corresponding fluctuating wall pressure.

6.8 Conclusions

The present study set out to utilise a novel flush-mounted fibre-optic sensor to explore the effect of turbulence on the wall pressure fluctuations in partially-filled pipe flows. The primary goal was to recover, from the strain response of the FBGs within the sensor, statistical descriptors of the wall pressure fluctuations, and changes thereof around the wetted perimeter, for two distinct flow regimes. Trends were validated against complementary LES data. Through a data-driven approach based on the Metropolis-Hastings scheme, narrowband frequency ranges were identified that preserved the expected correlation structure of bulk flow captured by the FOS data to the greatest extent. Correlation decay rates of the FOS response to wall pressure were subsequently identified and plotted as a function of angle around the wetted perimeter. Symmetry was observed about the pipe centreline for all quantities between regimes, with overall azimuthal trends captured well in both experimental and numerical cases. RMS pressure recovered from the LES data was maximised at angular locations corresponding to increased flow directed towards the wall from secondary flow motions. An inverse relationship with correlation decay parameters was observed, indicating that the streamwise coherence of wall pressure is influenced by the organisation of secondary flows.

Spectral analysis via frequency-wavenumber spectra allowed for the identification of all

resolvable scales, as compared to correlation, which inherently averages. Both FOS and LES datasets exhibited a dominant convective ridge corresponding to convection velocity U_c at a range of spatial and temporal scales. Values of U_c recovered from the gradient of the convective ridge showed good agreement between datasets in the range of $0.85 - 0.9U_b$ at the pipe bottom, considerably higher than values reported for turbulent boundary layer flows. Owing to the sparsity of studies on wall pressure fluctuations in partially-filled pipe flows, the frequency-wavenumber spectra presented here are expected to be different from those seen in previous literature due to the additional mechanisms driving the flow in this case. Spectral features recovered from the FOS data, absent in LES, were seen consistently around the whole wetted perimeter, highlighting contributions from additional scales not present in LES. Consistent fingerprints of LSMs were observed in both FOS flow regimes, with dominant wavelengths in the range of $0.8 < \lambda/h < 5$. Consistent VLSM signatures were only observed in flow regime 2; it is believed that disturbances caused by the honeycomb inlet reduced flow development in the case of flow regime 1, suppressing production of VLSMs at the FOS position. Consistent azimuthal patterns of the remaining two scales were consistent between regimes. It was found that energy contribution from larger scales dominated the FOS and LES time series towards the mixed corners, signalling larger structures dominating the wall pressure fluctuations at these positions.

Despite positive outcomes, limitations blurring the agreement between datasets must be recognised. Slight variations in azimuthal distributions between datasets were apparent, even with the same flow conditions. It was hypothesised that the experimental inlet conditions contributed to a less developed flow in the case of flow regime 1. In the LES case, the length of the streamwise domain was insufficient to accommodate the development of multi-scale behaviour, as is expected in partially-filled pipe flows. As it is understood that the driving force behind variability in hydraulic properties around the wetted perimeter is the organisation of secondary flow motions, potential future work could entail combining FOS measurements with S-PIV measurements taken in the spanwise plane at a streamwise position coinciding with the FOS. This would allow for direct correlation between FOS-recovered wall pressure distributions and the corresponding structure of the secondary flow, which may aid in addressing misalignments between LES and FOS results around the wetted perimeter. Further work could extend the LES domain to better capture VLSM behaviour, and a wider range of Reynolds numbers and filling ratios could be explored to match the flow conditions of previous studies. In summary, this study provides evidence that fibre-optic sensors can recover azimuthal variation in wall pressure statistics and fin-

gerprints of internal flow structures present in partially-filled pipes. This work is of great importance in the field of non-invasive sensing, and the results given here may be used to guide sensing methodologies applied to live hydraulic infrastructure.

Acknowledgements

The authors would like to acknowledge the support of the UK Engineering and Physical Sciences Research Council (EPSRC) Impact Acceleration Account (IAA), and through grant EP/S017283/1 on Distributed Fibre-optic Cable Sensing for Buried Pipe Infrastructure. This research made use of the Integrated Civil and Infrastructure Research Centre (ICAIR) which has been part funded by EPSRC, Grant UKCRIC-National Distributed Water Infrastructure Facility (EP/R010420/1) and the European Regional Development Fund (Project Number 28R15P00608). For the purpose of open access, the authors have applied a ‘Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising.

Chapter 7

Conclusions

The aim of this thesis was to develop a non-invasive flush-mounted fibre-optic sensing platform for turbulent partially-filled pipe flows, with the goal of recovering hydraulic conditions and multi-scale behaviour in these flows with strain data collected in response to fluctuating wall pressure at the pipe wall. This included azimuthal variations of the mentioned flow characteristics, in an effort to advance current understanding of partially-filled pipe flow dynamics. A discussion and summary of the outcomes of this thesis shall be given here, including avenues for potential future research.

7.1 Discussion

In Chapter 2, a comprehensive review of the current state of the art in the literature relevant to this thesis was presented. This includes a survey on partially-filled pipe flows and turbulence, leading to the presentation of additional flow structures that emerge in these flows due to the presence of a free surface. Existing methodologies used for measuring varied flow phenomena, particularly in the context of partially-filled pipe flows is given next. An introduction to fibre-optic sensing is given, highlighting current use cases for monitoring hydraulic infrastructure while drawing attention to the lack of non-intrusive applications of such measurement devices. The chapter is finalised with an introduction to inverse problems in engineering, with a focus on the Bayesian perspective of inferring unknown system parameters from measured experimental data. Together, the reviewed topics combine to present the motivation for this thesis: the development of a non-invasive sensing mechanism designed for recovering hydraulic and multi-scale behaviour of turbulent

partially-filled pipe flows. The reviewed literature also identifies a lack of knowledge concerning the spatial organisation of the fluctuating wall pressure in these flows, the link to overarching turbulent flow behaviour, and the lack of viable sensing methodologies capable of recording these phenomena.

In Chapter 3, I establish the core mathematical, statistical and experimental methodologies that are utilised throughout the analysis procedures of this thesis. From the Navier-Stokes equations, an expression for the wall pressure is given. Key statistical quantities of wall pressure used throughout this thesis, such as cross-correlation and PSD, are given next. Taylor’s frozen turbulence hypothesis is then introduced for a generic turbulent quantity, a crucial tool used to relate spatial and temporal information contained in the experimental data. Two Bayesian methods for parameter estimation tasks are introduced to provide information on the probabilistic frameworks utilised throughout this thesis. Finally, experimental methodologies used to set the flow conditions and data collection processes required to obtain the raw FOS data are also outlined. All of the above combine to provide the theoretical bridge between the underlying flow physics and the experimental sensing methodologies, from data collection to uncertainty quantification and physics-based inference of the FOS data. In particular, the statistical descriptors and inference frameworks introduced in this chapter underpin the later recovery of convection velocity, correlation structures, and spatial scales from FOS-recovered indirect strain measurements at the pipe wall.

Chapter 4 introduces a FOS utilised in the initial experimental phase of this thesis that was contained in a layered thin plate-like structure, clamped at a multitude of points along its boundary. A novel scheme coined combinatorial ABC-SMC was developed to resolve uncertainty in the state of the boundary conditions induced by the clamping system. This formally manifests as an inverse problem encompassing inference of a high-dimensional binary parameter. This method proved capable of uncovering boundary condition configurations of the FOS structure based on experimental displacement data of the plate (obtained by an LDV) when the structure was subjected to harmonic excitation at multiple points along the plate top boundary. In conjunction with an FEM model of the experimental procedure, the scheme was capable of uncovering the most probable clamping configurations of the FOS through comparison between FEM-generated and experimental data via custom-engineered metric functions. This foundational work provided sound inference for the problem in question. However, the scheme was unable to resolve boundary condition

uncertainty at locations far away from the excitation source, as the displacement profiles tested only covered a short range. Despite this limitation, the scheme successfully provided a sound calibration of the sensor response, which the authors believe may prove more informative if testing displacement profiles covering a greater length of the structure.

Although the sensor design and installation investigated throughout chapter 4 differs from the form utilised for hydraulic testing in later chapters, the work developed here establishes a necessary statistical-based inference procedure for resolving uncertainty in the sensor response during laboratory testing. In a wider context, the chapter demonstrates how Bayesian methodologies prove as sound tools for parameter estimation tasks in inverse problem settings, even when the parameters to be estimated are not common among the literature. This theme continues throughout the rest of the thesis, whereby turbulent flow properties are recovered indirectly from raw data of wall pressure-induced strain through physics-based inference.

