

University of Alberta

Modelling Concentration Fluctuations in Plumes Dispersing in
Urban Canopy Flows Within a Single-Particle Lagrangian
Description for Turbulent and Molecular Mixing

by

Jonathan Victor Postma

A thesis submitted to the Faculty of Graduate Studies and Research
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

Department of Earth and Atmospheric Sciences

© Jonathan Victor Postma
Spring 2010
Edmonton, Alberta

Permission is hereby granted to the University of Alberta Libraries to reproduce single copies of this thesis and to lend or sell such copies for private, scholarly or scientific research purposes only. Where the thesis is converted to, or otherwise made available in digital form, the University of Alberta will advise potential users of the thesis of these terms.

The author reserves all other publication and other rights in association with the copyright in the thesis and, except as herein before provided, neither the thesis nor any substantial portion thereof may be printed or otherwise reproduced in any material form whatsoever without the author's prior written permission.

Examining Committee

John D Wilson

Eugene Yee

Andrew Bush

Paul Myers

Gerhard Reuter

Brian Sawford

Lorenz Sigurdson

To my family.

Abstract

An interaction by exchange with the conditional mean (IECM) micromixing model is coupled to a three-dimensional single-particle Lagrangian stochastic (LS) model to estimate concentration fluctuations in plumes of a passive (i.e., non-buoyant), non-reactive (i.e., no chemistry) tracer dispersing from a variety of source configurations in four neutrally stratified flows: a horizontally-homogeneous wall shear layer flow; a horizontally-homogeneous representation of the Tombstone canopy flow; a three-dimensional inhomogeneous representation of the Tombstone canopy flow; and a three-dimensional inhomogeneous representation of the Mock Urban Setting Trials (MUST) canopy flow. The IECM micromixing model incorporates the combined effects of turbulent and molecular mixing on particle concentration. This allows the numerical estimation of all moments of the scalar concentration field, which is a significant advance over traditional LS models given that concentration fluctuations are a ubiquitous feature of a dispersing plume.

The single-particle implementation of the LS-IECM model is based upon a previously reported implementation that used simultaneously computed parti-

cle trajectories to estimate the conditional mean concentration field [Cassiani, M. A., Franzese, P. A. and Giostra, U. A.: 2005, A PDF micromixing model of dispersion for atmospheric flow. Part I: development of model, application to homogeneous turbulence and to a neutral boundary layer, *Atmospheric Environment* 39, 1457-1469]. The model used in this thesis pre-calculates the conditional mean concentration field with an LS model for use with the IECM model, which runs as a separate simulation. The principal advantage of this single-particle approach is the performance increase on parallel computer architecture, which scales directly with the number of processors. The simulations presented in this thesis go beyond those performed with the previous model by considering three-dimensional inhomogeneous flows, as well as one-dimensional horizontally-homogeneous flows.

The accuracy of the LS-IECM model was good for the flows with horizontal-homogeneity, and comparable to the results of previous simulations from older models. Rogue velocities in the simulations utilising inhomogeneous flow statistics resulted in acceptable to poor accuracy in these simulations. Suggestions for improvements to the model are made.

Acknowledgements

During my tenure as a Ph.D. student, I have been very fortunate in that I have had not only one, but two excellent supervisors. Both Dr John Wilson and Dr Eugene Yee have been exceedingly generous with their academic, technical, and financial support; readily available and quick to reply to emails, questions, and requests; meticulous editors with surprisingly short turn-around times, given the length of my thesis; and provided opportunities for me to gain valuable field experience, even though my project was strictly numerical. For all of these things I am deeply grateful and offer my sincere thanks.

I would also like to extend my appreciation and thanks to my supervisory and examination committee members: Dr Andrew Bush, Dr Paul Myers, Dr Gerhard Reuter, Dr Brian Sawford, and Dr Lorenz Sigurdson, for taking time out of their busy schedules to provide me with useful feedback and comments about my thesis. Together with suggestions from my supervisors, my thesis is more thorough, accurate, and interesting thanks to all of their efforts.

There were many difficulties during the writing of this thesis: segmentation faults, non-responsive programs, confusing papers, et cetera. At times like these the thought of turning quitting fantasies (which I suspect all graduate students experience) into reality seemed very seductive. It was in these times that I was most grateful to have a very close, supportive family, that I could call or visit to take my mind off of my thesis work until the quitting fantasies

subsided. The care packages filled with many delicious Dutch goodies (which are much harder to come by in Edmonton) sent by my wonderful parents also helped.

Thanks also to the EAS support staff in the main office and the DIF. Your contributions to the department are much appreciated. In particular, I would like to thank our fantastic Graduate Program Administrator: Marsha Boyd. If ever there was a question regarding the program requirements or an upcoming deadline, the quickest and most surefire way to learn more about it was to send Marsha an email.

Referring back to the quitting fantasies, thanks to my jovial lot of friends for all the ski trips, climbing trips, potluck dinners, games nights, and other events over the past few years. Stepping away from the computer every now and then certainly helped me complete this thesis.

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Previous Work	4
1.3	Outline	8
2	Background Theory	11
2.1	Turbulence	12
2.2	The Governing Equations	19
2.3	The Mean Flow Equations	22
2.4	Atmospheric Flows	27
2.4.1	Neutral Wall Shear Layer Flow	29
2.4.2	Canopy Flow	33
2.5	Turbulent Dispersion and Mixing	38
2.6	Probability Density Function Modelling	47
3	Closure and Numerical Modelling	50
3.1	Lagrangian-Stochastic Trajectory Modelling	50
3.2	The IECM Micromixing Model	55
3.2.1	The Micromixing Timescale	59

3.3	Description of the SPMMM Model Suite	64
3.3.1	The MEANS Pre-calculation Program	67
3.3.2	The SPMMM Micromixing Model	77
3.4	Model Evaluation Methodology	80
4	Dispersion in a Neutral Wall Shear Layer Flow	87
4.1	Experimental and Computational Setup	87
4.2	Parameter Calibration of MEANS	88
4.3	Model Consistency	96
4.4	Comparison to FR82 Experimental Data	104
4.4.1	Dispersion from an Elevated Point Source	106
4.4.2	Dispersion from a Ground-level Point Source	124
4.5	Error Analysis	127
4.6	Chapter Summary	129
5	Tombstone Canopy Simulations	132
5.1	Experimental and Computational Setup	132
5.2	Dispersion from a Continuous Line Source	148
5.2.1	Comparison of SPMMM to the Model of Cassiani et al.	156
5.3	Dispersion from a Continuous Point Source	159
5.4	Chapter Summary	187
6	MUST Canopy Simulations	192
6.1	Experimental and Computational Setup	193
6.2	Dispersion from a Continuous Point Source	209
6.2.1	Rogue Velocities Prevent Consistency	214
6.2.2	Above Canopy Concentration Statistics	225

6.3 Chapter Summary	230
7 Summary and Conclusions	235
7.1 Summary	235
7.2 Conclusions	244
Bibliography	250
A Mathematical and Statistical Tools	261
A.1 Eulerian and Lagrangian Coordinates	261
A.2 Tensor Notation	263
A.3 Probability Density Functions	266
A.3.1 Example: The Gaussian or Normal Distribution	269
A.4 Statistical Tools	270
A.4.1 The Unconditional and Conditional Means	270
A.4.2 Fluctuations and Moments	273

List of Tables

2.1	Summary of plume growth rates for absolute and relative dispersion	46
3.1	Expected behaviour of the micromixing timescale	60
3.2	Performance measures for a perfect and an acceptable model .	85
4.1	Consistency with the driving velocity statistics	99
5.1	Values of various physical and aerodynamic parameters from the Tombstone canopy experiments	140
5.2	Performance measures for the mean dimensionless concentration on the plume centreline from SPMMM simulations of an elevated, in-canopy point source in the Tombstone canopy flow .	164
5.3	Performance measures for the standard deviation of the dimensionless concentration on the plume centreline from SPMMM simulations of an elevated, in-canopy point source in the Tombstone canopy flow	170

5.4	Performance measures for the skewness of the dimensionless concentration on the plume centreline from SPMM simulations of an elevated, in-canopy point source in the Tombstone canopy flow	176
5.5	Performance measures for the kurtosis of the dimensionless concentration on the plume centreline from SPMM simulations of an elevated, in-canopy point source in the Tombstone canopy flow	182
6.1	The number of particle steps with rogue velocities in two MUST canopy simulations	216

List of Figures

2.1	Laminar and turbulent flows	13
2.2	Turbulent scales of motion and the energy cascade	15
2.3	A simplified atmosphere	27
2.4	Turbulent mixing and molecular diffusion	39
2.5	Stages of plume development	41
3.1	The time behaviour of the micromixing timescale	63
3.2	Dynamic discretisation of the plume	69
3.3	Discretisation of the velocity space	71
3.4	Accumulation of residence times	72
3.5	Determination of the conditional mean concentrations	75
3.6	Approximating the velocity PDF in a large spatial bin	76
3.7	Initialising SPMMM	79
3.8	Extracting data from SPMMM	81
4.1	Vertical profiles of FR82 wall shear layer velocity statistics and TKE dissipation rate	89
4.2	Vertical profiles of the normalised mean concentration for an ele- vated point source from the MEANS Kolmogorov constant tuning simulations in the FR82 flow	91

4.3	Performance measures from the MEANS Kolmogorov constant tuning simulations in the FR82 flow	92
4.4	Vertical profiles of the normalised mean concentration for an elevated point source from the MEANS timestep tuning simulations in the FR82 flow	94
4.5	Performance measures from the MEANS timestep tuning simulations in the FR82 flow	95
4.6	Vertical profiles of the normalised mean concentration for an elevated point source from the MEANS initial source distribution tuning simulations in the FR82 flow	97
4.7	MEANS and SPMMM first-order consistency check for simulations in the FR82 flow	100
4.8	Performance measures for the MEANS and SPMMM consistency check simulations in the FR82 flow	102
4.9	Inconsistency between MEANS and SPMMM due to $N_z^{\text{vel}} > N_z$	103
4.10	Smoothing a signal with the Savitzky-Golay filter	105
4.11	Comparison of the MEANS and SPMMM simulations with the FR82 vertical profiles of the normalised mean concentration for an elevated point source in the FR82 flow	107
4.12	Streamwise transect of the dimensionless maximum mean concentration for an elevated point source in the FR82 flow from MEANS and SPMMM simulations	108
4.13	Spanwise and vertical plume halfwidths from an elevated point source in the FR82 flow from MEANS and SPMMM simulations	109

4.14	Non-alteration of the first-order statistics by the IECM micromixing model.	111
4.15	Example of the effects of μ and C_r on the second-order concentration statistics	113
4.16	Contours in $\mu - C_r$ parameter space of the performance measures from SPMMM simulations of an elevated point source in the FR82 flow	114
4.17	Streamwise transects of the concentration fluctuation intensity from SPMMM simulations of an elevated point source in the FR82 flow	116
4.18	Performance measures resulting from the SPMMM initial source distribution tuning simulations of an elevated point source in the FR82 flow	118
4.19	Streamwise transects of the concentration fluctuation intensity from SPMMM simulations of elevated point sources of various sizes, and a ground-level source, in the FR82 flow	120
4.20	Vertical profiles of the normalised concentration variance at five downstream positions from SPMMM simulations of an elevated point source in the FR82 flow	121
4.21	Vertical profiles of the normalised mean concentration resulting from SPMMM simulations of a ground-level point source in the FR82 flow	125
4.22	Vertical profiles of the normalised concentration variance at five downstream positions from SPMMM simulations of a ground-level point source in the FR82 flow	126

4.23	Statistical error in the SPMMM simulated mean concentration and concentration variance	127
4.24	Bias error in the SPMMM simulated mean concentration and concentration variance	128
4.25	Discretisation error in the SPMMM simulated mean concentration and concentration variance	129
5.1	Tombstone canopy unit cell showing data extraction locations and the discretisation grid	133
5.2	Spatially-averaged Tombstone canopy velocity statistics from the wind-tunnel and water-channel experiments	136
5.3	Vertical profiles of the measured mean streamwise velocity from the Tombstone canopy water-channel experiments and the resulting interpolation.	141
5.4	Streamwise and spanwise transects of the interpolated mean streamwise velocity for the water-channel Tombstone canopy experiments.	143
5.5	Tombstone canopy TKE dissipation rate and the Lagrangian integral timescale measurements from wind-tunnel experiments	145
5.6	Vertical profiles of measured and modelled Tombstone canopy TKE dissipation rates	147
5.7	Vertical profiles of the dimensionless mean concentration resulting from SPMMM simulations of an elevated, in-canopy, cross-wind line source in the Tombstone canopy flow.	150

5.8	Vertical profiles of the dimensionless standard deviation of concentration from SPMMM simulations of an elevated, in-canopy, cross-wind line source in the Tombstone canopy flow.	152
5.9	Streamwise transect of the ground-level dimensionless concentration mean from SPMMM simulations of an elevated, cross-wind line source in the Tombstone canopy flow.	154
5.10	Streamwise transect of the concentration fluctuation intensity from SPMMM simulations of an elevated, in-canopy, cross-wind line source in the Tombstone canopy flow.	155
5.11	Comparison of the vertical profiles of the dimensionless mean concentration from an elevated, in-canopy, cross-wind line source in the Tombstone canopy flow, as simulated by SPMMM and Cassiani et al. (2007)	157
5.12	Comparison of the vertical profiles of the dimensionless standard deviation of concentration from an elevated, in-canopy, cross-wind line source in the Tombstone canopy flow as simulated by SPMMM and Cassiani et al. (2007)	158
5.13	Example of the consistency between the MEANS and SPMMM simulated vertical profile of mean concentration.	162
5.14	Streamwise transects of the simulated mean dimensionless concentration at six heights on the plume centreline from SPMMM simulations of an elevated, in-canopy point source in the Tombstone canopy flow	163

5.15	Spanwise transects of the dimensionless mean concentration at three heights, 1 and 2 rows downstream from the source from SPMMM simulations of an elevated, in-canopy point source in the Tombstone canopy flow	166
5.16	Spanwise transects of the dimensionless mean concentration at three heights, 4 and 8 rows downstream from the source from SPMMM simulations of an elevated, in-canopy point source in the Tombstone canopy flow	167
5.17	Spanwise transects of the dimensionless mean concentration at three heights, 12 and 16 rows downstream from the source from SPMMM simulations of an elevated, in-canopy point source in the Tombstone canopy flow	168
5.18	Streamwise transects of the standard deviation of the dimensionless concentration at six heights on the plume centreline resulting from SPMMM simulations of an elevated, in-canopy point source in the Tombstone canopy flow	169
5.19	Spanwise transects of the standard deviation of the dimensionless concentration at three heights, 1 and 2 rows downstream from the source from SPMMM simulations of an elevated, in-canopy point source in the Tombstone canopy flow	172
5.20	Spanwise transects of the standard deviation of the dimensionless concentration at three heights, 4 and 8 rows downstream from the source from SPMMM simulations of an elevated, in-canopy point source in the Tombstone canopy flow	173

5.21	Spanwise transects of the standard deviation of the dimensionless concentration at three heights, 12 and 16 rows downstream from the source from SPMMM simulations of an elevated, in-canopy point source in the Tombstone canopy flow	174
5.22	Streamwise transects of the skewness of the dimensionless concentration at six heights on the plume centreline from SPMMM simulations of an elevated, in-canopy point source in the Tombstone canopy flow	175
5.23	Spanwise transects of the skewness of the dimensionless concentration at three heights, 1 and 2 rows downstream from the source from SPMMM simulations of an elevated, in-canopy point source in the Tombstone canopy flow	178
5.24	Spanwise transects of the skewness of the dimensionless concentration at three heights, 4 and 8 rows downstream from the source from SPMMM simulations of an elevated, in-canopy point source in the Tombstone canopy flow	179
5.25	Spanwise transects of the skewness of the dimensionless concentration at three heights, 12 and 16 rows downstream from the source from SPMMM simulations of an elevated, in-canopy point source in the Tombstone canopy flow	180
5.26	Streamwise transects of the kurtosis of the dimensionless concentration at six heights on the plume centreline from SPMMM simulations of an elevated, in-canopy point source in the Tombstone canopy flow	181

5.27	Spanwise transects of the kurtosis of the dimensionless concentration at three heights, 1 and 2 rows downstream of the source from SPMMM simulations of an elevated, in-canopy point source in the Tombstone canopy flow	184
5.28	Spanwise transects of the kurtosis of the dimensionless concentration at three heights, 4 and 8 rows downstream of the source from SPMMM simulations of an elevated, in-canopy point source in the Tombstone canopy flow	185
5.29	Spanwise transects of the kurtosis of the dimensionless concentration at three heights, 12 and 16 rows downstream of the source from SPMMM simulations of an elevated, in-canopy point source in the Tombstone canopy flow	186
6.1	The MUST obstacle array	196
6.2	The reduced MUST obstacle array as used for the MEANS and SPMMM simulations	199
6.3	Vertical profiles of the measured mean streamwise velocity from the water-channel MUST canopy experiments and the resulting linear interpolation.	201
6.4	Streamwise transects of the interpolated mean streamwise velocity for the water-channel MUST canopy experiments.	202
6.5	Spanwise transects of the interpolated mean streamwise velocity for the water-channel MUST canopy experiments.	203
6.6	Modelled vertical profiles of the TKE dissipation rate for the water-channel MUST canopy experiments.	204

6.7	Vertical profiles of various velocity statistics from the water-channel MUST canopy experiments - part 1.	205
6.8	Vertical profiles of various velocity statistics from the water-channel MUST canopy experiments - part 2.	206
6.9	A suspected outlier in the Coanda σ_w measurements	209
6.10	Spanwise transects of the first four moments of the dimensionless concentration field at $z/h_c = 0.5$, 11 rows downstream of the source from SPMMM simulations of a ground-level point source in the MUST canopy flow.	212
6.11	Example of the first-order inconsistency between MEANS and SPMMM simulations in the MUST canopy flow	213
6.12	y - z scatter plot of rogue velocity positions	217
6.13	MEANS and SPMMM first-order consistency check for the MUST canopy simulations at $z/h_c = 7.0$	219
6.14	MEANS and SPMMM first-order consistency check for the MUST canopy simulations at $z/h_c = 4.0$	221
6.15	Comparison of the vertical profiles of $\langle \phi u, v, w \rangle$ from MEANS simulations for a velocity bin that receives re-initialised rogue particles.	223
6.16	Spanwise transects of the dimensionless mean concentration at six downstream positions from SPMMM simulations of an above canopy point source in the MUST canopy flow.	228

6.17	Spanwise transects of the standard deviation of the dimensionless concentration at six downstream positions from SPMMM simulations of an above canopy point source in the MUST canopy flow.	229
6.18	Spanwise transects of the skewness of the dimensionless concentration at six downstream positions from SPMMM simulations of an above canopy point source in the MUST canopy flow. . . .	231
6.19	Spanwise transects of the kurtosis of the dimensionless concentration at six downstream positions from SPMMM simulations of an above canopy point source in the MUST canopy flow. . . .	232
7.1	Contours of the percentage by volume of ethylene for a simulated accident at a petrochemical plant	249
A.1	The standardised Gaussian or normal distribution	271
A.2	Examples of skewness in probability density functions	275
A.3	Examples of excess kurtosis in probability density functions . . .	276

List of Symbols

Acronyms

ABL	atmospheric boundary layer
ASL	atmospheric surface layer
CBRN	chemical, biological, radiological, nuclear
CDF	cumulative distribution function
ES	elevated source
FAC2	factor of two from observations
FB	fractional bias
FB _{fn}	fractional bias (false negative)
FB _{fp}	fractional bias (false positive)
FP	Fokker-Planck
FR82	Fackrell and Robins (1982)
GLS	ground-level source
IDL	interactive data language
IEM	interaction by exchange with the mean
IECM	interaction by exchange with the conditional mean
LDV	laser Doppler velocimetry
LIF	laser induced fluorescence

LS	Lagrangian stochastic
MUST	mock urban setting trials
NMSE	normalised mean square error
PDF	probability density function
SDE	stochastic differential equation
RANS	Reynolds-averaged Navier-Stokes
RMS	root-mean-square
TKE	turbulent kinetic energy
WMC	well-mixed condition

Upper-case Roman Letters

A_f	frontal area of a canopy element
A_i	exact acceleration term
A_L	lot area
A_s	area of the source
B_Q	bias error in quantity Q
C	a constant
C_D	effective drag coefficient
C_0	Kolmogorov constant
C_{ij}	the covariance matrix
C_{ij}^{-1}	the inverse of the covariance matrix
C_r	Richardson constant
\mathcal{D}	molecular diffusivity
D_{ij}^L	second-order Lagrangian structure function
D_Q	deterministic error in quantity Q

E	energy spectrum function
$F_{\mathbf{u}}$	CDF for Eulerian velocity
\mathcal{G}	a domain in a turbulent flow
K_m	eddy viscosity for momentum
Ku_ϕ	excess kurtosis of ϕ
Ku_ϕ^*	excess kurtosis of dimensionless ϕ
L	length scale of the most energetic eddies
M	spatial resolution, $M = N_x \times N_y \times N_z$
N	number of particles used by SPMMM
N_p	number of computer processors
N_ϕ	number of particles used by MEANS
N_ϕ^v	number of particles that visit a velocity bin
N_u	number of streamwise velocity bins
N_v	number of spanwise velocity bins
N_w	number of vertical velocity bins
N_x	number of streamwise bins
N_x^{vel}	number of streamwise velocity statistic bins
N_y	number of spanwise bins
N_y^{vel}	number of spanwise velocity statistic bins
N_z	number of vertical bins
N_z^{vel}	number of vertical velocity statistic bins
P	the pressure
\mathcal{P}_s	shear production of TKE
\mathcal{P}_w	wake production of TKE
$Prob$	probability of an event occurring

Q	source strength
Q_o	an observed quantity (i.e., from experiment)
Q_p	a predicted quantity (i.e., from simulation)
R	Lagrangian velocity auto-correlation function
R_r	Lagrangian relative velocity auto-corr. function
R_{ij}	Reynolds stress tensor
Re	Reynolds number for the flow
Re_δ	Reynolds number at the top of the boundary layer
Re_l	Reynolds number for an eddy of size l
Re_η	Reynolds number for the Kolmogorov scales
$S(\phi)$	chemical reaction source term
S_Q	discretisation error in quantity Q
Sc	the Schmidt number
Sk_ϕ	skewness of ϕ
Sk_ϕ^*	skewness of dimensionless ϕ
T_L	Lagrangian integral timescale
T_r	integral timescale for relative velocities
\mathcal{T}	TKE transfer rate
\mathcal{T}_d	dispersive transport of TKE
\mathcal{T}_m	molecular transport of TKE
\mathcal{T}_p	pressure transport of TKE
\mathcal{T}_t	turbulent transport of TKE
\mathbf{U}	Lagrangian vector velocity $\mathbf{U} = (U, V, W)$
\mathbf{U}'	fluctuating component of above velocity
U	Lagrangian streamwise speed

U'	fluctuating Lagrangian streamwise speed
V	Lagrangian spanwise speed
V'	fluctuating Lagrangian spanwise speed
\mathbf{V}	sample space variable for \mathbf{U}
\mathcal{V}	volume
W	Lagrangian vertical speed
W_c	vertical speed of the plume centre of mass
W_r	Lagrangian relative vertical speed
W'	fluctuating Lagrangian vertical speed
\mathbf{X}	Lagrangian vector position, $\mathbf{X} = (X, Y, Z)$
X	Lagrangian streamwise position
Y	Lagrangian spanwise position
Z	Lagrangian vertical position
Z_c	vertical position of the plume centre of mass
Z_r	Lagrangian relative vertical position

Lower-case Roman Letters

a	element area density
a_i	deterministic drift term
b_{ij}	stochastic diffusion term
d	displacement height
d_r	relative plume width
d_s	source diameter
$d\xi_i$	incremental Wiener process

f_c	Coriolis parameter
f_{F_i}	form drag vector
$f_{\mathbf{u}}$	one-point, one-time PDF of \mathbf{u}
f_{V_i}	viscous drag vector
f_{ϕ}	one-point, one-time PDF of the ϕ
$f_{\phi \mathbf{u}}$	conditional PDF for ϕ and \mathbf{u}
$f_{\mathbf{u}\phi}$	joint PDF for \mathbf{u} and ϕ
\mathbf{g}	gravitational force per unit mass
g	local gravitational acceleration
g_a	density function of fluid particles
g_t	density function of tracer particles
h_c	canopy height
$h_c^* = A_f^{1/2}$	length scale for flow interaction
k	turbulent kinetic energy
k_v	von Kármán constant
l	characteristic length scale of an eddy
ℓ	Prandtl – von Kármán mixing length
p	modified pressure (includes gravitational terms)
t	time
t'	a dummy variable for time
t''	another dummy variable for time
t_0	source timescale scaled by $C_r^{1/3}$
t_m	micromixing timescale
t_s	characteristic timescale of the source
t_t	turbulence timescale

t_r	residence time in bin (I, J, L)
t_r^v	residence time in bin (I, J, K, L, M, N)
\mathbf{u}	Eulerian vector velocity, $\mathbf{u} = (u, v, w)$
u	Eulerian streamwise speed
\mathbf{u}'	fluctuating component of Eulerian velocity
u'	fluctuating Eulerian streamwise speed
u_L	velocity scale the most energetic eddies
u_ℓ	velocity scale of an eddy with size ℓ
\mathbf{u}_r	turbulent relative velocity
u_s	streamwise source speed
u_*	friction velocity
$\langle u \rangle_c$	mean streamwise speed at canopy top
$\langle u \rangle_\delta$	mean streamwise speed at $z = \delta$
$\langle u \rangle_s$	mean streamwise speed at source height
v	Eulerian spanwise speed
v'	fluctuating Eulerian spanwise speed
w	Eulerian vertical speed
w'	fluctuating Eulerian vertical speed
\mathbf{x}	Eulerian vector position, $\mathbf{x} = (x, y, z)$
x	Eulerian position on x -axis
y	Eulerian position on y -axis
z	Eulerian position on z -axis
z_0	roughness length
z_s	source height
z_{rft}	the reflection height

Upper-case Greek Letters

Δ	represents an interval, i.e Δx is the x bin width
Φ	Φ/g_a is the drift correction term in a_i
Ψ	the gravitational potential
Σ_Q	statistical error in quantity Q
Θ	exact diffusion-reaction term

Lower-case Greek Letters

β	extinction parameter
δ	boundary layer depth
δ_{ij}	Kronecker delta
δ_y	spanwise plume halfwidth
δ_z	vertical plume halfwidth
ϵ_{ijk}	alternating unit tensor
ϵ_Q	total error
ϵ_n	numerical error
ϵ_p	physical error
ϵ	turbulent kinetic energy dissipation rate
κ	a wavenumber
λ_B	Batchelor scale
λ_d	diffusion length scale
λ_f	frontal area index

λ_p	plan area index
μ	the micromixing constant
μ_n	n^{th} central moment
μ_r	scale factor for x, y, z space
μ_s	scale factor for the initial source distribution
μ_t	scale factor for the timestep
μ_v	scale factor for u, v, w space
ν	kinematic viscosity
ω	angular velocity of Earth
ρ	constant density
ϕ	scalar concentration
ϕ_s	scalar concentration at source
$\phi_* = \phi/\phi_s$	dimensionless scalar concentration
ϕ_*	temperature scale
ϕ'	fluctuating component of scalar concentration
φ	latitude
ψ	sample space variable for scalar concentration
σ	standard deviation
σ_ϕ	standard deviation of ϕ
σ_ϕ^*	standard deviation of dimensionless ϕ
σ_0	initial source distribution
σ_r	width of the instantaneous plume
σ_z	width of the absolute plume in z
σ^2	variance
σ_{u_i}	tensor form of the velocity variance

$\sigma_{\mathbf{U}_r}^2$	variance of \mathbf{U}_r
$\sigma_{w_r}^2$	variance of W_r
τ_a	an advection timescale
τ_L	turnover timescale of the most energetic eddies
τ_ℓ	turnover timescale of an eddy
θ	generic micromixing model

Chapter 1

Introduction

1.1 Motivation

Modern society produces many potentially harmful materials: the manufacturing sector requires and produces chemicals; power generation sometimes produces nuclear by-products; hospitals and research establishments frequently produce and use biohazardous and radiological materials. For the most part these chemical, biological, radiological, and nuclear (CBRN) materials are highly controlled and produced, used, and disposed of in a manner that is completely safe. However, in their transport from production sites to factories or storage facilities, there is a risk of release of one of these harmful substances into the environment. Alternatively, the material may end up in the hands of an ill-intentioned group or state. Given that transportation networks frequently run into the heart of cities, and that terrorists try to maximise terror, the possibility of an accidental or intentional release of a CBRN agent into a densely populated urban centre is real, albeit low. Given the potential consequences of such a release, an improved understanding of the transport and dispersion of the agent once in the environment is necessary (or certainly,

desirable) to predict and evaluate the possible impacts of such a release and to support the formulation of appropriate mitigation strategies.

Imagine that there was a train derailment and that a tanker car carrying a toxic chemical ruptured and was releasing its contents into the atmosphere. The subsequent exposure of the population in the vicinity of the spill can be assessed using the dosage, which is defined as the concentration integrated over time:

$$\text{Dosage} = \int_{t_1}^{t_2} \phi dt, \quad (1.1)$$

where the concentration is denoted by ϕ . Note that the dosage is not necessarily in linear relation with the total uptake by an organism, the latter being a potentially complex function of such factors as breathing rate, shelter, et cetera. Until recently the vast majority of computer models only predicted the mean concentration $\langle \phi \rangle$, ignoring all higher-order moments of the concentration field such as the concentration variance σ_ϕ^2 , the concentration skewness Sk_ϕ , the concentration kurtosis Ku_ϕ^* , and so on. Thus dosage estimates were based upon the integral of the mean concentration.

Clearly, the inclusion of these higher-order moments would allow various statistical characteristics of the dosage (other than simply its mean value) to be used in the hazard assessment. For example, if the concentration field were Gaussian then the prescription of the mean concentration and the concentration variance would describe the full concentration probability density function (PDF), thus providing a full statistical description of the dosage that the exposed population may have been subject to. Therefore, from an emergency preparedness and risk assessment standpoint, it is rational to pursue the devel-

*Note that in this thesis the term “kurtosis” refers to the excess kurtosis. See Section A.4.2 in Appendix A.

opment of models that predict the higher-order moments of the concentration field. Coupled to this, there must be development of a model that can accurately estimate the concentration field in the highly disturbed flow frequently found in urban environments where human habitation is concentrated. Again, until recently many models focused on highly idealised terrain.

From a scientific standpoint, modelling the concentration variance and the other higher-order moments of the concentration field is the next logical progression after modelling the mean concentration. It is also logical to pursue development of dispersion models for complex and disturbed flows such as those found in urban areas. It is worth noting that the complex problems involved with modelling the higher-order moments of the concentration field, and dispersion in highly disturbed flows, have thus far hindered progress in this research area. Theoretical developments have been slow and, until recently, comprehensive data sets of concentration fluctuations in canopy flows were unavailable. These facts, combined with a lack of computational resources, have resulted in limited advances in the development of numerical models of dispersion in a canopy environment.

In this work we aim to increase the understanding of dispersion of a passive scalar in complex canopy flows through the use of the interaction by exchange with the conditional mean (IECM) micromixing model coupled to a single-particle Lagrangian stochastic (LS) trajectory model. The LS trajectory model governs the evolution of the position and velocity of a tracer particle via specified equations. In addition to the position and velocity, the tracer particle also has a concentration variable assigned to it, the evolution of which is governed by the IECM micromixing model. Combined, these two models

allow the prediction of the higher-order moments of the scalar concentration field.

1.2 Previous Work

Before proceeding it is necessary to introduce the notation used throughout this thesis. Vectors shall be represented by ***bold*** typeface and both tensor notation and standard meteorological notation will be used, i.e.

$$\begin{aligned}\mathbf{u}(\mathbf{x}, t) &= u_i = (u_1, u_2, u_3) = (u, v, w), \\ \mathbf{x} &= x_i = (x_1, x_2, x_3) = (x, y, z).\end{aligned}$$

A vector field such as that of velocity is denoted as $\mathbf{u}(\mathbf{x}, t)$. Eulerian quantities shall be denoted by lower-case letters (i.e. \mathbf{u}) while Lagrangian quantities shall be denoted by upper-case letters (i.e., \mathbf{U}). Other notation will be introduced as needed. Appendix A contains a brief review of the mathematics and statistics necessary to fully appreciate this thesis. For more thorough descriptions of the statistics used in the atmospheric sciences see Van Kampen (1981), Gardiner (1983), Rodean (1996), or Pope (2000).

The second-order statistic of the scalar concentration field is the concentration variance. There are several approaches to modelling concentration variance. Direct numerical simulation could be used to solve the Navier-Stokes and advection-diffusion equations down to the smallest scales of space and time from which the concentration variance could be determined. At present, computational limitations make this approach infeasible for atmospheric science applications.

In Reynolds averaged form, the Eulerian scalar variance budget equation

is

$$\frac{\partial \langle \phi'^2 \rangle}{\partial t} + \langle u_i \rangle \frac{\partial \langle \phi'^2 \rangle}{\partial x_i} = -2 \langle u'_i \phi' \rangle \frac{\partial \langle \phi \rangle}{\partial x_i} - \frac{\partial}{\partial x_i} \langle u'_i \phi'^2 \rangle - 2\mathcal{D} \left\langle \frac{\partial \phi'}{\partial x_i} \frac{\partial \phi'}{\partial x_i} \right\rangle + \mathcal{D} \frac{\partial^2 \langle \phi'^2 \rangle}{\partial x_i \partial x_i}, \quad (1.2)$$

where \mathcal{D} is the molecular diffusivity of the scalar in air and u'_i and ϕ' are the fluctuating components of Eulerian velocity and scalar concentration, respectively. Equation (1.2) cannot be solved unless closure assumptions are made for the scalar flux $\langle u'_i \phi' \rangle$, the triple correlation $\langle u'_i \phi'^2 \rangle$, and the scalar dissipation $2\mathcal{D} \left\langle \frac{\partial \phi'}{\partial x_i} \frac{\partial \phi'}{\partial x_i} \right\rangle$. Closure assumptions that have been suggested in the literature range from being overly simplistic to being formidably complicated. In the former case, the result is often a less realistic model. In the latter case, the difficulty in calculating the numerical solution is increased. A thorough discussion of closure assumptions can be found in Launder (1978), and Chapter 6 of Stull (1988). The approach of directly solving the Reynolds-averaged Navier-Stokes (RANS) equations, along with various closure assumptions, was used with good success to model scalar dispersion in the idealised Mock Urban Setting Trials (MUST) Canopy (Yee and Biltoft, 2004) by Hsieh et al. (2007), and Wang et al. (2009).

Fluctuating plume models have been successfully used to calculate the concentration variance and other higher-order moments of concentration for non-reactive scalars (Yee et al., 1994; Yee and Wilson, 2000; Luhar et al., 2000; Franzese, 2003; Gailis et al., 2006). These models add parametrised relative in-plume concentration fluctuations to a meandering plume model (Gifford, 1959). This approach is not followed here as future work may involve working with reactive scalars, perhaps in the context of air-quality modelling, for example. As will be shown below, the IECM model handles chemical

reactions naturally.

Another group of models that have been used to compute concentration variance can be collectively referred to as “PDF modelling techniques”. The assumed PDF method specifies the form of the concentration PDF, which contains some unknown parameters. Using the transport equations for the mean concentration and the concentration variance, the values of these parameters can be determined thereby fully defining the concentration PDF which can be used to predict all higher-order moments of the concentration field. Yee et al. (2009) recently applied this technique with good success to an experimental urban canopy composed of rows of aligned cubes. Another promising approach is transported PDF modelling. One implementation of this approach numerically solves for the one-point, one-time joint velocity and concentration PDF, which is denoted by $f_{\mathbf{u}\phi}$. There are two main advantages to this approach: the results are approximations to the PDF from which all joint concentration-velocity statistics can be derived; and all chemical reactions are treated exactly with no closure assumptions.

However, the transport equation for the joint velocity and concentration PDF of a single scalar species in three spatial dimensions has at least eight degrees of freedom associated with it: three space variables, three velocity variables, a scalar concentration variable, and time. This fact alone makes an Eulerian approach impractical, unless symmetries of the experimental scenario reduce the dimensionality. Furthermore, numerical techniques must be employed to ensure that $f_{\mathbf{u}\phi}$ is well behaved (i.e., non-negative and normalised at all points in space and time) which increases the complexity of the code (Fox, 2003). Eulerian methods also suffer from numerical diffusion and alias-

ing. In contrast to Eulerian methods, Lagrangian stochastic (LS) methods are grid free and thus can handle many independent variables with ease, making them well suited for transported PDF methods. Since LS models use a finite number of particles, the simulation results will contain error. Provided that this error is understood (Xu and Pope, 1999), LS-PDF methods are a powerful tool for the determination of concentration statistics. In practice, a desired level of accuracy can be achieved by utilising more particles.