Chapter 5 details the design and implementation of a novel fibre-optic sensor engineered to respond to pressure fluctuations at the interior wall of a turbulent partially-filled pipe. Such a sensor was designed and installed into the pipe wall of an experimental pipe rig, such that the FBG elements within the sensor were aligned in the streamwise direction of flow. Its fundamental operating principle was to recover the axial strain of an embedded fibre-optic cable induced through dynamic plate bending. This plate bending in itself was a response to pressure fluctuations at the wall originating from turbulence. An automated filtering approach based on the Metropolis-Hastings scheme was developed and applied to the FOS data to infer what frequency range retains information on bulk flow velocity to the greatest extent. Inference on these appropriate filter ranges was achieved on data from two separate flow regimes, whereby a custom scoring function was developed by the authors to reward filter ranges yielding tighter distributions of bulk flow velocity estimates produced via TDE estimates from two-point measurements of the FBG time series. The analysis confirmed the FOS functionality by revealing consistent streamwise correlation structures between flow regimes, and identifying signatures of large-scale motions, the spatial scales of which aligned with established literature. Crucially, conclusions were drawn on data obtained after rotating the FOS section to an out-of-water azimuthal position on the pipe wall. These results proved uninformative in the case of velocity and correlation estimates, providing further reasoning to assert that the primary strain response of the FOS was

indeed due to fluctuations in wall pressure, and not other external processes that may contribute to a dynamic response of the FOS structure. Despite identifying some design limitations through outlier analysis, this chapter demonstrated the novel operating principle of the FOS, and its ability to capture turbulent information in a challenging hydraulic environment. The results of this chapter therefore provide evidence that meaningful hydraulic information is non-invasively recoverable solely from the fluctuating wall pressure. This achievement directly tackles one of the main research gaps identified in Chapter 2, namely the lack of sensing methodologies for non-invasively recovering turbulent flow behaviour in partially-filled pipe flows. The agreement between the recovered spatial scales and established scale ranges in the literature for LSM behaviour further support the efficacy of the novel FOS. Additionally, the ability of the sensing platform to distinguish between submerged and out-of-water positions demonstrates that the sensor response is genuinely flow-driven, strengthening confidence in the proposed sensing methodology.

In chapter 6, the rotatability of the FOS enabled data collection at different azimuths about the wetted perimeter was utilised for both flow regimes. This, along with supplementary LES wall pressure fluctuation data for equivalent flow regimes, allowed for a detailed experimental and numerical investigation into the azimuthal variation of wall pressure descriptors in turbulent partially-filled pipe flows. A methodological contribution was a further adaptation of the score-based Metropolis-Hastings scheme introduced in the previous chapter. By shifting attention to the streamwise correlation structure instead of velocity estimates, the scheme is able to identify frequency ranges at each azimuthal location that preserve the FOS-recovered streamwise correlation decay to the greatest extent. Correlation parameters L and τ_b governing the exponential decay of streamwise correlation of the FOS response were found to vary symmetrically about the wetted perimeter in a similar fashion for each regime. Local maxima of these parameters occur at different points on the wetted perimeter, indicating regions of the flow domain where the fluctuating wall pressure remains coherent over greater streamwise distances. In both regimes, the overall trend dictated a sharp decrease in these parameters at angles closer towards the free surface. LES-obtained wall pressure fluctuation data were filtered in the same way as FOS data, with equivalent correlation parameters produced for all azimuthal locations. Recovery of RMS pressure in the LES data displayed an inverse relation to the previously mentioned parameters, hinting that at azimuths with greater RMS pressure results in decorrelation over shorter streamwise distances as compared to azimuths with relatively low RMS. The

organisation of these wall pressure quantities elucidate at what azimuths secondary flow motions direct fluid towards the pipe wall, supporting conclusions from previous available numerical studies. While the LES-derived azimuthal trends differed quantitatively when compared to the FOS-derived equivalents, qualitatively similar variations were seen. Both were roughly symmetric and possessed minima and maxima at roughly similar azimuths. In the FOS data, however, local variations between minima and maxima were much less, which somewhat obscures definitive claims about the precise secondary flow structure. Frequency-wavenumber analysis reveals a dominant convective ridge in both FOS and LES datasets, yielding strong agreement of recovered convection velocity between data types. It was noted that FOS-obtained spectra contained signatures of multiple structures. As such, FOS pre-multiplied spectra identified the growing dominance of LSMs towards the free surface in both flow regimes. While the LES data was shown to contain only one main spatial scale, results agreed that the size of this scale shifted towards greater wavelengths for azimuths closer towards the free surface of the simulation, providing analogous conclusions. Recovered properties of spatial scales within the LES data provide similar azimuthal variations to RMS variations. Despite agreeable results between data types, the limitations of both data types were realised, particularly the FOS's dampened response and the LES's computational limitations induced by periodic boundary conditions, as well as the lack of multi-scale behaviour as seen in the experiment.

The findings of Chapter 6 support and develop upon those of Chapter 5, providing evidence that fluctuating wall pressure contains measurable signatures of the internal turbulent behaviour of partially-filled pipe flows. The recovered azimuthal variation in derived quantities such as correlation structure, RMS pressure, and dominant spatial scales support existing literature on the organisation of secondary flow motions, and the consequences they have on other flow aspects. Specifically, the qualitative agreement between FOS and LES datasets suggests that the proposed sensing methodology is capable of recovering physically meaningful flow characteristics despite the indirect nature of the measurements. At the same time, the quantitative differences between experimental and numerical datasets highlight practical limitations of running computational studies, as well as those involved in data collection from sensors. In particular, damping within the FOS, sensor length, undetermined sources of noise, and computational restrictions imposed by periodic LES domains all influence the recoverable turbulent information. These limitations are important for considerations of future sensor iterations, where environmental or vibrational noise, sensor durability, and installation challenges would likely be more detrimental to

sensor performance in the field, if not accounted for properly.

7.2 Final conclusions

This thesis presents the development of a non-invasive fibre-optic sensor system for the recovery of hydraulic and turbulent flow phenomena through its response to the fluctuating wall pressure in partially-filled pipe flows. Through a combination of experimental design, extensions to Bayesian inference methodologies, and application of signal processing techniques, this work demonstrates that meaningful global flow properties are indeed recoverable from information measured at the interior pipe wall. In doing so, this thesis has addressed a key gap in modern sensing methodologies for live hydraulic infrastructure.

A key contribution of this thesis is to provide new experimental confirmation that recovered flow properties, including bulk flow velocity, streamwise correlation structure, convection velocity, and signatures of large-scale motions, vary azimuthally about the wetted perimeter, advancing previous numerical findings concerning the influence of secondary flow motions. The qualitative agreement between experimental and numerical LES datasets further reinforces the conclusion that the sensing mechanism developed throughout this thesis captures genuine flow-induced excitation data, the results of which transferring between data types and flow regimes.

In addition to hydraulic findings, this thesis also demonstrates the utility of Bayesian inference methodologies in resolving uncertainty associated with fibre-optic sensing systems and indirect measurement problems, particularly in unseen scenarios where custom modifications to parameter estimation problems are required. Across the work presented, physics-based inference methodologies were shown to extract meaningful information from noisy and uncertain experimental FOS datasets. Finally, the practical limitations identified throughout this thesis, including sensor damping, installation constraints, external sources of noise, and numerical simulation limitations, provide important guidance for future development of fibre-optic sensing systems intended for deployment in operational partially-filled pipes.

7.3 Future work

There exist many promising avenues for future research that the work conducted in this thesis could form the base of. These can be categorised into three main areas: experimental sensor design choices, extended hydraulic testing, further data processing and analysis methods applied to the current experimental datasets, based on fluid dynamical principles.

7.3.1 Sensor design choices

Firstly, while consistent spatial and temporal structures were recoverable from analysis of the sensor time series data, it is believed that the physical design of the sensor could be improved. The spacings between individual FBG elements are given in Table 6.3, and their relative positions within the FOS are shown in Figure 5.2(b). FBGs 1 to 4 are located a short distance away from the plate ends in the streamwise orientation. The plate ends were not fixed in place, and as such, typically experience lower strain values than more central points along the plate length. Although this effect was not rigorously quantified throughout this study, this lack of constraint likely results in these FBGs experiencing lower strain transfer from the pressure fluctuations as compared to points closer towards the FOS centre. To address this, it is suggested future iterations be conducted on a FOS of greater streamwise length, with the individual FBG elements located at positions further into the FOS centre, away from influence of free end conditions.