The influential paper by Taylor (1921) gave LS methods their start. Since then they have become widely used. Examples of their application can be found in: Wilson and Sawford (1996); Pope (2000); Fox (2003). Single-particle LS models are straightforward and easy to implement; however, the higher-order moments of concentration cannot be determined from single-particle trajectories. To obtain the concentration variance through purely LS methods requires using a two-particle model (Durbin, 1980; Thomson, 1990). If the third-order concentration moment is desired, then a three-particle model must be used, and so on. Multi-particle LS models require multi-point relative velocity statistics to describe the flow. Unfortunately, except in idealised regimes of turbulence, these are difficult to obtain and therefore multi-particle models are not highly developed for use as predictive tools

By coupling a single-particle LS model to a micromixing model to describe molecular mixing, all higher-order concentration moments can be calculated without the need of multi-point relative velocity statistics. Dopazo et al. (1997) review many micromixing models. One of the earliest and simplest micromixing models is the interaction by exchange with the mean (IEM) model (Dopazo and O'Brien, 1974; Pope, 1985). It has some drawbacks, such as affecting first-

order statistics due to a spurious term in the scalar flux budget equation (Pope, 1998; Sawford, 2004a). In spite of this, due to its simplicity, the IEM model is still widely used in areas ranging from combustion (Pope, 1985) to atmospheric science (Sawford, 2004a,b). In Section 3.2 the scalar flux budget equation is presented and the spurious term discussed further.

The IECM model (Fox, 1996; Pope, 1998) does not suffer from the spurious flux and its implementation is not significantly more complicated than IEM models. Sawford (2004b) showed that the IECM model can be related to a meandering plume model and applied the technique successfully to simulate concentration statistics due to a continuous line source in grid turbulence. Recent applications of IECM models to atmospheric flows include simulations of dispersion of passive (i.e., non-buoyant), non-reactive (i.e., no chemistry) scalars within the neutral boundary layer (Cassiani, Franzese and Giostra, 2005a), within the convective boundary layer (Cassiani, Franzese and Giostra, 2005b; Luhar and Sawford, 2005) and within a canopy layer (Cassiani, Radicchi and Giostra, 2005; Cassiani et al., 2007).

1.3 Outline

The remainder of this thesis is laid out as follows. Chapter 2 covers the necessary mathematics and governing equations required to fully appreciate the models. It begins with some introductory terms and concepts of turbulence followed by a statement of the fundamental transport equations. The two flow types considered in this thesis (neutral wall shear layer flow, and neutral canopy flow) are then described in general. A mathematical description of dispersion, diffusion, and mixing follows. The chapter concludes with a pre-

sensation of the analytical PDF transport equation, highlighting the unclosed terms that require modelling. Chapter 3 presents the LS and IECM models used to close the unclosed terms in the PDF transport equation, and provides a detailed description on how they are implemented numerically. Chapter 4 opens with a description of the Fackrell and Robins (1982; FR82) wind-tunnel experiments of dispersion of a passive, non-reactive scalar from continuous elevated and ground-level point sources. A comparison of model predictions to experimental data follows, along with a discussion of the numerical error associated with the micromixing model. Dispersion of a passive, non-reactive scalar through a model plant canopy, consisting of a diamond array of rigid obstacles, is the topic of Chapter 5. Experimental data for this *Tombstone Canopy* (Raupach et al., 1986; Coppin et al., 1986; Legg et al., 1986; Raupach et al., 1987) are available from wind-tunnel experiments utilising an elevated, in-canopy, cross-wind continuous line source (Legg et al., 1986), and from water-channel experiments of a near ground-level continuous point source (Hilderman and Chong, 2007). The velocity statistics from the two data sets are compared to show their compatibility, and then the model predictions of the concentration fluctuations are presented and discussed. Two flow regimes are investigated for the point source simulations: a spatially averaged, horizontally-homogeneous flow, and a flow with inhomogeneities between the canopy elements. Chapter 6 discusses the results of simulations of dispersion of a passive, non-reactive scalar from a continuous ground-level point source in the Mock Urban Setting Trials (MUST) model urban canopy, consisting of a regular array of building-like obstacles (Yee and Biltoft, 2004). The velocity statistics for these simulations came from water-channel experiments of the

full-scale field experiments (Hilderman and Chong, 2007). Chapter 7 provides an overall discussion and summary of the work contained in this thesis, states conclusions, and makes suggestions for further advances.

Chapter 2

Background Theory

In atmospheric science we are interested in how the properties of the air (i.e. momentum and scalar concentration of a tracer in air) change in space and time. The vector field representing momentum varies in both space and time $\mathbf{u}(\mathbf{x}, t)$, as does the scalar field representing the tracer concentration $\phi(\mathbf{x}, t)$. Collectively we refer to these as random fields. Dealing with these four-dimensional random fields is difficult and frequently leads to unsolvable equations. Therefore assumptions and idealisations are frequently made when modelling these fields, to simplify the mathematics.

The remainder of the chapter is structured as follows. Some introductory terms and concepts of turbulence as they pertain to this work are introduced. The equations governing fluid flow are then stated and the mathematical properties of a wall shear layer flow and a canopy flow are discussed. Details of specific flows encountered in this thesis are discussed in later chapters. Next there is a mathematical description of dispersion, diffusion, and mixing. This chapter concludes with a presentation of the analytical PDF transport equation, highlighting the unclosed terms that require modelling.

2.1 Turbulence

Practically all flows of engineering interest involve and are affected by turbulence: through an exchange of momentum, turbulence increases drag on automobiles, planes, and ships; when designing cooling systems for electronics and other applications turbulent heat transfer is used to disperse the heat as rapidly as possible; and turbulent mixing is very important in chemical engineering applications. The study of turbulent flow is therefore important to many aspects of our society.

By examining fluid flow in a pipe, Reynolds (1894) discovered that a flow could be characterised by a single dimensionless number that is constructed from a length scale \mathcal{L} (the pipe diameter or radius), a velocity scale \mathcal{U} (typically the centreline mean velocity), and the kinematic viscosity ν . This is the so-called Reynolds number

$$\text{Re} = \frac{\mathcal{U}\mathcal{L}}{\nu}. \quad (2.1)$$

For the pipe flow examined by Reynolds, the flow was laminar for $\text{Re} \lesssim 2300$ and turbulent for $\text{Re} \gtrsim 4000$. Laminar flow is characterised by parallel streamlines and poor mixing as shown in the left panel of Figure 2.1. In contrast, turbulent flows are very efficient at mixing and are characterised by non-parallel streamlines with a chaotic appearance as shown in the right panel of Figure 2.1. The transition from laminar flow to turbulent flow depends on the setup of the system in question. Near ground, the atmosphere has characteristic scales of $\mathcal{L} \sim 100 \text{ m}$ and $\mathcal{U} \sim 5 \text{ m s}^{-1}$. The kinematic viscosity of air is $\nu_{\text{air}} = 1.5 \times 10^{-5} \text{ m}^2 \text{ s}^{-1}$. The corresponding Reynolds number for surface layer atmospheric flow is thus $\text{Re} \sim 10^7$, well within the turbulent regime.

As the winds blow across the landscape they exchange momentum with




Figure 2.1: Laminar and turbulent flow in a pipe. Reynolds (1894) found that the pipe flow was laminar for $Re \lesssim 2400$ and turbulent for $Re \gtrsim 4000$.

the surface features and vegetation producing turbulent kinetic energy (TKE; denoted k) via shear production and wake production (discussed more in Section 2.4). This energy is pumped into *eddies* that have a characteristic size L , comparable to the characteristic length scale of the flow domain. The term *eddy* is a somewhat vague conceptual term that can be thought of as all turbulent motions (i.e., blobs of vorticity, $\nabla \times \mathbf{u}$) occupying some region of space and time. Richardson (1922) proposed a largely heuristic model to explain the fate of the TKE contained in the large length scales. The model has two principal assumptions: that turbulence is composed of eddies of different sizes (eddy length scale l) which can occupy the same volume, hence a large eddy is coincident with several smaller eddies; and that as long as the Reynolds number for the eddy (Re_l) is large, viscous stresses acting on large eddies will be negligible.

Numerous studies have shown that eddies with a characteristic size $l \approx L$ contain the bulk of the TKE. This is therefore referred to as the *energy-containing range*. The eddies in this range are large, anisotropic, and unstable. Hence they break up and transfer their energy downwards in scale to smaller eddies. Eddies in the energy-containing range have Reynolds number compa-

rable to that of the bulk flow and thus an insignificant amount of energy is lost through viscous dissipation during the transfer. As this process continues, larger eddies breakup into smaller eddies and transfer their energy downwards in scale. At each step the Reynolds number of the eddies decreases, but is still sufficiently large to render the viscous stresses negligible. Physically, the eddies are being stretched and twisted by the velocity field into tubes and sheets, and becoming more localised. Their characteristic size is decreasing and energy is being transferred to smaller length scales.

Eventually the eddies are so small that they are no longer part of the energy-containing range and are part of the *universal equilibrium range*. This range is characterised by the fact that the eddies are isotropic (i.e. retain no memory of the geometry of the mean flow and the largest energy-containing eddies), and again in this range vortex interaction assures the transfer of eddy kinetic energy to smaller and smaller scales until eventually viscosity can act to dissipate the TKE. A turn-over timescale can be defined as the quotient of the characteristic length scale of an eddy divided by the characteristic velocity scale (u_l) for the eddy: $\tau_l = l/u_l$. Eddies in the universal equilibrium range are able to quickly adapt to the energy transfer rate from the energy-containing range since the turn-over time $\tau_l = l/u_l$ is small compared to the turn-over time of the largest eddies, $\tau_L = L/u_L$. This range is divided into two sub-ranges: the *inertial subrange* where energy transfer is dominated by inertia; and the *dissipation range*, where the eddies are very small (characteristic size comparable to the Kolmogorov microscale η) and the Reynolds number of the eddies is sufficiently small such that viscous forces dissipate the TKE. This process is referred to as the *energy cascade* and is illustrated in Figure 2.2.

Justification for the length scales demarcating the energy-containing range, the inertial subrange, and the dissipation range can be found in Pope (2000).

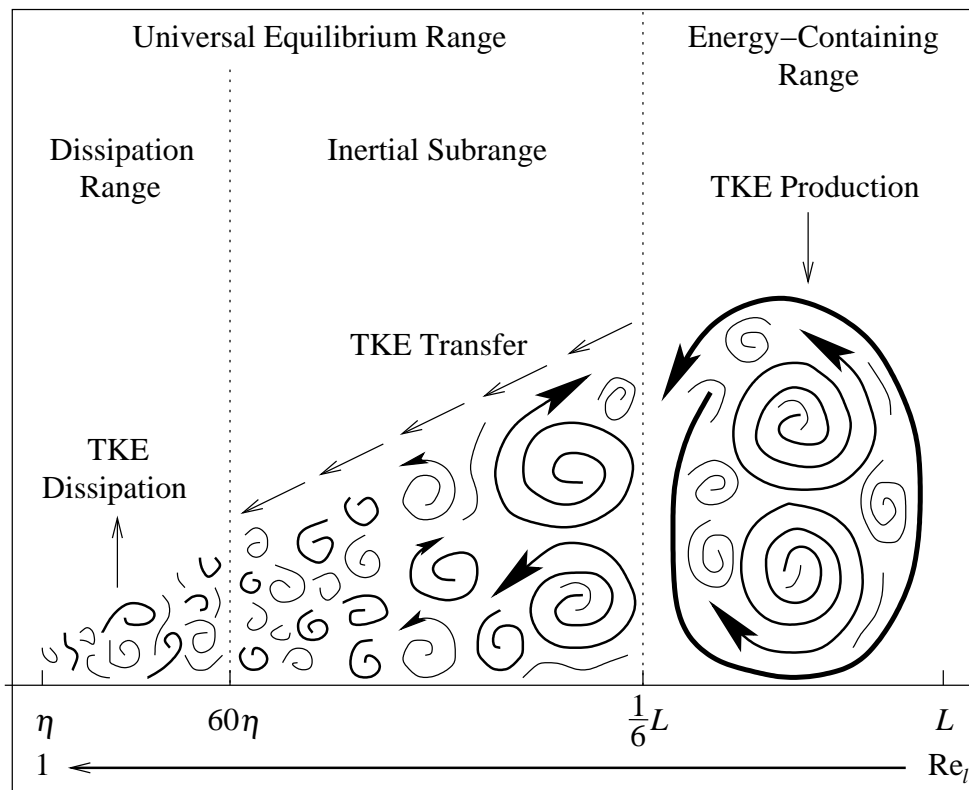


Figure 2.2: Turbulent scales of motion and the energy cascade for high Reynolds number turbulence. TKE enters the cascade in the energy-containing range and is transferred to smaller scales of motion in the inertial subrange. At each step in the cascade viscous losses are negligible until the dissipation range. At these small length scales $Re_l \rightarrow Re_\eta \rightarrow 1$, and viscous forces dissipate the TKE. Note the x -axis is not to scale.

The TKE contained within the eddy is proportional to the characteristic velocity squared, $k \propto u_l^2$. Based on dimensional arguments, an energy transfer rate from one eddy to the next can then be computed as $\mathcal{T} = u_l^2/\tau_l = u_l^3/l$. This TKE transfer rate must be the same across all eddy sizes otherwise there would be a build up of energy at some length scale. What goes in must come

out. This is true across all scales of motion, right down to the smallest scales where the actions of viscosity dissipate the TKE. Consequently, the turbulent kinetic energy dissipation rate must also be equal to the TKE transfer rate, $\varepsilon = \mathcal{T}$. Thus the rate of TKE production at the largest scales sets the rate of TKE transfer across scales at all scales of turbulent motion.

The energy cascade was put on a firmer mathematical footing by Kolmogorov (1941). An English translation can be found in Kolmogorov (1991). Thorough descriptions of turbulence can be found in Pope (2000) and Davidson (2004). Here we only briefly summarise some of the key findings. The length, velocity, and time scales of the smallest, dissipative scales of motion can be determined (by dimensional analysis) from ε and ν . They are

$$\eta \equiv (\nu^3/\varepsilon)^{1/4}, \quad (2.2)$$

$$u_\eta \equiv (\varepsilon\nu)^{1/4}, \quad (2.3)$$

$$\tau_\eta \equiv (\nu/\varepsilon)^{1/2}. \quad (2.4)$$

These are the Kolmogorov scales. Note that $\text{Re}_\eta = 1$ at this scale, consistent with the above assertion that the Reynolds number characterising eddy motion decreases at each step in the energy cascade. Kolmogorov (1941) proposed two similarity hypotheses for high Reynolds number turbulence. They are, loosely stated:

1. In the universal equilibrium range, the statistics of the small-scale motions for a high Reynolds number turbulent flow have a universal form that is determined uniquely by ε and ν .
2. In the inertial subrange, the statistics of the small-scale motions for

a high Reynolds number turbulent flow have a universal form that is determined uniquely by ε and independent of ν (in *effect* this conjecture defines the inertial subrange).

These hypotheses were refined (Kolmogorov, 1962) to account for intermittency in the rate of energy dissipation (i.e., the rate of energy dissipation is not constant, but exhibits spatial inhomogeneities).

A triumph of the Kolmogorov hypotheses is what is known as the “five-thirds law”. The wavenumber corresponding to an eddy with size l is $\kappa = 2\pi/l$. Let $E(\kappa, t)$ represent the energy spectrum function. The quantity $E(\kappa, t) d\kappa$ thus represents the contribution to the TKE of all modes in the range $[\kappa, \kappa + d\kappa]$. By integrating over all wavenumbers the TKE is recovered $k = \int_0^\infty E(\kappa, t) d\kappa$. The energy spectrum in the inertial subrange for homogeneous turbulence (according to Kolmogorov’s analysis) is

$$E(\kappa) = C\varepsilon^{2/3}\kappa^{-5/3}, \quad (2.5)$$

where C is a universal constant. Experimental evidence suggests that $C = 0.5$, with an uncertainty perhaps as small as 0.02 (Sreenivasan, 1995). The Kansas experiments produced a value of 0.52 ± 0.04 with atmospheric measurements (Wyngaard and Coté, 1971). The energy spectrum function in the inertial subrange has a five-thirds dependence on the wavenumber and is proportional only to ε as suggested in the second similarity hypothesis. We shall return to equation (2.5) when developing the micromixing model in Chapter 3.

The energy cascade is not the whole story of turbulence however. Of particular interest to some turbulence theorists are the so-called coherent structures, whose large scale and whose intensity recommend that their contribution to

transport and mixing be accounted explicitly, rather than submerged (and presumably mis-represented) in the statistical description (a visible example from the atmosphere would be the dust devil). They were first visualised in a mixing layer by Brown and Roshko (1974), whose results showed that coherent structures vary discontinuously in space and size (due to merging with neighbouring structures), and that they convect at a nearly constant speed. Their results also showed that the mixing layer is dominated by these structures. A brief review of coherent structures can be found in Lumley and Yaglom (2001) and references therein. The fact that they receive no further mention in this thesis is not to deny their potential importance in the context of concentration fluctuations. However for the present their explicit treatment is rather intractable.

Returning to the statistical theory and the tools provided by Kolmogorov, the second-order Lagrangian structure function,

$$D_{ij}^L(t') \equiv \langle [\mathbf{U}_i(t+t') - \mathbf{U}_i(t)][\mathbf{U}_j(t+t') - \mathbf{U}_j(t)] \rangle, \quad (2.6)$$

provides information regarding the structure of a turbulent flow, and will prove useful in the development of the models in the next chapter. For time offsets t' in the range $\tau_\eta \ll t' \ll T_L$ (where T_L is the Lagrangian integral time scale) the Kolmogorov hypotheses predict

$$D_{ij}^L(t') = \delta_{ij} C_0 \varepsilon t', \quad (2.7)$$

where C_0 is the Kolmogorov constant, and δ_{ij} is the Kronecker delta. To be consistent with the Kolmogorov hypotheses, the Lagrangian-stochastic (LS) model described in Chapter 3 must reproduce equation (2.7). The Kolmogorov constant is universal, although its value has yet to be determined with great

certainty. Values in the range of two to seven are considered acceptable.

2.2 The Governing Equations

An arbitrary volume of fluid (i.e., a fluid element) experiences two types of forces: those which are molecular in origin and arise through various stresses (denoted by $\tau_{ij} = \tau_{ji}$) are called surface forces; while those which act without contact are called body forces. Gravity is the body force of interest for the fluids encountered in this thesis. The gravitational acceleration (and also the field) remains approximately constant, with a value of $g \approx 9.81 \text{ m s}^{-2}$, for several kilometres above the surface. The gravitational potential per unit mass for a constant gravitational field is $\Psi = gz$, and the gravitational force per unit mass is $\mathbf{g} = -\nabla\Psi$. In atmospheric science one must also consider the effects of the fictitious Coriolis force, when sufficiently far from the surface. The Coriolis force is denoted by $f_c\epsilon_{ij3}u_j$, where $f_c = 2\omega \sin\varphi$ is the Coriolis parameter, the scalar ϵ_{ijk} is the alternating unit tensor, ω is the angular velocity of the Earth, and φ is the latitude. Considering the above forces, an application of Newton's second law results in the transport equation for momentum:

$$\frac{\partial u_i}{\partial t} + u_j \frac{\partial u_i}{\partial x_j} = \frac{1}{\rho} \frac{\partial \tau_{ij}}{\partial x_j} - \frac{\partial \Psi}{\partial x_i} - f_c \epsilon_{ij3} u_j, \quad (2.8)$$

where ρ is the mass density. The data sets used to validate the models in this thesis are from wind-tunnel and water-channel experiments. At the relatively small time and velocity scales of these experiments, the contribution of the Coriolis force to the momentum budget can be neglected, and we do so, in all but one instance. However, if the models developed in this thesis are used for full-scale atmospheric applications, then the Coriolis effects must be included.

The continuity or mass-conservation equation is

$$\frac{\partial \rho}{\partial t} + \frac{\partial(\rho u_i)}{\partial x_i} = 0, \quad (2.9)$$

which, assuming the mass density of the fluid to be constant, simplifies to

$$\frac{\partial u_i}{\partial x_i} = \nabla \cdot \mathbf{u} = 0, \quad (2.10)$$

demonstrating that the velocity field of a constant density fluid is *non-divergent*.

The fluids of interest for the work described in this thesis are air and water, both of which are Newtonian fluids. A Newtonian fluid is one whose shear stress is linearly proportional to the velocity gradient perpendicular to the shear plane. The constant of proportionality is called the viscosity (μ). The stress tensor for a Newtonian fluid is

$$\tau_{ij} = -P\delta_{ij} + \mu \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right), \quad (2.11)$$

where P is the pressure. Substituting equation (2.11) into equation (2.8), and using the non-divergence of the velocity field (equation (2.10)) to eliminate some terms, results in the constant density Navier-Stokes equation for a non-divergent flow:

$$\frac{\partial u_i}{\partial t} + u_j \frac{\partial u_i}{\partial x_j} = -\frac{1}{\rho} \frac{\partial P}{\partial x_i} - \frac{\partial \Psi}{\partial x_i} + \nu \frac{\partial^2 u_i}{\partial x_j \partial x_j}, \quad (2.12)$$

where $\nu = \mu/\rho$ is the kinematic viscosity. Henceforth equation (2.12) shall simply be referred to as the Navier-Stokes equation. Following Pope (2000) we define the modified pressure to be $p = P + \rho\Psi$ (herein referred to as pressure), which simplifies the Navier-Stokes equation to

$$\frac{\partial u_i}{\partial t} + u_j \frac{\partial u_i}{\partial x_j} = -\frac{1}{\rho} \frac{\partial p}{\partial x_i} + \nu \frac{\partial^2 u_i}{\partial x_j \partial x_j}. \quad (2.13)$$

From left to right in equation (2.13) the terms represent the storage of momentum, advection of momentum, pressure-gradient forces, and viscous transport of momentum. Note that owing to the (assumed) non-divergence of the velocity field, the advection term may equivalently be written in “transport” or “flux” form as a flux divergence,

$$u_j \frac{\partial u_i}{\partial x_j} = \frac{\partial}{\partial x_j} u_i u_j. \quad (2.14)$$

In constant density flow, there are no buoyancy forces and no relationship between density and pressure (Pope, 2000). To determine what (if any) the restrictions upon p are, take the divergence of the Navier-Stokes (equation (2.13)) and simplify,

$$\left(\frac{\partial}{\partial t} + u_j \frac{\partial}{\partial x_j} - \nu \frac{\partial^2}{\partial x_j \partial x_j} \right) \left(\frac{\partial u_i}{\partial x_i} \right) = -\frac{1}{\rho} \frac{\partial^2 p}{\partial x_i \partial x_i} - \frac{\partial u_i}{\partial x_j} \frac{\partial u_j}{\partial x_i}. \quad (2.15)$$

If the initial and boundary conditions to the above equation were $\partial u_i / \partial x_i = 0$, then the only solution possible is if, and only if, the right-hand side equals zero everywhere (Pope, 2000). Therefore for a non-divergent velocity field to remain non-divergent it must satisfy the following Poisson equation:

$$\frac{\partial^2 p}{\partial x_i \partial x_i} = \nabla^2 p = -\rho \frac{\partial u_i}{\partial x_j} \frac{\partial u_j}{\partial x_i}. \quad (2.16)$$

This equation demonstrates the complex, non-linear coupling between the velocity and pressure fields, a coupling which gives rise to well known pitfalls in the numerical procedure for solving finite difference approximations to these differential equations. Analytical solutions to equation (2.16) can only be found in the simplest case of highly symmetric laminar flow.

The transport equation for a reactive scalar, denoted by $\phi(\mathbf{x}, t)$, is

$$\frac{\partial \phi}{\partial t} + u_j \frac{\partial \phi}{\partial x_j} = \mathcal{D} \frac{\partial^2 \phi}{\partial x_j \partial x_j} + S(\phi(\mathbf{x}, t)), \quad (2.17)$$

where \mathcal{D} is the molecular diffusivity of the scalar in the bulk fluid (here considered to be constant and uniform across the flow), and $S(\phi)$ is the chemical reaction source term. From left to right in equation (2.17), the terms represent the storage of the scalar, advection of the scalar, molecular diffusion (transport) of the scalar, and the chemical reaction source term. We shall leave the chemical source term in the equations for now to highlight one of the advantages of the IECM modelling approach, even though all scalars considered in this work are non-reactive. If the scalar was not reactive then $S(\phi) = 0$ and the conservation equation in Lagrangian form becomes

$$\frac{d\phi}{dt} = \mathcal{D}\nabla^2\phi. \quad (2.18)$$

In this Lagrangian form, we see that diffusion is solely responsible for the evolution of the scalar concentration. Note that since ϕ does not appear in the Navier-Stokes equation it is a passive scalar (viz. it has no effect on the flow).

2.3 The Mean Flow Equations

The problem of turbulence is so difficult that it has evaded an exact mathematical description for hundreds of years. The majority of the progress made in understanding turbulence has been through a statistical approach. Much of what follows in this section comes from Gardiner (1983) and Pope (2000), although many other good books describing these statistics exist.

The descriptive power of the Navier-Stokes equation is also its shortcoming. Equation (2.13) describes all scales of motion from the smallest to the largest. Laminar flow solutions to the Navier-Stokes equation can be found or approxi-

mated numerically. As the Reynolds number increases, the flow becomes more and more turbulent and the number of available degrees of freedom increases. In other words, there are more and more scales of motion that need to be simulated. The number of degrees of freedom is proportional to the ratio of the characteristic sizes of the largest eddies to the smallest eddies, L/η . This ratio increases drastically with the Reynolds number.

With direct numerical simulation (DNS) all scales of motion can be computed. However, this approach is extremely computationally intensive and, given the limitation of computational resources, not yet practical for the atmospheric sciences. Recall that the Reynolds number for a typical atmospheric flow is around $\text{Re} \sim 10^7$. By using $\eta = (\nu^3/\varepsilon)^{1/4}$ (equation (2.2)) and the relation $\varepsilon \sim u_L^3/L$, the Reynolds number dependence between the smallest and largest scales of motion is found to be $\eta/L \sim \text{Re}^{-3/4}$ (Pope, 2000). If we take $L \sim 100$ m, then $\eta \sim 10^{-4}$ m for atmospheric flow. To simulate all scales of motion of this flow with DNS would require computational resources to deal with a range of spatial scales that covers six orders of magnitude. The computational resources required to undertake such a simulation do not yet exist.

Large eddy simulation (LES) reduces the computational complexity (i.e., the number of degrees of freedom) by resolving only the flow dependent, large scales. This approach is also computationally intensive (although not as intensive as DNS) but, due to recent (and expected future) advances in computational power, stands at the forefront of promising modelling efforts. For example, Yue et al. (2007) used an LES model to simulate turbulent flow in a corn canopy. Their model even resolved gusts that penetrated deeply into the

canopy, as did the model of Finnigan et al. (2009) (this paper has beautiful, colour visualisations of eddy evolution). LES models will undoubtedly benefit future dispersion models.

The most severe reduction in the number of degrees of freedom can be achieved by statistical averaging (instead of filtering), leading to the Reynolds-averaged Navier-Stokes (RANS) equations. An instantaneous signal can be thought of as being composed of a mean component and a fluctuating component. For example, a velocity measurement from a turbulent flow can be written as

$$u_i(\mathbf{x}, t) = \langle u_i(\mathbf{x}, t) \rangle + u'_i(\mathbf{x}, t), \quad (2.19)$$

where $\langle u_i \rangle$ is the mean component of the velocity and u'_i is the fluctuating component. Taking the derivative of equation (2.19) and using equation (2.10) gives

$$\frac{\partial}{\partial x_i} (\langle u_i \rangle + u'_i) = \frac{\partial \langle u_i \rangle}{\partial x_i} + \frac{\partial u'_i}{\partial x_i} = 0. \quad (2.20)$$

Taking the mean of equation (2.10) leads to the conclusion that $\partial \langle u_i \rangle / \partial x_i = 0$, and since \mathbf{u} is non-divergent, it follows from this that the velocity fluctuations are also non-divergent, $\partial u'_i / \partial x_i = 0$.

The RANS equation can be derived by utilising Reynolds decomposition, taking the mean of the Navier-Stokes equation (2.13), assuming non-divergence, utilising Reynolds decomposition and some algebra. Doing so results in the RANS equation

$$\frac{\partial \langle u_i \rangle}{\partial t} + \langle u_j \rangle \frac{\partial \langle u_i \rangle}{\partial x_j} = -\frac{1}{\rho} \frac{\partial \langle p \rangle}{\partial x_i} + \nu \frac{\partial^2 \langle u_i \rangle}{\partial x_j \partial x_j} - \underbrace{\frac{\partial}{\partial x_j} \langle u'_i u'_j \rangle}_{\text{new term}}. \quad (2.21)$$

Comparing equations (2.13) and (2.21), we see that the third term on the right-hand side of equation (2.21) does not appear in equation (2.13). This

new term represents the contributions of the Reynolds stresses, which are momentum fluxes. To conceptualise the Reynolds stresses imagine a cube within a fluid flow. Momentum flux from the fluid can interact with the cube, via turbulent eddies, through any one of the six faces. If, for example, it enters (or leaves) only one face, the momentum on that side of the cube will begin to increase (or decrease), causing it to accelerate (or decelerate). Meanwhile, the motion of the other sides of the cube, whose momentum is not changing, remain the same. The end result is deformation of the cube.

The eddy momentum flux density, which has up to three spatial components (i.e., x , y , z directions), can enter any of the three directional faces of the cube; therefore the Reynolds stresses have nine components. It is convenient to represent them in tensor notation,

$$\begin{aligned} \langle u'_i u'_j \rangle = R_{ij} &= \begin{bmatrix} \langle u'_1 u'_1 \rangle & \langle u'_1 u'_2 \rangle & \langle u'_1 u'_3 \rangle \\ \langle u'_2 u'_1 \rangle & \langle u'_2 u'_2 \rangle & \langle u'_2 u'_3 \rangle \\ \langle u'_3 u'_1 \rangle & \langle u'_3 u'_2 \rangle & \langle u'_3 u'_3 \rangle \end{bmatrix} \\ &= \begin{bmatrix} \langle u' u' \rangle & \langle u' v' \rangle & \langle u' w' \rangle \\ \langle v' u' \rangle & \langle v' v' \rangle & \langle v' w' \rangle \\ \langle w' u' \rangle & \langle w' v' \rangle & \langle w' w' \rangle \end{bmatrix}. \end{aligned} \quad (2.22)$$

This is the Reynolds stress tensor. It is symmetric, $\langle u'_i u'_j \rangle = \langle u'_j u'_i \rangle$. The three diagonal elements of the above array are called the *normal* Reynolds stresses which are in fact variances

$$\langle u'_1 u'_1 \rangle = \sigma_u^2, \quad (2.23)$$

$$\langle u'_2 u'_2 \rangle = \sigma_v^2, \quad (2.24)$$

$$\langle u'_3 u'_3 \rangle = \sigma_w^2. \quad (2.25)$$

By definition, the TKE is one half of the trace of the Reynolds stress tensor,

$$k \equiv \frac{1}{2} \langle u'_i u'_i \rangle = \frac{1}{2} (\sigma_u^2 + \sigma_v^2 + \sigma_w^2). \quad (2.26)$$

The six off-diagonal elements of the above array are called the *shear* (or tangential) Reynolds stresses and are simply velocity covariances with the following form:

$$\langle u'_1 u'_2 \rangle = \langle u' v' \rangle, \quad (2.27)$$

$$\langle u'_1 u'_3 \rangle = \langle u' w' \rangle, \quad (2.28)$$

$$\langle u'_2 u'_3 \rangle = \langle v' w' \rangle. \quad (2.29)$$

The process of deriving the mean scalar equation is the same as deriving the RANS equation. We begin by decomposing the scalar into mean and fluctuating component

$$\phi(\mathbf{x}, t) = \langle \phi(\mathbf{x}, t) \rangle + \phi'(\mathbf{x}, t) \quad (2.30)$$

and taking the mean of the conservation equation (2.17). Again non-divergence is assumed and Reynolds decomposition is used on the velocity. After some algebra we arrive at

$$\frac{\partial \langle \phi \rangle}{\partial t} + \langle u_j \rangle \frac{\partial \langle \phi \rangle}{\partial x_j} = \mathcal{D} \frac{\partial^2 \langle \phi \rangle}{\partial x_j \partial x_j} - \underbrace{\frac{\partial \langle u'_j \phi' \rangle}{\partial x_j}}_{\text{new term}} + \langle S(\phi) \rangle. \quad (2.31)$$

Upon comparison of equations (2.17) and (2.31) we see that once again the result of the Reynolds averaging is a new term in $\langle u'_j \phi' \rangle$. This quantity, called the scalar flux (strictly the scalar eddy flux density), represents the flux of the scalar due to the fluctuating velocity field. To simplify the wording in the remainder of this work we hereafter will refer to ϕ as the concentration.

2.4 Atmospheric Flows

We now have sufficient background knowledge to discuss the types of atmospheric flows of interest to us in this thesis, and the associated mathematics. Two types of flows will be considered: wall shear layer flow, and canopy flow. Before these considerations, let us briefly discuss the region of the atmosphere that we are interested in. For a more thorough coverage the reader is directed to Stull (1988) or Garratt (1992).

Figure 2.3 presents a highly simplified view of the lower portion of the atmosphere. The layers shown in the figure are all part of the troposphere which extends from ground level to the tropopause, roughly eleven kilometres above the surface. The *atmospheric boundary layer* (ABL) is the region of the




Figure 2.3: A highly simplified diagram of the atmosphere highlighting the region of interest to this work (in grey). The boundary layer depth is denoted by δ .

atmosphere that is in contact with, and directly influenced by, the surface of the Earth on a timescale of roughly one hour. Its thickness (denoted by δ) is approximately 100 to 3000 metres and varies in both space and time in response to forcings such as: terrain-induced frictional drag, heat transfer, thermal stratification of the layer (i.e., stability), and evaporation and transpiration. It undergoes diurnal variation and is generally thicker during the day than during the night. Radiative cooling at cloud tops or radiative heating at the ground, results in convectively driven turbulence which mixes momentum, moisture, heat, and pollutants throughout the ABL.

The bottom 10% (nominally) of the ABL is called the *atmospheric surface layer* (ASL). The primary feature of the ASL is that the turbulent fluxes and stresses vary by less than 10% relative to their surface values and, as a consequence, the ASL is sometimes referred to as the constant stress layer. This is rather ironic given that the gradients of the turbulent fluxes and stresses are actually steepest in the bottom 10% of the ABL (as consideration of the RANS equations easily shows). The ASL is commonly modelled as being horizontally-homogeneous. In parts of the bottom portion of the ASL there are plants, trees, and buildings. Collectively, if in some sense uniform in their distribution, and of sufficiently high density on the landscape, these objects are referred to as a *canopy*, and the layer they produce is called the *canopy layer*. This layer of strong mean wind shear and mixing has a drastically different state of motion than would occur in the absence of a canopy. Since practically all of land-based biological organisms live in the ABL, our work is focused on this layer. Above the ABL there is a region, known as the *free atmosphere*, that does not respond to forcing caused by terrain features, other

than on longer timescales.

2.4.1 Neutral Wall Shear Layer Flow

This simplest real world turbulent flow that can be imagined is when wind blows over a flat, uniform surface. Friction at the surface extracts momentum from the winds and alters the flow. The situation can be further simplified if we assume that there are no buoyancy forces (i.e., the flow is neutral). If we choose the Cartesian coordinate system such that the x -axis is aligned along the mean wind direction, assume stationarity and horizontal-homogeneity, neglect the effects of viscosity, and include the effects of the Coriolis force, the RANS equation (2.21) becomes

$$\begin{aligned} \frac{\partial \langle u'w' \rangle}{\partial z} &= -\frac{1}{\rho} \frac{\partial \langle p \rangle}{\partial x} + f_c \langle v \rangle \\ &= f_c \left(-\frac{1}{\rho f_c} \frac{\partial \langle p \rangle}{\partial x} + \langle v \rangle \right). \end{aligned} \quad (2.32)$$

At mid-latitudes the Coriolis parameter is on the order of $f_c \sim 10^{-4}$. The first term on the right hand side of equation (2.32) defines the y -component of the *geostrophic wind* and is on the order of $\sim 10 \text{ m s}^{-1}$. In the surface layer (where $\langle v \rangle = 0$, by choice of the orientation of the coordinate system) we therefore have

$$\frac{\partial \langle u'w' \rangle}{\partial z} \approx 10^{-3} \approx 0, \quad (2.33)$$

supporting the above assertion of the surface being a constant stress layer. We note, however, that the stress gradient is at its *maximum* at ground. Nevertheless, the approximation of the surface layer as a constant stress layer is a convenient fiction. Similar arguments can be made for the other stresses.

An important scaling variable can be defined from the magnitude of the Reynolds stresses at the surface. The *friction velocity* is defined in relation to the total vertical flux of horizontal momentum,

$$u_*^2 \equiv (\langle u'w' \rangle^2 + \langle v'w' \rangle^2)^{1/2}. \quad (2.34)$$

With the appropriate choice of coordinates this expression may be simplified to

$$u_*^2 \equiv (\langle u'w' \rangle^2)^{1/2} \quad (2.35)$$

Note that in this work it will be assumed that values of u_* are for the surface unless otherwise stated. When scaled by u_* , velocity statistics from flows of different scales take on a similar form, thus allowing us to simulate a neutral ABL in a wind-tunnel or water-channel with good confidence that the turbulence is being well represented.

The prognostic equation for $\langle u'w' \rangle$ contains triple correlation terms such as $\langle w'u'w' \rangle$. The prognostic equations for these triple correlations in turn have quadruple correlation terms in them. This problem continues as we move to higher orders. This proliferation of unknowns is the turbulence *closure problem*. To get around this closure problem there is no alternative but to introduce a model to provide additional relationships between the unknowns, the simplest such closure being the gradient-diffusion paradigm, known by many other names, including K -theory. It is fundamentally incorrect but in certain circumstances can produce accurate results. One such circumstance is when seeking an expression for the vertical profile of the mean streamwise velocity for a neutral wall shear layer flow under the horizontally homogeneous and stationary assumptions.

Assume that there is a simple relation between the shear stress and vertical gradient of mean streamwise velocity

$$\langle u'w' \rangle = -K_m \frac{\partial \langle u \rangle}{\partial z}, \quad (2.36)$$

where K_m is the eddy viscosity, or eddy diffusivity for momentum. It has the dimensions of a diffusivity (i.e., $\text{m}^2 \text{s}^{-1}$). Heuristically, equation (2.36) is related to how efficiently turbulence transports momentum down the gradient. On dimensional grounds we may assume that K_m is the product of an appropriate velocity fluctuation scale (habitually taken as u_* , although a priori \sqrt{k} would at first appear equally legitimate) and an appropriate length scale. It turns out to “work” if the latter is specified as the distance z to the boundary, such that in short $K_m = k_v u_* z$ where $k_v = 0.4$ is a constant of proportionality known as the von Kármán constant. Thus by equations (2.35) and (2.36) we have

$$u_*^2 = k_v u_* z \frac{\partial \langle u \rangle}{\partial z}, \quad (2.37)$$

which rearranged gives

$$\frac{\partial \langle u \rangle}{\partial z} = \frac{u_*}{k_v z}. \quad (2.38)$$

Integration of this equation with respect to z yields the log-law for streamwise velocity

$$\langle u \rangle = \frac{u_*}{k_v} \ln \left(\frac{z}{z_0} \right). \quad (2.39)$$

The constant of integration z_0 is called the *roughness length*. There is a relationship between the roughness length and the heights of the individual roughness elements (i.e. grasses, trees, buildings, etc.) but they are not equal. The roughness length is smaller than the height of the roughness elements. A typical value for a flat, grass covered prairie is on the order of 0.01 m, for a city

centre $z_0 \sim 1$ m (Stull, 1988). To determine the roughness length in neutral flows one may plot $\ln(z)$ versus $\langle u \rangle$ and extrapolate the plot to where $\langle u \rangle = 0$ on the ordinate axis. The intercept on the abscissa axis allows the roughness length to be determined.

The TKE budget equation for a neutral wall shear flow is

$$\begin{aligned} \frac{\partial k}{\partial t} + \langle u_j \rangle \frac{\partial k}{\partial x_j} = & -\langle u'_i u'_j \rangle \frac{\partial \langle u_i \rangle}{\partial x_j} - \frac{1}{2} \frac{\partial \langle u'_i u'_i u'_j \rangle}{\partial x_j} - \frac{1}{\rho} \frac{\partial \langle u'_j p' \rangle}{\partial x_j} \\ & - \nu \left\langle \frac{\partial u'_i}{\partial x_j} \frac{\partial u'_i}{\partial x_j} \right\rangle + \nu \frac{\partial^2 k}{\partial x_j \partial x_j}, \end{aligned} \quad (2.40)$$

(see Stull, 1988 for a derivation). The terms of the right-hand-side represent: shear production (\mathcal{P}_s), turbulent transport (\mathcal{T}_t), pressure correlation, dissipation ($-\varepsilon$), and viscous diffusion. Under the assumptions stated throughout this section, neglecting the viscous diffusion (suitable for high Reynolds number flows), and with the further assumption of local equilibrium (viz., turbulent transport can be ignored), production and dissipation terms balance and equation (2.40) simplifies to

$$\varepsilon = -\langle u'w' \rangle \frac{\partial \langle u \rangle}{\partial z}. \quad (2.41)$$

Substituting equations (2.35) and (2.38) in the above expression yields the TKE dissipation rate for a neutral wall shear flow, subject to the stated assumptions

$$\varepsilon = u_*^2 \frac{u_*}{k_v z} = \frac{u_*^3}{k_v z}. \quad (2.42)$$

Numerous measurements of wall shear layer flows in the atmosphere, wind-tunnels, and water-channels have shown equation (2.39) and, with slightly lower fidelity, equation (2.42) to be valid, regardless of the assumptions made in their derivation. One example of their validity will be shown in Chapter 4.

Both equations can be derived using Monin-Obukhov similarity theory as in Stull (1988); Garratt (1992); Kaimal and Finnigan (1994). Since it does not add to this thesis, we do not cover it here.

2.4.2 Canopy Flow

As conveniently simple and comprehensible as a wall shear layer is, its occurrence in the natural world is rare. The surface of the earth is for the most part neither flat nor uniform. It is covered with different types of vegetation and buildings of different shapes. These plants and buildings are collectively known as a canopy. In the case of a city, the term mixed canopy is appropriate as over small distances the individual canopy elements change rapidly as we pass over parks, downtown cores, and residential areas. There are however also large areas where the canopy is approximately uniform. Modern suburban areas for example have row upon row of similarly shaped and sized houses, and many forests contain very similarly aged and sized trees. Perhaps the most uniform canopy can be found in crops, a monoculture of equally aged and uniformly spaced plants, such as a cornfield.

Regardless of their composition canopies share a common feature: through momentum exchange they strongly alter the flow. Canopies can be characterised in several ways. The frontal area index is calculated as the ratio of the frontal (streamwise) area of a canopy element to the lot area (the area on which a single obstacle sits in the array), $\lambda_f \equiv A_f/A_L$. The plan area index is defined as $\lambda_p \equiv A_p/A_L$, where A_p is the plan (floor) area of a canopy element. Finally, the element area density (denoted by a), or surface area density, is calculated as the frontal area per unit volume, where the volume is calculated

as the product of the lot area and the canopy height.

Provided the canopy is sufficiently dense, equations (2.39) and (2.42) can still be used in the region above the canopy, with a minor modification. In this situation the tops of the canopy elements act as one and displace the flow upwards. The minor modification to the equations is the addition of a displacement height d ,

$$\langle u \rangle = \frac{u_*}{k_v} \ln \left(\frac{z - d}{z_0} \right), \quad (2.43)$$

$$\varepsilon = \frac{u_*^3}{k_v(z - d)}. \quad (2.44)$$

The displacement height is generally around seventy-five percent of the height of the canopy elements, $d \sim 0.75h_c$ (Kaimal and Finnigan, 1994). It can be determined in the neutral case by plotting $\ln(z - d)$ versus $\langle u \rangle$ over a range of different values for d . When d is correctly set the line should have no curvature, a slope of u_*/k_v , and have an intercept with the abscissa axis that allows z_0 to be inferred. When utilising the above equations it is important to note that the values for u_* and z_0 are for the displacement height, not the ground level.

Analytical expressions for the mean streamwise velocity and the TKE dissipation are harder to determine inside the canopy. Given present mathematical tools an exact treatment of the problem is impossible. Real world canopies are simply too complex, with too many length scales. Consider the fractal nature of trees, or the architectural subtleties of many buildings. To approach this problem requires simplifications and assumptions. Given the complex nature of canopies, the governing equations must be spatially-averaged (represented by the square brackets in this section only), shown here for $\langle \mathbf{u} \rangle$,

$$[\langle \mathbf{u} \rangle](\mathbf{x}, t) = \frac{1}{\mathcal{V}} \iiint_{\mathcal{V}} \langle \mathbf{u} \rangle(\mathbf{x} + \mathbf{r}, t) d\mathbf{r}. \quad (2.45)$$

The averaging volume is denoted by \mathcal{V} . The departure from the spatially-averaged quantity is denoted by a double prime

$$\langle u \rangle = [\langle u \rangle] + \langle u \rangle'', \quad (2.46)$$

and satisfies $[\langle u \rangle''] = 0$.

The spatially-averaged transport equation for mean momentum within a canopy is

$$\begin{aligned} \frac{\partial[\langle u_i \rangle]}{\partial t} + [\langle u_j \rangle] \frac{\partial[\langle u_i \rangle]}{\partial x_j} = & -\frac{1}{\rho} \frac{\partial[\langle p \rangle]}{\partial x_i} - \frac{\partial[\langle u'_i u'_j \rangle]}{\partial x_j} - \frac{\partial[\langle u_i \rangle'' \langle u_j \rangle'']}{\partial x_j} \\ & + \nu \frac{\partial^2[\langle u_i \rangle]}{\partial x_j \partial x_j} + f_{F_i} + f_{V_i}, \end{aligned} \quad (2.47)$$

where f_{F_i} and f_{V_i} are the form and viscous drag force vectors respectively (Raupach et al., 1986) and formally represent the flux of momentum to solid surfaces within the averaging volume. The third term on the right-hand-side is new, and represents the dispersive flux.

Within the canopy, momentum absorption rapidly attenuates the mean streamwise velocity. A popular model for the in-canopy mean streamwise velocity can be found, once again, via K -theory. Consider a canopy composed of rigid elements of height h_c , and assume that the turbulence is stationary and the flow is neutral with no streamwise evolution within or above the canopy. To a good approximation the spatially averaged velocity statistics within the canopy are horizontally-homogeneous (Raupach et al., 1986). Under these assumptions, the streamwise component of equation (2.47) is

$$\frac{\partial[\langle u'w' \rangle]}{\partial z} + \frac{\partial[\langle u \rangle'' \langle w \rangle'']}{\partial z} - \nu \frac{\partial^2[\langle u \rangle]}{\partial z^2} = f_{F_x} + f_{V_x}. \quad (2.48)$$

This equation can be further simplified by considering only the dominant

terms,

$$\frac{\partial[\langle u'w' \rangle]}{\partial z} = -\frac{1}{2}C_D a[\langle u \rangle]^2, \quad (2.49)$$

where the right-hand-side is the parametrisation of the form drag with C_D being the drag coefficient. Using equation (2.36) to parametrise the Reynolds stress term, and the Prandtl-von Kármán mixing-length theory to specify the form of the eddy diffusivity,

$$K_m = \ell^2 \frac{d[\langle u \rangle]}{dz} \quad (2.50)$$

(where the mixing length ℓ is assumed to be constant within the canopy), and solving gives

$$[\langle u \rangle](z) = \langle u \rangle_c \exp\left(-\beta \left(1 - \frac{z}{h_c}\right)\right). \quad (2.51)$$

Here, $\langle u \rangle_c = \langle u \rangle(h_c)$ is the mean streamwise velocity at the canopy height. The extinction parameter β can be connected with the eddy diffusivity of the K -theory model but is frequently left as a fitting parameter. Even though this equation violates the no-slip condition (i.e., $\langle u \rangle(0) \neq 0$), it has been shown to capture the essence of the mean streamwise velocity within a canopy reasonably well. Equations (2.43) and (2.51) can be combined (with small modifications to ensure that the velocity is continuous across the $z = h_c$ interface) to parametrise the mean streamwise velocity at any location within or above the canopy,

$$[\langle u \rangle] = \langle u \rangle_c \exp\left(-\beta \left(1 - \frac{z}{h_c}\right)\right), \quad \text{for } z \leq h_c \quad (2.52)$$

$$\langle u \rangle = \langle u \rangle_c + \frac{u_*}{k_v} \ln\left(\frac{z/h_c - d/h_c}{1 - d/h_c}\right), \quad \text{for } z > h_c \quad (2.53)$$

where u_* is calculated based on the shear stresses at the canopy top.

As Raupach et al. (1986) displayed, the volume averaged TKE budget equation under the above assumptions is

$$\begin{aligned} \frac{\partial [k]}{\partial t} = & -[\langle u'w' \rangle] \frac{\partial [\langle u \rangle]}{\partial z} - \left[\langle u'_i u'_j \rangle'' \frac{\partial \langle u_i \rangle''}{\partial x_j} \right] - \frac{1}{2} \frac{\partial [\langle w' u'_i u'_i \rangle]}{\partial z} \\ & - \frac{1}{2} \frac{\partial [\langle w \rangle'' \langle u'_i u'_i \rangle'']}{\partial z} - \frac{\partial [\langle p' w' \rangle]}{\partial z} + \nu \frac{\partial^2 [k]}{\partial z^2} - [\varepsilon], \end{aligned} \quad (2.54)$$

The rigidity of the canopy elements allows one to neglect the contributions of canopy waving to the TKE budget. On the right-hand side of equation (2.54) the terms, from left to right, represent: shear production (\mathcal{P}_s), wake production (\mathcal{P}_w) above and around the canopy elements, turbulent transport of TKE (\mathcal{T}_t), dispersive transport (\mathcal{T}_d), pressure transport (\mathcal{T}_p), molecular transport (\mathcal{T}_m), and TKE dissipation. Raupach et al. (1986) showed how to evaluate the wake production in terms of measurable quantities. By assuming that dispersive and molecular contributions to the shear stress are negligible, the wake production term can be written as

$$\mathcal{P}_w = [\langle u \rangle] \frac{\partial [\langle u'w' \rangle]}{\partial z}. \quad (2.55)$$

For the Tombstone Canopy (Raupach et al., 1986; Coppin et al., 1986; Legg et al., 1986; Raupach et al., 1987), it was shown that the TKE dissipation rate can be calculated as

$$\begin{aligned} [\varepsilon] = & [\langle u'w' \rangle] \frac{\partial [\langle u \rangle]}{\partial z} + \left[\langle u'_i u'_j \rangle'' \frac{\partial \langle u_i \rangle''}{\partial x_j} \right] + \frac{1}{2} \frac{\partial [\langle w' u'_i u'_i \rangle]}{\partial z} \\ = & \mathcal{P}_s + \mathcal{P}_w + \mathcal{T}_t, \end{aligned} \quad (2.56)$$

by assuming that the production of TKE equaled the dissipation of TKE, and that most of the transport terms were negligible. The shear production and wake production are in closed form and are readily calculable. The turbulent transport term is unclosed and needs to be modelled. The particulars of the Tombstone Canopy and flow will be more thoroughly described in Chapter 5.

2.5 Turbulent Dispersion and Mixing

When a scalar is released into a flow it spreads out and mixes with the ambient environment. We are familiar with this phenomenon by observing pollutants emitted from smokestacks, exhaust from automobile tailpipes, and smoke from the end of a cigarette, to give some examples. Mixing of the scalar into the ambient environment is caused either by the effects of the turbulent flow or the effects of molecular diffusion. The left panel of Figure 2.4 shows alternating layers of pure air (white) and a scalar (black). This setup will be taken as the initial conditions for the discussion below.

The ratio of the kinematic viscosity to the molecular diffusivity is called the Schmidt number

$$\text{Sc} \equiv \frac{\nu}{\mathcal{D}}. \quad (2.57)$$

The Schmidt number relates the smallest length scale of the scalar field to the Kolmogorov scale as

$$\lambda_B \equiv \text{Sc}^{-1/2} \eta, \quad (2.58)$$

for fluids with $\text{Sc} \geq 1$. This is called the Batchelor scale and it represents the smallest scale where turbulent diffusion is balanced by molecular mixing (see Fox (2003)). For fluids with $\text{Sc} \ll 1$, a different diffusion scale applies,

$$\lambda_d \equiv \text{Sc}^{-3/4} \eta. \quad (2.59)$$

For gases $\text{Sc} \approx 1$ and so $\lambda_B \approx \eta$. For scalar eddies much larger than the Batchelor scale, molecular diffusion can be ignored and the transport equation for the passive, non-reactive scalar becomes

$$\frac{\partial \phi}{\partial t} + u_j \frac{\partial \phi}{\partial x_j} = \frac{d\phi}{dt} = 0. \quad (2.60)$$

This shows that at these larger scales the scalar will be redistributed but not mixed via molecular diffusion towards the mean concentration. This is illustrated in the centre panel of Figure 2.4 . After larger-scale mixing, the regions of pure scalar and pure air have been stretched and deformed but they have not mixed to produce regions of intermediate scalar concentration. When the scalar eddies are comparable in size to the Batchelor scale, molecular diffusion cannot be neglected and the mixing towards the mean concentration occurs. This is shown in the right panel of Figure 2.4 by the grey scale shading. There are no longer areas of completely pure air or completely pure scalar. Diffusion has resulted in regions with scalar concentrations somewhere between the two extremes. Given enough time the entire box would reach a more or less uniform scalar concentration.




Figure 2.4: This figure illustrates two mixing scenarios for the initial slab conditions shown in (a). The first is when the scalar eddies are much larger than the Batchelor scale and thus molecular diffusion is ignored and the scalar is only redistributed (b). The result is regions of pure air and regions of pure scalar. When the scalar eddies are comparable in size to the Batchelor scale, molecular diffusion cannot be ignored which results in the scalar mixing with the pure air to give intermediate concentrations (c). As time progresses molecular diffusion makes it increasingly difficult to find areas of pure air or pure scalar. The eventual end state in this scenario is a more or less uniform scalar concentration. Figures after Fox (2003).

The relative sizes of the contaminant plume and the eddies have a marked

effect on the development of a plume. Consider the train derailment shown in Figure 2.5. One of the cars has ruptured and is releasing a contaminant into the atmosphere. The figure shows the *instantaneous plume* which is a snapshot of the plume. The *average plume* is the area that the instantaneous plume covers in some large amount of time. To view the average plume one could take a long, open exposure photograph of the instantaneous plume.

When the plume is small relative to the characteristic eddy size it is simply pushed back and forth by the flow in a process called *meandering*. A detector on the centreline of the time-averaged plume would detect concentration fluctuations as the plume sweeps back and forth over it. When the plume grows to a size comparable to the characteristic eddies it incorporates uncontaminated air into it in a process called *entrainment*. Some meandering of the plume still occurs but its effects are becoming less significant. Concentration fluctuations observed by a detector in this region of the plume are a result of alternating regions of relatively pure, freshly entrained air, and impure contaminated material passing the detector. Other than at the source of a very small release, the plume will always be larger than the Batchelor scale. Molecular diffusion will become important when material lines in the plume have been folded and stretched by the turbulent motions as to be sufficiently thin (comparable to Batchelor scale) to allow molecular diffusion to act. This generally occurs when the plume is much larger than the characteristic eddies. At this point concentration fluctuations are generated principally through dissipation and internal fluctuations. Small-scale structure in the plume is fine enough to allow molecular diffusion to be significant, resulting in the dissipation of the scalar fluctuations.




Figure 2.5: The three stages of plume development. Initially the plume is smaller than the characteristic eddies and it meanders back and forth. When the plume is comparable in size to the characteristic eddies entrainment incorporates uncontaminated air into contaminated air. Once the plume is much larger than the characteristic eddies, molecular diffusion acts to homogenise the fluctuations in scalar concentration.

There are two views on dispersion: *absolute dispersion*, where we are interested in the properties of the average plume; and *relative dispersion*, where we are interested the statistical properties of the plume relative to its instantaneous centreline (i.e., the instantaneous plume). In both cases mathematical analysis can be applied to determine the rate of growth of the plume. We begin first with absolute dispersion. This problem, first given a mathematical foundation by Taylor (1921), considers the average distance that a single particle moves from its source in some amount of time as a result of turbulent motion.

Let us assume that a passive, non-reactive scalar is being released into homogeneous, isotropic, and stationary turbulence. For simplicity we only consider motion on the z -axis. At $t = 0$ a tracer particle* is released from $Z(0) = 0$ with a vertical velocity of $W(t)$. Lagrangian notation is used since we are following the particle. The particle's position at time t is computed as

$$Z(t) = \int_0^t W(t') dt'. \quad (2.61)$$

If we considered an ensemble of tracer particles, then the spread of the plume could be defined as

$$\sigma_z^2 \equiv \langle Z^2 \rangle, \quad (2.62)$$

and the time evolution of this spread calculated as

$$\frac{d\sigma_z^2}{dt} = 2 \left\langle Z \frac{dZ}{dt} \right\rangle, \quad (2.63)$$

where we have made use of the property that the operation of ensemble averaging commutes with differentiation. Utilising equation (2.61), the preceding

*A tracer particle is one which follows the flow exactly; it does not have buoyancy and does not undergo gravitational settling.

equation can be expressed as

$$\begin{aligned}\frac{d\sigma_z^2}{dt} &= 2 \left\langle \int_0^t W(t') dt' \frac{d}{dt} \int_0^t W(t'') dt'' \right\rangle \\ &= 2 \left\langle W(t) \int_0^t W(t') dt' \right\rangle,\end{aligned}\tag{2.64}$$

where t' and t'' are dummy variables for time. Changing notation, let $\tau = t - t'$ be a time lag and define

$$R(\tau) \equiv \frac{\langle W(t)W(t - \tau) \rangle}{\sigma_w^2}\tag{2.65}$$

as the Lagrangian velocity auto-correlation function, which will be discussed below. Note that since the turbulence is stationary $R(\tau)$ is not a function of time. The vertical velocity variance is denoted by σ_w^2 . We now have

$$\begin{aligned}\frac{d\sigma_z^2}{dt} &= 2 \int_0^t \langle W(t)W(t - \tau) \rangle d\tau \\ &= 2 \int_0^t \sigma_w^2 R(\tau) d\tau\end{aligned}\tag{2.66}$$

which brings us to Taylor's result,

$$\frac{d\sigma_z^2}{dt} = 2\sigma_w^2 \int_0^t R(\tau) d\tau.\tag{2.67}$$

To proceed further, knowledge of the Lagrangian velocity auto-correlation function is needed. The *Lagrangian integral timescale* is defined as the area beneath the Lagrangian velocity auto-correlation function

$$T_L = \int_0^\infty R(\tau) d\tau.\tag{2.68}$$

It can be thought of as the characteristic persistence time of a vertical velocity fluctuation. At very small time lags ($\tau \simeq 0$), we would expect $W(t) \simeq W(t - \tau)$ and thus $R \simeq 1$ as the two velocities are highly correlated. Hence, upon

integration of equation (2.67), the evolution of σ_z^2 in the *near field* (i.e., close to the source) is

$$\sigma_z^2 = \sigma_w^2 t^2. \quad (2.69)$$

Alternatively, in the *far field* ($\tau \gg T_L$) we expect the plume to grow as

$$\sigma_z^2 = 2\sigma_w^2 T_L t. \quad (2.70)$$

This result is consistent with pure Fickian diffusion which predicts $\sigma_z \propto \sqrt{t}$.

Now let us consider relative dispersion. To find the growth rates of the instantaneous plume relative to its centreline, let us consider the rate of separation of two tracer particles released simultaneously into homogeneous, isotropic, and stationary turbulence as Richardson (1926) did. Let Z_c represent position of the centre of mass for the two particles and W_c represent the vertical velocity of the centre of mass. The position and velocity of a particle relative to the centre of mass are respectively $Z_r = Z - Z_c$ and $W_r = W - W_c$. Modifying equation (2.62) to deal with relative positions we have

$$\sigma_r^2 \equiv \langle Z_r^2 \rangle. \quad (2.71)$$

At small times, close to the source, relative dispersion proceeds by diffusion only and the evolution equation for σ_r^2 can be written as

$$\frac{d\sigma_r^2}{dt} = 2\mathcal{D}, \quad (2.72)$$

the integration of which yields $\sigma_r^2 = \sigma_0^2 + 2\mathcal{D}t$, where the initial source distribution σ_0 accounts for a finite sized source.

The above result for plume growth should be valid to times that are small with respect to the Kolmogorov time scale τ_η . Now consider the situation when relative dispersion is accomplished by convection. In intermediate times

when $\eta \ll \sigma_r \ll L$ (i.e., the inertial subrange), eddies with a characteristic size comparable to the size of the instantaneous plume ($l \sim \sigma_r$) are responsible for plume growth. Recall from Section 2.1 that the energy transfer rate down the cascade to smaller length scales was shown to be $\mathcal{T} \sim \varepsilon \sim u_l^3/l$, which is rearranged to give

$$u_l \sim (\varepsilon l)^{1/3} \sim (\varepsilon \sigma_r)^{1/3}. \quad (2.73)$$

Since only eddies with a characteristic size comparable to the size of the instantaneous plume are responsible for plume growth we have

$$\frac{d\sigma_r}{dt} \sim u_l \sim (\varepsilon \sigma_r)^{1/3}. \quad (2.74)$$

This is rearranged and integrated to give

$$\sigma_r^2 - \sigma_0^2 = C_r \varepsilon t^3, \quad (2.75)$$

where C_r is a constant of proportionality called the Richardson constant and σ_0 is a constant of integration which accounts for the finite source size. Upon defining $t_0 = (\sigma_0^2/C_r \varepsilon)^{1/3}$ we arrive at the Richardson Law

$$\sigma_r^2 = C_r \varepsilon (t + t_0)^3. \quad (2.76)$$

The Richardson Law is sometimes called the Richardson-Obukhov Law, and the Richardson constant the Richardson-Obukhov constant, in honour of the man who gave the Richardson Law a solid mathematical footing, Obukhov (1941a,b). The details of Obukhov's work can be found in English in Monin and Yaglom (1975).

The value of the Richardson constant is not yet known with great certainty. Values found in the literature are in the range $0.1 \lesssim C_r \lesssim 6$, depending on

the type of model used to estimate it. Franzese and Cassiani (2007) and references therein provide many examples of Richardson constant value estimates from various models. In addition, they relate the Richardson constant to the Kolmogorov constant for homogeneous isotropic turbulence. They find that $C_r \approx C_0/11$. This will be a useful guide for tuning the micromixing model in later chapters.

At large times the centre of mass of the plume is no longer subject to displacement as it has grown large relative to the eddies. In this time regime the evolution of the relative plume is the same as the evolution of the absolute plume,

$$\frac{d\sigma_r^2}{dt} = 2\sigma_{w_r}^2 T_r, \quad (2.77)$$

which upon integration yields

$$\sigma_r^2 = \sigma_0^2 + 2\sigma_{w_r}^2 T_r t, \quad (2.78)$$

where $\sigma_{w_r}^2 = \langle W_r^2 \rangle$ is the variance of the relative vertical velocity. Again there is Fickian diffusion type behaviour. Table 2.1 summarises the growth rates of the absolute plume and the instantaneous plume at various times.

Time	Absolute Dispersion (Average Plume)	Relative Dispersion (Instantaneous Plume)
$t \rightarrow 0$	$\sigma_z \propto t$	$\sigma_r \propto t^{1/2}$
$0 \ll t \ll T_L$		$\sigma_r \propto t^{3/2}$
$t \gg T_L$	$\sigma_z \propto t^{1/2}$	$\sigma_r \propto t^{1/2}$

Table 2.1: Summary of the growth rate of a plume for absolute dispersion and relative dispersion.

2.6 Probability Density Function Modelling

The one-point, one-time, joint velocity and concentration probability density function ($f_{\mathbf{u}\phi} = f_{\mathbf{u}\phi}(\mathbf{v}, \psi; \mathbf{x}, t)$) can provide all one-point and one-time statistical information about the underlying velocity and concentration fields. The PDF is a density with respect to the quantities to the left of the semi-colon, and a function with respect to the quantities to the right of the semi-colon. To the left of the semi-colon there are: \mathbf{v} , the sample space (i.e., dummy) variable for \mathbf{u} ; and ψ , the sample space variable for ϕ . Arguments of $f_{\mathbf{u}\phi}$ will be omitted (but implied) for notational simplicity. As the velocity and concentration fields have associated transport equations, so too does the joint velocity and concentration PDF,

$$\frac{\partial f_{\mathbf{u}\phi}}{\partial t} + v_i \frac{\partial f_{\mathbf{u}\phi}}{\partial x_i} = -\frac{\partial}{\partial v_i} [\langle A_i | \mathbf{v}, \psi \rangle f_{\mathbf{u}\phi}] - \frac{\partial}{\partial \psi} [\langle \Theta | \mathbf{v}, \psi \rangle f_{\mathbf{u}\phi}]. \quad (2.79)$$

There are multiple ways to derive this equation, but all are quite lengthy, and would not add much to the present discussion. The interested reader is directed to Pope (1985, 2000) or Fox (2003) for its derivation. The coefficient A_i comes from the Navier-Stokes equation (2.13),

$$A_i = \nu \frac{\partial^2 u_i}{\partial x_j \partial x_j} - \frac{1}{\rho} \frac{\partial p}{\partial x_i}. \quad (2.80)$$

The coefficient Θ comes from the transport equation for a scalar (2.17) and is a combination of molecular diffusion and a chemical source term which represents chemical reactions,

$$\Theta = \mathcal{D} \frac{\partial^2 \phi}{\partial x_j \partial x_j} + S(\phi). \quad (2.81)$$

Equation (2.79) has two unclosed terms in it: the conditional mean accel-

eration

$$\langle A_i | \mathbf{v}, \psi \rangle = \left\langle \nu \frac{\partial^2 u_i}{\partial x_j \partial x_j} - \frac{1}{\rho} \frac{\partial p}{\partial x_i} \middle| \mathbf{v}, \psi \right\rangle, \quad (2.82)$$

which corresponds to viscous dissipation and pressure fluctuations; and the conditional mean diffusion-reaction

$$\langle \Theta | \mathbf{v}, \psi \rangle = \left\langle \mathcal{D} \frac{\partial^2 \phi}{\partial x_j \partial x_j} + S(\phi) \middle| \mathbf{v}, \psi \right\rangle = \left\langle \mathcal{D} \frac{\partial^2 \phi}{\partial x_j \partial x_j} \middle| \mathbf{v}, \psi \right\rangle + S(\psi), \quad (2.83)$$

which corresponds to molecular diffusion and chemical reactions (Fox, 2003). In equations (2.82) and (2.83) the means are conditioned upon the velocity and the concentration respectively. To better understand terms such as these consider the question, “what is the mean acceleration (diffusion-reaction) of a fluid element given its velocity (concentration)?” Fluid elements with a different velocity and/or concentration will likely (but not necessarily) have different mean accelerations and diffusion-reactions.

Note that the chemical source term appears in closed form in equation (2.83). Exact representation of $S(\phi)$ is a major advantage of the IEM and the IECM approach as the chemical source term is often a complicated, non-linear function. Since this project pertains to non-reactive scalars, we now set $S(\phi) = 0$, and the conditional mean diffusion-reaction is now simply the conditional mean diffusion.

From equations (2.79), (2.82), and (2.83) it is seen that the evolution of the joint PDF of velocity and concentration is driven by transport in: real space through v_i ; velocity phase space through $\langle A_i | \mathbf{v}, \psi \rangle$; and in composition phase space through $\langle \Theta | \mathbf{v}, \psi \rangle$ (Fox, 2003). If we wish to simulate the evolution of the joint velocity-concentration PDF then models must be prescribed for the conditional mean acceleration and the conditional mean diffusion-reaction. In

the next chapter, a single-particle LS scheme will be coupled to the IECM micromixing model to close these terms.

Chapter 3

Closure and Numerical Modelling

Chapter 2 concluded with the presentation of the unclosed terms in the transport equation of the one-point, one-time joint velocity and concentration PDF. In this chapter, models to close these conditional fluxes will be presented and their implementation discussed.

3.1 Lagrangian-Stochastic Trajectory Modelling

To close the conditional mean acceleration (equation (2.82)) a LS trajectory model is employed. Under the assumption that the velocity and position of a fluid element are jointly a continuous Markov process, and assuming validity of the Kolmogorov similarity relationship for the Lagrangian second-order structure function (equation (2.7)), it follows that the motion of N independent tracer particles is governed by the following stochastic equations:

$$dU'_i = a_i(\mathbf{X}, \mathbf{U}', t)dt + b_{ij}(\mathbf{X}, \mathbf{U}', t)d\xi_j(t), \quad (3.1)$$

$$dX_i = (\langle u_i \rangle + U'_i) dt, \quad (3.2)$$

where U'_i is the Lagrangian velocity fluctuation relative to the Eulerian mean, dt is a small timestep and $d\xi_j(t)$ represents an incremental Wiener process with

zero mean and variance dt . On the right-hand-side of equation (3.1) we have the deterministic term $a_i dt$, and the stochastic diffusion term $b_{ij} d\xi_j$. Thomson (1987) considered five criteria for LS models of particle trajectories which aid in the determination of a_i and b_{ij} :

1. “The well-mixed condition” (WMC): if the tracer particles are initially well-mixed in position and velocity space, will they remain so?
2. “The small-time behaviour of the velocity distribution of particles from a point source”: is it correct?
3. “The requirement of compatibility with the Eulerian equations”: are the true Eulerian equations and those derived from the Lagrangian model compatible?
4. “Forward and reverse dispersion”: is there consistency between the forward and reverse formulations of the dispersion?
5. “The small-timescale limit”: as the Lagrangian timescale tends towards zero does the model reduce to a diffusion equation?

He then went on to show that criteria 2-4 are mathematically equivalent to criterion 1, and that criterion 5 is a weaker condition than criterion 1. Let $g_t(\mathbf{x}, \mathbf{u}', t)$ represent the density function of the $(\mathbf{x}, \mathbf{u}')$ phase space distribution of the tracer particles and $g_a(\mathbf{x}, \mathbf{u}', t)$ represent the density function of the distribution of fluid elements. If the tracer is *well-mixed*, and there are no sources or sinks for the tracer, then by definition g_t is proportional to g_a , or alternatively, the ratio g_t/g_a is constant. Thomson (1987) showed that satisfaction of the WMC ensures satisfaction of the other criteria.

Solutions for the deterministic coefficient a_i will be discussed briefly, but first we discuss the form of the stochastic diffusion coefficient. An expression for b_{ij} can be found by assuring consistency with Kolmogorov's theory of local isotropy for the Lagrangian structure function (equation (2.7)) as outlined in Thomson (1987). For $dt \ll T_L$, the diffusion process in equation (3.1) produces the following expression for the second-order Lagrangian structure function,

$$D_{ik}^L = \langle b_{ij}b_{jk} \rangle dt. \quad (3.3)$$

Equating the above equation with equation (2.7) gives

$$\langle b_{ij}b_{jk} \rangle = \delta_{ij}C_0\varepsilon. \quad (3.4)$$

Consistency of the LS model and Kolmogorov's theory of local isotropy is assured by the specification

$$b_{ij} = \delta_{ij}(C_0\varepsilon)^{1/2}. \quad (3.5)$$

Solutions for the deterministic coefficient a_i are found by imposing the WMC. Equations (3.1) & (3.2) imply that g_a (or g_t , since the two are proportional in the well-mixed state) satisfies the following Fokker-Planck equation:

$$\frac{\partial g_a}{\partial t} + \frac{\partial}{\partial x_i}[\langle v_i \rangle + V'_i]g_a = -\frac{\partial}{\partial V'_i}(a_i g_a) + \frac{\partial^2}{\partial V'_i \partial V'_k} \left(\frac{1}{2} b_{ij} b_{jk} g_a \right). \quad (3.6)$$

By isolating $a_i g_a$ from the above equation we arrive at a mathematical expression of the well-mixed condition,

$$a_i g_a = \frac{\partial}{\partial V'_k} \left(\frac{1}{2} b_{ij} b_{jk} g_a \right) + \Phi_i, \quad (3.7)$$

$$\frac{\partial \Phi_i}{\partial V'_i} = -\frac{\partial g_a}{\partial t} - \frac{\partial}{\partial x_i}[\langle v_i \rangle + V'_i]g_a, \quad (3.8)$$

with $\Phi(\mathbf{x}, \mathbf{U}', t) \rightarrow 0$ as $\mathbf{U}' \rightarrow \infty$. For more details on this derivation the interested reader is directed to Thomson (1987) and Rodean (1996). Before continuing further in the determination of the deterministic coefficient in equation (3.1), we must first specify the form of the PDF of the Eulerian velocity fluctuations. For many cases of practical importance and interest a Gaussian form is sufficient:

$$g_a(\mathbf{x}, \mathbf{u}') = \frac{[\det(R^{-1})]^{1/2}}{(2\pi)^{3/2}} \exp\left(-\frac{1}{2}u'_i R_{ij}^{-1} u'_j\right), \quad (3.9)$$

with $R_{ij}^{-1}(\mathbf{x})$ being the inverse Reynolds stress tensor, where stationarity has been assumed (but is not necessary), therefore the Reynolds stresses are independent of time. Even under these conditions the WMC (equations (3.7) and (3.8)) cannot produce a unique solution for a_i for two or three-dimensional turbulence. In fact, a unique solution is only possible in the one-dimensional case.