It would also be of interest to utilise a fibre-optic cable possessing a greater number of FBGs. By increasing the number of FBG elements of unique Bragg wavelengths, in conjunction with the considerations outlined in the above paragraph, greater streamwise distances of flow could be covered. This would allow for a better understanding of the temporal evolution of and convection of flow structures over the sensor. For n gratings of non-uniform spacing, there exists $\frac{n(n-1)}{2}$ unique separations. Non-uniform spacings in this extended FBG array would allow a greater number of unique two-point separations to be calculated. This would extend the streamwise correlation profile, and ensure an enhanced resolution of the data in the wavenumber domain, providing a more complete spectral description of the turbulent energy contributions in $f - k$ space.

To ensure measured wavelength shifts are related solely to axial strain rather than temperature fluctuations, a temperature compensation scheme would be desired in future sensor iterations. A method as simple as utilising a reference FBG free from influences of axial

strain has been used to estimate the temperature effects in the functional FBG elements. By subtracting the temperature FBG measurements from the functional FBG measurements, an estimate of Bragg wavelength shift due solely to strain can be recovered [168]. While a similar approach was utilised in preliminary testing of the FOS under near-identical conditions, it was found that wavelength shifts due to temperature were on several orders of magnitude lower than those due to strain, justifying their omission in the present study. However, for future sensor iterations, especially those with FBGs or deployment in environments with larger thermal fluctuations, a temperature compensation scheme would be desirable.

A fundamental avenue for future work involves advanced calibration procedures applied to the FOS structure. Where in this work, the focus was on temporal and spatial changes of FOS response to fluctuating wall pressure, and as such conversion to pressure was not deemed necessary. Future iterations would benefit from a dedicated experimental study assessing the transfer function between the pressure field applied to the plate and the resulting strain recovered by individual gratings. Rigorous assessment of this transfer would enable sound quantification of FOS spatial sensitivity to better understand sensor limitations. Testing of different gel and plate material properties has the potential to alter FOS sensitivity, which may prove useful depending on future FOS use-cases.

7.3.2 Further hydraulic testing

Two distinct flow regimes were considered in this study, given in Table 6.1, which is relatively few compared to other studies recovering flow dynamics of partially-filled pipe flows. A more comprehensive exploration of flow conditions would prove useful in fully characterising changes in turbulent phenomena. Indeed, the authors of [190, 17] considered four flow depths to isolate the influence of filling ratio on flow components. The authors of [19] considered six flow conditions, although these conditions were varied based on Reynolds number, all of which were low relative to the conditions in this thesis. Even so, the broader range of flows tested allowed for a more robust analysis of how the wall shear stress distribution changes with Reynolds number.

A primary goal of this work was to recover azimuthal variations in wall pressure fluctuation-induced FOS statistics and identify the overarching flow structures imprinting this fluctuating pressure onto the pipe wall. As such, it would be of interest to test more flow depths. This would not only allow us to observe relative changes in azimuthal variations in

quantities derived from the wall pressure fluctuations in a greater number of flow depths, but more importantly, it would enable these changes to be directly related back to the changes of the underlying flow structures themselves over varied flow depths.

Out of the few studies available on partially-filled pipe flows, most contain information on the whole flow field, whether that be through numerical simulation, or optical methods [3, 1, 47]. The results presented here are the first directly on the fluctuating wall pressure in partially-filled pipe flows. While qualitative trends on FOS response to wall pressure (and azimuthal variations) were recovered, there exists no velocity distribution data obtained from the same experimental rig to uncover the true nature of secondary currents or LSM/VLSM behaviour for the particular setup. As such, properties of these flow structures are inferred through their imprint onto the pipe wall, and knowledge obtained from the literature. For future iterations of this work, it may be helpful to conduct similar experiments on additional flow conditions while simultaneously recovering additional measurements with S-PIV methods to complement the FOS data. This combined approach could unambiguously combine spanwise velocity distribution data on secondary currents to the wall pressure, moving from inference to a confirmed relation to flow structure behaviour in the given experimental testing rig.

7.3.3 Further analyses of FOS data

While Taylor’s hypothesis has been utilised in this work to equate relations between temporal and spatial scales in the flow, the author believes there are additional analyses that can be conducted through leveraging this assumption. As this hypothesis assumes all scales within the flow travel at the same local velocity U_c , the phase difference φ between two signals obtained by sensors of streamwise separation ξ in the direction of mean flow is given by [123, 124]

$$\varphi(f) = 2\pi\tau^*f = 2\pi f\xi/U_c, \quad (7.1)$$

where τ^* is given in Equation 5.4. Differentiating with respect to frequency, I get

$$\frac{d\varphi}{df} = 2\pi\tau^* = 2\pi\xi/U_c. \quad (7.2)$$

This tells us that under Taylor’s hypothesis, the derivative of phase with respect to frequency obtained from the CSD of two-point measurements in Equation 3.7 will be a con-

stant. I can compute the above quantities to observe at what frequency ranges this is obeyed.




Figure 7.1: *Plots of Equation 7.2 for varied FBG combinations at the pipe bottom position in flow regime 1. Darker lines denote increased streamwise separation. Vertical dashed lines denote approximate frequency limits where Taylor’s hypothesis is satisfied.*

It can be seen there exists a slim frequency range when all two-point measurements have a constant phase derivative, this constant (time lag) increases for greater streamwise separations, as to be expected. The frequency range over which the hypothesis holds is notably wider for smaller sensor separations, extending into higher frequencies. For the greatest separations, this constant section is constrained more tightly to lower frequencies. This is due to the evolution of structures as they travel along the FOS. Structures lose coherence over greater distances, violating the ‘frozen’ assumption for such large spacings. The smaller (higher frequency) structures decay along this distance, whereas only the larger (lower frequency) structures remain coherent over longer distances. This range spans roughly $3.5 \sim 8\text{Hz}$ in the case shown in Figure 7.1, which is roughly the same frequency bounds identified through the MCMC-based filter identification scheme as seen in Algorithm 4. Thus, it is understood that this optimal filter range is based on the maximum streamwise separation of the FBG elements within the sensor, and by use of the analysis in this subsection, larger frequency ranges in which Taylor’s hypothesis is valid may be identified for two-point measurements, depending on the streamwise separation.

As such, further work could quantify how the frequency range on which Taylor’s hypothesis is valid changes with streamwise separation. While the bounds in Figure 7.1 were input manually, data-driven approaches could be explored that are capable of identifying the valid range of Taylor’s hypothesis for a given FBG pair, which may prove efficient when resolving future datasets. This has the potential to uncover additional higher frequency information on smaller scale separations. Building on the foundation of this thesis, the azimuthal variation of this relationship should also be explored to determine if the convective properties of turbulent structures change significantly around the pipe azimuth, and identify wall locations where higher frequency components (smaller spatial scales) dominate.

Another avenue for future work is continued exploration of the multi-scale behaviour of the current FOS data. By separating the pre-multiplied spectra into isolated ‘zones’ where each spatial scale dominates, an estimate can be made of what wavelength range contributes to each spatial scale. This can be visualised by separating the pre-multiplied spectra in Figure 6.20 at points corresponding to local minima. Converting these wavelength ranges back to a frequency range with Taylor’s hypothesis, an estimate of the frequency range at which each spatial scale operates can be recovered. Following the approach of Bailey & Smits [210], the CSD of two-point measurements can be converted back to a filtered version of the space-time correlation function (Equation (3.4)) corresponding to contributions from an individual spatial scale. This is achieved through computing

$$R_{\varepsilon\varepsilon}^{ab}(\xi, \tau) = \int_a^b [C_1(\xi) \cos(2\pi f\tau) + C_2(\xi) \sin(2\pi f\tau)] df. \quad (7.3)$$

Here, a, b correspond to the lower and upper frequency limits of the identified spatial scale I wish to isolate the contribution of in $R_{\varepsilon\varepsilon}$. C_1 and C_2 denote the real and imaginary parts of the cross-spectral density $\Gamma(\xi, 2\pi f)$. This technique effectively isolates the space-time correlation, $R_{\varepsilon\varepsilon}^{ab}$, associated with a specific range of eddy sizes. These filtered correlation functions could subsequently be transformed into the frequency-wavenumber domain. This final step would reveal where dominant contributions to energy for each isolated scale lie in $f - k$ space, allowing for further decomposition of this multi-scale behaviour recovered by the FOS.