The “simplest” solution for a_i in three-dimensions for Gaussian turbulence, which is attributed to Thomson (1987), is

$$a_i = T_i^{(0)} + T_{ij}^{(1)} U'_j + T_{ijk}^{(2)} U'_j U'_k, \quad (3.10)$$

where

$$T_i^{(0)} \equiv \frac{1}{2} \frac{\partial R_{i\ell}}{\partial x_\ell}, \quad (3.11)$$

$$\begin{aligned} T_{ij}^{(1)} &\equiv -\frac{1}{2}(C_0\varepsilon)R_{ij}^{-1} + \frac{1}{2}R_{j\ell}^{-1} \frac{\partial R_{i\ell}}{\partial x_k} \langle u_k \rangle, \\ &= -\frac{1}{2}(C_0\varepsilon)R_{ij}^{-1} + T_{ijk}^{(2)} \langle u_k \rangle, \end{aligned} \quad (3.12)$$

$$T_{ijk}^{(2)} \equiv \frac{1}{2}R_{j\ell}^{-1} \frac{\partial R_{i\ell}}{\partial x_k}, \quad (3.13)$$

where $\langle u_k \rangle$ is the mean Eulerian velocity and R_{ij} is the Reynolds stress tensor (equation (2.22)). The first term in equation (3.12) is called the *fading memory*

term, and all of the other terms relate to *drift correction* for inhomogeneous turbulence. With equations (3.5 & 3.10), along with equations (3.11 – 3.13), the evolution equations for the Lagrangian velocity and position (equations (3.1) & (3.2)) can be rewritten

$$dU'_i = \left(T_i^{(0)} + T_{ij}^{(1)}U'_j + T_{ijk}^{(2)}U'_jU'_k \right) dt + (C_0\varepsilon)^{1/2}d\xi_i(t), \quad (3.14)$$

$$dX_i = (\langle u_i \rangle + U'_i) dt. \quad (3.15)$$

In homogeneous-isotropic turbulence the Reynolds stress tensor is diagonal and equation (3.14) simplifies to

$$dU'_i = -\frac{C_0\varepsilon}{2\sigma^2}U'_i dt + (C_0\varepsilon)^{1/2}d\xi_i(t), \quad (3.16)$$

where the velocity variance σ^2 is the same in all directions, due to isotropy. These are the familiar Langevin equations; the stochastic differential equations associated with the simplest, statistically stationary diffusion process, the Ornstein-Uhlenbeck process (Pope, 2000).

If the fields of the mean velocities and the Reynolds stresses possess steep and/or erratic gradients or local distortions, such as those found in canopies, then the Thomson model described above may produce unrealistic velocities, commonly referred to as *rogue trajectories* or *rogue velocities*. Yee and Wilson (2007) suggested that rogue velocities arise due to dynamical and/or numerical instabilities within the Langevin equations, and proposed a method to prevent them. Unfortunately, the prevention algorithm proposed by those authors requires significant computational overhead and is not employed in the model presented here. Whenever a rogue velocity is produced by the models used in this thesis (here detected when the velocity fluctuation is greater than six

times the local standard deviation), the three components of the velocity are re-initialised based upon the local velocity statistics.

In practice, the derivatives in equations (3.14) and (3.15) are replaced by forward difference equations so that the updated velocity and positions are given as

$$U'_i(t + \Delta t) = U'_i(t) + a_i \Delta t + (C_0 \varepsilon)^{1/2} \Delta \xi_i(t), \quad (3.17)$$

$$X_i(t + \Delta t) = X_i(t) + (\langle u_i \rangle + U'_i) \Delta t, \quad (3.18)$$

where equation (3.10) was used to simplify equation (3.17). In this thesis, the timestep is chosen to be a fraction of the Lagrangian integral timescale

$$\Delta t = \mu_t \min[T_{L_u}, T_{L_v}, T_{L_w}] \quad (3.19)$$

where $\mu_t \ll 1$ is the timestep constant and the Lagrangian integral timescales associated with the u , v , and w velocities are calculated as

$$T_{L_u} = \frac{2\sigma_u^2}{C_0 \varepsilon}, \quad T_{L_v} = \frac{2\sigma_v^2}{C_0 \varepsilon}, \quad T_{L_w} = \frac{2\sigma_w^2}{C_0 \varepsilon}. \quad (3.20)$$

Together, equations (3.14) and (3.15) can only provide information about first-order statistics such as the mean concentration. A micromixing model may be used to calculate the higher-order moments of the concentration field, in which case the compound Markovian state variable is enlarged to (U'_i, X_i, ϕ) .

3.2 The IECM Micromixing Model

The rate of change in concentration as calculated by the IECM model is

$$\frac{d\phi}{dt} = \theta(\mathbf{x}, \psi, t) = -\frac{1}{t_m}(\phi - \langle \phi | \mathbf{u} \rangle), \quad (3.21)$$

where ψ is the sample space variable for ϕ , t_m is the mixing timescale, and $\langle\phi|\mathbf{u}\rangle$ is the scalar concentration conditioned on the local velocity (also called the conditional concentration). To better understand the conditional concentration consider the question, “what is the concentration of the scalar given its velocity?” The mathematically simpler IEM model can be attained by substituting the unconditional mean concentration $\langle\phi\rangle$ for $\langle\phi|\mathbf{u}\rangle$ in equation (3.21). The unconditional and conditional mean concentrations are related through

$$\langle\phi\rangle = \int_{\mathbf{v}} \langle\phi|\mathbf{v}\rangle f_{\mathbf{u}}(\mathbf{v}) d\mathbf{v}, \quad (3.22)$$

where (once again) \mathbf{v} is the sample space variable for \mathbf{u} . In the IECM model, the particle’s concentration will relax to the local mean concentration conditioned on velocity, via equation (3.21). The rate at which this mixing occurs is inversely proportional to the micromixing timescale t_m . Heuristically, conditional scalar mixing can be thought of as occurring between all fluid elements that occupy the same eddy. Furthermore, fluid elements with the same velocity are more likely to remain together for times comparable to the Lagrangian integral time and therefore are more likely to mix (Fox, 1996; Pope, 1998). Plume flapping or meandering contributes to the scalar variance but not to the scalar dissipation (Sawford, 2004a). It is important for a mixing model to properly represent the bulk motions and the in-plume, variance-dissipating motions. By conditioning on velocity, the IECM model achieves this segregation (Sawford, 2004a).

The Fokker-Planck (FP) equation corresponding to equations (3.1), (3.2)

and (3.21) is

$$\begin{aligned} \frac{\partial f_{\mathbf{u}\phi}}{\partial t} + (\langle v_i \rangle + V'_i) \frac{\partial f_{\mathbf{u}\phi}}{\partial x_i} = & - \frac{\partial}{\partial v_i} [a_i(\mathbf{V}', \psi) f_{\mathbf{u}\phi}] - \frac{\partial}{\partial \psi} [\theta(\psi) f_{\mathbf{u}\phi}] \\ & + \frac{1}{2} C_0 \varepsilon \frac{\partial^2 f_{\mathbf{u}\phi}}{\partial v_i \partial v_i}, \end{aligned} \quad (3.23)$$

(Pope, 1985, 2000; Fox, 2003). Comparing the exact equation (2.79) with the modelled equation (3.23), we see that the closure for the conditional mean acceleration is

$$\begin{aligned} \frac{\partial}{\partial v_i} \left[\left\langle \nu \frac{\partial^2 u_i}{\partial x_j \partial x_j} - \frac{1}{\rho} \frac{\partial p}{\partial x_i} \middle| \mathbf{v}, \psi \right\rangle f_{\mathbf{u}\phi} \right] = & \frac{\partial}{\partial v_i} [a_i(\mathbf{V}', \psi) f_{\mathbf{u}\phi}] \\ & - \frac{1}{2} C_0 \varepsilon \frac{\partial^2 f_{\mathbf{u}\phi}}{\partial v_i \partial v_i}, \end{aligned} \quad (3.24)$$

and the closure for the conditional mean diffusion is

$$\left\langle \mathcal{D} \frac{\partial^2 \phi}{\partial x_j \partial x_j} \middle| \mathbf{v}, \psi \right\rangle = \theta(\psi) = -\frac{1}{t_m} (\phi - \langle \phi | \mathbf{v} \rangle). \quad (3.25)$$

The budget equation for the mean concentration can be retrieved by multiplying equation (3.23) by ψ and integrating over (\mathbf{u}, ϕ) phase space. The result is

$$\frac{\partial \langle \phi \rangle}{\partial t} + \langle u_i \rangle \frac{\partial \langle \phi \rangle}{\partial x_i} + \frac{\partial}{\partial x_i} \langle U'_i \phi' \rangle = \langle \theta \rangle = 0. \quad (3.26)$$

This expression involves the scalar flux $\langle U'_i \phi' \rangle$ which should not be affected by any aspect of the micromixing model. To see if this is the case derive the budget equation for the scalar flux by multiplying equation (3.23) by $v_j \psi$ and integrating over (\mathbf{u}, ϕ) phase space, using the mean momentum and mean concentration equations to eliminate the extra terms. The resulting equation is

$$\begin{aligned} \frac{\partial \langle U'_j \phi' \rangle}{\partial t} + \langle u_i \rangle \frac{\partial}{\partial x_i} \langle U'_j \phi' \rangle + \langle U'_i U'_j \rangle \frac{\partial \langle \phi \rangle}{\partial x_i} + \langle U'_i \phi' \rangle \frac{\partial \langle u_j \rangle}{\partial x_i} \\ + \frac{\partial}{\partial x_i} \langle U'_i U'_j \phi' \rangle = \langle a_j \phi \rangle + \langle u_j \theta \rangle. \end{aligned} \quad (3.27)$$

The first term on the right-hand-side is due to the LS closure for the conditional mean acceleration, and is equal to zero. The effects of the micromixing model are manifested through the last term on the right-hand-side. In order to have no effect on the scalar flux $\langle u_i \theta \rangle = 0$ is required. For the IEM model we have

$$\langle u_j \theta \rangle = -\frac{1}{t_m} \langle u_j (\phi - \langle \phi \rangle) \rangle = -\frac{1}{t_m} \langle u_j \phi' \rangle \neq 0. \quad (3.28)$$

From this equation we see that the IEM model has a spurious term in the scalar flux budget equation. In the literature this extra term is referred to as the “spurious flux”.

In contrast, the IECM model does not have a spurious flux,

$$\begin{aligned} \langle u_j \theta \rangle &= -\frac{1}{t_m} \langle u_j (\phi - \langle \phi | \mathbf{v} \rangle) \rangle \\ &= -\frac{1}{t_m} \langle u_j \phi - u_j \langle \phi | \mathbf{v} \rangle \rangle \\ &= -\frac{1}{t_m} (\langle u_j \phi \rangle - \langle u_j \langle \phi | \mathbf{v} \rangle \rangle) = 0. \end{aligned} \quad (3.29)$$

In the above demonstration, the definitions of the mean and the conditional PDF were used:

$$\begin{aligned} \langle u_j \langle \phi | \mathbf{v} \rangle \rangle &= \left\langle u_j \int_{\psi} \psi f_{\phi | \mathbf{u}} d\psi \right\rangle \\ &= \left\langle \int_{\psi} v_j \psi f_{\phi | \mathbf{u}} d\psi \right\rangle \\ &= \int_{\mathbf{v}} \int_{\psi} v_j \psi f_{\phi | \mathbf{u}} d\psi f_{\mathbf{u}} d\mathbf{v} \\ &= \int_{\mathbf{v}} \int_{\psi} v_j \psi f_{\mathbf{u} \phi} d\psi d\mathbf{v} \\ &= \langle u_j \phi \rangle. \end{aligned} \quad (3.30)$$

Therefore the IECM model does not alter first-order statistics of the concentration field while the IEM model does. The spurious flux acts as a source,

or more commonly a sink, of the scalar flux (Pope, 1998). The significance of not having a spurious flux is that single-particle trajectories can be used to pre-calculate the conditional mean concentrations, resulting in greatly reduced computational loads.

3.2.1 The Micromixing Timescale

Equation (3.21) has one free parameter, the micromixing timescale, t_m . Through this parameter all of the processes that contribute to mixing are heuristically included into the model. Some early micromixing models (Pope, 1985) assumed the micromixing timescale to be equal to the turbulence timescale, commonly modelled as

$$t_t \equiv k/\varepsilon. \quad (3.31)$$

When the length scale of the turbulence is much smaller than the length scale of the source, then this assumption produces reasonable results. However, this is not the case for most atmospheric science applications. Consequently, for problems involving atmospheric dispersion the relationship between the micromixing timescale and the turbulence timescale is more complicated, and furthermore is time dependent (Sykes et al., 1984).

Mixing is intimately related to the growth of the plume about its instantaneous centreline and thus to relative dispersion (Sawford, 2004b). Therefore, it is reasonable to expect the mixing timescale to evolve with the instantaneous plume size as

$$t_m \sim \left(\frac{\sigma_r^2}{\sigma_{\mathbf{U}_r}^2} \right)^{1/2}, \quad (3.32)$$

where $\sigma_{\mathbf{U}_r}^2 = \langle \mathbf{U}_r^2 \rangle$ is the variance of the Lagrangian relative velocity fluctuations (Luhar and Sawford, 2005). To calculate \mathbf{U}_r , take the difference between

the turbulent velocity and the velocity of the instantaneous centre of mass. To determine the evolution of the mixing timescale, consider dispersion from a localised point source of diameter σ_0 for which one may define a characteristic timescale $t_s = (\sigma_0^2/\varepsilon)^{1/3}$. At early times $t \ll t_s$, when the plume is small, the source characteristics dominate the evolution of the plume and thus we expect $t_m \propto t_s$. As discussed in Section 2.5, Richardson's t^3 law (equation (2.76)) applies in the inertial subrange $t_s \ll t \ll T_L$. The instantaneous plume grows as $\sigma_r^2 \propto t^3$ and $\sigma_{U_r}^2$ grows as t . Therefore, by equation (3.32), we expect $t_m \propto t$. At late times, $t \gg T_L$, the instantaneous plume has grown to fill the time-averaged envelope of the absolute plume. During this regime of plume development, σ_r^2 grows proportionally to t and $\sigma_{U_r}^2$ is constant. We therefore expect $t_m \propto t^{1/2}$. Table 3.1 summarises the expected behaviour of the micromixing timescale in the three time regimes discussed above. The parametrisation of t_m used in the IECM model should reproduce these expectations.

Regime	Time	Behaviour of t_m
1	$t \rightarrow 0$	$t_m \propto t_s$
2	$t_s \ll t \ll T_L$	$t_m \propto t$
3	$t \gg T_L$	$t_m \propto t^{1/2}$

Table 3.1: Summary of the expected time behaviour of the micromixing timescale in three regimes of plume development.

We begin by describing the parametrisation of the micromixing timescale for homogeneous and isotropic turbulence. Following Sykes et al. (1984); Sawford (2004b); and Cassiani, Franzese and Giostra (2005a), for short and

medium times we assume that

$$t_m = \mu \left(\frac{\sigma_r^2}{\sigma_{\mathbf{U}_r}^2} \right)^{1/2}. \quad (3.33)$$

The micromixing constant (μ) is empirically determined and depends upon the type of turbulence, the source configuration, and the stage of development of the plume. It is treated as a “tuning” parameter. The variance of the Lagrangian relative velocity fluctuations ($\sigma_{\mathbf{U}_r}^2$) is in general very difficult to calculate. However, it represents the fraction of the energy responsible for expansion of the plume about its instantaneous centreline. It is possible to estimate an expression for it in the inertial subrange by integrating Kolmogorov’s 5/3 law (equation (2.5)) with respect to the wavenumber

$$\begin{aligned} \int_{\kappa}^{\infty} E_{\kappa'} d\kappa' &= \int_{\kappa}^{\infty} C\varepsilon^{2/3} \kappa'^{-5/3} d\kappa' \\ &= -\frac{3}{2} C\varepsilon^{2/3} \kappa^{-2/3}. \end{aligned} \quad (3.34)$$

Casting this result in terms of the eddy sizes with $l = 2\pi/\kappa$, we find that the energy available for plume expansion is proportional to $l^{2/3}$. The largest, energy-containing eddies have size L , so the total amount of energy available for plume expansion is proportional to $L^{2/3}$.

By constructing a ratio of the instantaneous plume width to the largest eddies Franzese (2003) and Cassiani, Franzese and Giostra (2005a) modelled $\sigma_{\mathbf{U}_r}^2$ as

$$\sigma_{\mathbf{U}_r}^2 = \sigma^2 \left(\frac{\sigma_r}{L} \right)^{2/3}, \quad (3.35)$$

where $\sigma^2 = \frac{2}{3}k$ is the turbulent velocity variance (energy) in a given direction (assumed to be the same in every direction because of isotropy). When $\sigma_r > L$ the constraint $\sigma_{\mathbf{U}_r}^2 = \sigma^2$ is imposed. The lengthscale of the most energetic

eddies is calculated as

$$L = \frac{(3\sigma^2/2)^{3/2}}{\varepsilon}. \quad (3.36)$$

The instantaneous plume spread is modelled as

$$\sigma_r^2 = \frac{d_r^2}{1 + (d_r^2 - \sigma_0^2)/(\sigma_0^2 + 2\sigma^2 T_L t)}, \quad (3.37)$$

where d_r is the separation between two particles in the instantaneous plume and calculated with the Richardson Law

$$d_r^2 = C_r \varepsilon (t + t_0)^3. \quad (3.38)$$

The constant $t_0 = t_s/C_r^{1/3}$ (where $t_s = (\sigma_0^2/\varepsilon)^{1/3}$ is the characteristic timescale of the source) ensures that $t_m \rightarrow t_s$ as $t \rightarrow 0$. The time evolution of the mixing timescale scaled by the source timescale, as predicted by equations (3.33)–(3.38) for homogeneous-isotropic turbulence, is shown in Figure 3.1. We see that this parametrisation for the micromixing timescale does indeed result in the proper time behaviour. The asymptotic values for t_m are:

$$\begin{aligned} t_m &= \mu(3/2)^{1/2} t_s & t \rightarrow 0; \text{ Regime 1,} \\ t_m &= \mu(3/2)^{1/2} C_r^{1/3} t & t_s \ll t \ll T_L; \text{ Regime 2,} \\ t_m &= \mu\sqrt{2T_L t} & t \gg T_L; \text{ Regime 3,} \end{aligned} \quad (3.39)$$

all of which have the expected time dependence (Cassiani, Franzese and Giostra, 2005a). They are represented by the dashed lines in Figure 3.1.

In non-homogeneous, non-isotropic turbulence the parametrisation of t_m is slightly modified. As in Cassiani, Franzese and Giostra (2005a), local equilibrium and local isotropy are assumed and the variance used in equation (3.35) is replaced by a local mean variance

$$\sigma^2 = \frac{\sigma_u^2 + \sigma_v^2 + \sigma_w^2}{3}. \quad (3.40)$$




Figure 3.1: Evolution of the dimensionless mixing timescale in homogeneous-isotropic turbulence. The short, dashed lines represent the asymptotic values given by equation grouping (3.39).

Equation (3.38) is discretised by invoking linearisation, viz.

$$d_r^2(t + \Delta t) = d_r^2(t) + 3C_r\varepsilon(t_0 + t)^2\Delta t, \quad (3.41)$$

and enforcing the constraint $\sigma_r^2(t + \Delta t) \geq \sigma_r^2(t)$.

3.3 Description of the SPMMM Model Suite

The micromixing model used in this thesis is called **SPMMM**. This is an acronym for Single Particle MicroMixing Model. In contrast to the simultaneous release of N particles to determine the concentration statistics, as used by Cassiani, Franzese and Giostra (2005a), **SPMMM** releases a single particle at a time. The particle samples a pre-calculated conditional mean concentration field that is provided by a program called **MEANS**. This single-particle, pre-calculation architecture of the model is permitted since the IECM model has no spurious flux, as shown in Section 3.2. Collectively **MEANS** and **SPMMM** shall be referred to as the **SPMMM** model suite. In theory, these models can provide all moments of the scalar field. In practice, however, this is limited by the availability of the computational resources. The use of single-particle trajectories results in computational simplicity, but comes at the cost of not being able to calculate the concentration field of reactive species, as would be the case if simultaneous particle trajectories were used. Since we only consider non-reactive species in this thesis, this is not a cause of concern.

While the use of simultaneous trajectories allows the incorporation of chemical reaction, it is much more difficult to parallelise since the particles are interactive in that at each time step the conditional mean concentration must be calculated based on the particles that occupy a particular region of space. If the

particles within a region are being processed by different computer processors (which is very likely), then there will be a large computational communication overhead since at each time step, the processors will have to pause, share particle data, then continue. In contrast, the single-particle trajectory framework allows for trivial parallelisation and a direct increase in performance; the time required to run SPMMM on N_p computer processors is $\sim 1/N_p$ the time required to run SPMMM on one computer processor. The only time the processors must communicate is when sharing the plume extent data, at the beginning of the simulation.

Before describing the models some terminology must be introduced. For both models, space is divided into three spatial dimensions $\mathbf{x} = (x, y, z)$ and three velocity dimensions $\mathbf{u} = (u, v, w)$. The spatial dimensions shall at times be referred to as the *physical space* or *physical domain* and the velocity dimensions shall at times be referred to as the *velocity space* or *velocity domain*. Collectively, the six spatial and velocity dimensions shall sometimes be referred to as the *simulation domain*. The user-specified extent of the physical domain has dimensions of $x_{\min} \leq x \leq x_{\max}$, $y_{\min} \leq y \leq y_{\max}$ and $z_{\min} \leq z \leq z_{\max}$. These physical domain boundaries are set at the beginning of simulation. The velocity domain boundaries: $u_{\min} \leq u \leq u_{\max}$, $v_{\min} \leq v \leq v_{\max}$, and $w_{\min} \leq w \leq w_{\max}$, are determined by the program MEANS as will be described below. The simulation domain is discretised into N_x streamwise bins, N_y spanwise bins, N_z vertical bins, N_u streamwise velocity bins, N_v spanwise velocity bins, and N_w vertical velocity bins. Upper-case subscripts (and indices) will be used to denote this discretisation; $(x, y, z) \Rightarrow (x_I, y_J, z_K)$ and $(u, v, w) \Rightarrow (u_L, v_M, w_N)$. The *source region* refers to where the emissions

originate.

Perfectly reflective boundary conditions are used on the upstream face and at the bounding horizontal planes. That is, if $X < x_{\min}$ then

$$\begin{aligned} X &= 2x_{\min} - X, \\ U' &= -U', \\ W' &= -W'; \end{aligned} \tag{3.42}$$

if $Z < z_{\text{rft}}$, where z_{rft} is the height of the reflection surface near the bottom of the physical domain then,

$$\begin{aligned} Z &= 2z_{\text{rft}} - Z, \\ U' &= -U', \\ W' &= -W'; \end{aligned} \tag{3.43}$$

and if $Z > z_{\max}$ then

$$\begin{aligned} Z &= 2z_{\max} - Z, \\ U' &= -U', \\ W' &= -W'. \end{aligned} \tag{3.44}$$

Periodic boundary conditions are used on the lateral faces, if $Y < y_{\min}$ then

$$Y = y_{\max} - y_{\min} + Y, \tag{3.45}$$

and if $Y > y_{\max}$ then

$$Y = y_{\min} - y_{\max} + Y. \tag{3.46}$$

The velocity statistics used to drive the models are in discretised form. These velocity statistics can be obtained from discretisations of analytical

equations, from interpolations of experimentally measured flow fields, or provided by another model. The discretisation of the velocity statistics is not necessarily the same as the discretisation of the physical space. We therefore let N_x^{vel} , N_y^{vel} , N_z^{vel} represent the number of velocity statistics bins in physical space in the streamwise, spanwise, and vertical directions respectively.

The descriptions of the **MEANS** and **SPMMM** models below will be for the full six-dimensional simulation domain, (x, y, z, u, v, w) . For illustrative purposes, a simplified three-dimensional domain (x, z, w) will be used. This simplified domain will be discretised into $N_x = 5$, $N_z = 3$, and $N_w = 4$ bins. The source region will be represented by an arrow.

3.3.1 The MEANS Pre-calculation Program

MEANS is an implementation of equations described in Section 3.1. The purpose of **MEANS** is to pre-calculate the unconditional mean concentration field $\langle \phi \rangle$ and the conditional mean concentration field $\langle \phi | \mathbf{u} \rangle$ for use by **SPMMM**. To maximise the spatial resolution of the model, the calculation of the mean concentrations is carried out on a dynamic grid that encompasses the plume. The approximate spatial extent of the plume is determined from the release of a small sub-ensemble of particles, released from the source region, into the domain. The fetch (i.e., $x_{\text{max}} - x_{\text{min}}$) is divided into N_x extraction planes resulting in a streamwise bin width of

$$\Delta x = \frac{x_{\text{max}} - x_{\text{min}}}{N_x}. \quad (3.47)$$

As the particles' trajectories cross these extraction planes, their spanwise and vertical positions are recorded. Once all the particles in the sub-ensemble

have completed their trajectories, the extrema in the physical space at the I^{th} extraction plane ($1 \leq I \leq N_x$) can be determined as

$$y_{\min}^I = \langle y \rangle - \mu_r \sigma_y, \quad (3.48)$$

$$y_{\max}^I = \langle y \rangle + \mu_r \sigma_y, \quad (3.49)$$

$$z_{\min}^I = \langle z \rangle - \mu_r \sigma_z, \quad (3.50)$$

$$z_{\max}^I = \langle z \rangle + \mu_r \sigma_z, \quad (3.51)$$

where μ_r is a scale factor that is set by the user and σ_y and σ_z are the root mean square values (over the sub-ensemble) of particle displacements from the the centroids $\langle y \rangle$ and $\langle z \rangle$ at sample plane I . For all simulations reported in this thesis we set $\mu_r = 6.0$.

Once the spatial extrema at each extraction plane have been determined, the spanwise and vertical bin widths at the I^{th} plane are calculated as

$$\Delta y^I = \frac{y_{\max}^I - y_{\min}^I}{N_y}, \quad (3.52)$$

$$\Delta z^I = \frac{z_{\max}^I - z_{\min}^I}{N_z}. \quad (3.53)$$

A two-dimensional representation of the resulting physical grid is shown in Figure 3.2. The extraction planes are represented by dotted lines and the sub-domain where $\phi \neq 0$ (i.e. the tracer concentration is non-zero) is shaded.

The extent of the velocity domain is determined from the driving velocity statistics. For each position in the discretised physical space the mean velocities and the velocity variances are known. From these the global velocity

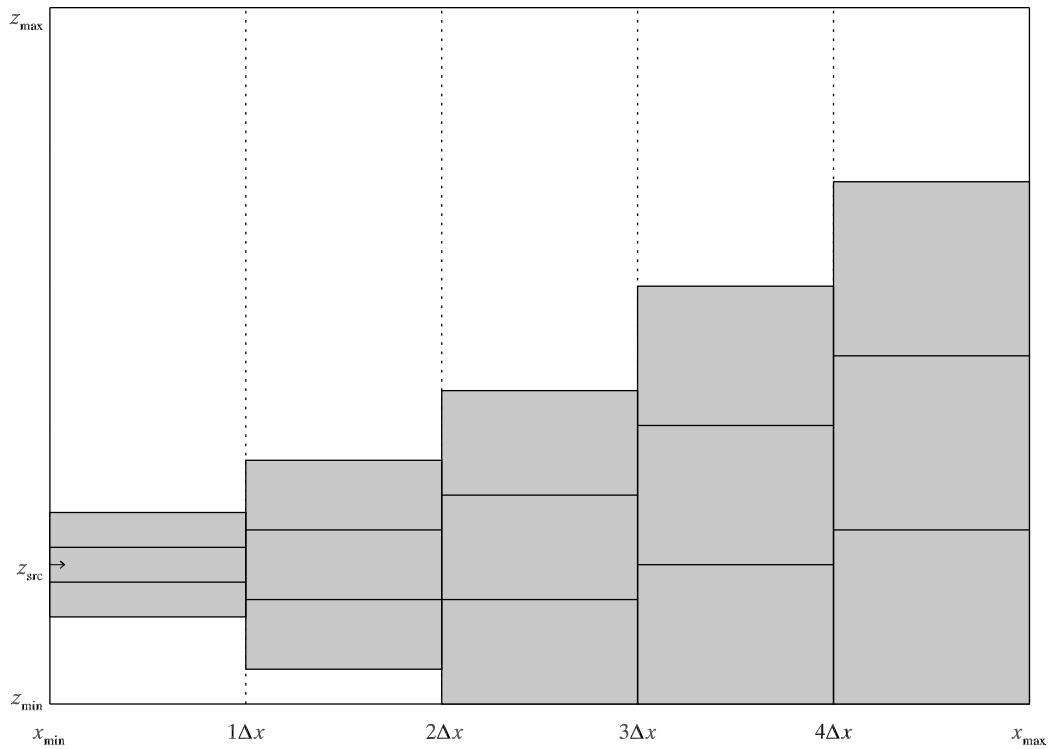


Figure 3.2: A two-dimensional representation showing the physical grid resulting from the dynamic discretisation procedure with $N_x = 5$ and $N_z = 3$. The extraction planes are represented by dotted lines and the sub-domain where $\phi \neq 0$ is shaded. Note how the vertical bin width expands with the plume.

extrema can be computed as

$$u_{\min} = \min[\langle u(x_I, y_J, z_K) \rangle - \mu_v \sigma_u(x_I, y_J, z_K)], \quad (3.54)$$

$$u_{\max} = \max[\langle u(x_I, y_J, z_K) \rangle + \mu_v \sigma_u(x_I, y_J, z_K)], \quad (3.55)$$

$$v_{\min} = \min[\langle v(x_I, y_J, z_K) \rangle - \mu_v \sigma_v(x_I, y_J, z_K)], \quad (3.56)$$

$$v_{\max} = \max[\langle v(x_I, y_J, z_K) \rangle + \mu_v \sigma_v(x_I, y_J, z_K)], \quad (3.57)$$

$$w_{\min} = \min[\langle w(x_I, y_J, z_K) \rangle - \mu_v \sigma_w(x_I, y_J, z_K)], \quad (3.58)$$

$$w_{\max} = \max[\langle w(x_I, y_J, z_K) \rangle + \mu_v \sigma_w(x_I, y_J, z_K)], \quad (3.59)$$

where the scale factor μ_v is set by the user. The global velocity bin widths are then calculated as

$$\Delta u = \frac{u_{\max} - u_{\min}}{N_u}, \quad (3.60)$$

$$\Delta v = \frac{v_{\max} - v_{\min}}{N_v}, \quad (3.61)$$

$$\Delta w = \frac{w_{\max} - w_{\min}}{N_w}. \quad (3.62)$$

A physical domain is created for each velocity domain, as shown in Figure 3.3.

For all simulations reported in this thesis we set $\mu_v = 6.0$.

The boundaries that partition both physical and velocity space are now known. To determine the unconditional and conditional mean concentrations, N_ϕ particles are released from the source region and tracked downstream until they leave the domain. The amount of time that a particle spends in each physical and velocity bin is related to the concentration in that bin, and is recorded as shown in Figure 3.4. For clarity, these *residence times* are only shown for a single physical bin in each of the four velocity bins in the figure. In actuality the residence times are being accumulated for every bin. From




Figure 3.3: The resulting physical and velocity grids resulting from the dynamic discretisation procedure with $N_x = 5$, $N_z = 3$, and $N_w = 4$. Note that the physical grid is the same for each velocity bin.

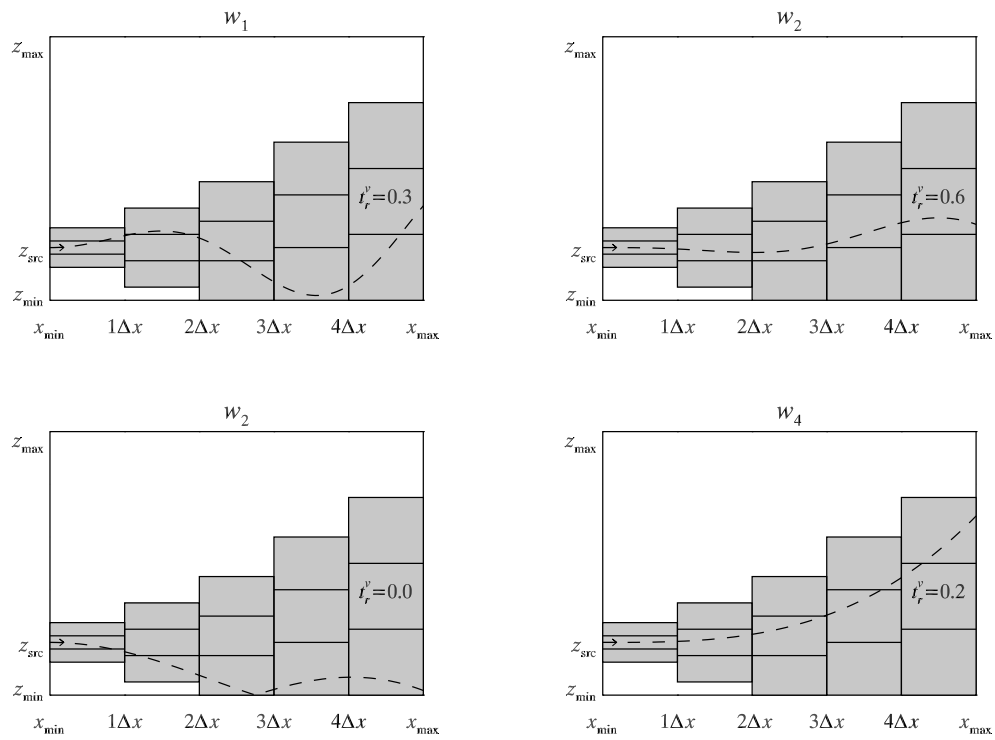


Figure 3.4: The accumulation of residence times in the simulation domain. The dashed lines represent particle trajectories. Note that for clarity the residence times are only shown for a single physical bin in each of the four velocity bins.

these residence times

$$t_r^v = t_r^v(x_I, y_J, z_K, u_L, v_M, w_N), \quad (3.63)$$

the conditional and unconditional mean concentration fields (for a given source configuration and flow) can be determined. A residence time approach was chosen over a flux method (where the concentration at a position is proportional to the number of particles that cross through a rectangular bin perpendicular to the mean streamwise flow at that location) for determining the mean concentrations as the intended use for the model is with highly disturbed flows where there is the potential for significant upstream travel. A residence time approach can capture the effects of this upstream travel on the mean fields better than a flux approach, since in the flux approach the upstream travel may occur between two of the extraction planes and thus never get “noticed” by the model.

To calculate the unconditional mean concentration for each physical bin, we need to sum over all velocity bins to compute the unconditional residence time as

$$t_r = t_r(x_I, y_J, z_K) = \sum_{L,M,N} t_r^v(x_I, y_J, z_K, u_L, v_M, w_N). \quad (3.64)$$

The unconditional mean concentration in physical bin (x_I, y_J, z_K) is then computed as

$$\langle \phi \rangle = \langle \phi(x_I, y_J, z_K) \rangle = \frac{Qt_r}{\mathcal{V}N_\phi}, \quad (3.65)$$

where Q is the source strength and $\mathcal{V} = \mathcal{V}(x_I, y_J, z_K)$ is the volume of the physical bin. The conditional mean concentration in physical bin (x_I, y_J, z_K) and velocity bin (u_L, v_M, w_N) is computed as

$$\langle \phi | \mathbf{u} \rangle = \langle \phi | \mathbf{u} \rangle(x_I, y_J, z_K, u_L, v_M, w_N) = \frac{Qt_r^v}{\mathcal{V}N_\phi^v}, \quad (3.66)$$

where

$$N_\phi^v = N_\phi^v(x_I, y_J, z_K, u_L, v_M, w_N) \quad (3.67)$$

is the number of particles that visit bin $(x_I, y_J, z_K, u_L, v_M, w_N)$ during the simulation. Due to computer memory limitations N_ϕ^v is not tracked explicitly. Instead it is estimated from the velocity PDF, $f_{\mathbf{u}} = f_{\mathbf{u}}(\mathbf{v}; \mathbf{x})$ as

$$N_\phi^v = N_\phi f_{\mathbf{u}} \Delta u \Delta v \Delta w, \quad (3.68)$$

giving

$$\langle \phi | \mathbf{u} \rangle = \frac{Q t_r^v}{\mathcal{V} N_\phi f_{\mathbf{u}} \Delta u \Delta v \Delta w}. \quad (3.69)$$

The result of the **MEANS** pre-calculation program is a file containing the conditional mean concentrations to be used by **SPMMM**. The unconditional mean concentrations are used solely to confirm or investigate consistency between **MEANS** and **SPMMM** in the first-order statistics. Figure 3.5 shows the resulting conditional mean concentrations for the small illustrative simulation domain. The different shades of grey represent different conditional mean concentrations.

The estimation of N_ϕ^v via equation (3.69) allows **MEANS** and **SPMMM** to be run at double the resolution they could be if N_ϕ^v was being accumulated numerically. However there is a potential source of error involved with determining N_ϕ^v in this manner. In many situations, the spatial resolution of the conditional concentration field may be coarser than the spatial resolution of the driving velocity statistics. Alternatively a conditional concentration bin may span two driving velocity statistic bins. The value of $f_{\mathbf{u}}$ varies in each velocity statistic bin thus **MEANS** computes an average velocity PDF for each of its conditional concentration bins. Figure 3.6 demonstrates this situation for the




Figure 3.5: The conditional mean concentration fields that result from a MEANS simulation. The different shades of grey represent different conditional mean concentrations. Note that each velocity bin has a different conditional mean concentration field. This data is used by SPMM to allow the calculation of the higher-order moments of the concentration field.

case of horizontally homogeneous velocity statistics and a single, large conditional concentration bin which spans three velocity statistic bins (shown in grey). In this case, MEANS will estimate the velocity PDF in the conditional concentration bin as $\langle f_{\mathbf{u}} \rangle = \frac{1}{3}[f_{\mathbf{u}}(z_2) + f_{\mathbf{u}}(z_3) + f_{\mathbf{u}}(z_4)]$. In many cases this mean value well approximates the actual PDF. However, if the driving velocity statistics are discretised into more bins than the spatial domain (i.e. $N_x^{\text{vel}} > N_x$, $N_y^{\text{vel}} > N_y$, $N_z^{\text{vel}} > N_z$) then this approximation can lead to error, as will be demonstrated in Section 4.3.




Figure 3.6: The approximation of the velocity PDF in the situation where the resolution of the velocity statistics is greater than the resolution of the conditional concentration field as determined by MEANS. In this situation the velocity PDF in the conditional concentration bin (shown in grey) will be estimated to be $\langle f_{\mathbf{u}} \rangle = \frac{1}{3}[f_{\mathbf{u}}(z_2) + f_{\mathbf{u}}(z_3) + f_{\mathbf{u}}(z_4)]$.

3.3.2 The SPMMM Micromixing Model

Once the unconditional and conditional mean concentrations have been computed by **MEANS**, the next step is to use the **SPMMM** micromixing model to allow mixing to occur in order to provide information on the higher-order moments of the concentration field. The first step is to compute the micromixing timescales on the physical grid. Again a small sub-ensemble of particles are released from the source region one at a time and tracked downstream. As the particles pass through each bin in physical space, t_m is computed according to the equations in Section 3.2.1 and recorded. Since the particles will all follow different trajectories, the computed value of t_m for a given physical bin may differ greatly from particle to particle. Once all the particles in the sub-ensemble have exited the physical domain through $x \geq x_{\max}$, the mean micromixing timescale in bin (x_I, y_J, z_K) is calculated for use in the next stage of the simulation. If the micromixing timescale is larger than the turbulence timescale, then t_m is reset to t_t . Furthermore, in regions outside the plume, mixing still occurs and does so at a rate governed by the turbulence timescale, hence for these regions $t_m = t_t$. For the remainder of the **SPMMM** simulation, the timestep as shown in equation (3.19) is modified to include the micromixing timescale

$$\Delta t = \mu_t \min[T_{L_u}, T_{L_v}, T_{L_w}, t_m]. \quad (3.70)$$

During the micromixing stage of the **SPMMM** model suite, N particles are uniformly released, one at a time, anywhere on the upstream face of the simulation domain. If the particle originates outside of the source region it is given an initial concentration of $\phi_0 = 0$, as shown in the left panel of Figure 3.7. Otherwise it is given an initial concentration of $\phi_0 = \phi_{\text{src}}$, where ϕ_{src} is the

source concentration which varies according to the source configuration and type. This is shown in the right panel of Figure 3.7 where $\phi_{\text{src}} = 1.7$. The source configuration in **SPMMM** can be either a circular point or a cross-wind line. The initial distribution of particles about this source is Gaussian or tophat. The source concentration profiles for the Gaussian sources are,

$$\phi_{\text{src}} = \frac{Q}{2\pi\sigma_0^2U} \exp\left(\frac{-r^2}{2\sigma_0^2}\right) \quad \text{Gaussian point source,} \quad (3.71)$$

$$\phi_{\text{src}} = \frac{Q}{\sqrt{2\pi}\sigma_0U} \exp\left(\frac{-r^2}{2\sigma_0^2}\right) \quad \text{Gaussian line source.} \quad (3.72)$$

For tophat sources they are,

$$\phi_{\text{src}} = \frac{4Q}{\pi d_s^2U} \quad \text{tophat point source,} \quad (3.73)$$

$$\phi_{\text{src}} = \frac{Q}{d_sU} \quad \text{tophat line source.} \quad (3.74)$$

The initial source distribution is $\sigma_0 = \mu_s d_s$, where d_s is the physical diameter of the source and μ_s is the source constant, a tunable parameter ($\sigma_0 = d_s$ for tophat sources). For point sources, the distance from the particle position to the source centre in the yz -plane is $r^2 = (y - y_s)^2 + (z - z_s)^2$. For cross-wind line sources, the distance from the particle position to the line source is $r^2 = (z - z_s)^2$. For Gaussian sources, a particle is considered to be in the source region if it is initialised within five initial source distributions of the source centre, $r \leq 5\sigma_0$. For tophat sources, a particle is considered to a source particle if it falls within the area occupied by the source.

As the particle travels downstream, its concentration is compared to the conditional mean concentration of the physical bin and velocity bin that the particle occupies and the micromixing timescale for the physical bin is retrieved. Mixing then occurs according to equation (3.21). Under the assump-

tion that $\langle\phi|\mathbf{u}\rangle$ is approximately constant (it is in fact very slowly varying for sufficiently small timestep Δt) equation (3.21) can be solved to give

$$\phi(t + \Delta t) = \phi(t) \exp(-\Delta t/t_m) + \langle\phi|\mathbf{u}\rangle(1 - \exp(-\Delta t/t_m)), \quad (3.75)$$

which is the equation used by **SPMMM** to update the particle's concentration. Equation (3.75) always results in the concentration of the particle mixing towards the conditional mean concentration. This is illustrated in both panels of Figure 3.7. In the left panel $\phi_0 = 0.0$ and $\langle\phi|w\rangle = 0.9$. After mixing, $\phi_1 = 0.5$, having mixed towards $\langle\phi|w\rangle = 0.9$. In the right panel $\phi_0 = 1.7$ and $\langle\phi|w\rangle = 1.1$. After mixing, $\phi_1 = 1.3$, again having mixed towards the conditional mean concentration of $\langle\phi|w\rangle = 1.1$.

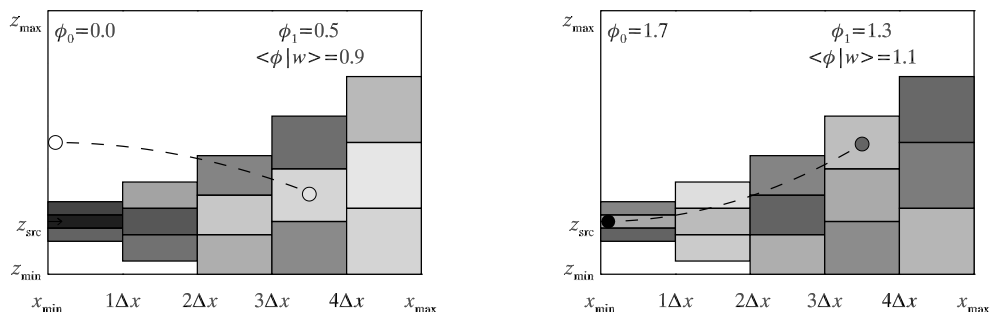


Figure 3.7: The two possible initial particle positions for **SPMMM**. The panel on the left shows a particle initialised outside the source region and thus given an initial concentration of $\phi_0 = 0$. The panel on the right shows a particle initialised in the source region and thus given an initial concentration of $\phi_0 = \phi_{\text{src}}$, set here to be 1.7. Note that mixing via equation (3.75) results in ϕ moving towards $\langle\phi|w\rangle$.

As a particle travels downstream, it passes through user-specified extraction planes and its position and concentration are saved to file for future processing. The extraction planes need not be uniformly spaced and they are not necessarily the same extraction planes used to determine the extent of the

plume in MEANS. From this output data, another grid can be imposed and the higher order moments of the concentration field calculated, as shown in Figure 3.8. This processing grid is frequently the same size as the grid used in MEANS to determine the mean concentrations. However it need not be, in the event of a low particle number simulation for example.

The decision to process the raw data outside of the SPMMM program was made to allow the user to decide which statistics to view after the simulation was completed. To decide beforehand would limit the order of the statistic to view as the ability to calculate the statistic would have to be coded into the program. If, for whatever reason, the user wished to view the seventeenth moment of the concentration field after the simulation had finished the option still exists if the data is processed outside of the program. The Interactive Data Language (IDL) by ITT Visual Information Solutions was used to process the data. Given the stochastic nature of the models described above the data contains statistical noise. A Savitzky-Golay smoothing filter was used to remove this noise while maintaining the signal. An example of the statistical noise and the smoothing filter is presented in Chapter 4 (see Figure 4.10).

3.4 Model Evaluation Methodology

Model evaluation is best performed both qualitatively and quantitatively. The qualitative assessment will be carried out through comparison of the model predictions with the experimental observations. These comparisons will be in the form of plots of various profiles showing the observed and predicted values. The quantitative assessment will utilise the statistical, bias, and discretisation errors as defined in Xu and Pope (1999), and some of the performance measures




Figure 3.8: Data extraction and the determination of concentration statistics for SPMMM. The panel on the left shows an extraction plane (dotted line) and three particle trajectories (propagated individually, not simultaneously). As a particle crosses through the plane its position and concentration is saved to file. From this data, concentration statistics can be determined, as shown in the right panel for the mean concentration and concentration variance.

outlined in Hanna (1989), Hanna et al. (1993) and Chang and Hanna (2004). In the following discussion, \mathcal{Q}_o shall represent an *observed* quantity, which in our case is a wind-tunnel, water-channel, or atmospheric measurement, and \mathcal{Q}_p shall represent a quantity *predicted* by the model. The model predictions are functions of the number of particles used in the simulations (N), the spatial discretisation of the simulation domain ($M = N_x \times N_y \times N_z$), and the timestep (Δt). Since the simulations in this thesis were run at the smallest practical timestep, we do not consider the contribution of the timestep to the error.

The error between a predicted quantity and an observed quantity is defined as the difference between them. It is composed of a physical error, which is a result of the models not representing the physical processes exactly, and a numerical error:

$$\epsilon_{\mathcal{Q}} \equiv \mathcal{Q}_p^{N,M} - \mathcal{Q}_o = \epsilon_p + \epsilon_n, \quad (3.76)$$

where ϵ_p and ϵ_n are the physical and numerical errors respectively. The numerical error can be further decomposed into statistical error ($\Sigma_{\mathcal{Q}}$) and de-

terministic error ($D_{\mathcal{Q}}$), which is the combination of the bias error ($B_{\mathcal{Q}}$) and discretisation error ($S_{\mathcal{Q}}$),

$$\epsilon_n = \Sigma_{\mathcal{Q}} + D_{\mathcal{Q}} = \Sigma_{\mathcal{Q}} + B_{\mathcal{Q}} + S_{\mathcal{Q}}. \quad (3.77)$$

The statistical error results from utilising a finite number of particles in the simulations. It can be identified as

$$\Sigma_{\mathcal{Q}} = \mathcal{Q}_p^{N,M} - \langle \mathcal{Q}_p^{N,M} \rangle_E, \quad (3.78)$$

where $\langle \cdot \rangle_E$ denotes the ensemble average. It is expected to scale as

$$\Sigma_{\mathcal{Q}} \propto N^{-1/2}, \quad (3.79)$$

(Xu and Pope, 1999). The bias error is a deterministic error that results from statistical error and is identified as

$$B_{\mathcal{Q}} = \langle \mathcal{Q}_p^{N,M} \rangle_E - \mathcal{Q}_p^{\infty,M}, \quad (3.80)$$

where $\mathcal{Q}_p^{\infty,M} = \lim_{N \rightarrow \infty} \mathcal{Q}_p^{N,M}$. The bias is expected to scale as

$$B_{\mathcal{Q}} \propto N^{-1}, \quad (3.81)$$

(Xu and Pope, 1999). Terms such as the deterministic coefficient in equation (3.1) and the conditional mean concentration in equation (3.21) involve velocity fluctuations or the velocity. If these quantities were computed by the model then they would have an associated statistical error (due to the use of a finite number of particles to calculate them) that would feedback into the model and induce bias (Xu and Pope, 1999). In the SPMMM micromixing model, the velocity statistics are externally supplied and the conditional concentrations are pre-calculated by MEANS. Therefore, even though the conditional concentrations probably have a bias error associated with them (due

to their determination from a finite number of particles) no further bias error should be introduced into the SPMMM micromixing model as the conditional concentrations are not re-computed at each timestep in SPMMM. We expect SPMMM to therefore have minimal bias error (only that which is contained in the conditional mean concentration field). The discretisation error is identified as

$$S_Q = Q_p^{\infty, M} - Q_o, \quad (3.82)$$

and is expected to decrease with an increasing spatial resolution (i.e., more bins). The use of more bins requires the use of more particles to maintain the number of particles per bin to keep the magnitude of the statistical error constant. Increasing the spatial resolution therefore greatly increases the computational costs.

We now introduce the performance measures used in this thesis. The first performance measure to consider is the fractional bias,

$$\text{FB} = \frac{(\overline{Q_o} - \overline{Q_p})}{0.5(\overline{Q_o} + \overline{Q_p})}, \quad (3.83)$$

where the over-bar indicates an arithmetic mean of all (or some subset of) the available observations or predictions. The FB is a measure of the systematic bias of the model, the difference between the observed and predicted quantities $Q_o - Q_p$. As such, if the model both under-predicts and over-predicts the results, it is possible for $\text{FB} = 0.0$ due to the cancellation of errors. To this end we can decompose the fractional bias into two parts, one part which considers over-predicted results, or false positives,

$$\text{FB}_{\text{fp}} = \frac{0.5[|\overline{Q_o} - \overline{Q_p}| + (\overline{Q_p} - \overline{Q_o})]}{0.5(\overline{Q_o} + \overline{Q_p})}, \quad (3.84)$$

and another part which considers under-predicted results, or false negatives,

$$\text{FB}_{\text{fn}} = \frac{0.5[|\overline{Q_o} - \overline{Q_p}| + (\overline{Q_o} - \overline{Q_p})]}{0.5(\overline{Q_o} + \overline{Q_p})}. \quad (3.85)$$

The original fractional bias is recovered through $\text{FB} = \text{FB}_{\text{fn}} - \text{FB}_{\text{fp}}$.

To better understand the FB, rearrange equation (3.83) to give

$$\frac{\overline{Q_p}}{\overline{Q_o}} = \frac{1 - \frac{1}{2}\text{FB}}{1 + \frac{1}{2}\text{FB}}. \quad (3.86)$$

From this equation, we can infer that FB should be zero for a perfect model as then the ratio of the mean predicted quantity and the mean observed quantity would be unity, implying that they are identical. If $\text{FB} = 2/3$ then $\overline{Q_p}/\overline{Q_o} = 1/2$ implying a factor of two under-prediction and if $\text{FB} = -2/3$ then $\overline{Q_p}/\overline{Q_o} = 2$ implying a factor of two over-prediction. Chang and Hanna (2004) suggest $-0.3 < \text{FB} < 0.3$ as an acceptable range for the FB.

The normalised mean square error is defined by

$$\text{NMSE} = \frac{\overline{(\overline{Q_o} - \overline{Q_p})^2}}{\overline{Q_o Q_p}}. \quad (3.87)$$

It is a measure of the mean relative scatter of the model results. A perfect model would have no scatter and thus have $\text{NMSE} = 0$. For the sake of understanding let $\text{NMSE} = 1$, implying that $\overline{(\overline{Q_o} - \overline{Q_p})^2} = \overline{Q_o Q_p}$. Now assume that the observed data and predicted results are uncorrelated (this of course means that the model is very poor at predicting the observed results) allowing $\overline{Q_o Q_p} = \overline{Q_o} \overline{Q_p}$. If we further assume that $\overline{Q_p} = \overline{Q_o}$ then we have $\overline{(\overline{Q_o} - \overline{Q_p})^2} = \overline{Q_o}^2$. That is, the root-mean-square error is equal to the mean. As the NMSE gets larger, so too does the scatter in the predicted results. Chang and Hanna (2004) suggest an acceptable value of $\text{NMSE} < 4$. Note that since the error in equation (3.87) is squared the NMSE is susceptible to outliers.

The fraction of data within a factor of 2 of the observations is determined by

$$\text{FAC2} = \text{fraction of data that satisfy: } 0.5 \leq \frac{Q_p}{Q_o} \leq 2.0. \quad (3.88)$$

It is a robust performance measure as it is not susceptible to outliers in the data. A perfect model would have $\text{FAC2} = 1$, that is one hundred percent of all data would be within a factor of two of the observations. Chang and Hanna (2004) suggest $\text{FAC2} > 0.5$ for an acceptable model.

The last performance measure to be used is the normalised absolute error,

$$\text{NAE} = \frac{|\overline{Q_o - Q_p}|}{0.5(\overline{Q_o} + \overline{Q_p})} = \text{FB}_{\text{fp}} + \text{FB}_{\text{fm}}. \quad (3.89)$$

Like the FAC2 it is not susceptible to outliers. A perfect model would have $\text{NAE} = 0$. No value for the NAE for an acceptable model could be found in the literature. As with all error though we seek to minimise the NAE. Table 3.2 summarises the values of the four main statistical evaluations used in this work for a perfect model and for an acceptable model.

Performance Measure	Perfect Model	Acceptable Model
FB	0	-0.3 - 0.3
NMSE	0	< 4
FAC2	1	> 0.5
NAE	0	minimised

Table 3.2: Values of the performance measures for a perfect model and an acceptable model.

The acceptable values of the FB, NMSE, and FAC2 reported by Chang and Hanna (2004) were determined by evaluating the performance measures reported in many field experiments and model simulations, and determining which of these experiments and simulations had “good” results (quotation

marks are also in their paper). While “good” is a very subjective term, the simulation results of **MEANS** and **SPMMM** presented in the next three chapters support the acceptable values listed in Table 3.2 put forth by Chang and Hanna (2004) in that the profiles and transects of various quantities (e.g., mean concentration, concentration variance, etc.) for simulations which produced performance measures in the acceptable ranges looked qualitatively similar to the observed profiles and transects.

Chapter 4

Dispersion in a Neutral Wall Shear Layer Flow

4.1 Experimental and Computational Setup

To evaluate the SPMMM model suite we simulated the concentration fluctuations due to a continuous point source emitting into a neutral wall shear layer flow, which is arguably the simplest regime of atmospheric turbulence. In particular, in the constant stress layer well-known analytic profiles of the key observables (mean wind speed, shear stress, TKE and its dissipation rate) satisfy appropriately simplified governing equations. The Fackrell and Robins (1982; FR82) experiments were designed to investigate the effects of the source size on the concentration fluctuations. They were carried out in the Marchwood Engineering Laboratories' open-circuit wind-tunnel, measuring $24 \times 9.1 \times 2.7$ metres. A neutral wall shear layer, corresponding to a natural atmospheric boundary layer, was grown within the wind-tunnel. The boundary layer depth was $\delta = 1.2$ m, the mean streamwise velocity at the top of the boundary layer (the free-stream velocity) was $\langle u \rangle_\delta = 4.0 \text{ m s}^{-1}$, resulting in a Reynolds number of $\text{Re}_\delta \approx 320\,000$ (based on δ and $\langle u \rangle_\delta$). The friction velocity was reported as $u_* = 0.188 \text{ m s}^{-1}$, the roughness length as $z_0 = 2.88 \times 10^{-4} \text{ m}$.

A neutrally buoyant mixture of propane and helium was released isokinetically from horizontally-oriented ground-level and elevated point sources. The ground-level sources had diameters of 3, 9, and 15 mm, while the elevated sources had diameters of: 3, 8.5, 9, 15, 25 and 35 mm. The height of the elevated source was $z_s = 0.19\delta = 0.228$ m. Propane concentration measurements were extracted at several downstream locations with a modified flame-ionisation detection system described in Fackrell (1980). These data were used to evaluate the **SPMMM** model suite.

The measured velocity statistics and the TKE dissipation rate for the flow are shown as symbols in Figure 4.1. The solid lines are the profiles used to drive the **SPMMM** model suite. For the mean streamwise velocity and the TKE dissipation rate, the solid lines represent equations (2.39) and (2.42) respectively. These show that the assumptions made in deriving these two equations are valid. For the velocity variances and the covariance, the solid lines represent the best fit to the measured data. Since the fetch available for measurements within the wind-tunnel was relatively short there was no significant streamwise evolution of the velocity statistics where the data were extracted. The flow was therefore assumed to be fully-developed, as well as being horizontally-homogeneous and stationary.

4.2 Parameter Calibration of **MEANS**

Since the **SPMMM** micromixing model and the **MEANS** pre-calculation model must agree at the level of first-order statistics, we can use **MEANS** for the initial calibration of the model suite. The tunable parameters in **MEANS** are: the Kolmogorov constant, the timestep, and the initial source distribution. In




Figure 4.1: Dimensionless velocity statistics and TKE dissipation rate for the FR82 neutral wall shear layer flow. The symbols are extracted from Figure 1 of Fackrell and Robins (1982). The lines are the fitted profiles used to drive the **MEANS** and **SPMMM** models. In the case of the mean streamwise velocity and the TKE dissipation rate, the solid lines represent equations (2.39) and (2.42) respectively. For the stresses, the solid lines represent the best fit to the experimental data.

the simulations used in the calibration of the model, the spatial resolution was set higher than a typical simulation to minimise as much as possible the discretisation error caused by the spatial resolution. This was possible because the conditional mean concentrations were not needed and thus velocity space was treated as a single large bin ($N_u = N_v = N_w = 1$), which freed computational resources for physical space discretisation. The numbers of bins in the streamwise, spanwise, and vertical directions were 100, 60, and 60 respectively. The particle number in the simulations was $N_\phi = 2 \times 10^7$. One hundred vertical bins ($N_z^{\text{vel}} = 100$) were used to discretise the velocity statistics. The elevated point source had a diameter of $d_s = 8.5$ mm and the initial source distribution was set to $\sigma_0 = 0.5d_s$, with the exception of the initial source distribution tuning simulations, where σ_0 varied.

Of all of the free parameters in the model, the Kolmogorov constant arguably has the greatest influence on the model since it appears in both the stochastic term (equation (3.5)) and the deterministic drift term (equation (3.12)) of the LS model. We therefore begin by tuning it. Several runs of the **MEANS** pre-calculation program were carried out with differing values of C_0 . The results from three of these simulations can be seen in Figure 4.2. This figure shows the vertical profiles of the normalised mean concentration on the plume centreline at five downstream positions. The symbols represent the FR82 wind-tunnel measurements and the lines are the **MEANS** results. Note that at each position the data has been normalised by maximum mean concentration at that position, hence in each panel the data ranges from zero to one. The advection timescale is calculated as $\tau_a = x/(\langle u \rangle(z_s))$. The ratio $\tau_a(z_s)/T_L(z_s)$ (displayed as τ_a/T_L on the figure) is an estimation of how many

Lagrangian integral timescales have passed since a particle's release. Values less than one suggest that the particle remembers the conditions of its release. Of the three simulations shown in Figure 4.2, the results from the $C_0 = 6.0$




Figure 4.2: Vertical profiles of the normalised mean concentration on the plume centreline for an elevated point source at five downstream locations in the FR82 flow. The open circles are from the FR82 wind-tunnel experiments. The lines are simulation results from MEANS with differing values of the Kolmogorov constant. The source diameter was 8.5 mm and the source height was 0.19δ .

simulations fit the FR82 the best, particularly near ground-level.

By comparing the vertical profiles of the normalised mean concentration from MEANS with those from the FR82 experiments, the various performance measures can be calculated. Figure 4.3 shows the behaviour of four performance measures (FAC2, FB, NMSE, NAE), based upon the 37 FR82 data points in Figure 4.2, in response to altering the Kolmogorov constant. To re-




Figure 4.3: Behaviour of the performance measures for MEANS simulations of dispersion from an 8.5 mm elevated point source in the FR82 flow in response to altering the Kolmogorov constant. The performance measures were calculated by comparing the profile of normalised mean concentration from MEANS with those from the FR82 experiments.

duce clutter on the plots (both below and for the remainder of the thesis), only the FB is shown. In situations where $\text{FB} \approx 0$, the values of FB_{fp} and FB_{fn} will be reported. The FB is near its optimal value of zero for $C_0 = 4.5 - 5.0$. However, this is due to cancellation of FB_{fp} and FB_{fn} . For example, for $C_0 = 5.0$, $\text{FB}_{\text{fp}} = 0.0530$ and $\text{FB}_{\text{fn}} = 0.0531$. The NMSE and the NAE are optimised for $C_0 \approx 6.0$ and FAC2 is closest to its optimal value of one for $C_0 = 4.0 - 5.0$. The FAC2 value at $C_0 = 6.0$ is 0.892 compared with 0.946 at $C_0 = 4.5$. This difference arises from the top two data points in the fifth panel of Figure 4.2. Given that the $C_0 = 6.0$ simulations produced the best fit to the FR82 data (particularly near ground-level) and the performance measures are optimised (or near to it) at this value of C_0 , the Kolmogorov constant will be set to this value for the remainder of the simulations of the FR82 flow, and for most of the other simulations in later chapters.

The effects of altering the timestep were investigated by running several MEANS simulations with various timesteps, achieved by altering the timestep constant. Figure 4.4 shows the vertical profiles of normalised mean concentration for three timesteps, with timestep constants of: $\mu_t = 0.02$, $\mu_t = 0.10$, and $\mu_t = 0.20$. There is little difference amongst the profiles in the left two panels (nearer to the source) before there has been significant contact between the plume and the ground. In the right three panels, the ground-level performance of the simulations with smaller timesteps is seen to be better, suggesting that smaller timesteps improve the performance of the reflection algorithm. In the last panel on the right, the performance of the simulation with $\mu_t = 0.02$ shows the best performance over all heights.

Figure 4.5 shows the performance measures from the timestep tuning sim-




Figure 4.4: Vertical profiles of the normalised mean concentration on the plume centreline for an elevated point source at five downstream locations in the FR82 flow. The open circles are from the FR82 wind-tunnel experiments. The lines are simulation results from **MEANS** with differing values of the timestep. The source diameter was 8.5 mm and the source height was 0.19δ .

ulations. The performance of MEANS can be seen to improve with decreasing timestep. The FB, NMSE, and NAE steadily decrease while the FAC2 remains approximately constant. Despite the fact that better performance is




Figure 4.5: Behaviour of the performance measures for MEANS simulations of dispersion from an 8.5 mm elevated point source in the FR82 flow in response to altering the timestep. Recall that $\Delta t = \mu_t \min(T_{L_u}, T_{L_v}, T_{L_w})$ in the MEANS pre-calculation program.

realised with a smaller timestep there are also practical considerations; smaller timesteps lead to longer simulations. We therefore set $\mu_t = 0.02$ to provide accurate results with reasonable run times. With this timestep, a 10^7 particle MEANS simulation of the FR82 elevated source experiments takes approximately two hours on a 3.0 GHz Pentium 4 processor.

Figure 4.6 shows that the vertical profiles of the normalised mean concen-

tration produced by **MEANS** are not sensitive to the initial source distribution, at least at these five downstream locations. Cassiani, Franzese and Giostra (2005a) also noted that the source size “does not significantly influence the mean field”. However, it is expected that the initial source distribution would have a great effect on the concentration fluctuations, particularly near the source, where meandering of the instantaneous plume is principally responsible for giving rise to concentration fluctuations. **MEANS** however cannot predict concentration fluctuations and thus the tuning of initial source distribution must be performed using **SPMMM**.

4.3 Model Consistency

There are two types of consistency expected from the **SPMMM** micromixing model. First, the velocity statistics of the N particles should reproduce the driving velocity statistics. Second, there should be a consistency between the first-order concentration statistics produced by **MEANS** and the first-order concentration statistics produced by **SPMMM**. For these consistency simulations the Kolmogorov constant was set to $C_0 = 6.0$, the initial source distribution was $\sigma_0 = 0.5d_s$, and the timestep constant was $\mu_t = 0.02$. Since the IECM model does not suffer from a spurious flux (as displayed by equation (3.29)), the micromixing timescale parameters should not affect the first-order consistency between **MEANS** and **SPMMM**. For these simulations we set the micromixing model’s free parameters to $C_r = 0.30$, the value used by Borgas and Sawford (1994) and Cassiani, Franzese and Giostra (2005a), and $\mu = 0.70$. However, as stated above, the micromixing model should not affect first-order concentration statistics, so these are arbitrary values. In order to preserve the computational




Figure 4.6: Vertical profiles of the normalised mean concentration on the plume centreline for an elevated point source at five downstream locations in the FR82 flow. The open circles are from the FR82 wind-tunnel experiments. The lines are simulation results from MEANS with differing values of the initial source distribution (σ_0). The source diameter was 8.5 mm and the source height was 0.19δ .

resources for the velocity space resolution, the physical space resolution was reduced to $N_x = 40$, $N_y = 40$ and $N_z = 40$. The driving velocity statistics had the same vertical resolution $N_z^{\text{vel}} = 40$. No velocity component was favoured, thus $N_u = N_v = N_w$. Six simulations with 5^3 , 10^3 , 15^3 , 20^3 , 25^3 , and 30^3 velocity bins were carried out.

To produce equivalent statistics, the number of particles used in **SPMMM** must be greater than the number used in **MEANS**. This is because in **MEANS** the particles all originate in the source region whereas in **SPMMM** the particles originate anywhere on the upstream face of the simulation domain. Thus, the probability of a particle passing through a specific region of the plume in **SPMMM** is much lower than the corresponding probability for **MEANS**. The magnitude of this effect is much larger near the source, where the plume dimensions are much smaller. To this end, 2×10^7 particles were used in the **MEANS** simulations and 10 ensembles of 10^7 particles were used for the **SPMMM** simulations.

The consistency of **SPMMM** with the velocity statistics was checked at two locations: a point away from the reflection boundary, $(x, y, z) = (0.6, -0.5, 0.3)$; and a point near the reflection boundary, $(x, y, z) = (5.0, 0.0, 0.05)$. This point was chosen to evaluate the effect of the reflection algorithm on the consistency of the model. Table 4.1 displays the results. At both locations **SPMMM** reproduced very well the driving velocity statistics, as required.

As discussed in Section 3.2, the IECM model does not suffer from a spurious flux and thus does not alter first-order concentration statistics. Therefore, the mean concentration profiles from **MEANS** and **SPMMM** should (in theory) be identical. In practice, the agreement of the first-order statistics will depend on the resolution of the model. Recall from Chapter 3 the relationship between

Statistic	Input Value	SPMMM Value	Input Value	SPMMM Value
$\langle u \rangle$	3.29	3.30	2.18	2.19
$\langle v \rangle$	0.00	0.0002	0.00	0.0002
$\langle w \rangle$	0.00	0.0045	0.00	-0.001
σ_u^2	0.1042	0.1052	0.191	0.194
σ_v^2	0.0619	0.0624	0.0955	0.0961
σ_w^2	0.0460	0.0465	0.0534	0.0540
$\langle u'w' \rangle$	0.0289	0.0281	0.0345	0.0343
	$(x, y, z) = (0.6, -0.5, 0.3)$		$(x, y, z) = (5.0, 0.0, 0.005)$	

Table 4.1: Comparison of the input driving velocity statistics with the velocity statistics produced by SPMMM at two locations, one far away from the reflection boundary (columns 2 and 3) and one just above the reflection boundary (columns 4 and 5).

the unconditional and conditional mean concentrations

$$\langle \phi \rangle = \int_{\mathbf{v}} \langle \phi | \mathbf{v} \rangle f_{\mathbf{u}} d\mathbf{v}. \quad (3.22)$$

The velocity space in SPMMM is discretised into $N_u \times N_v \times N_w$ bins. As a particle propagates downstream in the simulation domain it samples the discretised conditional mean concentration field, rather than the continuous space conditional mean concentration field as shown in the above equation. This results in numerical error, an investigation of which is the main focus of this section.

Figure 4.7 shows the vertical profiles of the mean concentration on the plume centreline at five downstream locations for the 10^3 and 30^3 velocity bin simulations, as produced by MEANS and SPMMM. Note that the data have not been normalised. The consistency between the two models is clearly better for the 30^3 velocity bin simulation than for the 10^3 velocity bin simulation. The discrepancy between MEANS and SPMMM in the left-most panel of both plots is due to the number of particles used in the SPMMM simulations, and thus to the

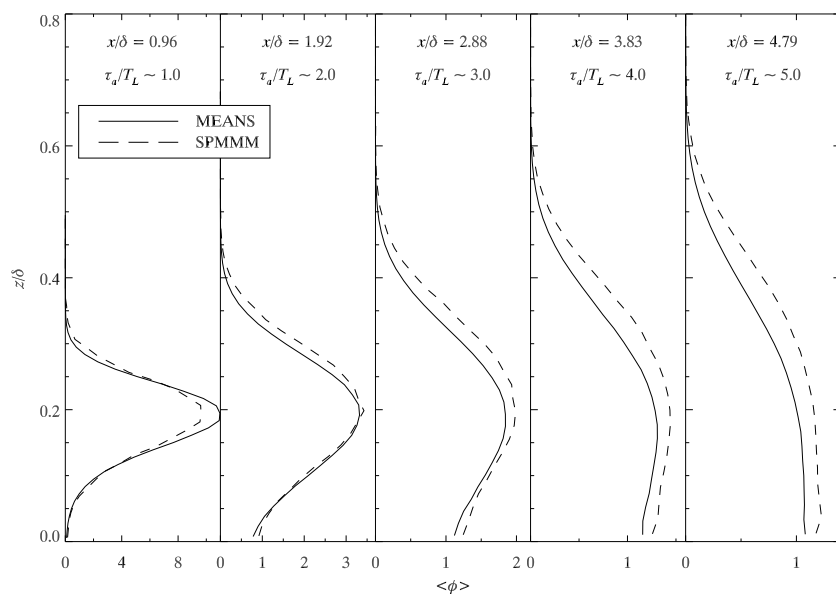
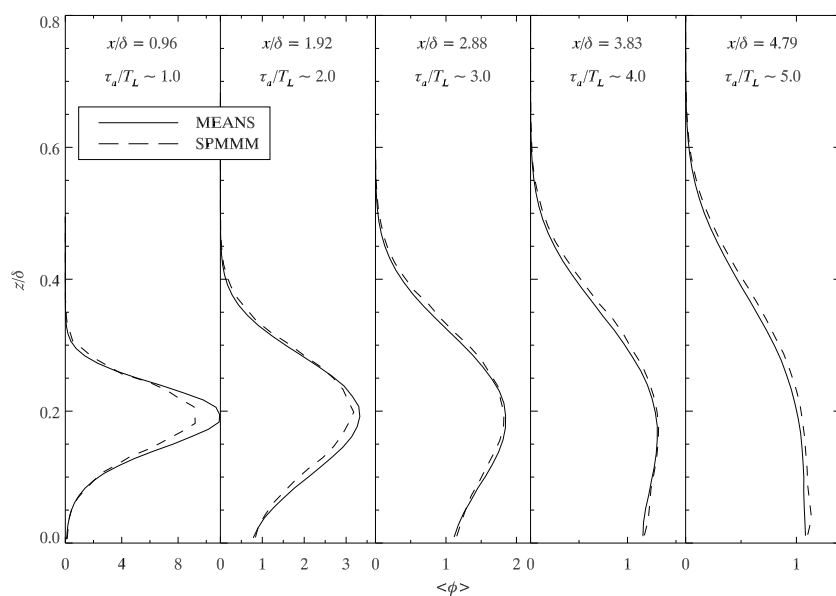
(a) 10^3 velocity bins(b) 30^3 velocity bins

Figure 4.7: Vertical profiles of the mean concentration on the plume centreline for an elevated point source at five downstream positions in the FR82 flow produced by MEANS and SPMMM for the 10^3 and 30^3 velocity bin simulations. The source diameter was 8.5 mm and the source height was 0.19δ . The consistency between MEANS and SPMMM increases with the number of velocity bins used in the simulations.

number of particles that originate in the source region. When fewer particles are used this discrepancy is worse.

Figure 4.8 shows four performance measures for six **MEANS** and **SPMMM** consistency check simulations. The performance measures were calculated from the data in Figure 4.7 and the equivalent data from other simulations. In order to show the variation of the performance measures for larger numbers of velocity bins, the performance measures for the 5^3 velocity bin simulations were left off of the figure as they were orders of magnitude different than the values from the other simulations. They were: $\text{FB} = -1.68$, $\text{NMSE} = 29.31$, $\text{FAC2} = 0.00535$, and $\text{NAE} = 1.69$. When there are fewer than 15^3 velocity bins the first-order consistency between **MEANS** and **SPMMM** worsens. At 20^3 the four performance measures level off and hold approximately the same values for larger numbers of velocity bins. This is because at approximately 20^3 velocity bins the numerical integral of the velocity PDF is unity (to within the numerical precision of the computer),

$$\sum_{L=1}^{N_u} \sum_{M=1}^{N_v} \sum_{N=1}^{N_w} f_{\mathbf{u}}(x_I, y_J, z_K, u_L, v_M, w_N) \Delta u \Delta v \Delta w = f_{\mathbf{u}}(x_I, y_J, z_K) = 1. \quad (4.1)$$

Therefore the normalisation constants for the conditional residence times (N_{ϕ}^v) computed by equation (3.67) are accurate. This is fortunate as it allows **SPMMM** to be run at a reasonably low velocity space resolution thereby allowing a greater spatial resolution. However, a more distorted velocity field (e.g., canopy flow) may require higher velocity space resolution. For the remainder of the FR82 simulations 20^3 velocity bins are utilised. We found that up to five percent error in equation (4.1) (i.e., 0.95–1.05) had little effect on the results of the models.