Bibliography

- [1] D. Knight and M. Sterling. Boundary shear in circular pipes running partially full. *Journal of Hydraulic Engineering*, 126(4):263–275, 2000.
- [2] K. C. Kim and R. J. Adrian. Very large-scale motion in the outer layer. *Physics of Fluids*, 11(2):417–422, 1999.
- [3] J. F. Brosda. *Numerical Investigation of Secondary Currents in Partially Filled Pipe Flow*. PhD thesis, Technical University of Munich, 2022.
- [4] A. Krynkin, K. V. Horoshenkov, A. Nichols, and S. J. Tait. A non-invasive acoustical method to measure the mean roughness height of the free surface of a turbulent shallow water flow. *Review of Scientific Instruments*, 85(11), 2014.
- [5] N. Hu. Sensor-size-related attenuation correction of wall pressure spectra measurements. *Physics of Fluids*, 34(6), 2022.
- [6] J. Prisutova, A. Krynkin, S. Tait, and K. Horoshenkov. Use of fibre-optic sensors for pipe condition and hydraulics measurements: A review. *CivilEng*, 3(1):85–113, 2022.
- [7] N. J. Yang, T. Li, X. Xiao, Z. T. H. Tse, C. M. Lim, and H. Ren. Chapter 14 - force sensing in compact concentric tube mechanism with optical fibers. In Hongliang Ren, editor, *Flexible Robotics in Medicine*, pages 327–347. Academic Press, 2020.
- [8] Q. Zhao, H. K. Zheng, R. Q. Lv, Y. F. Gu, Y. Zhao, and Y. Yang. Novel integrated optical fiber sensor for temperature, pressure and flow measurement. *Sensors and Actuators, A: Physical*, 280:68–75, 2018.
- [9] W. McMahon, M. Evans, M. Burtwell, and J. Parker. Minimising Street Works Disruption: The Real Costs of Street Works to the Utility Industry and Society. Technical Report UKWIR Report 05/WM/12/8, UKWIR, London, UK, 2006.

- [10] Water and treated water - GOV.uk. <https://www.gov.uk/government/publications/water-and-treated-water/water-and-treated-water>, May 2015.
- [11] J. Bevan. Escaping the jaws of death: Ensuring enough water in 2050. <https://www.gov.uk/government/speeches/escaping-the-jaws-of-death-ensuring-enough-water-in-2050>, 2019.
- [12] D. Carrington. Water companies losing vast amounts through leakage, as drought fears rise. <https://www.theguardian.com/environment/2017/may/11/water-companies-losing-vast-amounts-through-leakage-raising-drought-fears>, 2017.
- [13] Ofwat. Service delivery report 2020-2021. Technical Report November, Ofwat, Birmingham, 2021.
- [14] Water UK. Water companies record lowest leakage levels from pipes, 2020.
- [15] H. C. H. Ng, H. L. F. Cregan, J. M. Dodds, R. J. Poole, and D. J. C. Dennis. Partially filled pipes: Experiments in laminar and turbulent flow. *Journal of Fluid Mechanics*, 848:467–507, 2018.
- [16] H. C. H. Ng, E. Collignon, R. J. Poole, and D. J. C. Dennis. Energetic motions in turbulent partially filled pipe flow. *Physics of Fluids*, 33(2), feb 2021.
- [17] Y. Liu, T. Stoesser, and H. Fang. Impact of turbulence and secondary flow on the water surface in partially filled pipes. *Physics of Fluids*, 34(3), 2022.
- [18] Jiayi Wu. *Free surface dynamics and non-contact radar sensing in a partially filled pipe*. PhD thesis, University of Sheffield, 2023.
- [19] J. Brosda and M. Manhart. Numerical investigation of semifilled-pipe flow. *Journal of Fluid Mechanics*, 932:A25, feb 2022.
- [20] N. Sabri, S. A. Aljunid, M. S. Salim, and S. Fouad. Fiber optic sensors: Short review and applications. *Recent Trends in Physics of Material Science and Technology*, pages 299–311, 2014.
- [21] M. Borecki. Intelligent fiber optic sensor for estimating the concentration of a mixture-design and working principle. *Sensors*, 7(3):384–399, 2007.

- [22] Y. Yang, Z. Wang, T. Chang, M. Yu, J. Chen, G. Zheng, and H.-L. Cui. Seismic observation and analysis based on three-component fiber optic seismometer. *IEEE Access*, 8:1374–1382, 2020.
- [23] T. M. Daley, B. M. Freifeld, J. Ajo-Franklin, S. Dou, R. Pevzner, V. Shulakova, S. Kashikar, D. E. Miller, J. Goetz, J. Henningses, and S. Lueth. Field testing of fiber-optic distributed acoustic sensing (DAS) for subsurface seismic monitoring. *The Leading Edge*, 32(6):699–706, June 2013.
- [24] K. N. Choi, J. C. Juarez, and H. F. Taylor. Distributed fiber optic pressure/seismic sensor for low-cost monitoring of long perimeters. *Ground Sensor Technologies and Applications V*, 5090:134, 2003.
- [25] T. Wu, G. Liu, S. Fu, and F. Xing. Recent progress of fiber-optic sensors for the structural health monitoring of civil infrastructure. *Sensors*, 20(16):1–25, 2020.
- [26] J. M. López-Higuera, L. R. Cobo, A. Q. Incera, and A. Cobo. Fiber optic sensors in structural health monitoring. *Journal of Lightwave Technology*, 29(4):587–608, 2011.
- [27] N. Vahabi, E. Willman, H. Baghsiahi, and D. R. Selviah. Fluid flow velocity measurement in active wells using fiber optic distributed acoustic sensors. *IEEE Sensors Journal*, 20(19):11499–11507, 2020.
- [28] W. Peng, G. R. Pickrell, Z. Huang, J. Xu, D. W. Kim, B. Qi, and A. Wang. Self-compensating fiber optic flow sensor system and its field applications. *Applied Optics*, 43(8):1752–1760, 2004.
- [29] M. P. Lipus, S. Kranz, T. Reinsch, C. Cunow, J. Henningses, and M. Reich. Distributed viscosity and flow velocity measurements using a fiber-optic shear stress sensor. *Sensors and Actuators A: Physical*, 345:113760, 2022.
- [30] S. C. Jain. *Open-Channel Flow*. John Wiley & Sons, 2000.
- [31] V. T. Chow. *Open-Channel Hydraulics*. McGraw-Hill Civil Engineering Series. Blackburn Press, 2009.
- [32] P. K. Kundu, I. M. Cohen, D. R. Dowling, and J. Capecelatro. *Fluid Mechanics*. Elsevier, 2024.

- [33] F. M. White. *Fluid Mechanics*. McGraw-Hill Series in Mechanical Engineering. McGraw-Hill, seventh edition, 1979.
- [34] M. H. Chaudhry. *Open-Channel Flow*. Springer Science+Business Media, LLC, Columbia, second edition, 2008.
- [35] K. B. Ranger and A. M. J. Davis. Steady pressure driven two-phase stratified laminar flow through a pipe. *The Canadian Journal of Chemical Engineering*, 57(6):688–691, 1979.
- [36] J. Guo and R. N. Meroney. Theoretical solution for laminar flow in partially-filled pipes. *Journal of Hydraulic Research*, 51(4):408–416, 2013.
- [37] M. Brocchini and D. H. Peregrine. The dynamics of strong turbulence at free surfaces. Part 1. description. *Journal of Fluid Mechanics*, 449:225–254, 2001.
- [38] F. Muraro, G. Dolcetti, A. Nichols, S. J. Tait, and K. V. Horoshenkov. Free-surface behaviour of shallow turbulent flows. *Journal of Hydraulic Research*, 59(1):1–20, 2021.
- [39] J.-I. Yoon, J. Sung, and M. H. Lee. Velocity profiles and friction coefficients in circular open channels. *Journal of Hydraulic Research*, 50(3):304–311, 2012.
- [40] L. Prandtl. Über die ausgebildete turbulenz. In *Verhandlungen Des II. Internationalen Kongresses Für Technische Mechanik*, pages 62–75, Zurich, Switzerland, 1926.
- [41] I. Nezu. Open-channel flow turbulence and its research prospect in the 21st century. *Journal of Hydraulic Engineering*, 131(4):229–246, 2005.
- [42] Y. Sakai. *Coherent Structures and Secondary Motions in Open Duct Flow*. PhD thesis, Karlsruhe Institute of Technology, 2016.
- [43] I. Nezu and W. Rodi. Experimental study on secondary currents in open channel flow. In *21st International Association for Hydraulic Research*, volume 7, pages 62–63, Melbourne, Australia, 1985.
- [44] Y. Liu, T. Stoesser, and H. Fang. Effect of secondary currents on the flow and turbulence in partially filled pipes. *Journal of Fluid Mechanics*, 938:1–35, 2022.