Recall in Section 3.3.1 it was mentioned that the process of estimating the




Figure 4.8: Behaviour of the performance measures for the MEANS and SPMM first-order consistency check simulations in response to the velocity space resolution. The simulations were for an 8.5 mm elevated point source in the FR82 flow. The performance measures were calculated from the data in Figure 4.7 and the equivalent data from other simulations.

value of the velocity PDF in a situation where the resolution of the driving velocity statistics is greater than the resolution of the conditional concentration bin field (see Figure 3.6) can lead to error. Figure 4.9 demonstrates this error for the case where $N_z = 40$ and $N_z^{\text{vel}} = 100$. Away from ground-level,




Figure 4.9: Vertical profiles of the mean concentration on the plume centreline for an elevated point source at five downstream locations in the FR82 flow. The source diameter was 8.5 mm and the source height was 0.19δ . This figure shows the inconsistency between MEANS and SPMMM that arises as a result of a poor estimation of the velocity PDF in bins near ground-level, where the driving velocity statistics vary rapidly in the vertical direction. This inconsistency arises due to the poor estimation of $f_{\mathbf{u}}$, and thus N_{ϕ}^v , when $N_z^{\text{vel}} > N_z$.

where the velocity statistics vary relatively slowly, there are no major inconsistencies between MEANS and SPMMM. In contrast, near ground-level, there are major inconsistencies between MEANS and SPMMM, as a result of a poor estimation of $f_{\mathbf{u}}$ in these lower bins. A poor estimation of $f_{\mathbf{u}}$ in turn results in

a poor estimation of the conditional residence time normalising constant, N_ϕ^v (see equation (3.68)). The inconsistency is worse at the farther downstream locations for two reasons. First, as the plume grows, so too does the size of the individual conditional mean concentration bins and they thus span more velocity statistic bins resulting in a poorer estimate of $f_{\mathbf{u}}$. Second, once a particle enters the ground-level region its concentration will mix toward an improperly normalised conditional concentration. Since, in Figure 4.9, the SPMMM profiles show greater mean concentrations near ground-level than the MEANS profiles we can conclude that the conditional mean concentrations were too high, and thus the normalising constants were too low. If the particle eventually travels upwards it will take this excess concentration (and inconsistency) up with it. Fortunately, this error is easy to avoid by ensuring that $N_z^{\text{vel}} \leq N_z$.

4.4 Comparison to FR82 Experimental Data

Having identified optimal values of the free parameters for both MEANS and SPMMM, and having shown that they are consistent with one another, we now compare the results of the SPMMM micromixing model with the FR82 experimental data. For all simulations in this section, the Kolmogorov constant was set to $C_0 = 6.0$ and the timestep constant was $\mu_t = 0.02$. To determine the conditional mean concentration field, MEANS used $N_\phi = 2 \times 10^7$ particles. Ten ensembles of two million particles ($N = 10 \times (2 \times 10^6)$) were used in SPMMM simulations. Physical space was discretised into forty bins in each of the three spatial dimensions ($N_x = N_y = N_z = 40$). The driving velocity statistics were discretised into forty bins as well ($N_z^{\text{vel}} = 40$). Velocity space was discretised into twenty bins in each direction ($N_u = N_v = N_w = 20$).

Given the stochastic nature of the models, the resulting profiles sometimes contain statistical noise. This noise was removed with a fourth-order Savitzky-Golay smoothing filter. This filter was chosen as it minimises the bias introduced by the filter, preserves higher moments within the data, and in general, preserves heights and widths of the profiles being smoothed. Figure 4.10 displays a vertical profile of concentration variance, calculated with the raw data from a 10^7 particle SPMM simulation, and the corresponding profile calculated with the same data smoothed with the Savitzky-Golay filter. Also shown is a profile of vertical concentration variance calculated with the raw data from a 10^8 particle SPMM simulation. By comparing the smoothed 10^7 particle profile




Figure 4.10: This figure shows the vertical profiles of concentration variance from two SPMM simulations (raw data), and the corresponding profile smoothed with a fourth-order Savitzky-Golay filter.

with the raw 10^8 particle profile, it can be seen that the Savitzky-Golay filter is effective at reducing the noise in a signal while retaining the general shape and amplitude of the signal.

4.4.1 Dispersion from an Elevated Point Source

We begin the comparison of SPMMM with the FR82 experimental data by comparing concentration statistics from simulations of dispersion from an 8.5 mm elevated point source at a height of 0.19δ . The vertical profiles of the normalised mean concentration for an elevated point source at five downstream locations are shown in Figure 4.11. Also shown are the corresponding profiles from MEANS. Good agreement is seen overall between the predictions from the two models and the FR82 observations, at all downstream locations. By increasing the Kolmogorov constant this far-field ground-level fit can be improved, but at the cost of a poorer fit overall (see Figure 4.3). Since the data at each location has been scaled by the maximum mean concentration at the location, it can only be used to evaluate the shape of the profiles and not the magnitude of the mean concentration.

To examine the magnitude of the mean concentration, Fackrell and Robins (1982) provided the streamwise transect of the dimensionless maximum mean concentration, as displayed in Figure 4.12. The maximum mean concentration was made dimensionless by multiplying by the mean streamwise velocity at source height ($\langle u \rangle_s$) and the boundary layer depth squared, then dividing by the source strength. Good agreement between the FR82 results and the SPMMM simulations can be seen in the figure. Furthermore, we see that SPMMM is consistent with MEANS, as it should be since $\max(\langle \phi \rangle)$ is a first-order statistic.




Figure 4.11: Vertical profiles of the normalised mean concentration on the plume centreline for an elevated point source at five downstream locations in the FR82 flow. The open circles are from the FR82 wind-tunnel experiments. The lines are simulation results from MEANS and SPMMM. The source diameter was 8.5 mm and the source height was 0.19δ .




Figure 4.12: The streamwise transect of the dimensionless maximum mean concentration for an elevated point source in the FR82 flow. The open circles are from the FR82 wind-tunnel experiments. The lines are simulation results from MEANS and SPMMM. The source diameter was 8.5 mm and the source height was 0.19δ .

Figure 4.13 shows the plume halfwidths in the spanwise direction (δ_y) and the vertical direction (δ_z). Fackrell and Robins (1982) defined the plume halfwidth as the distance between the location of the maximum mean concentration and the location where the mean concentration is one-half its maximum value. The data were made dimensionless by scaling with the boundary layer depth. Fackrell and Robins (1982) observed that the transition of the plume




Figure 4.13: The streamwise transects of the spanwise halfwidth (δ_y) and the vertical halfwidth (δ_z) for an elevated point source in the FR82 flow as simulated by MEANS and SPMMM. The halfwidths have been scaled by the boundary layer depth (δ). The open symbols are from the FR82 wind-tunnel experiments. The lines are simulation results from MEANS and SPMMM. The source diameter was 8.5 mm and the source height was 0.19δ .

from an elevated form (i.e., $\max(\langle\phi\rangle)$ above ground-level) to a ground-level form (i.e., $\max(\langle\phi\rangle)$ at ground-level) “actually develops in a rather complex

manner”. Although not shown, the SPMMM simulations displayed this complex behaviour as well. Consequently, FR82 fitted the vertical profiles of mean concentration with a reflected Gaussian profile of the form

$$\phi(z) \sim \exp\left(-0.693\frac{(z+z_s)^2}{\delta_z^2}\right) + \exp\left(-0.693\frac{(z-z_s)^2}{\delta_z^2}\right), \quad (4.2)$$

where the vertical plume halfwidth was left as a fitting parameter, resulting in “a straightforward growth of the vertical spread with distance”. The same process was carried out with the SPMMM simulation results, and therefore the vertical halfwidth data and simulation results in Figure 4.13 represent the best fit to equation (4.2). The spanwise halfwidth profiles were computed according to the FR82 definition mentioned above, near the beginning of this paragraph.

Good agreement between SPMMM, MEANS, and the FR82 data is realised for the both the spanwise and vertical halfwidths. Since the halfwidths are related to the mean concentration, we once again expect MEANS and SPMMM to produce consistent results, as they did. For $x/\delta \lesssim 2.0$ the simulated plume grew more slowly than the experimental plume in the spanwise direction. Lower values of the Kolmogorov constant resulted in a better fit to the FR82 data in the near field at the cost of a poorer fit in the far field. After $x/\delta \gtrsim 2.0$ the agreement between the experimental and simulated spanwise plume halfwidths is excellent. There is good agreement between the simulated transect of the vertical plume halfwidth and the FR82 experimental observations. We note that equation (4.2) is sensitive to the value of δ_z , and that small changes in its value resulted in a better or worse fit to the experimental transect. However, the fit was always qualitatively acceptable.

To check that SPMMM does not affect the first-order concentration statistics, three simulations with varying values of the IECM model’s two free parameters

(μ and C_r) were performed. Figure 4.14 shows vertical profiles of mean concentration at five downstream positions for the three simulations. The shaded region in the figure represents plus or minus one standard error. As all the curves lie within a standard error of each other it is evident that the IECM model does not affect the first-order statistics.




Figure 4.14: Vertical profiles of the normalised mean concentration on the plume centreline for an elevated point source at five downstream locations in the FR82 flow. The shaded region represents plus or minus one standard error. The source diameter was 8.5 mm and the source height was 0.19δ . Differing values of IECM model's two free parameters have been used in each simulation to illustrate that the IECM model does not affect first-order statistics.

Having shown that SPMMM can accurately simulate the first-order concentration statistics from the FR82 experiments, we now move on to evaluate its performance in simulating the second-order concentration statistics. To

determine the optimal values of μ and C_r we compared the SPMMM simulation results with the streamwise transect of the concentration fluctuation intensity, defined by FR82 as $\max(\sigma_\phi)/\max(\langle\phi\rangle)$, where $\sigma_\phi = \sigma_\phi(x, y, z)$ is the standard deviation of the concentration. FR82 provided profiles for 3 mm, 9 mm, 15 mm, 25 mm, and 35 mm elevated sources and for 3, 9 and 15 mm ground-level sources. To calibrate μ and C_r , we utilised only the 9 mm data as it was closest to the 8.5 mm elevated source used to determine the vertical profiles of normalised concentration and normalised variance. The initial source distribution was set to $\sigma_0 = 0.7d_s$ for the tuning of the IECM models' free parameters. Once μ and C_r were tuned then σ_0 was tuned. Figure 4.15 shows an example of the influence μ and C_r can have over the SPMMM simulation results. From this figure it is clear that increasing the micromixing timescale results in larger concentration fluctuations.

Numerous SPMMM simulations were performed while varying the μ and C_r . Performance measures were calculated by comparing model results with the FR82 experimental results and then contouring in $\mu - C_r$ parameter space. The performance measures were based upon six FR82 data points shown in Figure 4.15. The available FR82 vertical profile of concentration variance data has 38 data points, but unfortunately it is normalised, with no information on the maximum concentration variance to denormalise it. As a consequence, it can only be used to evaluate relative magnitudes of the concentration variance, and therefore is not a suitable data set with which to determine μ and C_r . Figure 4.16 shows the contours of the FB, NMSE, FAC2, and NAE. In these figures lighter shades of grey signify better performance of the model.

Examining the FB contours, we see that for C_r in the range 0.30 – 0.60, FB




Figure 4.15: Streamwise transects of the concentration fluctuation intensity. The open circles are from the FR82 wind-tunnel experiments. The lines are simulation results from SPMM. The source diameter was 9 mm and the source height was 0.19δ . This figure shows the marked effects that the IECM model's two free parameters can have on the simulation results of second-order concentration statistics.




Figure 4.16: Contours in $\mu - C_r$ parameter space of the performance measures for the concentration fluctuation intensity. The data used to calculate the performance measures were from SPMM simulations of an elevated point source in the FR82 flow. The source diameter was 9 mm and the source height was 0.19δ . Lighter shades of grey signify better performance.

is optimised for $\mu \approx 0.77 - 0.67$ respectively. For $C_r = 0.30 - 0.60$, the NMSE is optimised for $\mu \approx 0.85 - 0.66$. The FAC2 contours do not provide much information regarding the optimal values of μ and C_r , since for many values of these two free parameters the resulting model profiles were with a factor of two of the FR82 observations. We can see that if $C_r \gtrsim 0.40$ and $\mu \gtrsim 0.65$ then $\text{FAC2} = 1$. The NAE contours perhaps provide the most definitive information on the optimal values of the free parameters as its contours have an island of optimisation, as opposed to the bands of optimisation seen in the FB and NMSE contour plots. The NAE is minimised for $\mu \approx 0.72 - 0.77$ and $C_r \approx 0.40 - 0.50$.

There was no single simulation whose values of μ and C_r simultaneously optimised the FB, NMSE, and the NAE. The FB was optimised for $(\mu, C_r) = (0.70, 0.50)$, the NMSE was optimised for $(\mu, C_r) = (0.75, 0.50)$, and the NAE was optimised for $(\mu, C_r) = (0.75, 0.45)$. The resulting streamwise transects of the concentration fluctuation intensity for these three simulations are shown in Figure 4.17. The results from the three simulations look similar. However, the run which optimised the NAE appears to have the best overall fit to the FR82 data and we therefore set $\mu = 0.75$ and $C_r = 0.45$ for the remainder of the simulations in this thesis (except for one in Chapter 5 where SPMMM will be compared with a previously reported model). These values differ slightly from the values used by Cassiani, Franzese and Giostra (2005a). For their model, the latter authors used $\mu \approx 0.65$ and $C_r = 0.30$. A reason for this difference in model tuning is suggested below.

With the optimised values of the micromixing model parameters set, the initial source distribution can be tuned. Six simulations with differing values




Figure 4.17: Streamwise transects of the concentration fluctuation intensity for an elevated point source in the FR82 flow. The open circles are from the FR82 wind-tunnel experiments. The lines are simulation results from SPMMM. The source diameter was 9 mm and the source height was 0.19δ . The SPMMM profiles shown are for the three runs that optimised the FB ($\mu = 0.70, C_r = 0.50$), the NAE ($\mu = 0.75, C_r = 0.45$), and the NMSE ($\mu = 0.75, C_r = 0.50$).

of σ_0 were performed. The resulting streamwise transects of concentration fluctuation intensity varied only close to the source, where meandering of the plume is the principal source of concentration fluctuations. The performance measures from these simulations are displayed in Figure 4.18. As the initial source distribution increased in size, the model results displayed more bias towards under-prediction (a positive FB). The near zero value of FB at $\sigma_0 = 0.9d_s$ is due to a cancellation of FB_{fp} and FB_{fn} . Both the NMSE and the NAE are optimised at $\sigma_0 = 0.8d_s$. The FAC2 is unity for all σ_0 investigated here. Based on these results we set the initial source distribution to $\sigma_0 = 0.8d_s$. For comparison, Cassiani, Franzese and Giostra (2005a) used $\sigma_0 \approx 0.82d_s$.

With the free parameters optimised, simulations of dispersion from elevated sources (ES) with diameters of 3 mm, 9 mm, 15 mm, 25 mm, 35 mm, and from a 15 mm ground-level source (GLS) were carried out and the resulting streamwise concentration fluctuation intensity profiles constructed. They are shown in Figure 4.19. As it was shown in FR82 that the relative fluctuation intensity for ground-level sources was not sensitive to the source size, we only display the results for a 15 mm ground-level source. The simulation results for the 3 and 9 mm ground-level sources agreed with the findings of FR82. The symbols in Figure 4.19 are FR82 observations and the lines are SPMMM simulation results. The profiles for the elevated sources all display the same trend, an initial rise followed by a slow decay. Since the model was tuned to the 9 mm profile it is not surprising that the corresponding profile displays the best fit to the FR82 observations. The initial rise of the concentration fluctuation intensity is too low for the 3 mm elevated source but captured reasonably well for the other sources. Farther downstream, it appears that the mixing is too vigorous for




Figure 4.18: Behaviour of the performance measures for SPMM simulations of dispersion from an 8.5 mm elevated point source in the FR82 flow in response to altering the initial source distribution. The performance measures were calculated by comparing the streamwise transect of concentration fluctuation intensity from SPMM with those from the FR82 experiments. The diameter of the source is denoted by d_s .

$0.80 \lesssim x/\delta \lesssim 2.0$. In this range the modelled fluctuations from the 15 mm, 25 mm, and 35 mm sources are approximately ten to fifteen percent too low. It is interesting to compare the resulting profiles from the 15 mm elevated source simulations to those from the 15 mm ground-level source simulations. Both profiles start at approximately the same value but the profile for the elevated source increased substantially before undergoing a slow decay whereas the profile for the ground-level source exhibits a very small rise and then stays approximately constant. Physically this is caused by increased stretching, twisting, and folding of the material lines due to increased velocity shear near ground-level. This figure shows that the parametrisation for the micromixing timescale captures reasonably well this effect. The **SPMMM** results displayed in this figure are very similar to the results of similar simulations performed by Cassiani, Franzese and Giostra (2005a).

Figure 4.20 shows vertical profiles of the normalised concentration variance on the plume centreline for an 8.5 mm elevated point source at five downstream locations from an **SPMMM** simulation with the optimised values of μ , C_r , and σ_0 . The area shaded grey represents plus or minus one standard error. Also shown for comparison are the corresponding results from Cassiani, Franzese and Giostra (2005a; CASS). The performance measures for this simulation are: $\text{FB} = 0.166$, $\text{NMSE} = 0.108$, $\text{FAC2} = 0.816$, and $\text{NAE} = 0.254$. In the first panel from the left ($x/\delta = 0.96$) the agreement between **SPMMM** and the FR82 measurements is generally quite good. The height of the maximum variance occurs a bit too low in the **SPMMM** results and the modelled variance near the very top of the plume is too low as well. Farther downstream, the modelled maximum variance is always at a greater height than the FR82 measurements,




Figure 4.19: Streamwise transects of the concentration fluctuation intensity for elevated point sources of various sizes, and a ground-level source, in the FR82 flow. The open circles are from the FR82 wind-tunnel experiments. The lines are simulation results from SPMMM. The source height was 0.19δ .

as seen in the last four panels. The modelled variance close to the ground is low compared with the FR82 measurements. The general shapes of the modelled profiles are very similar to the experimental profiles but they appear to be shifted upwards in height.




Figure 4.20: Vertical profiles of the normalised concentration variance on the plume centreline for an elevated point source at five downstream locations in the FR82 flow. The source diameter was 8.5 mm and the source height was 0.19δ . The open circles are from the FR82 wind-tunnel experiments. The lines are simulation results from SPMMM utilising optimised values of the free parameters: $\mu = 0.75$, $C_r = 0.45$, and $\sigma_0 = 0.8d_s$. The area shaded grey represents plus or minus one standard error. The CASS profiles are from Cassiani, Franzese and Giostra (2005a), whose model was tuned differently: $\mu \approx 0.65$, $C_r = 0.3$, and $C_0 = 5.0$

From Figure 4.20, it appears that mixing near the ground-level is too vigorous resulting in the concentration fluctuations being dissipated too rapidly.

Increasing the micromixing timescale (to slow the rate of mixing) resulted in a much better fit to the FR82 vertical profile of normalised concentration variance data, but at the cost of a poorer fit to the reported streamwise transect of concentration fluctuation intensity data. For example, **SPMMM** simulations with $\mu = 1.60$ and $C_r = 0.30$ produced an excellent fit to the observed vertical profiles of normalised concentration variance (FB = 0.010, NMSE = 0.014, FAC2 = 0.919, NAE = 0.086), but a very poor fit to the streamwise transect of concentration fluctuation intensity. If simulating normalised quantities then this is an acceptable solution. However, normalised data is rarely as useful of non-normalised data. It was thought that the poor ground-level performance may have been related to the reflection algorithm, but numerous investigations into it revealed no problems.

The Cassiani, Franzese and Giostra (2005a) vertical profiles of concentration variance displayed a better agreement with the FR82 measurements for $z/\delta \lesssim 0.2$. One possible explanation for this lies in the nature of the two models. The Cassiani, Franzese and Giostra (2005a) model utilised simultaneous particle trajectories, and the conditional mean concentration in a bin was computed at each time step by considering the concentrations and velocities of the particles that lay within that bin at a particular timestep. It is therefore conceivable that conditional mean concentration in the bin would vary slightly from step to step. Mixing towards this varying conditional mean concentration may result in increased concentration fluctuations. In contrast, the conditional mean concentrations in **SPMMM** are pre-calculated and remain the same throughout the simulation. For each step in a particular spatial bin a particle's concentration mixes with unchanging conditional mean concentra-

tion of that bin. If the particle's concentration reaches the conditional mean concentration then no more mixing will occur. In practise, this will not happen due to the exponential nature of the relaxation towards the conditional mean, however the particle's concentration will closely approximate the conditional mean concentration. If many particles that travel through this bin reach the conditional mean concentration then concentration fluctuations will decrease.

This hypothesis is supported by two pieces of evidence. First, the fit between SPMMM and the FR82 measurements in the left-most panel of Figure 4.20 is quite good. This panel corresponds to the earliest available travel time, and therefore the particles' concentrations have not yet had much time to mix towards the conditional mean concentration. Second, increasing the mixing timescale resulted in a better fit to the normalised data in the figure. A longer mixing timescale results in the particles' concentrations mixing more slowly to the conditional mean concentrations. This hypothesis may also explain why Cassiani, Franzese and Giostra (2005a) used smaller values of μ and C_r than were used with SPMMM.

Another possible source of the discrepancy (again related to the implementation of the models) is that by calculating the conditional mean concentration on the fly the Cassiani, Franzese and Giostra (2005a) model contains more bias error than SPMMM, which has the pre-calculated conditional mean concentration supplied to it. A bias towards over-prediction would improve the results of the SPMMM simulations, although it would be due to error and not model accuracy. Given that both the Cassiani, Franzese and Giostra (2005a) model and SPMMM utilise the same micromixing model, it is possible that the better results obtained by Cassiani, Franzese and Giostra (2005a) for the vertical profiles of

normalised concentration variance are due to bias and may be fortuitous.

4.4.2 Dispersion from a Ground-level Point Source

We now consider dispersion from a 15 mm ground-level source. As for the elevated sources, various profiles from SPMM will be compared with their counterparts from the FR82 experiments. Figure 4.21 shows the vertical profile of the normalised mean concentration on the plume centreline at five downstream locations. The filled circles are FR82 data in the downstream region $1.67 \leq x/\delta \leq 5.92$. The other symbols correspond to SPMM simulation results at the specified extraction locations. The vertical coordinate has been scaled by the vertical plume halfwidth to display the self-preserving nature of the plume (i.e., the shape of the plume is invariant). Excellent agreement between SPMM and the FR82 experiments is realised in this figure. The self-preserving nature of the plume is evident as the profiles from the five extraction locations lie over one another.

Similarly, the vertical profile of the normalised concentration variance from the SPMM simulations shows excellent agreement with the FR82 experimental data, as seen in Figure 4.22. The variance profile for the ground-level source fits the observations much better than the variance profile from the elevated source. Recall that the TKE production and dissipation rates increase towards the ground. This leads to increased mixing and dissipation of the concentration fluctuations near the ground. Evidently, the parametrisation of the micromixing timescale used in this thesis captures this effect.




Figure 4.21: Vertical profile of the normalised mean concentration on the plume centreline for a ground-level point source at five downstream locations in the FR82 flow in the downstream region $1.67 \leq x/\delta \leq 5.92$. The filled circles are from the FR82 wind-tunnel experiments. The other symbols are SPMMM simulation results. The source diameter was 15 mm. The vertical coordinate has been scaled by the vertical plume halfwidth to display the self-preserving nature of the plume.




Figure 4.22: Vertical profiles of the normalised concentration variance on the plume centreline for a ground-level point source at five downstream locations in the FR82 flow in the downstream region $1.67 \leq x/\delta \leq 5.92$. The filled circles are from the FR82 wind-tunnel experiments. The other symbols are SPMMM simulation results. The source diameter was 15 mm. The vertical coordinate has been scaled by the vertical plume halfwidth to display the self-preserving nature of the plume.

4.5 Error Analysis

In this section we investigate the statistical, bias, and discretisation errors in the SPMMM simulation results. These errors are investigated for both mean concentration and concentration variance for an 8.5 mm elevated point source utilising the vertical profiles of the respective quantity (i.e., Figures 4.11 and 4.20). As there are 37 data points for the mean concentration data, and 38 for the concentration variance data, only a subset of these will be shown. Figure 4.23 shows the statistical error (as the root-mean-square (RMS) of Σ) in mean concentration and concentration variance at three positions on logarithmic axes. The expected behaviour of the statistical error ($\Sigma \propto N^{-1/2}$) is

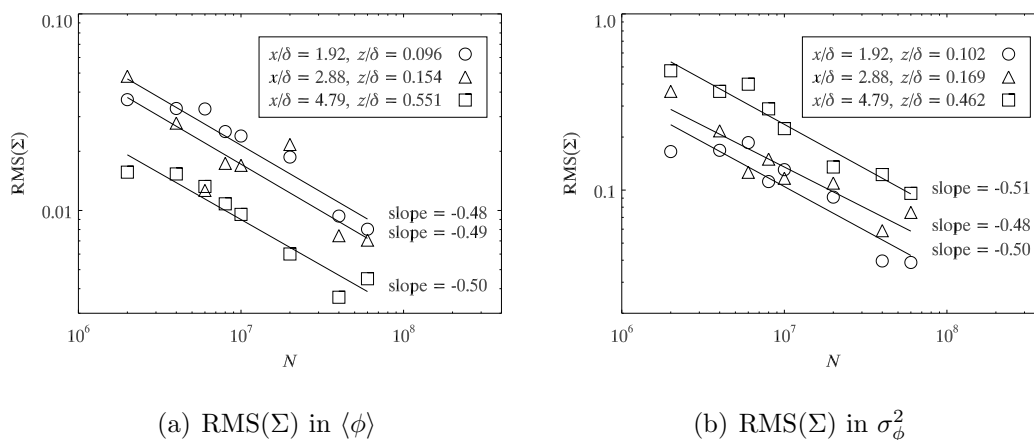


Figure 4.23: Statistical error in the mean concentration and the concentration variance at three positions from SPMMM simulations.

realised as all the slopes are close to $-1/2$ on the log-log plots. The spatial resolution was held constant at $M = 40^3$ for these simulations.

As discussed in Section 3.4, SPMMM should have minimal bias error due to the fact that the velocity statistics and the conditional mean concentration are supplied or calculated externally. Figure 4.24 shows the bias error for the

mean concentration and the concentration variance to demonstrate that it is approximately zero. It is believed that the small non-zero values represent the bias error contained in the conditional mean concentration field, or are a result of using a very large particle simulation ($N = 10^8$) to represent $Q_p^{\infty, M}$ in equation (3.80). Regardless, the bias error is one to three orders of magnitude smaller than the statistical error.

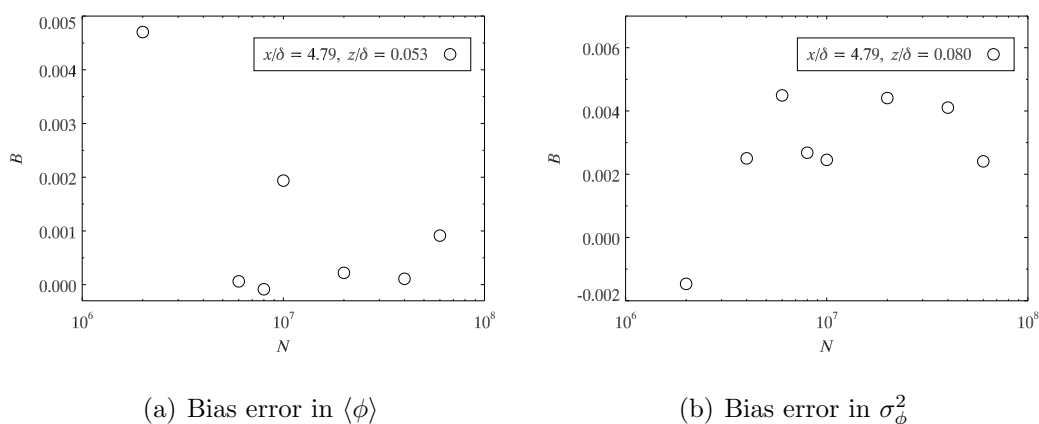


Figure 4.24: Bias error in the mean concentration and the concentration variance from SPMMM simulations.

Lastly we investigate the scaling behaviour of the discretisation error with respect to the number of bins used to discretise the physical space. To produce a better estimate, the discretisation error was calculated as the difference between the total error and the statistical error ($S + \epsilon_p = \epsilon - \Sigma$; see equations (3.76) and (3.77)). This estimation was possible due to the very small bias error. Note that we have assumed that the physical error is constant and therefore would not alter the scaling behaviour of the discretisation error. The simulations were conducted with different spatial resolutions while maintaining the number of particles per bin to keep the statistical error constant between simulations. Simulations with $M = 10^3$, 20^3 , 30^3 , and 40^3 spatial bins

were carried out. However, the available FR82 data is for the plume centre-line, so in reality we investigated the discretisation error along the z -axis only. Figure 4.25 displays the absolute value of the discretisation error in mean concentration and concentration variance plotted against the inverse number of vertical bins at three positions. From left to right, the four data points in each line of the plots correspond to $N_z^{-1} = 40^{-1}$, 30^{-1} , 20^{-1} , and 10^{-1} . In both mean concentration and concentration variance we see that the discretisation error decreases linearly with increasing resolution along the z -axis.

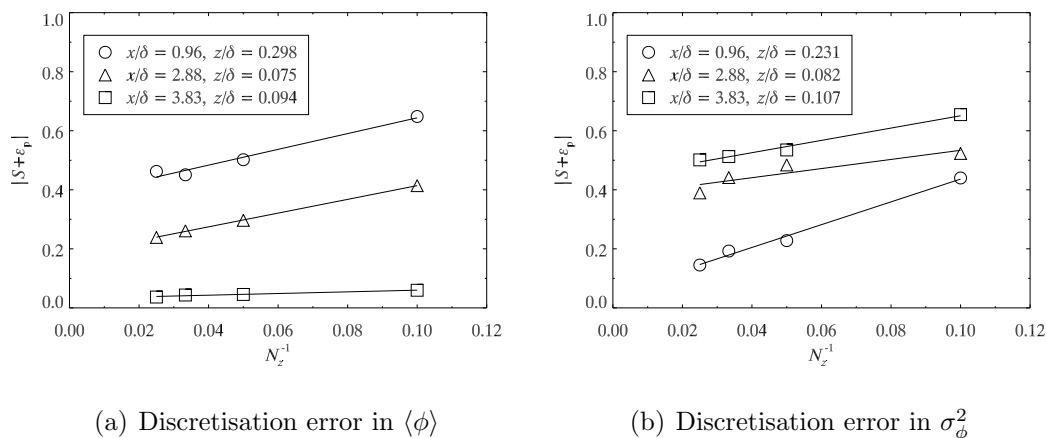


Figure 4.25: Scaling behaviour of the discretisation error in the mean concentration and the concentration variance from SPMM simulations.

4.6 Chapter Summary

In this chapter the ability of SPMM to simulate dispersion from elevated and ground-level point sources in the FR82 wall shear layer flow was investigated. By comparing vertical profiles of normalised mean concentration from MEANS simulations with those from the FR82 experiments, it was determined that $C_0 = 6.0$ produced the best fit to the experimental data and that the

performance of **MEANS** improved with decreasing timestep. However, based on the practical consideration of having reasonable simulation times, the timestep constant was set to $\mu_t = 0.02$.

MEANS and **SPMMM** have been shown to be consistent by comparing first-order concentration statistics. The modelled profiles of the normalised mean concentration, plume halfwidths, and the maximum mean concentrations for plumes resulting from dispersion of an 8.5 mm elevated point source were in very good agreement with the FR82 measurements. We found that **MEANS** and **SPMMM** were consistent provided that at least 15^3 velocity bins were used in the simulations. We therefore and thereafter discretised the velocity space into 20^3 bins to ensure consistency and to allow simulations with a higher spatial resolution. The modelled ground-level plume also had a vertical profile of normalised mean concentration that matched the FR82 measurements very well.

By comparing the streamwise transect of the concentration fluctuation intensity from **SPMMM** simulations to the corresponding FR82 measurements for a 9 mm elevated point source, the free parameters of the micromixing model have been optimised to $\mu = 0.75$ and $C_r = 0.45$. The optimal initial source distribution was found to be $\sigma_0 = 0.8d_s$. These optimised values resulted in an acceptable agreement with the measured concentration fluctuation intensity profiles for other elevated source sizes (3 mm 15 mm, 25 mm, 35 mm), as well as for a 15 mm ground-level source. The modelled profiles for the elevated sources were ten to fifteen percent lower than the measured values in the downstream range $0.80 \lesssim x/\delta \lesssim 2.0$.

The modelled vertical profiles of normalised concentration variance from an

8.5 mm elevated source were in reasonable agreement with the FR82 measurements, although the SPMMM produced variance levels near ground-level were too low compared to the measurements. It was hypothesised that this may be related to the fact that SPMMM mixes a particle's concentration towards the unchanging, pre-calculated conditional concentrations. Much better SPMMM profiles could be produced by setting $\mu = 1.60$ and $C_r = 0.30$, but at the cost of an unacceptably poor modelled streamwise transect of concentration fluctuation intensity. The modelled ground-level vertical profile of normalised concentration variance was much better than the elevated source results.

The error in the SPMMM micromixing model results for mean concentration and concentration variance behaved as expected. The statistical error went as $\Sigma \propto N^{-1/2}$, the bias error was approximately zero, and the discretisation error along the z -axis decreased with increasing spatial resolution as $|S| \propto N_z^{-1}$.

The SPMMM micromixing model displayed a good overall agreement with the FR82 measurements for dispersion of a passive scalar in a wall shear layer flow for elevated and ground-level sources. We now move on to simulating dispersion of a passive scalar in a more complicated canopy flow.

Chapter 5

Tombstone Canopy Simulations

5.1 Experimental and Computational Setup

The Tombstone Canopy consisted of a regular diamond-shaped array of thin billboard-like obstacles intended (by the originators, i.e., Raupach, Coppin, Legg) to represent a uniform plant canopy. Each obstacle was a rectangular aluminum tab measuring 10 mm in width, 1 mm in streamwise thickness, and $h_c = 60$ mm in height. The centre to centre streamwise and spanwise spacings of the tabs were 44 mm and 60 mm, respectively. This distribution of canopy elements resulted in a frontal area index of $\lambda_f = 0.23$ and a plan area index of $\lambda_p \approx 0$. As the canopy is regular, it can be geometrically represented by tiling a *unit cell* in the streamwise and spanwise directions, although the flow itself was not exactly periodic in the sense of repeating from cell to cell (i.e. there were flow gradients in the horizontal, albeit quite weak, on a scale exceeding the unit cell size). Figure 5.1 shows a unit cell in plan view. The lettered points in the figure correspond to velocity statistic data extraction locations used by Raupach et al. (1986) and Hilderman and Chong (2007). The dotted lines show the grid we have used to discretise the data. Both the points

and the grid are discussed below. Experimental data sets from the original wind-tunnel experiments and the subsequent water-channel experiments are considered in this thesis.




Figure 5.1: Plan view of a Tombstone canopy unit cell. The location and size of the canopy obstacles are shown by the thin, orange rectangles. Points A – J (blue points) are velocity statistic data extraction locations used by Raupach et al. (1986). Hilderman and Chong (2007) extracted velocity statistics not only at A-J, but also at the remaining points shown (viz. K-S; green points). Primed points share the u and w velocity statistics of the corresponding unprimed points, but have their v velocity statistics reflected about zero (e.g., $\langle v \rangle_{C'} = -\langle v \rangle_C$). The dotted lines show the grid used to discretise the data.

The results of the wind-tunnel experiments are contained in three papers: Raupach et al. (1986) describes the turbulence structure within and above the canopy*; Coppin et al. (1986) describes the results of scalar disper-

*There were two errors in the axes labels of Figure 6 in Raupach et al. (1986). The figure

sion experiments from an elevated plane source (which was in reality several line sources); and Legg et al. (1986) describes the results of scalar dispersion experiments from an elevated line source. The experiments were carried out at the Commonwealth Scientific and Industrial Research Organisation Pye Laboratory open-return blower tunnel with a working section measuring $10.6 \text{ m} \times 1.78 \text{ m} \times 0.65 \text{ m}$ in length, width, and height respectively. A neutral boundary layer was grown over a rough surface (consisting of gravel) and then encountered the Tombstone Canopy, which extended for three metres in the streamwise direction and covered the entire span of the wind-tunnel. Velocity statistics for the u and w components were extracted from the flow at points A – J using a three element hot-wire anemometer and a miniature one-component sonic anemometer. However, these measurements were restricted to heights $z < \frac{5}{3}h_c = 100 \text{ mm}$. At point D, velocity statistics were extracted up to $z = \frac{20}{3}h_c = 400 \text{ mm}$. These data are shown in Figure 5.2 as dashed lines. Also shown are the spatially-averaged velocity statistics from the water-channel experiments (circles and solid lines), which will be described in detail below.

For now we will focus only on the wind-tunnel profiles. The displayed profiles are polynomial interpolations to the data presented in Figure 6 of Raupach et al. (1986) and Raupach et al. (1987), as used by Cassiani et al. (2007) to drive their model. We chose to use these profiles as we wish to make a direct comparison of the **SPMMM** simulation results to the results of the Cassiani et al. (2007) model, and thus identical velocity statistics must be used. Since the Raupach et al. (1986) velocity data did not include any

is corrected in Raupach et al. (1987).

v -component measurements, the standard deviation of the spanwise velocity was estimated as $\sigma_v^2 = (\sigma_u^2 \sigma_w^2)^{1/2}$, as suggested by Brunet et al. (1994), and used by Cassiani et al. (2007).