- [45] A. Bose, D. J. Borman, T. N. Hunter, J. T. Spencer, and C. J. Cunliffe. Numerical investigation of fill height and secondary currents in an inclined partially filled pipe flow. In *Journal of Physics: Conference Series*, volume 2899, page 012017, 2024.
- [46] S. P. Clark and N. Kehler. Turbulent flow characteristics in circular corrugated culverts at mild slopes. *Journal of Hydraulic Research*, 49(5):676–684, 2011.
- [47] J. E. Berlamont, K. Trouw, and G. Luyckx. Shear stress distribution in partially filled pipes. *Journal of Hydraulic Engineering*, 129(9):697–705, 2003.
- [48] I. Nezu and H. Nakagawa. *Turbulence in Open Channel Flows*. Turbulence in Open Channel Flows. CRC Press, 1993.
- [49] H. C.-H. Ng, E. Collignon, R. J. Poole, and D. J. C. Dennis. Energetic motions in turbulent partially filled pipe flow. *Physics of Fluids*, 33(2):025101, 02 2021.
- [50] B. Gibeau and S. Ghaemi. Low- and mid-frequency wall-pressure sources in a turbulent boundary layer. *Journal of Fluid Mechanics*, 918:1–45, 2021.
- [51] G. M. Corcos and Berkeley. Institute of Engineering Research University of California. *Pressure Fluctuations in Shear Flows*. IER reports. University of California, Institute of Engineering Research, 1962.
- [52] B. M. Efimtsov. Similarity criteria for the spectra of wall pressure fluctuations in a turbulent boundary layer. *Soviet Physics Acoustics*, 30:33–35, February 1984.
- [53] A. V. Smol'yakov. Calculation of the spectra of pseudosound wall-pressure fluctuations in turbulent boundary layers. *Acoustical Physics*, 46(3):342–347, May 2000.
- [54] M. Goody. Empirical spectral model of surface pressure fluctuations. *AIAA Journal*, 42(9):1788–1794, 2004.
- [55] Reg Herschy. The velocity-area method. *Flow Measurement and Instrumentation*, 4(1):7–10, 1993. Special Issue - Open Channel Measurements.
- [56] G. Gomit, L. Chatellier, and L. David. Free-surface flow measurements by non-intrusive methods: A survey. *Experiments in Fluids*, 63(6):1–25, 2022.
- [57] G. Dolcetti, K. V. Horoshenkov, A. Krynkin, and S. J. Tait. Frequency-wavenumber spectrum of the free surface of shallow turbulent flows over a rough boundary. *Physics of Fluids*, 28(10), 2016.

- [58] K. V. Horoshenkov, A. Nichols, S. J. Tait, and G. A. Maximov. The pattern of surface waves in a shallow free surface flow. *Journal of Geophysical Research: Earth Surface*, 118(3):1864–1876, 2013.
- [59] A. Nichols. Free surface dynamics in shallow turbulent flows. *Ph.D. thesis*, pages 1–266, 2013.
- [60] Andrew Nichols, Simon J. Tait, Kirill V. Horoshenkov, and Simon J. Shepherd. A model of the free surface dynamics of shallow turbulent flows. *Journal of Hydraulic Research*, 54(5):516–526, 2016.
- [61] A. Nichols, S. Tait, K. Horoshenkov, and S. Shepherd. A non-invasive airborne wave monitor. *Flow Measurement and Instrumentation*, 34:118–126, 2013.
- [62] F. G. Bass and I. M. Fuks. Wave scattering from statistically rough surfaces, 1979.
- [63] M.-D. Johnson, J. Cuenca, T. Lähivaara, G. Dolcetti, M. Alkmim, L. De Ryck, and A. Krynkina. Bayesian reconstruction of surface shape from phaseless scattered acoustic data. *The Journal of the Acoustical Society of America*, 156(6):4024–4036, 2024.
- [64] M.-D. Johnson, F. Muraro, S. J. Tait, G. Dolcetti, and A. Krynkina. Recovering the free-surface dynamics of shallow turbulent flows using digital image correlation (dic). *Journal of Hydraulic Research*, 63(5):578–593, 2025.
- [65] Jiayi Wu, Andrew Nichols, Anton Krynkina, and Martin Croft. Digital Image Correlation for stereoscopic measurement of water surface dynamics in a partially filled pipe. *Acta Geophysica*, 70(5):2451–2467, 2022.
- [66] Andrew Nichols, Matteo Rubinato, Yun Hang Cho, and Jiayi Wu. Optimal use of titanium dioxide colourant to enable water surfaces to be measured by kinect sensors. *Sensors (Switzerland)*, 20(12):1–17, 2020.
- [67] Shiyao Wang, Jesus Leonardo Corredor Garcia, Jonathan Davidson, and Andrew Nichols. Conductance-based interface detection for multiphase pipe flow. *Sensors (Switzerland)*, 20(20):1–18, 2020.
- [68] M. Raffel, C. E. Willert, F. Scarano, C. J. Kähler, S. T. Wereley, and J. Kompenhans. *Particle Image Velocimetry*. Springer International Publishing, Cham, 2018.

- [69] R.J. Adrian and J. Westerweel. *Particle Image Velocimetry*. Cambridge Aerospace Series. Cambridge University Press, 2011.
- [70] R. J. Adrian. Multi-point optical measurements of simultaneous vectors in unsteady flow-a review. *International Journal of Heat and Fluid Flow*, 7(2):127–145, 1986.
- [71] R. Meynart. Instantaneous velocity field measurements in unsteady gas flow by speckle velocimetry. *Applied Optics*, 22(4):535, 1983.
- [72] P. Jacquot and P. K. Rastogi. Influence of out-of-plane deformation and its elimination in white light speckle photography. *Optics and Lasers in Engineering*, 2(1):33–55, 1981.
- [73] P. Arroyot and C. A. Greated. Stereoscopic particle image velocimetry. *Measurement Science and Technology*, 2:1181–1186, 1991.
- [74] A. K. Prasad. Stereoscopic particle image velocimetry. *Experiments in Fluids*, 29(2):103–116, 2000.
- [75] G. E. Elsinga, F. Scarano, B. Wieneke, and B. W. Van Oudheusden. Tomographic particle image velocimetry. *Experiments in Fluids*, 41(6):933–947, 2006.
- [76] N. A. Malik, T. Dracos, and D. A. Papantoniou. Particle tracking velocimetry in three-dimensional flows. *Experiments in Fluids*, 15-15(4-5):279–294, 1993.
- [77] V. Weitbrecht, G. Kühn, and G. H. Jirka. Large scale piv-measurements at the surface of shallow water flows. *Flow Measurement and Instrumentation*, 13(5-6):237–245, 2002.
- [78] A. S. Thompson, D. Maynes, and J. D. Blotter. Internal turbulent flow induced pipe vibrations with and without baffle plates. In *American Society of Mechanical Engineers, Fluids Engineering Division (Publication) FEDSM*, volume 3, pages 649–659, 2010.
- [79] P. Martinez-Vazquez, S. Sharifi, and S. Soroosh. Modelling boundary shear stress distribution in open channels using a face recognition technique. *Journal of Hydroinformatics*, 19(2):10, 2016.
- [80] W. K. Blake. Turbulent boundary-layer wall-pressure fluctuations on smooth and rough walls. *Journal of Fluid Mechanics*, 44:637–660, 2070.