The boundary layer depth at point D was $\delta = 540$ mm, and the free-stream velocity was $\langle u \rangle_\delta = 11.25$ m s⁻¹. The Reynolds number for the flow (based on δ and $\langle u \rangle_\delta$) was $\text{Re}_\delta \approx 4.0 \times 10^5$. Raupach et al. (1986) reported the friction velocity (based on the spatially-averaged[†] mean shear stress just above the canopy) as $u_* = 1.03$ m s⁻¹, the roughness length (deduced from the spatially-averaged mean streamwise wind profile well above the canopy) as $z_0 = 8.3$ mm, and the spatially-averaged mean streamwise velocity at canopy height as $\langle u \rangle_c = 3.40$ m s⁻¹. Based on $\langle u \rangle_c$ and h_c the Reynolds number was $\text{Re}_c \approx 1.4 \times 10^4$.

The source in the continuous, cross-wind, elevated line source experiments was a 0.9 mm diameter hot wire at a height of $z_s = 0.85h_c = 51$ mm, placed midway between two rows of obstacles. The source strength varied from $Q = 60 - 350$ W m⁻¹ from experiment to experiment. The heat generated by the wire was sufficiently low such that it did not alter the turbulence structure. This was ensured and confirmed by: insulating the floor of the wind-tunnel and measuring the heat flux through the floor; allowing the experimental apparatus to establish thermal equilibrium before measurements were made by leaving it running overnight; and by measuring the heat flux within the canopy elements and through planes normal to the mean flow. The measurements showed: the heat flux through the floor to be negligible; the integrated heat flux through planes normal to the mean flow to be conserved; that at a measurement loca-

[†]Spatial-averaging was performed over the unit cell area.

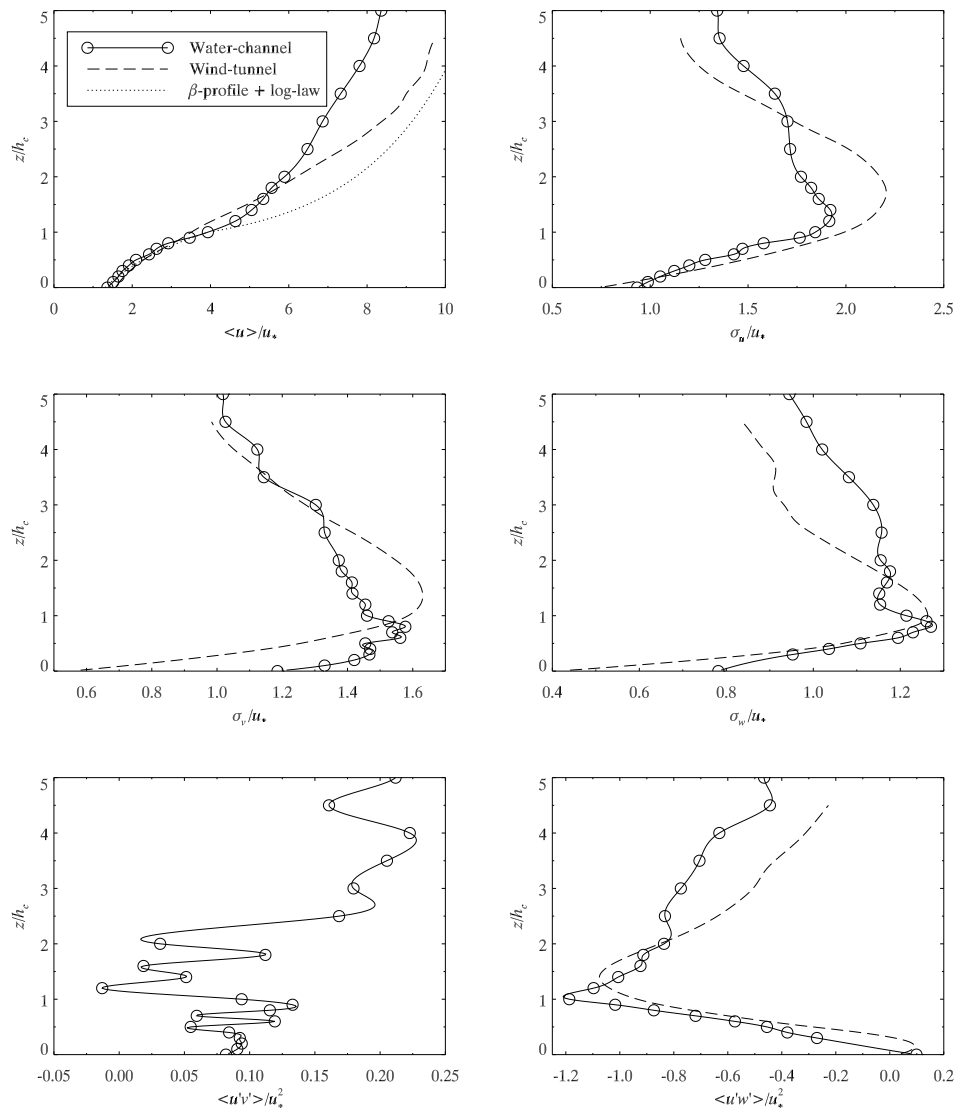


Figure 5.2: Vertical profiles of the spatially-averaged Tombstone canopy velocity statistics from water-channel experiments and the corresponding measurements from the wind-tunnel experiments. For the water-channel experiments, the symbols are flow measurements and the lines are the cubic spline interpolations used to drive the models. The dotted line in the vertical profile of the streamwise velocity represents equations (2.52) and (2.53), with $\beta = 1.1$ and a displacement height of $d = 43 \text{ mm} \approx 0.72h_c$, the value reported by Raupach et al. (1986). Note that for the wind-tunnel the σ_v profile was estimated as $\sigma_v^2 = (\sigma_u^2 \sigma_w^2)^{1/2}$ as no v -component velocity measurements were made.

tion 0.023 m from the source (i.e., the closest measurement location); the heat flux per unit ground area within the canopy elements was less than 5% the vertical heat flux in the air (Coppin et al., 1986; Legg et al., 1986).

The ratio of the buoyancy forces to the inertial forces (calculated as the Grashof Number/(Reynolds Number)²) close to the source was 0.0015, and the maximum gradient Richardson number 0.023 m from the source was -0.001 below the plume, and 0.005 above the plume (Legg et al., 1986). Therefore it is concluded with good certainty that the heat from the source was non-buoyant and a passive tracer, allowing the temperature fluctuations to be interpreted as concentration fluctuations. Concentration statistics were extracted with a cold-wire resistance thermometer and made dimensionless with

$$\phi_* = \frac{Q}{\rho c_p z_s \langle u \rangle_s}, \quad (5.1)$$

where c_p is the heat capacity of air at constant pressure and $\langle u \rangle_s$ is the mean streamwise velocity at source height.

Hilderman and Chong (2007) revisited scalar dispersion in the Tombstone Canopy using a near ground-level point source, a tracer dispersion configuration that was not examined by the original wind-tunnel study. The experiments were performed in the Coanda Research and Development Corporation (Burnaby, British Columbia) recirculating water-channel, with a test section measuring 10.0 m \times 1.5 m \times 1.0 m in length, width, and depth respectively. There were 210 rows of canopy elements (which were identical to those used in the wind-tunnel experiments) within the test section of the water-channel. As for the wind-tunnel experiments, the neutrally-stratified wall shear layer was fully-developed in the upstream fetch, before interacting with the canopy in the test section of the tunnel.

The flow was seeded with titanium dioxide particles to allow the extraction of velocity statistics with laser Doppler velocimetry (LDV). Measurements of the u , v , and w velocity components at the nineteen points A – S were made in the unit cell centred on the seventh tab of row 160, $x = 7.04$ m into the canopy and $y = 0.42$ m from the sidewall of the water-channel. At all points, velocity measurements were made up to a height of $3h_c = 180$ mm. For points D and J, velocity measurements were made up to a height of $12h_c = 720$ mm. The u and w velocity statistics of the primed points are identical to their corresponding unprimed points, but the v statistics have to be reflected about zero (e.g., $\langle v \rangle_{C'} = -\langle v \rangle_C$). For the work in this thesis, the Coanda velocity statistic data from the ten points A – J were spatially-averaged as

$$\frac{(A + B + 2C + D + E + (F + G)/2 + 2H + 2I + J)}{12}, \quad (5.2)$$

where each point has been weighted by the total fraction of the unit cell area occupied by the lot in which it sits[‡]. The spatially-averaged velocity statistics from the water-channel experiments are shown in Figure 5.2. The symbols are flow measurements and the lines are the cubic spline interpolations used to drive the models. From this figure we can see that the wind-tunnel and water-channel statistics are similar to one another within the canopy, but this similarity decreases above the canopy.

In the vertical profile of the mean streamwise velocity, the analytical profile as predicted by equations (2.52) and (2.53) is also shown for $\beta = 1.1$ and $d = 43$ mm $\approx 0.72h_c$ (the value reported by Raupach et al. (1986)). Within the canopy, the β -profile fits the experimental data quite well. The

[‡]To determine the lot areas, draw lines midway between all neighbouring points (only A – J). The area of the resulting rectangles is easily calculated.

above canopy fit, in the wall shear layer region, is poorer. For both data sets the streamwise velocity profile is attenuated more rapidly than predicted by equation (2.53). However, this discrepancy can be eliminated by tuning the displacement height. A value of $d = 10 \text{ mm} \approx 0.17h_c$ results in an excellent fit to the water-channel data, whereas a value of $d = 30 \text{ mm} = 0.5h_c$ results in an improved fit to the wind-tunnel data. These values for d are both less than the reported value, a finding also noted by Wilson et al. (1998).

From the spatially-averaged velocity statistics, the friction velocity at the top of the canopy for the water-channel simulations was determined to be $u_* = 0.018 \text{ m s}^{-1}$, and the roughness length was $z_0 = 9 \text{ mm}$. The boundary layer depth was $\delta = 540 \text{ mm}$ and the free-stream velocity was $\langle u \rangle_\delta = 0.175 \text{ m s}^{-1}$, giving a Reynolds number (based on δ and $\langle u \rangle_\delta$) of $\text{Re}_\delta = 9.4 \times 10^4$. At the top of the canopy the mean streamwise velocity was $\langle u \rangle_\delta = 0.07 \text{ m s}^{-1}$, giving $\text{Re}_c = 4.2 \times 10^3$. When scaled by their respective friction velocities, the wind-tunnel and water-channel velocity statistics are similar. Table 5.1 summarises physical and aerodynamic parameters from the wind-tunnel and water-channel Tombstone canopy experiments.

To drive simulations with locally inhomogeneous velocity statistics (where the velocity statistics change from one location to another within the unit cell but have no large-scale inhomogeneities), the point data from A–S have been interpolated onto a grid consisting of $N_x^{\text{vel}} = 22$ bins in the streamwise direction, $N_y^{\text{vel}} = 15$ bins in the spanwise direction (as shown in Figure 5.1 by the grid of dotted lines), and $N_z^{\text{vel}} = 40$ bins in the vertical direction. For each point, vertical profiles of the statistics have been produced using cubic spline interpolation. Below $3h_c$ each point has its own unique velocity statis-

Parameter	Symbol	Wind-tunnel	Water-channel
Source configuration		line	point
Source diameter	d_s	0.9 mm	15 mm
Source height	z_s	$0.85h_c$	$0.2h_c$
Friction velocity	u_*	1.03 m s^{-1}	0.018 m s^{-1}
Roughness length	z_0	8.7 mm	9 mm
Boundary layer depth	δ	540 mm	540 mm
Free-stream velocity	$\langle u \rangle_\delta$	11.25 m s^{-1}	0.175 m s^{-1}
Reynolds numbers	Re_δ	4.0×10^5	9.4×10^4
	Re_c	1.4×10^4	4.2×10^3

Table 5.1: Physical and aerodynamic parameters from the Tombstone canopy wind-tunnel and water-channel experiments. The Reynolds number was based upon the boundary layer depth and the free-stream velocity ($\text{Re}_\delta = \langle u \rangle_\delta \delta / \nu$). The canopy was identical for both the wind-tunnel and water-channel experiments: a canopy height of $h_c = 60$ mm; a streamwise obstacle spacing of 44 mm; and a centre to centre spanwise obstacle spacing of 60 mm

tics. Above, the average of the full height profiles (up to $12h_c$) from points D and J was used for each point. Therefore, below $3h_c$ the flow is locally inhomogeneous, and above $3h_c$ the flow quickly attains horizontal-homogeneity. Results from the interpolation procedure described above are shown for the mean streamwise velocity at points A and R in Figure 5.3. The transition from a locally inhomogeneous flow to a horizontally-homogeneous flow occurs between $z = 3h_c$ and $z = 3.5h_c$

The vertical profiles were then interpolated onto the $N_x^{\text{vel}} \times N_y^{\text{vel}}$ grid for each vertical level using thin plate splines. Thin plate splines were used because they: pass through the experimentally measured data points; grow linearly when far away from the control points (i.e., the experimental data); are smooth functions with continuous first derivatives; are able to model functions with complex local distortions, such as those which may be found within




Figure 5.3: Vertical profiles of the mean streamwise velocity from the Tombstone canopy water-channel experiments and the resulting interpolation.

a canopy. Figure 5.4 displays four streamwise transects of the mean streamwise velocity for $y = 0$ mm, and four spanwise transects for $x = 2$ mm, just behind the central obstacle, through point A. The horizontal axes are unscaled to facilitate comparison with Figure 5.1. From Figure 5.4(a) it is evident that $\langle u \rangle \neq 0$ inside the obstacle, located at $x = 0$ mm. This is true for the other statistics as well and is due to interpolation through the obstacle. This was permitted to avoid singularities in the Reynolds stress tensor and to prevent a particle getting “stuck” in the simulation domain (i.e., not propagating downstream to x_{\max}). For future obstacle resolving versions of SPMMM, the obstacles will be boundaries and reflection will occur when a particle encounters them.

In the Coanda water-channel experiments, a continuous point source with a diameter of $d_s = 15$ mm was positioned immediately behind the centre tab of row 116 at a height of $z_s = 12$ mm $= 0.2h_c$. Since the sharp sided, aluminum tabs are very efficient turbulence generators, Hilderman and Chong (2007) found that the resulting concentration field resulting from dispersion from a very small source was extremely sensitive to the exact placement within the canopy, and that the resulting plume could easily be biased to one side of the channel or the other by a slight change in the release angle or velocity. As this is undesirable from the standpoint of being able to reproduce experimental results, a relatively large source was used. The source was designed to have minimal momentum and had a fine mesh over the end, to ensure a uniform distribution of the dye over the source area. Due to the physical size of the tube that composed the source, the front of the source was 26 mm downstream from row 116. Sodium fluorescein dye was used as the dynamically passive tracer. The velocity of the dye at the source was $u_s = 2.264 \times 10^{-3} \text{ m s}^{-1}$, which



(a) Streamwise transects of $\langle u \rangle / u_*$ along $y = 0$ mm(b) Spanwise transects of $\langle u \rangle / u_*$ along $x = 2$ mm

Figure 5.4: Streamwise and spanwise transects of the interpolated mean streamwise velocity at four heights for the water-channel Tombstone canopy experiments. The horizontal axes are unscaled to facilitate comparison with Figure 5.1.

corresponds to a release rate of $Q = 24 \text{ mL min}^{-1}$. Sodium fluorescein is a weakly diffusive scalar, with a (molecular) Schmidt number of 1920 in water.

Laser induced fluorescence (LIF) techniques were used to extract concentration data 1, 2, 4, 7, 8, 12, and 16 rows downstream from the source at 7 to 10 heights, depending on the location. The measurements were made dimensionless by dividing by the source concentration, i.e., $\phi^* = \phi/\phi_s$ where

$$\phi_s = \frac{Q}{A_s u_s}, \quad (5.3)$$

with A_s representing the area of the source. Once dimensionless, higher-order statistical quantities were calculated. In this thesis, the dimensionless quantities are denoted with a superscript asterisk: $\langle \phi^* \rangle$ for the mean dimensionless concentration, σ_ϕ^* for the standard deviation of the dimensionless concentration, Sk_ϕ^* for the skewness of the dimensionless concentration, and Ku_ϕ^* for the kurtosis of the dimensionless concentration.

In addition to the driving velocity statistics, the TKE dissipation rate is needed to drive the models (see equations (3.17), (3.20), (3.36), and (3.38)). The left panel of Figure 5.5 shows three different determinations of the TKE dissipation rate from Figure 15 of Raupach et al. (1986): from the u -velocity spectra; from the w -velocity spectra; and from a residual calculation according to

$$\varepsilon = \mathcal{P}_s + \mathcal{P}_w + \mathcal{I}_t, \quad (2.56)$$

which was first encountered in Section 2.4.2. The lines in the figure are to aid the eye. The three measurements agree reasonably well within the canopy but display a factor of two discrepancy between the spectral and residual methods above the canopy. The right panel of Figure 5.5 shows the polynomial fit to

the far field Lagrangian integral timescale from Figure 10 of Legg et al. (1986), used by Cassiani et al. (2007) to infer the TKE dissipation rate to drive their model.

In their determination of the T_L profile, Legg et al. (1986) assumed the Lagrangian vertical velocity autocorrelation function to be exponential, and that the variances of Eulerian and Lagrangian vertical velocities were equal, which is strictly true only in homogeneous turbulence. Consequently, the profile is meant to provide some indication of the behaviour of T_L , not represent it exactly. No measurements of T_L were made above $3h_c$; above this height the value is held constant.

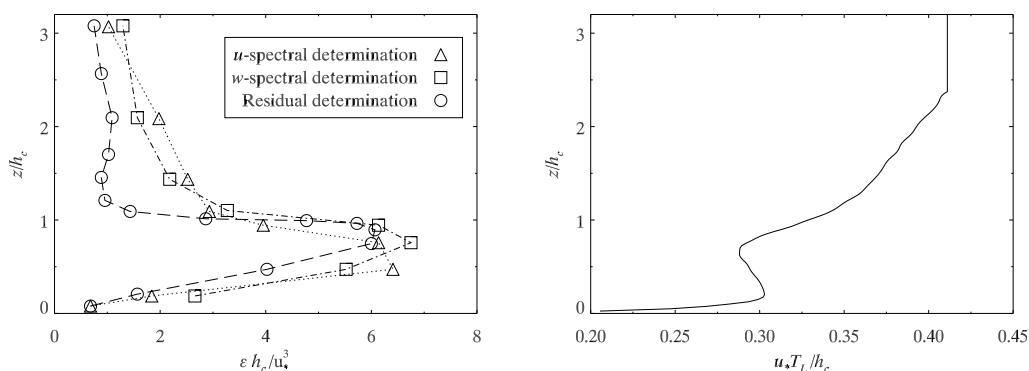


Figure 5.5: Vertical profiles of the TKE dissipation rate (left; lines added to aid the eye) and the far field Lagrangian integral timescale (right) from the Tombstone canopy wind-tunnel experiments. The TKE dissipation data are from three different determinations: from the u -velocity spectra; from the w -velocity spectra; and from a residual calculation according to equation (2.56). Above $3h_c$ there were no measurements of the Lagrangian integral timescale so the value was held constant. The data in the left hand panel are also shown in Figure 5.6 along with various formulae.

Figure 5.6 shows modelled, vertical profiles of the TKE dissipation rate, along with the analytical profile for a displaced wall shear layer (equation (2.44)) above the canopy. The u -spectral determination and (to a lesser extent)

the w -spectral determination agree quite well with the displaced wall shear layer profile for the range $z/h_c = 2-3$, lending support to the assumptions made in deriving equation (2.44). The residual determination does not show good agreement with the analytical profile. Cassiani et al. (2007) interpreted the Lagrangian integral timescale measurements of Legg et al. (1986) to be T_{L_w} and used the relation $T_{L_w} = 2\sigma_w^2/C_0\varepsilon$ to determine the TKE dissipation rate. They found that $C_0 = 2.0$ produced an acceptable fit to the measured TKE dissipation data, as shown in Figure 5.6. While at the lower end of the currently entertained range of C_0 , this value is substantially different than $C_0 = 6.0$ used in the FR82 simulations. However, the TKE dissipation rate implied from T_L with $C_0 = 6.0$ does not fit any of the in-canopy TKE dissipation measurement very well, as seen in the figure. A second method of calculating the TKE dissipation rate is with equation (2.56).

As shown in Chapter 2, the shear production and wake production of TKE are in closed form and are readily calculated from the available velocity statistics. However, the turbulent transport of TKE is unclosed and must be modelled, which has the undesirable side-effect of introducing more arbitrary constants into the model. We therefore chose simplicity and assumed local equilibrium and calculated the TKE dissipation rate as $\varepsilon = P_s + P_w$. Overall this method produced a better fit to the experimental data, and to the displaced wall shear layer profile, than the method of inferring a TKE dissipation rate from a profile of the Lagrangian integral timescale, as seen in Figure 5.6. Furthermore, it does so with no reference to (or dependence upon a chosen value of) the Kolmogorov constant. A Savitzky-Golay filter was used to remove noise from the profile. The assumption of local equilibrium within the canopy




Figure 5.6: Measured and modelled vertical profiles of the Tombstone canopy TKE dissipation rate.

is certainly incorrect. Raupach et al. (1986) note that turbulent transport was a major loss in the TKE budget just above the canopy and was the principal gain within the upper canopy. Thus, the inclusion of T_t in the calculation of the TKE dissipation rate would probably result in better fit to the experimental data. However, as shown by the resulting profile in Figure 5.6, the assumption of local equilibrium can be used to determine a reasonably accurate TKE dissipation rate, without the need of introducing further tunable parameters into the model.

5.2 Dispersion from a Continuous Line Source

In this section we report on results of simulations of dispersion from a continuous cross-wind elevated line source with the characteristics discussed above. The horizontally-homogeneous, spatially averaged, wind-tunnel measured, velocity statistics shown in Figure 5.2 were used to drive SPMMM. These were the same as the velocity statistics used by Cassiani et al. (2007). The vertical profile of TKE dissipation determined by assuming local equilibrium was used.

The use of a continuous line source effectively reduces the simulations to two dimensions allowing the use of fewer particles. To pre-calculate the conditional mean concentration field, we used $N_\phi = 2 \times 10^6$ particles. Similarly, since in a line source simulation there is a much greater chance of a particle being a source particle in SPMMM, we used $N = 2 \times 10^6$ particles. The simulations were performed at the same spatial, velocity, and temporal resolution as the FR82 simulations: 40^3 spatial bins, 20^3 velocity bins (which was enough to ensure that the velocity PDF integrates to unity; equation (4.1)), and $\mu_t = 0.02$. The

driving velocity statistics were discretised with forty bins as well, $N_z^{\text{vel}} = 40$. An initial source distribution of $\sigma_0 = 2.0d_s$ was found to produce the best fit (albeit not a great fit) to the experimental measurements of the standard deviation of concentration close to the source.

In Section 4.4.1, the micromixing model parameters had been optimised to $\mu = 0.75$ and $C_r = 0.45$ by simulating dispersion from an elevated point source in the FR82 wall shear layer flow. As discussed in Thomson (1996), Sawford (2004b), and Luhar and Sawford (2005), the micromixing constant will be larger for a line source than for a point source since mixing is more efficient for a point source plume. This is because for a line source plume entrainment only occurs in the vertical direction as opposed to the spanwise and vertical directions for a point source plume. Investigation showed that increasing the micromixing model constant to $\mu = 0.80 - 0.85$ resulted in a slightly better fit to the experimental data above the height of the maximum concentration standard deviation and a slightly poorer fit below. Overall, the performance measures did improve slightly upon increasing μ , but the difference was not marked, so we used $\mu = 0.75$ and $C_r = 0.45$ for the majority of the simulations to avoid arbitrarily re-tuning of the model. The Kolmogorov constant was specified as $C_0 = 6.0$, which is unchanged in view of its earlier “tuning” relative to the FR82 experiment.

Figure 5.7 displays the vertical profile of the dimensionless mean concentration at three downstream positions. As for the FR82 results, an estimation of the characteristic timescale τ_a/T_L is displayed in each panel of the figure. The grey shaded area represents plus or minus one standard error (note: in this plot the error is too small to see). The symbols are the experimental re-

sults from Legg et al. (1986; TOMB3) and the lines are the SPMMM results. The




Figure 5.7: Vertical profiles of the dimensionless mean concentration from SPMMM simulations of an elevated, cross-wind line source ($z_s/h_c = 0.85$) in the Tombstone canopy flow. Horizontally-homogeneous velocity statistics were used to drive this simulation. The symbols are wind-tunnel experimental results and the lines are SPMMM predictions.

performance measures for the data displayed in the figure are: $FB = 0.295$, $NMSE = 0.491$, $FAC2 = 0.729$, and $NAE = 0.415$, all within their respective limits for an acceptable model.

Qualitatively, the simulation matches the experimental data reasonably well. At all three locations the model produced profiles with a similar shape and magnitude to the experimental data. The fit is better above the canopy. This may be related to the Kolmogorov constant, as it is the only tunable

parameter that affects the mean concentration, as shown in Chapter 4. Poggi et al. (2008) investigated the behaviour of C_0 inside dense canopies. They found that C_0 is reduced by a factor of ~ 5 relative to its ASL value due to wake production, and a factor of ~ 1.5 due to finite Reynolds number effects. These reductions were offset by an increase in C_0 , due to a short circuiting of the energy cascade within the canopy, but not enough to compensate for the reductions. Although not shown, a re-tuning of the Kolmogorov constant to $C_0 = 4.0$ does improve the agreement between the experimental data and the present simulations.

The discrepancy may also be related to the driving velocity statistics, or the modelled TKE dissipation rate. An earlier single-particle LS model (Flesch and Wilson, 1992) produced more accurate results than those shown in Figure 5.7 (particularly below canopy at $x/h_c = 2.78$), probably because those authors had formulated their model in terms of the Lagrangian integral timescale (T_{L_w} ; as opposed to the TKE dissipation rate), and used the profile of T_{L_w} from Legg et al. (1986). As shown in Figure 5.6, in order to get an acceptable fit to the measured TKE dissipation rate data the Kolmogorov constant must be $C_0 \approx 2.0$, at the low end of its recently reported estimates. Since the water-channel Tombstone canopy measurements do not include any data on the TKE dissipation rate or the Lagrangian integral timescale, and for the reasons mentioned in Section 5.1 above, we opted to calculate the TKE dissipation rate as the sum of shear production and wake production. However, below we show the results of a simulation carried out using the TKE dissipation rate inferred from the Lagrangian integral timescale reported by Legg et al. (1986), when comparing our model to the model of Cassiani et al.

(2007).

Figure 5.8 displays three vertical profiles of the dimensionless standard deviation of concentration. Experimental data for the standard deviation




Figure 5.8: Vertical profiles of the dimensionless standard deviation of concentration from SPMMM simulations of an elevated, cross-wind, line source ($z_s/h_c = 0.85$) in the Tombstone canopy flow. Horizontally-homogeneous velocity statistics were used to drive this simulation. The symbols are wind-tunnel experimental results and the lines are SPMMM predictions. The grey shading (very small to invisible) represents plus or minus one standard error.

at $x/h_c = 0.38$ and $z/h_c \approx 1$ (left-most panel) was not presented in the original paper (Legg et al., 1986), and therefore cannot be presented here. The performance measures for these data are: $FB = 0.650$, $NMSE = 1.487$, $FAC2 = 0.625$, and $NAE = 0.680$. The FB is outside of its acceptable range due to the simulated standard deviation of concentration being significantly

under-predicted at $x/h_c = 0.38$. No amount of tuning of σ_0 could eliminate the under-prediction from the model at this location. This seems to suggest that the present parametrisation of the micromixing timescale cannot capture near source phenomena accurately, or that the velocity statistics are not sufficiently detailed or realistic. The other performance measures are within their respective acceptable ranges. Qualitatively, the model predictions and the experimental data show good agreement, especially for $z/h_c \lesssim 2$ at $x/h_c = 2.78$ and $x/h_c = 11.6$.

The streamwise transect of the ground-level dimensionless mean concentration is shown in Figure 5.9. The performance measures for these data are: $\text{FB} = 0.728$, $\text{NMSE} = 1.141$, $\text{FAC2} = 0.444$, and $\text{NAE} = 0.803$. For $x/h_c \lesssim 6$ the fit is poor with the mean concentration under-predicted. The $x/h_c = 2.78$ panel in Figure 5.7 displays the cause of these poor results. At this location **SPMMM** does not accurately simulate the ground-level mean concentration. For $x/h_c \gtrsim 6$ the simulations agree better with the experimental data.

The **SPMMM** predictions for the streamwise transect of the concentration fluctuation intensity (defined here as $\max(\sigma_\phi)/\max\langle\phi\rangle$) are much better, as seen in Figure 5.10. For the data shown in the figure, the performance measures are: $\text{FB} = -0.030$, $\text{NMSE} = 0.034$, $\text{FAC2} = 1.000$, and $\text{NAE} = 0.149$. In Figure 5.8 we saw that **SPMMM** accurately simulates the magnitude of the concentration standard deviation away from the source (the centre and right panels). Thus for $x/h_c \lesssim 6$, the observed over-prediction of the concentration fluctuation intensity is due to the under-prediction of the mean concentration. Likewise, the observed under-prediction of the concentration fluctuation intensity for $x/h_c \gtrsim 6$ is due to the over-prediction of the mean concentration.




Figure 5.9: Streamwise transect of the ground-level dimensionless mean concentration from SPMMM simulations of an elevated, cross-wind, line source ($z_s/h_c = 0.85$) in the Tombstone canopy flow. Horizontally-homogeneous velocity statistics were used to drive this simulation.




Figure 5.10: Streamwise transect of the concentration fluctuation intensity from SPMMM simulations of an elevated line, cross-wind line source ($z_s/h_c = 0.85$) in the Tombstone canopy flow. Horizontally-homogeneous velocity statistics were used to drive this simulation.

As for the simulations of the FR82 flow in Chapter 4, we see that SPMMM does not capture fully the large initial rise in the concentration fluctuation intensity close to the source. The modelled value is approximately twenty percent too low, equal to the discrepancy noted in the simulations of the FR82 flow shown in Figure 4.19.

5.2.1 Comparison of SPMMM to the Model of Cassiani et al.

For comparison with the Cassiani et al. (2007; CASS) results, a simulation was performed using the TKE dissipation profile inferred from the vertical profile of T_L shown in Figure 5.6 with a Kolmogorov constant of $C_0 = 2.0$. As briefly mentioned in Chapter 2, Franzese and Cassiani (2007) derived the relation $C_r \approx C_0/11$ for homogeneous isotropic turbulence. Cassiani et al. (2007) mention that for a smaller C_0 a proportionally smaller C_r must be used. Although not explicitly stated in their paper, we assumed that they used $C_r = 2/11 \approx 0.18$, as did we for the comparison simulation.

The vertical profiles of the dimensionless mean concentration at three locations for this simulation are shown in Figure 5.11. The results from the two models are very similar, but the magnitudes of the SPMMM results are slightly smaller than the corresponding CASS results. Figure 5.12 compares the dimensionless concentration standard deviation. Again the two models produced similar results. The differences are likely due to the implementation of the two models as discussed in Chapter 4 (SPMMM used single-particle trajectories with a pre-calculated conditional mean concentration field, whereas model of Cassiani et al. (2007) used simultaneous trajectories and calculated the conditional




Figure 5.11: Comparison of the vertical profiles of the dimensionless mean concentration from an elevated, cross-wind line source ($z_s/h_c = 0.85$) in the wind-tunnel tombstone canopy flow as simulated by SPMMM and Cassiani et al. (2007). Horizontally-homogeneous velocity statistics were used to drive these simulations.




Figure 5.12: Comparison of the vertical profiles of the dimensionless standard deviation of concentration from an elevated, cross-wind line source ($z_s/h_c = 0.85$) in the Tombstone canopy flow as simulated by SPMMM and Cassiani et al. (2007). Horizontally-homogeneous velocity statistics were used to drive these simulations.

mean concentration field at each timestep).

Lastly, we note that the **SPMMM** results shown in Figure 5.11 resemble the Flesch and Wilson (1992; Figure 6) simulation results more than the **SPMMM** results shown in Figure 5.7. This is probably due to using a TKE dissipation profile inferred from the vertical profile of T_L , which is more in line with the Lagrangian integral timescale formulation of the Flesch and Wilson (1992) model.

5.3 Dispersion from a Continuous Point Source

The spatial, temporal, and velocity resolution for the simulations of the continuous point source in the Tombstone canopy flow were the same as for the continuous line source simulations above. The free parameters, with the exception of σ_0 , were assigned the same values as above: $C_0 = 6.0$, $C_r = 0.45$, and $\mu = 0.75$. Since the experimental source was designed to have a uniform concentration over its outlet area, in **SPMMM** it was modelled as a tophat source. Due to the increased dimensionality of the problem, the number of particles used was increased to 2×10^7 for both the **MEANS** pre-calculation simulations and the **SPMMM** micromixing simulations. Simulations were performed using the spatially averaged, horizontally-homogeneous velocity statistics (denoted as **SPMMM-HH** in the figures below) and using the locally inhomogeneous A – S point data (denoted as **SPMMM-LI** in the figures below). The experimental measurements of Hilderman and Chong (2007) are denoted as **COANDA**. The next several pages display the results of the simulations, arranged in increasing order of the moment of the concentration field: mean, standard deviation, skewness, and kurtosis. Note that in the following figures the scales on the

vertical axes change between panels.

Before examining the simulation results, a brief discussion on the consistency between the **MEANS** and **SPMMM** simulated mean concentration is required. The discretised unit cell shown in Figure 5.1 measures $88 \text{ mm} \times 66 \text{ mm}$, and is divided into $N_x^{\text{vel}} = 22$ by $N_y^{\text{vel}} = 15$ velocity statistic bins in the streamwise and spanwise directions, respectively. Within the simulation domain there are 8 unit cells in the streamwise direction and 20 in the spanwise direction, for a total of 176×300 velocity statistic bins in the xy -plane in the simulation domain. In order to accurately estimate the velocity PDF in each velocity bin, and thus provide an accurate estimate of the number of normalising particles (recall equation (3.68): $N_\phi^v = N_\phi f_{\mathbf{u}} \Delta u \Delta v \Delta w$) for each residence time bin, **MEANS** should be run with a horizontal resolution of at *least* $N_x = 176$ and $N_y = 300$ (refer back to the end of Section 3.3.1 and Section 4.3 for more information). In this way each conditional concentration bin will contain only a single value of the velocity PDF when the plume is fully expanded, and equation (3.68) will produce accurate values. However, with $N_x = 176$, $N_y = 300$, $N_z = N_z^{\text{vel}} = 40$, and 20^3 velocity bins, the residence time array in **MEANS** alone would require 62.9 GB of memory[§], well beyond the amount that is available. To get around this limitation, the velocity PDF in equation (3.68) was replaced by a horizontally-homogeneous velocity PDF that resulted from taking a spatial average of the varying velocity PDF across the unit cell. This method resulted in acceptable mean concentration consistency between **MEANS** and **SPMMM** and allowed the models to be run with the available computational resources.

[§]At single precision the number of bytes is $4 \times N_x \times N_y \times N_z \times N_u \times N_v \times N_w$.

Figure 5.13 shows the vertical profiles of the cross-wind integrated mean concentration from MEANS and SPMMM simulations utilising locally inhomogeneous velocity statistics, 16 rows downstream from the source. While acceptable, the consistency is not as good as for the horizontally-homogeneous simulation (not shown, but are as good as the horizontally-homogeneous FR82 results shown in Figure 4.7(b)). The major differences between the SPMMM-HH and the SPMMM-LI simulations are: the use of the spatially-averaged velocity PDF in equation (3.68) for the SPMMM-LI simulations; and the occurrence of rogue velocities in the SPMMM-LI simulations, there were none in the SPMMM-HH simulations. These two factors are believed to have caused the discrepancy seen in Figure 5.13. The effects of rogue velocities on the conditional concentration field and first-order consistency are more thoroughly examined in Chapter 6.

Figure 5.14 shows the streamwise transects of the mean dimensionless concentration at six heights on the plume centreline. Close to the source ($x/h_c \lesssim 2.0$), the SPMMM-LI results fit the experimental data better than the SPMMM-HH results, at all heights. This is likely due to the relatively large positive mean vertical velocity found in the regions where the flow went up and over an obstacle (i.e., near point A in Figure 5.1). Farther from the source ($x/h_c \gtrsim 2.0$), there is little difference between the SPMMM-LI results and the SPMMM-HH results within the canopy ($z/h_c < 1.0$). At the top of the canopy ($z/h_c = 1.0$), the SPMMM-LI results are slightly poorer than the SPMMM-HH results. Above the canopy ($z/h_c > 1.0$), the SPMMM-LI results are consistently higher than the SPMMM-HH results and therefore closer to the experimental measurements, but the agreement between experiment and simulation is still marginal at best. For both the SPMMM-HH and SPMMM-LI simulations it appears that the parti-




Figure 5.13: The vertical profiles of the cross-wind integrated mean concentration from MEANS and SPMMM simulations utilising locally inhomogeneous velocity statistics 16 rows downstream from the source. The discrepancies are due to the use of a spatially-averaged velocity PDF in the normalising particle count calculation (see equation (3.68)), and the presence of rogue velocities in the SPMMM-LI simulations.