- [81] K. Selvam, E. Öngüner, J. Peixinho, E. S. Zanoun, and C. Egbers. Wall pressure in developing turbulent pipe flows. *Journal of Fluids Engineering, Transactions of the ASME*, 140(8):1–7, 2018.
- [82] J. M. Clinch. Measurements of the wall pressure field at the surface of a smooth-walled pipe containing turbulent water flow. *Journal of Sound and Vibration*, 9(3):398–419, 1969.
- [83] J. M. Clinch. Miniature transducer assembly for measuring the properties of the wall-pressure field in turbulent flows. *Journal of the Acoustical Society of America*, 40(1):254–255, 1966.
- [84] B. M. Abraham and W. L. Keith. Direct measurements of turbulent boundary layer wall pressure wavenumber-frequency spectra. *Journal of Fluids Engineering*, 1996-V(March):177–187, 1996.
- [85] R. M. Lueptow. Transducer resolution and the turbulent wall pressure spectrum. *The Journal of the Acoustical Society of America*, 97(1):370–378, jan 1995.
- [86] W. K. Blake. *Mechanics of Flow-Induced Vibration*. Elsevier Inc., London, second edition, 2017.
- [87] G. Schewe. On the structure and resolution of wall-pressure fluctuations associated with turbulent boundary-layer flow. *Journal of Fluid Mechanics*, 134:311–328, 1983.
- [88] J. Huang, L. Duan, K. M. Casper, R. M. Wagnild, and N. P. Bitter. Transducer resolution effect on pressure fluctuations beneath hypersonic turbulent boundary layers. *AIAA Journal*, 62(3):882–895, 2024.
- [89] E. Hecht. *Optics*. Pearson, 2017.
- [90] K. T. V. Grattan and T. Sun. Fiber optic sensor technology: An overview. *Sensors and Actuators, A: Physical*, 82(1):40–61, 2000.
- [91] A. Rogers. Distributed optical-fibre sensing. *Measurement Science and Technology*, 10(8):R75, 1999.
- [92] E. Udd. Overview of fiber optic sensors. *Fiber Optic Sensors, Second Edition*, 4030:1–34, 2017.

- [93] T. G. Giallorenzi. *Fiber Optic Sensors*, volume 2. Marcel Dekker Inc., 1980.
- [94] A. Masoudi and T. P. Newson. Contributed review: Distributed optical fibre dynamic strain sensing. *Review of Scientific Instruments*, 87(1), 2016.
- [95] M. K. Barnoski and S. M. Jensen. Fiber waveguides: A novel technique for investigating attenuation characteristics. *Applied Optics*, 15(9):2112–2115, 1976.
- [96] W. Eickhoff and R. Ulrich. Optical frequency domain reflectometry in single-mode fiber. *Applied Physics Letters*, 39(9):693–695, 1981.
- [97] G. Bolognini and A. Hartog. Raman-based fibre sensors: Trends and applications. *Optical Fiber Technology*, 19(6, Part B):678–688, 2013. Optical Fiber Sensors.
- [98] Dao Yuan Tan, Zhen Yu Tang, Zhen Rui Yan, Jing Wang, Wei Zhang, Jing Wu Huang, Peng Wang, Zhiguo Yuan, Huan Feng Duan, Bin Shi, and Hong Hu Zhu. Real-time monitoring of water states in large-diameter aqueducts – learning from distributed acoustic sensing signals. *Communications Engineering*, 4(1), 2025.
- [99] K. O. Hill, Y. Fujii, D. C. Johnson, and B. S. Kawasaki. Photosensitivity in optical fiber waveguides: Application to reflection filter fabrication. *Applied Physics Letters*, 32(10):647–649, 1978.
- [100] M. M. Werneck, R. C. S. B. Allil, B. A. Ribeiro, and F. V. B. de Nazaré. A guide to fiber bragg grating sensors. In Christian Cuadrado-Laborde, editor, *Current Trends in Short- and Long-period Fiber Gratings*, chapter 1. IntechOpen, London, 2013.
- [101] M. Majumder, T. K. Gangopadhyay, A. K. Chakraborty, K. Dasgupta, and D. K. Bhattacharya. Fibre Bragg gratings in structural health monitoring-Present status and applications. *Sensors and Actuators, A: Physical*, 147(1):150–164, sep 2008.
- [102] B. Glisic and D. Inaudi. *Fibre Optic Methods for Structural Health Monitoring*. John Wiley & Sons, 2007.
- [103] M. F. Bado and J. R. Casas. A review of recent distributed optical fiber sensors applications for civil engineering structural health monitoring. *Sensors*, 21(5):1818, 2021.
- [104] T. Jiang, L. Ren, Z. Jia, D. Li, and H. Li. Application of fbg based sensor in pipeline safety monitoring. *Applied Sciences*, 7(6), 2017.

- [105] L. Pereira, I. Sousa, E. Mesquita, A. Cabral, N. Alberto, C. Diaz, H. Varum, and P. Antunes. Fbg-based accelerometer for buried pipeline natural frequency monitoring and corrosion detection. *Buildings*, 14(2):456, 2024.
- [106] S. Takashima, H. Asanuma, and H. Niitsuma. A water flowmeter using dual fiber bragg grating sensors and cross-correlation technique. *Sensors and Actuators, A: Physical*, 116(1):66–74, 2004.
- [107] J. Huang, D. T. Pham, C. Ji, Z. Wang, and Z. Zhou. Multi-parameter dynamical measuring system using fibre bragg grating sensors for industrial hydraulic piping. *Measurement*, 134:226–235, 2019.
- [108] D. Creighton, P. P. Kirwan, T. P. O’Brien, C. Costello, and K. W. Moloney. Temperature compensated toroidal centripetal flowmeter. *IEEE Sensors Journal*, 16(18):6874–6878, 2016.
- [109] S. A. Sisson and Y. Fan. Likelihood-free markov chain monte carlo, 2010.
- [110] A. Tarantola. *Inverse Problem Theory*. Society for Industrial and Applied Mathematics, 2004.
- [111] J. Adler and O. Öktem. Deep bayesian inversion. In *Data-Driven Models in Inverse Problems*, number section 2, pages 359–412. De Gruyter, 2024.
- [112] C. Andrieu, N. De Freitas, A. Doucet, and M. I. Jordan. An introduction to mcmc for machine learning. *Machine Learning*, 50(1):5–43, 2003.
- [113] H. Ait Saadi, F. Ykhlef, and A. Guessoum. Mcmc for parameters estimation by bayesian approach. In *Eighth International Multi-Conference on Systems, Signals & Devices*, pages 1–6, 2011.
- [114] G. I. Valderrama-Bahamóndez and H. Fröhlich. Mcmc techniques for parameter estimation of ode based models in systems biology. *Frontiers in Applied Mathematics and Statistics*, 5(November):1–10, 2019.
- [115] J. S. Teixeira, L. T. Stutz, D. C. Knupp, and A. J. Silva Neto. A new adaptive approach of the metropolis-hastings algorithm applied to structural damage identification using time domain data. *Applied Mathematical Modelling*, 82:587–606, 2020.

- [116] M. A. Beaumont, W. Zhang, and D. J. Balding. Approximate bayesian computation in population genetics. *Genetics*, 162(4):2025–2035, 2002.
- [117] K. Csilléry, M. G. B. Blum, O. E. Gaggiotti, and O. François. Approximate bayesian computation (abc) in practice. *Trends in Ecology and Evolution*, 25(7):410–418, 2010.
- [118] S. A. Sisson, Y. Fan, and M. M. Tanaka. Sequential monte carlo without likelihoods. *Proceedings of the National Academy of Sciences of the United States of America*, 104(6):1760–1765, 2007.
- [119] S. Filippi, C. P. Barnes, J. Cornebise, and M. P. H. Stumpf. On optimality of kernels for approximate bayesian computation using sequential monte carlo. *Statistical Applications in Genetics and Molecular Biology*, 12(1):87–107, 2013.
- [120] Y. Deng, Y. Pei, C. Li, and B. Zhu. Model selection and parameter estimation for an improved approximate bayesian computation sequential monte carlo algorithm. *Discrete Dynamics in Nature and Society*, 2022, 2022.
- [121] K. N. Dinh, Z. Xiang, Z. Liu, and S. Tavaré. Approximate bayesian computation sequential monte carlo via random forests. *Statistics and Computing*, 35(6), Oct 2025.
- [122] G. I. Taylor. The spectrum of turbulence. *Proceedings of the Royal Society of London. Series A - Mathematical and Physical Sciences*, 164(919):476–490, 1938.
- [123] A. Cenedese, G. P. Romano, and F. Di Felice. Experimental testing of taylor’s hypothesis by l.d.a. in highly turbulent flow. *Experiments in Fluids*, 11(6):351–358, 1991.
- [124] D. Schlipf, D. Trabucchi, and O. Bischoff. Testing of Frozen Turbulence Hypothesis for Wind Turbine Applications with a Scanning LIDAR System. In *10th Symposium for the International Society of Acoustic Remote Sensing of the Atmosphere and Oceans*, volume 12, page 5410, 2010.
- [125] P. MOIN. Revisiting taylor’s hypothesis. *Journal of Fluid Mechanics*, 640:1–4, 2009.
- [126] J. C. Del Álamo and J. Jiménez. Estimation of turbulent convection velocities and corrections to taylor’s approximation. *Journal of Fluid Mechanics*, 640:5–26, 2009.

- [127] M. Guala, S. E. Hommema, and R. J. Adrian. Large-scale and very-large-scale motions in turbulent pipe flow. *Journal of Fluid Mechanics*, 554:521–542, 2006.
- [128] J. M. Marin, P. Pudlo, C. P. Robert, and R. J. Ryder. Approximate bayesian computational methods. *Statistics and Computing*, 22(6):1167–1180, 2012.
- [129] J. L. Beck and L. S. Katafygiotis. Updating models and their uncertainties. i: Bayesian statistical framework. *Journal of Engineering Mechanics*, 124(4):455–461, 1998.
- [130] K.-V. Yuen. *Bayesian methods for structural dynamics and civil engineering*. John Wiley & Sons, 2010.
- [131] S. H. Cheung and J. L. Beck. Bayesian model updating using hybrid monte carlo simulation with application to structural dynamic models with many uncertain parameters. *Journal of Engineering Mechanics*, 135(4):243–255, 2009.
- [132] J. Ching and Y.-C. Chen. Transitional markov chain monte carlo method for bayesian model updating, model class selection, and model averaging. *Journal of Engineering Mechanics*, 133(7):816–832, 2007.
- [133] A. Lye, A. Cicirello, and E. Patelli. Sampling methods for solving bayesian model updating problems: A tutorial. *Mechanical Systems and Signal Processing*, 159:107760, 2021.
- [134] S.-K. Au and J. L. Beck. Importance sampling in high dimensions. *Structural Safety*, 25(2):139–163, 2003.
- [135] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *Journal of Machine Learning Research*, 3(Jan):993–1022, 2003.
- [136] Y. Teh, D. Newman, and M. Welling. A collapsed variational bayesian inference algorithm for latent dirichlet allocation. *Advances in Neural Information Processing Systems*, 19, 2006.
- [137] Z. Ghahramani and G. E. Hinton. Variational learning for switching state-space models. *Neural Computation*, 12(4):831–864, 2000.

- [138] A. Tran, M. Tran, and Y. Wang. Constrained mixed-integer Gaussian mixture Bayesian optimization and its applications in designing fractal and auxetic meta-materials. *Structural and Multidisciplinary Optimization*, 59(6):2131–2154, jun 2019.
- [139] O. Cappé, A. Guillin, J.-M. Marin, and C. P. Robert. Population monte carlo. *Journal of Computational and Graphical Statistics*, 13(4):907–929, 2004.
- [140] A. Gelman, J. B. Carin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin. *Bayesian Data Analysis*. Chapman and Hall/CRC, 2013.
- [141] D. B. Rubin. Bayesianly Justifiable and Relevant Frequency Calculations for the Applied Statistician. *The Annals of Statistics*, 12(4):1151–1172, 1984.
- [142] P. Marjoram, J. Molitor, V. Plagnol, and S. Tavaré. Markov chain monte carlo without likelihoods. *PNAS December*, 23(26):15324–15328, 2003.
- [143] S. A. Sisson, Y. Fan, and M. M. Tanaka. Sequential Monte Carlo without likelihoods. *Proceedings of the National Academy of Sciences*, 104(6):1760–1765, 2007.
- [144] R. D. Wilkinson. *Bayesian Inference of Primate Divergence Times*. PhD thesis, University of Cambridge, 2007.
- [145] E. O. Buzbas and N. A. Rosenberg. Aabc: Approximate approximate bayesian computation for inference in population-genetic models. *Theoretical Population Biology*, 99:31–42, 2015.
- [146] A. Ben Abdesslem, N. Dervilis, D. Wagg, and K. Worden. Model selection and parameter estimation in structural dynamics using approximate bayesian computation. *Mechanical Systems and Signal Processing*, 99:306–325, 2018.
- [147] T. Toni, D. Welch, N. Strelkowa, A. Ipsen, and M. P. H. Stumpf. Approximate bayesian computation scheme for parameter inference and model selection in dynamical systems. *Journal of The Royal Society Interface*, 6(31):187–202, 2009.
- [148] P. J. Wangersky. Lotka-volterra population models. *Annual Review of Ecology and Systematics*, 9:189–218, 1978.
- [149] M. A. Beaumont. Approximate bayesian computation. *Annual Review of Statistics and Its Application*, 6:379–403, 2019.

- [150] J. K. Pritchard, M. T. Seielstad, A. Perez-Lezaun, and M. W. Feldman. Population Growth of Human Y Chromosomes: A Study of Y Chromosome Microsatellites. *Molecular Biology and Evolution*, 16(12):1791–1798, 12 1999.
- [151] A. Doucet, N. Freitas, and N. Gordon. Sequential monte carlo methods in practice. *Sequential Monte Carlo Methods in Practice*, 2001.
- [152] M. A. Beaumont, J.-M. Cornuet, J.-M. Marin, and C. P. Robert. Adaptive approximate bayesian computation. *Biometrika*, 96(4):983–990, 2009.
- [153] N. Chopin. A sequential particle filter method for static models. *Biometrika*, 89(3):539–552, 2002.
- [154] N. Chopin, O. Papaspiliopoulos, et al. *An Introduction to Sequential Monte Carlo*. Springer, 2020.
- [155] Schalk Willem Jacobsz and Sebastian Ingo Jahnke. Leak detection on water pipelines in unsaturated ground by discrete fibre optic sensing. *Structural Health Monitoring*, 19(4):1219–1236, 2020.
- [156] Comsol multiphysics[®] v. 6.2. www.comsol.com. COMSOL AB, stockholm, sweden, 2023.
- [157] D. Jeanbourquin, D. Sage, L. Nguyen, B. Schaeli, S. Kayal, D. A. Barry, and L. Rossi. Flow measurements in sewers based on image analysis: Automatic flow velocity algorithm. *Water Science and Technology*, 64(5):1108–1114, 2011.
- [158] M. K. Bull. Wall-pressure fluctuations beneath turbulent boundary layers: Some reflections on forty years of research. *Journal of Sound and Vibration*, 190(3):299–315, 1996.
- [159] D. M. Chase. Modeling the wavevector-frequency spectrum of turbulent boundary layer wall pressure. *Journal of Sound and Vibration*, 70(1):29–67, 1980.
- [160] D. A. Jackson and J. D. C. Jones. Fibre optic sensors. *Optica Acta*, 33(12):1469–1503, 1986.
- [161] J. Prisutova, A. Krynkin, S. Tait, and K. Horoshenkov. Use of fibre-optic sensors for pipe condition and hydraulics measurements: A review. *CivilEng*, 3(1):85–113, 2022.

- [162] N. S. Nosseir and B. Dagan. Tests of a fiber-optic velocity transducer. *Experiments in Fluids*, 3(4):239–243, 1985.
- [163] M. Bertulesi, D. F. Bignami, I. Boschini, M. Longoni, G. Menduni, and J. Morosi. Experimental Investigations of Distributed Fiber Optic Sensors for Water Pipeline Monitoring. *Sensors*, 23(13):1–18, 2023.
- [164] L. Wong, R. Deo, S. Rathnayaka, B. Shannon, C. Zhang, W. K. Chiu, J. Kodikara, and H. Widyastuti. Leak detection in water pipes using submersible optical optic-based pressure sensor. *Sensors (Switzerland)*, 18(12):1–17, 2018.
- [165] K. O. Hill and G. Meltz. Fiber Bragg grating technology fundamentals and overview. *Journal of Lightwave Technology*, 15(8):1263–1276, 1997.
- [166] C. E. Campanella, A. Cuccovillo, C. Campanella, A. Yurt, and V. M. N. Passaro. Fibre Bragg Grating based strain sensors: Review of technology and applications. *Sensors (Switzerland)*, 18(9), 2018.
- [167] A. Othonos, K. Kalli, D. Pureur, and A. Mugnier. Fibre Bragg gratings. *Springer Series in Optical Sciences*, 123:189–269, 2006.
- [168] Y. J. Rao. In-fibre Bragg grating sensors. *Measurement Science and Technology*, 8(4):355–375, 1997.
- [169] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953.
- [170] M A Beaumont, W Y Zhang, and D J Balding. Approximate bayesian computation in population genetics. *Genetics*, 162(4):2025–2035, 2002. ID number: ISI:000180502300043.
- [171] E. Salze, C. Bailly, O. Marsden, and D. Juvé. Investigation of the wall pressure wavenumber-frequency spectrum beneath a turbulent boundary layer with pressure gradient. In *International Symposium On Turbulence and Shear Flow Phenomena (TSFP-9)*, Melbourne, Australia, June 2015.
- [172] D. Palumbo. Determining correlation and coherence lengths in turbulent boundary layer flight data. *Journal of Sound and Vibration*, 331(16):3721–3737, 2012.

- [173] N. Lakshmanan, S. Arunachalam, S. Selvi Rajan, G. Ramesh Babu, and J. Shanmugasundaram. Correlations of aerodynamic pressures for prediction of across wind response of structures. *Journal of Wind Engineering and Industrial Aerodynamics*, 90(8):941–960, 2002.
- [174] L. Dechant, J. Smith, and M. Barone. Approximate analytical models for turbulent boundary layer wall pressure and wall shear fluctuation spectra and coherence functions. In *AIAA SciTech Forum - 55th AIAA Aerospace Sciences Meeting*, pages 1–14, 2017.
- [175] R. E. Falco. Coherent motions in the outer region of turbulent boundary layers. *The Physics of Fluids*, 20(10):S124–S132, 10 1977.
- [176] A. B. Shvidchenko and G. Pender. Macroturbulent structure of open-channel flow over gravel beds. *Water Resources Research*, 37(3):709–719, 2001.
- [177] A. G. Roy, T. Buffin-Bélanger, H. Lamarre, and A. D. Kirkbride. Size, shape and dynamics of large-scale turbulent flow structures in a gravel-bed river. *Journal of Fluid Mechanics*, 500:1–27, 2004.
- [178] Y. Duan, Q. Chen, D. Li, and Q. Zhong. Contributions of very large-scale motions to turbulence statistics in open channel flows. *Journal of Fluid Mechanics*, 892:A3, jun 2020.
- [179] B. Cerino, A. Zampiron, S. Proust, C. Berni, and V. Nikora. Large and very large scale motions in open-channel flows over rough bed in the presence and absence of a mixing layer. *Environmental Fluid Mechanics*, 24(6):1421–1442, dec 2024.
- [180] A. Zampiron, B. Cerino, C. Berni, S. Proust, and V. Nikora. Very-large-scale motions in open-channel flow: Insights from velocity spectra, correlations, and structure functions. *Physics of Fluids*, 36(4), apr 2024.
- [181] K. C. Kim and R. J. Adrian. Very large-scale motion in the outer layer. *Physics of Fluids*, 11(2):417–422, feb 1999.
- [182] A. E. Perry and C. J. Abell. Scaling laws for pipe-flow turbulence. *Journal of Fluid Mechanics*, 67(2):257–271, jan 1975.

- [183] K. J. Bullock, R. E. Cooper, and F. H. Abernathy. Structural similarity in radial correlations and spectra of longitudinal velocity fluctuations in pipe flow. *Journal of Fluid Mechanics*, 88(3):585–608, oct 1978.
- [184] Sébastien Debert, Marc Pachebat, Vincent Valeau, and Yves Gervais. Ensemble-Empirical-Mode-Decomposition method for instantaneous spatial-multi-scale decomposition of wall-pressure fluctuations under a turbulent flow. *Experiments in Fluids*, 50(2):339–350, 2011.
- [185] Hassan Berro and Pierre Moussou. Vibration spectrum of a water pipe under the effect of unsteady fluid loading. *American Society of Mechanical Engineers, Pressure Vessels and Piping Division (Publication) PVP*, 4(July 2015), 2015.
- [186] P. D. Lysak. Modeling the Wall Pressure Spectrum in Turbulent Pipe Flows. *Journal of Fluids Engineering*, 128(2):216–222, mar 2006.
- [187] D. M. Chase. The character of the turbulent wall pressure spectrum at subconvective wavenumbers and a suggested comprehensive model. *Journal of Sound and Vibration*, 112(1):125–147, 1987.
- [188] K. Zhao, Y. Li, R. Pei, H. Li, and G. J. Bennett. Review of wavenumber-frequency spectrum models of turbulent boundary-layer wall pressure fluctuations. *AIAA Journal*, 63(3):1140–1166, 2025.
- [189] W. W. Willmarth and C. E. Wooldridge. Measurements of the fluctuating pressure at the wall beneath a thick turbulent boundary layer. *Journal of Fluid Mechanics*, 14(2):187–210, oct 1962.
- [190] Y. Liu, T. Stoesser, and H. Fang. Effect of secondary currents on the flow and turbulence in partially filled pipes. *Journal of Fluid Mechanics*, 938:A16, 2022.
- [191] M. Uhlmann. An immersed boundary method with direct forcing for the simulation of particulate flows. *Journal of Computational Physics*, 209(2):448–476, 2005.
- [192] F. Nicoud and F. Ducros. Subgrid-scale stress modelling based on the square of the velocity gradient tensor. *Flow, Turbulence and Combustion*, 62(3):183–200, 1999.
- [193] P. Ouro, B. Fraga, U. Lopez-Novoa, and T. Stoesser. Scalability of an eulerian-lagrangian large-eddy simulation solver with hybrid mpi/openmp parallelisation. *Computers & Fluids*, 179:123–136, 2019.

- [194] P. Ouro and T. Stoesser. Impact of environmental turbulence on the performance and loadings of a tidal stream turbine. *Flow, Turbulence and Combustion*, 102(3):613–639, 2019.
- [195] S. Osher and J. A. Sethian. Fronts propagating with curvature-dependent speed: Algorithms based on hamilton-jacobi formulations. *Journal of Computational Physics*, 79(1):12–49, 1988.
- [196] J. A. Sethian and P. Smereka. Level set methods for fluid interfaces. *Annual Review of Fluid Mechanics*, 35(1):341–372, 2003.
- [197] S. O. R. Fedkiw and S. Osher. Level set methods and dynamic implicit surfaces. *Surfaces*, 44:77, 2002.
- [198] D. H. Zhao, H. W. Shen, J. S. Lai, and G. Q. T. III. Approximate riemann solvers in fvm for 2d hydraulic shock wave modeling. *Journal of Hydraulic Engineering*, 122(12):692–702, 1996.
- [199] M. Sussman, P. Smereka, and S. Osher. A level set approach for computing solutions to incompressible two-phase flow. *Journal of Computational Physics*, 114(1):146–159, 1994.
- [200] S. Kara, M. C. Kara, T. Stoesser, and T. W. Sturm. Free-surface versus rigid-lid les computations for bridge-abutment flow. *Journal of Hydraulic Engineering*, 141(9):04015019, 2015.
- [201] M. V. Lowson, United States. National Aeronautics, Space Administration, and George C. Marshall Space Flight Center. *Pressure Fluctuations in Turbulent Boundary Layers*. NASA technical note. National Aeronautics and Space Administration, 1965.
- [202] T. M. Farabee and M. J. Casarella. Spectral features of wall pressure fluctuations beneath turbulent boundary layers. *Physics of Fluids A*, 3(10):2410–2420, oct 1991.
- [203] H. P. Bakewell, G. F. Carey, J. J. Libuka, H. H. Schloemer, and W. A. Von Vinkle. Wall pressure correlations in turbulent pipe flow. Technical report, Defence Technical Information Center, 1962.

- [204] C. J. Geoga, C. L. Haley, A. R. Siegel, and M. Anitescu. Frequency–wavenumber spectral analysis of spatio-temporal flows. *Journal of Fluid Mechanics*, 848:545–559, aug 2018.
- [205] A. Zampiron, S. Cameron, and V. Nikora. Secondary currents and very-large-scale motions in open-channel flow over streamwise ridges. *Journal of Fluid Mechanics*, 887, 2020.
- [206] S. M. Cameron, V. I. Nikora, and M. T. Stewart. Very-large-scale motions in rough-bed open-channel flow. *Journal of Fluid Mechanics*, 814:416–429, 2017.
- [207] Otis M Solomon Jr. Psd computations using welch’s method. *NASA STI/Recon Technical Report N*, 92:23584, 1991.
- [208] N. R. M. Schipper, J. G. M. Kuerten, and J. C. H. Zeegers. Reducing honeycomb-generated turbulence with a passive grid. *Experiments in Fluids*, 64(7):1–15, 2023.
- [209] L. C. Thijs, R. A. Dellaert, S. Tajfirooz, J. C. H. Zeegers, and J. G. M. Kuerten. Honeycomb-generated Reynolds-number-dependent wake turbulence. *Journal of Turbulence*, 22(9):535–561, 2021.
- [210] S. C. C. Bailey and A. J. Smits. Experimental investigation of the structure of large- and very-large-scale motions in turbulent pipe flow. *Journal of Fluid Mechanics*, 651:339–356, 2010.