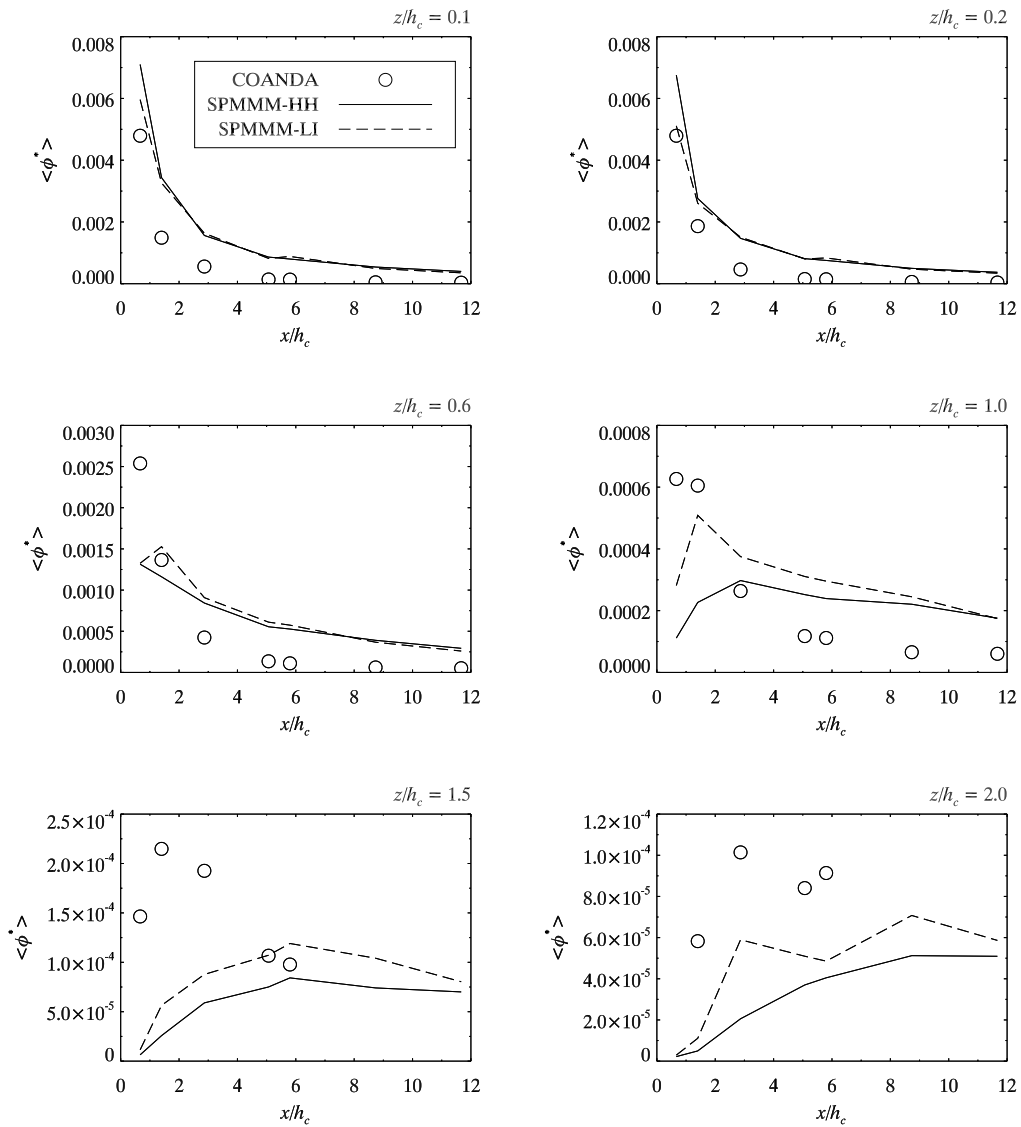


Figure 5.14: Streamwise transects of the mean dimensionless concentration at selected heights on the plume centreline from SPMMM simulations of an elevated point source ($z_s/h_c = 0.2$) in the Tombstone canopy flow.

cles are not escaping the canopy as rapidly as they should have, resulting in over-predicted concentrations within the canopy, and under-predicted concentrations above the canopy.

Table 5.2 presents the performance measures for the Figure 5.14 data. From these performance measures we can see that the use of locally inhomogeneous

z/h_c	flow	FB	NMSE	FAC2	NAE
0.1	HH	-0.686	0.760	0.143	0.686
	LI	-0.607	0.511	0.143	0.607
0.2	HH	-0.565	0.471	0.286	0.565
	LI	-0.436	0.234	0.286	0.436
0.6	HH	-0.082	0.656	0.429	0.666
	LI	-0.230	0.556	0.286	0.637
1.0	HH	0.194	1.196	0.143	0.866
	LI	-0.321	0.403	0.429	0.578
1.5	HH	1.007	1.959	0.400	1.007
	LI	0.311	0.364	0.600	0.522
2.0	HH	1.352	3.525	0.250	1.352
	LI	0.640	0.591	0.500	0.640

Table 5.2: Performance measures for the mean dimensionless concentration on the plume centreline from SPMMM simulations of an elevated point source ($z_s/h_c = 0.2$) in the Tombstone canopy flow. The performance measures were calculated from the data displayed in Figure 5.14.

geneous velocity statistics leads to a generally better fit to the experimental data overall. We note however that many of these performance measures lie outside of their acceptable ranges (e.g., FAC2).

Figures 5.15 – 5.17 show the spanwise transects of the mean dimensionless concentration 1, 2, 4, 8, 12, and 16 rows downstream of the source ($x/h_c = 0.67, 1.40, 2.87, 5.80, 8.73, 11.7$) at three heights: within the canopy at source height ($z/h_c = 0.2$); at the canopy top ($z/h_c = 1.0$); and above

the canopy ($z/h_c = 1.5$ for rows 1,2,4,8; $z/h_c = 3.0$ for rows 12 and 16, as this is the lowest above canopy measurement location for these rows). Note once again that the scales on the vertical axes may vary between panels. The water-channel experimental data sets are very large, and so to reduce clutter on the graphs every twentieth Coanda experimental data point is plotted in the spanwise transects. These figures support the statements made above: close to the source (Figure 5.15), at all heights the SPMMM-LI simulations match the experimental data better than the SPMMM-HH data; farther from the source (Figures 5.16 & 5.17), both simulations produced comparable results within the canopy; at the canopy top SPMMM-HH outperformed SPMMM-LI, and above the canopy SPMMM-LI outperformed SPMMM-HH. From these figures it appears that using locally inhomogeneous velocity statistics results in more lateral dispersion of the particles and thus a better fit to the experimental measurements.

With respect to the standard deviation of the dimensionless concentration on the plume centreline (Figure 5.18), the SPMMM-LI results are more accurate than the SPMMM-HH results close to the source ($x/h_c \lesssim 2.0$), and less accurate farther from the source ($x/h_c \gtrsim 2.0$). Table 5.3 presents the performance measures for this figure. With respect to the FB, the SPMMM-LI simulations performed better than the SPMMM-HH simulations, and were within the acceptable range. This is due to the better near field predictions of the SPMMM-LI simulations. This better near field prediction, and reasonable far field predictions, also resulted in NMSE values within acceptable limits for the SPMMM-LI simulations. The NMSE values of the SPMMM-HH simulations are consistently greater than those of the SPMMM-LI simulations, and above the

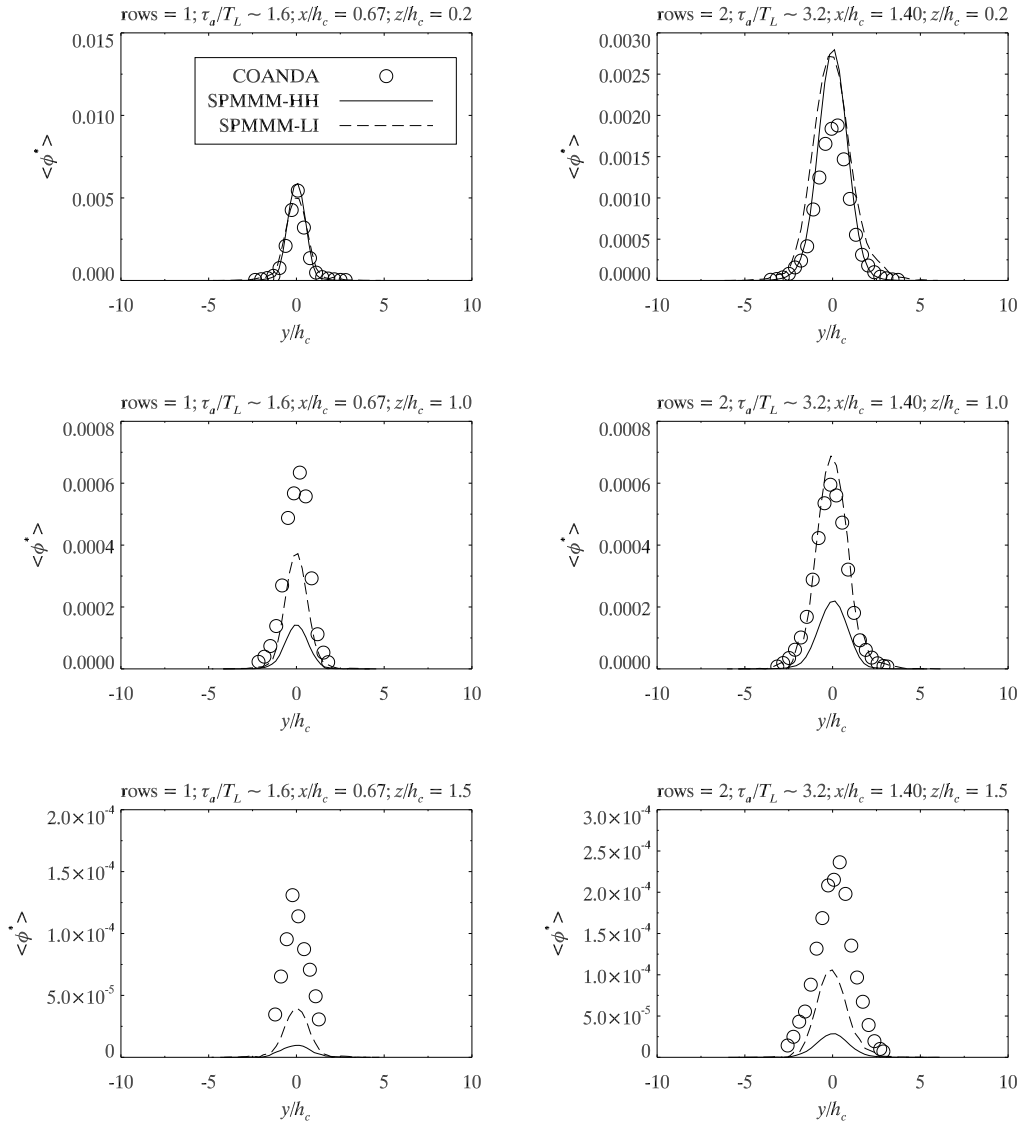


Figure 5.15: Spanwise transects of the dimensionless mean concentration at selected heights, one row ($x/h_c = 0.67$; left column) and two rows ($x/h_c = 1.40$; right column) downstream from the source from SPMMM simulations of an elevated point source ($z_s/h_c = 0.2$) in the Tombstone canopy flow.

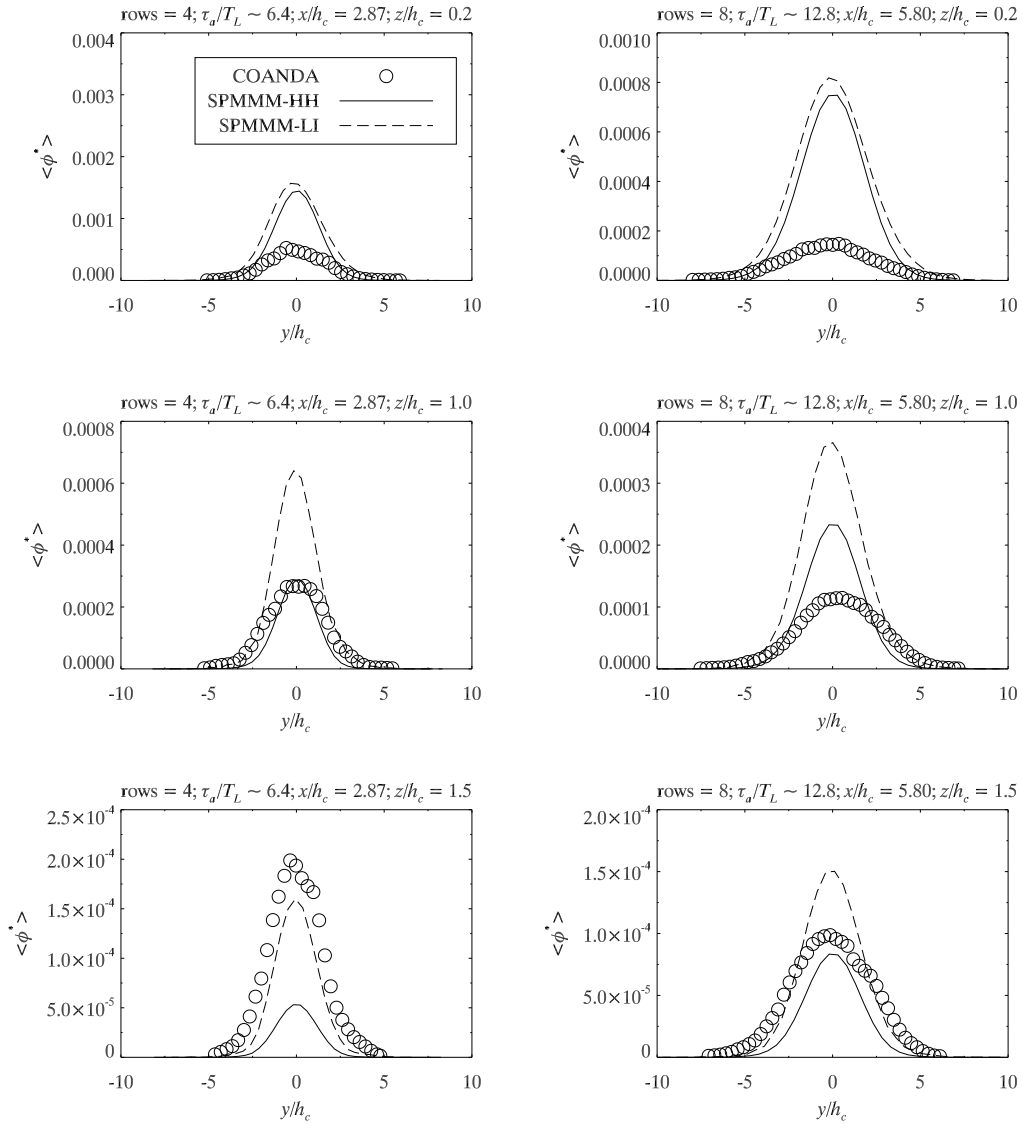


Figure 5.16: Spanwise transects of the dimensionless mean concentration at selected heights, four rows ($x/h_c = 2.87$; left column) and eight rows ($x/h_c = 5.80$; right column) downstream from the source from SPMMM simulations of an elevated point source ($z_s/h_c = 0.2$) in the Tombstone canopy flow.

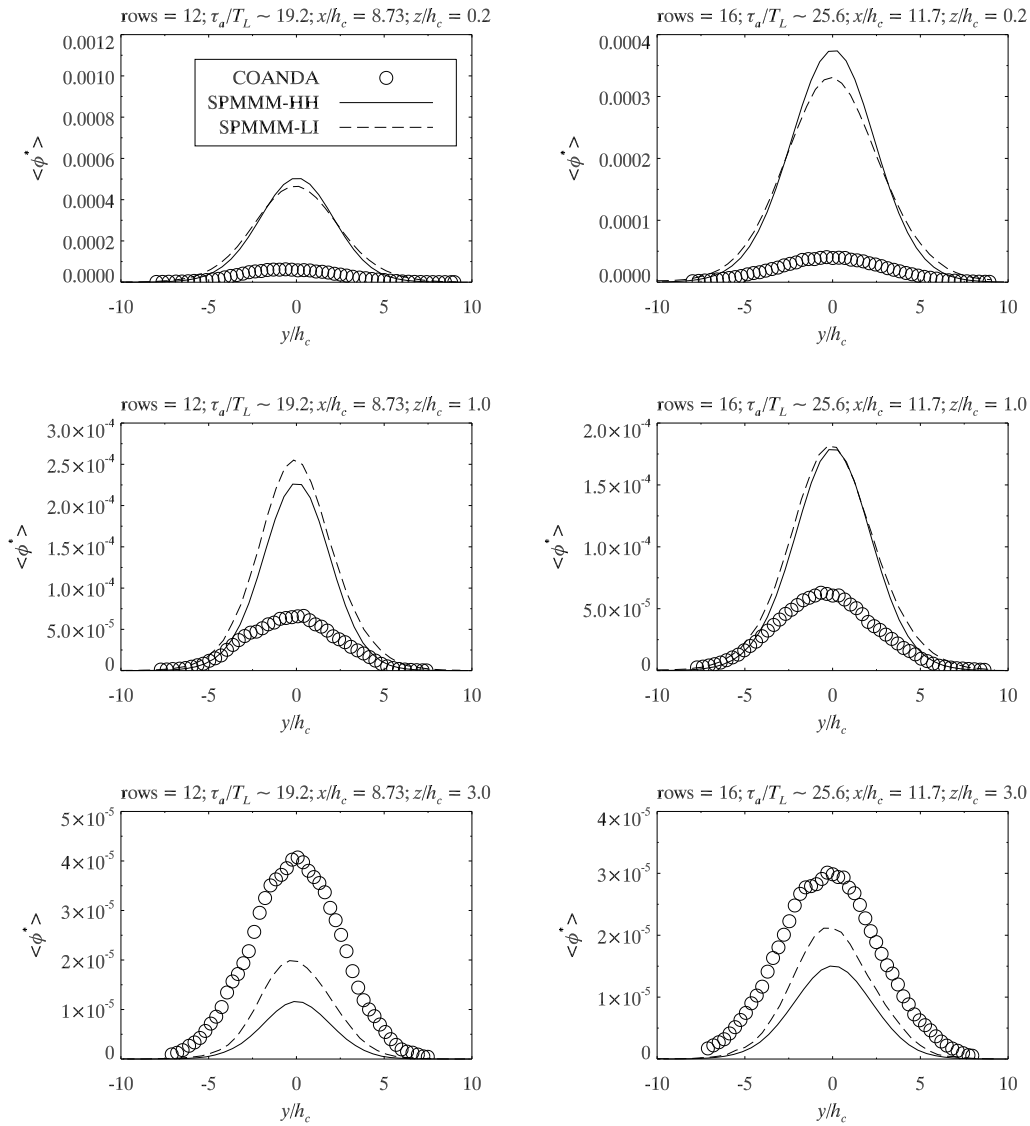


Figure 5.17: Spanwise transects of the dimensionless mean concentration at selected heights, twelve rows ($x/h_c = 8.73$; left column) and sixteen rows ($x/h_c = 11.7$; right column) downstream from the source from SPMMM simulations of an elevated point source ($z_s/h_c = 0.2$) in the Tombstone canopy flow.

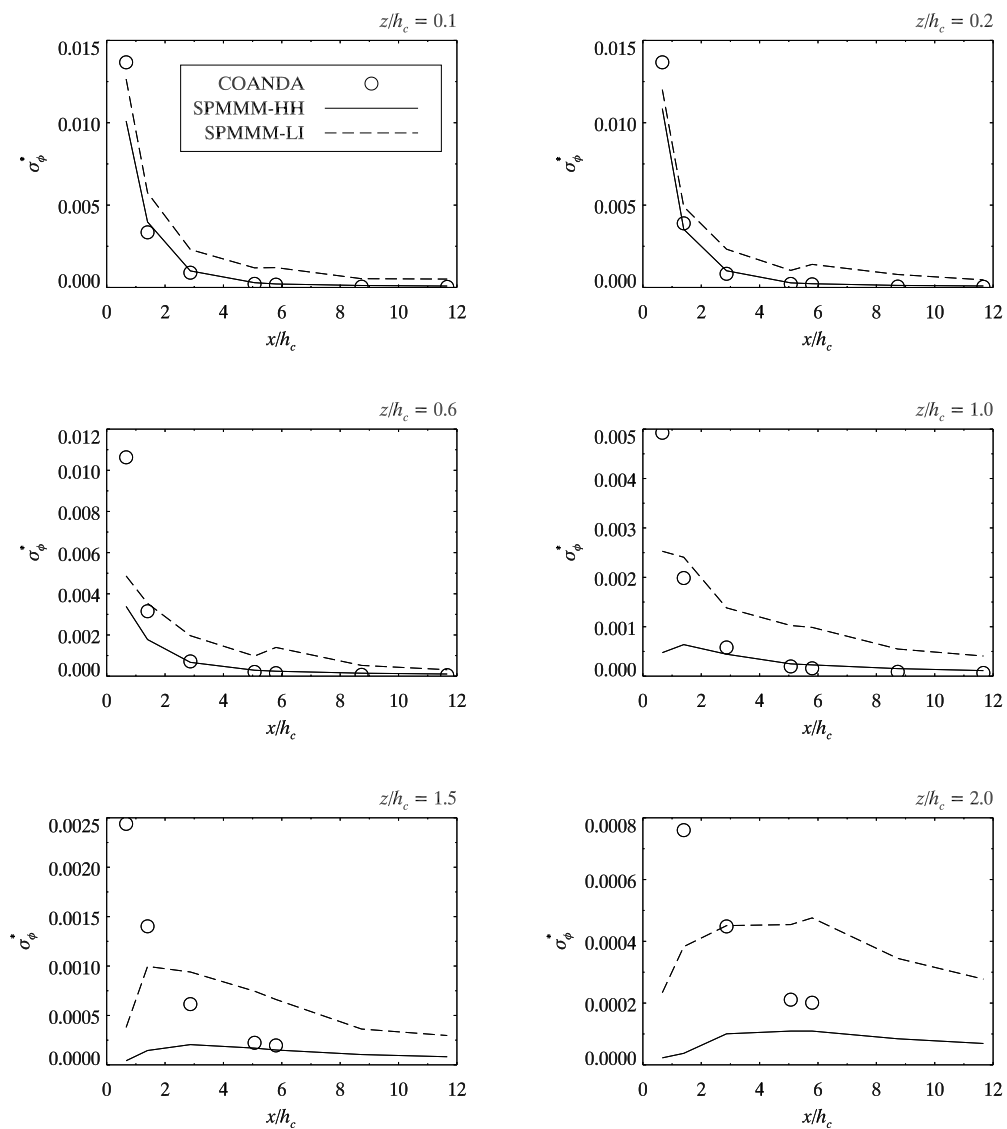


Figure 5.18: Streamwise transects of the standard deviation of the dimensionless concentration at selected heights on the plume centreline from SPMMM simulations of an elevated point source ($z_s/h_c = 0.2$) in the Tombstone canopy flow.

z/h_c	flow	FB	NMSE	FAC2	NAE
0.1	HH	0.152	0.317	0.714	0.266
	LI	-0.271	0.178	0.286	0.367
0.2	HH	0.158	0.186	0.714	0.204
	LI	-0.193	0.143	0.286	0.352
0.6	HH	0.774	3.851	0.571	0.832
	LI	0.098	1.288	0.143	0.710
1.0	HH	1.107	8.242	0.714	1.198
	LI	-0.148	0.779	0.286	0.703
1.5	HH	1.491	10.81	0.400	1.491
	LI	0.268	1.378	0.400	0.879
2.0	HH	1.429	6.703	0.250	1.429
	LI	0.063	0.653	0.250	0.690

Table 5.3: Performance measures for the standard deviation of the dimensionless concentration on the plume centreline from SPMM simulations of an elevated point source ($z/h_c = 0.2$) in the Tombstone canopy flow. The performance measures were calculated from the data displayed in Figure 5.18.

canopy they are outside of the acceptable range, again due to large under-predictions of the near-source standard deviation. The same arguments apply to the NAE, but the discrepancy is not as great since the NAE does not square the error like the NMSE. The **SPMMM-HH** results do show better FAC2 values than the corresponding **SPMMM-LI** results. This is due to more accurate far field predictions, particular within the canopy.

Figures 5.19 – 5.21 show the spanwise transects of the standard deviation of the dimensionless concentration at the same locations as for the mean dimensionless concentration above. As for the streamwise transects shown above, we see more accurate predictions from **SPMMM-LI** close to the source (Figure 5.19), and more accurate predictions from **SPMMM-HH** farther from the source (Figures 5.20 & 5.21), in terms of both the magnitude of the standard deviation and the spanwise spread.

The streamwise transects of the skewness of the dimensionless concentration at six heights on the plume centreline can be seen in Figure 5.22. The **SPMMM-HH** and **SPMMM-LI** results are comparable close to the source but quickly diverge. At and below the canopy height, the **SPMMM-HH** results are in fair agreement with the experimental results until $x/h_c \sim 3.0$, after which the model under-predicts the skewness and at $x/h_c \sim 8.0$ begins to predict the skewness to be approximately zero. Above the canopy the fit is fair. In contrast, the **SPMMM-LI** results over-predicts the skewness at all downstream locations $x/h_c \gtrsim 1.0$.

The more accurate performance with regards to the skewness of the **SPMMM** run with horizontally-homogeneous velocity statistics compared to the **SPMMM** run with locally inhomogeneous statistics is evident in Table 5.4, which shows

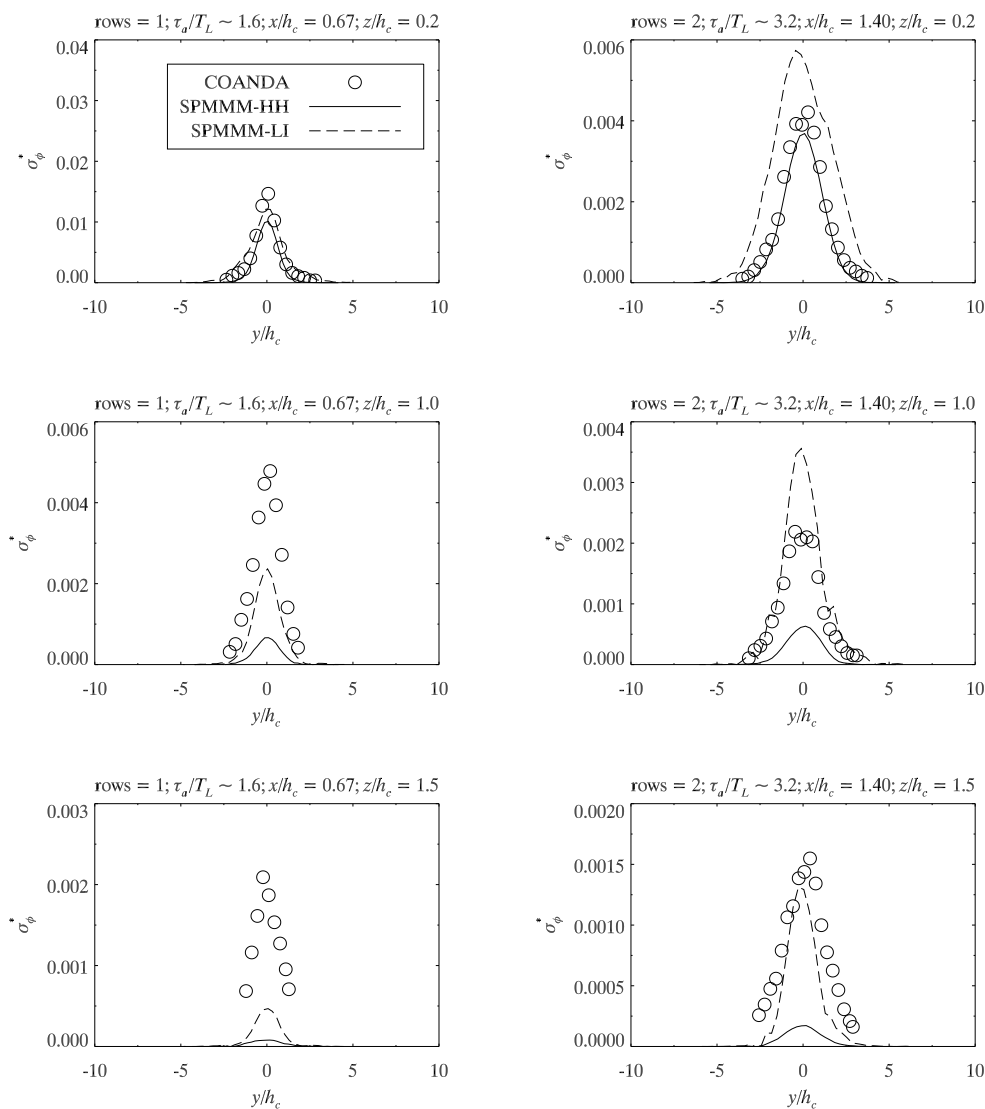


Figure 5.19: Spanwise transects of the standard deviation of the dimensionless concentration at three heights, one row ($x/h_c = 0.67$; left column) and two rows ($x/h_c = 1.40$; right column) downstream from the source from SPMMM simulations of an elevated point source ($z_s/h_c = 0.2$) in the Tombstone canopy flow.

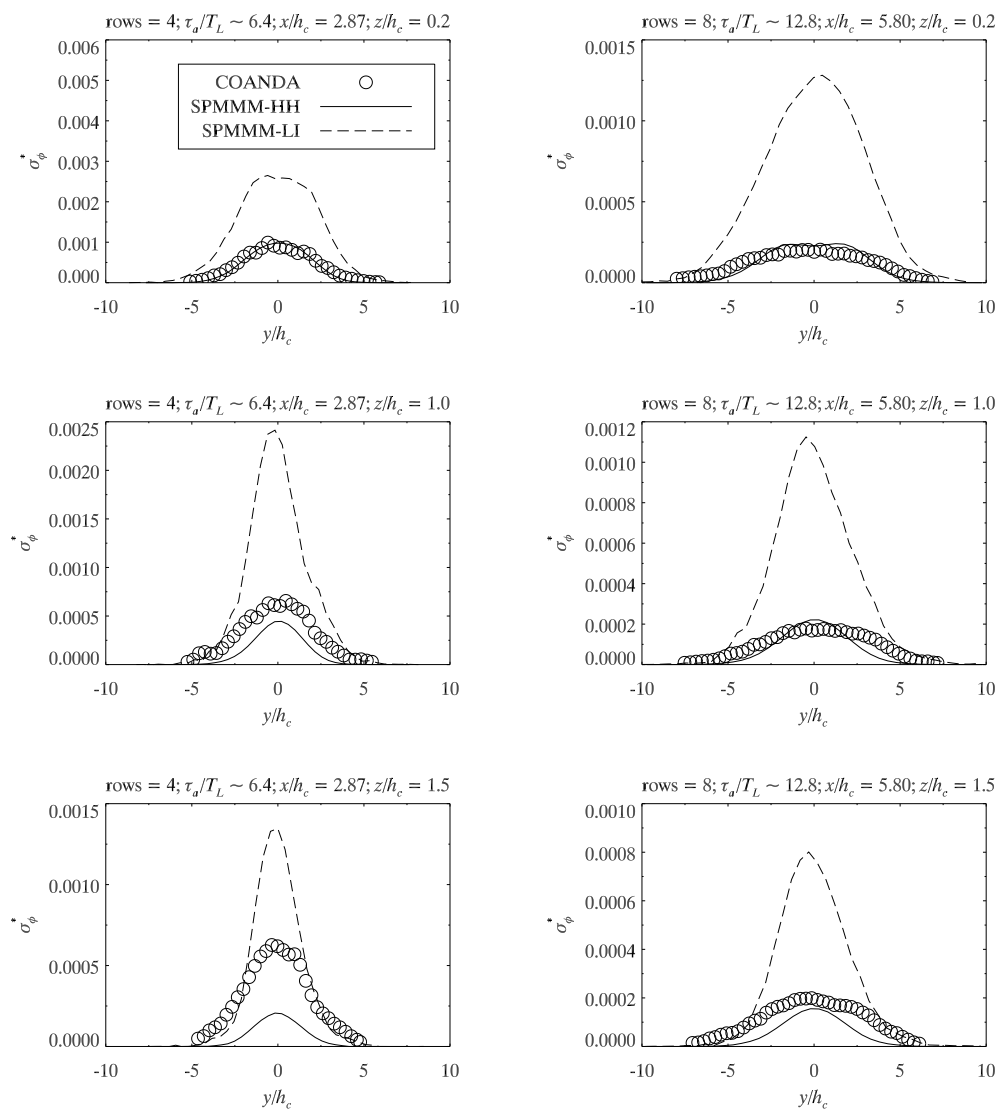


Figure 5.20: Spanwise transects of the standard deviation of the dimensionless concentration at three heights, four rows ($x/h_c = 2.87$; left column) and eight rows ($x/h_c = 5.80$; right column) downstream from the source from SPMMM simulations of an elevated point source ($z_s/h_c = 0.2$) in the Tombstone canopy flow.

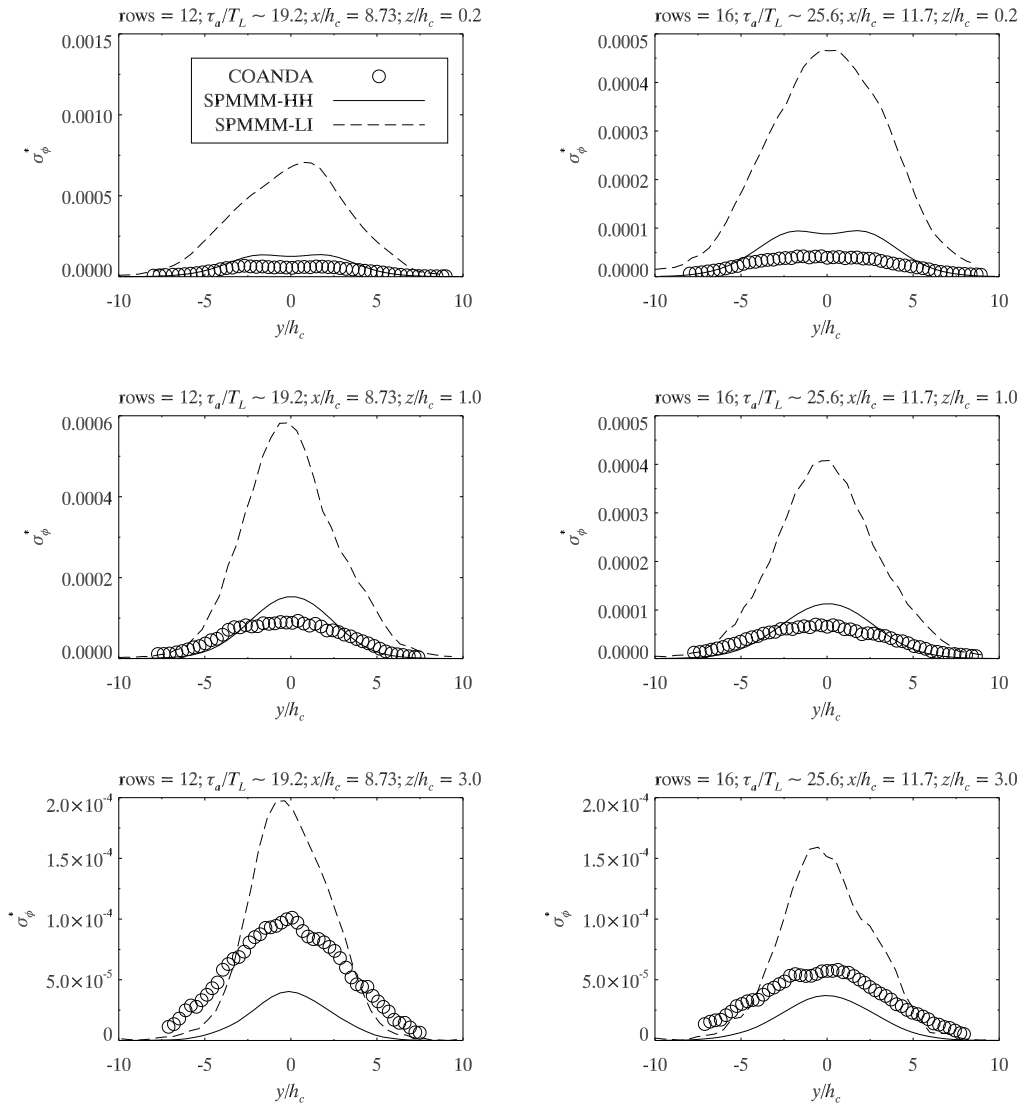


Figure 5.21: Spanwise transects of the standard deviation of the dimensionless concentration at three heights, twelve rows ($x/h_c = 8.73$; left column) and sixteen rows ($x/h_c = 11.7$; right column) downstream from the source from SPMMM simulations of an elevated point source ($z_s/h_c = 0.2$) in the Tombstone canopy flow.

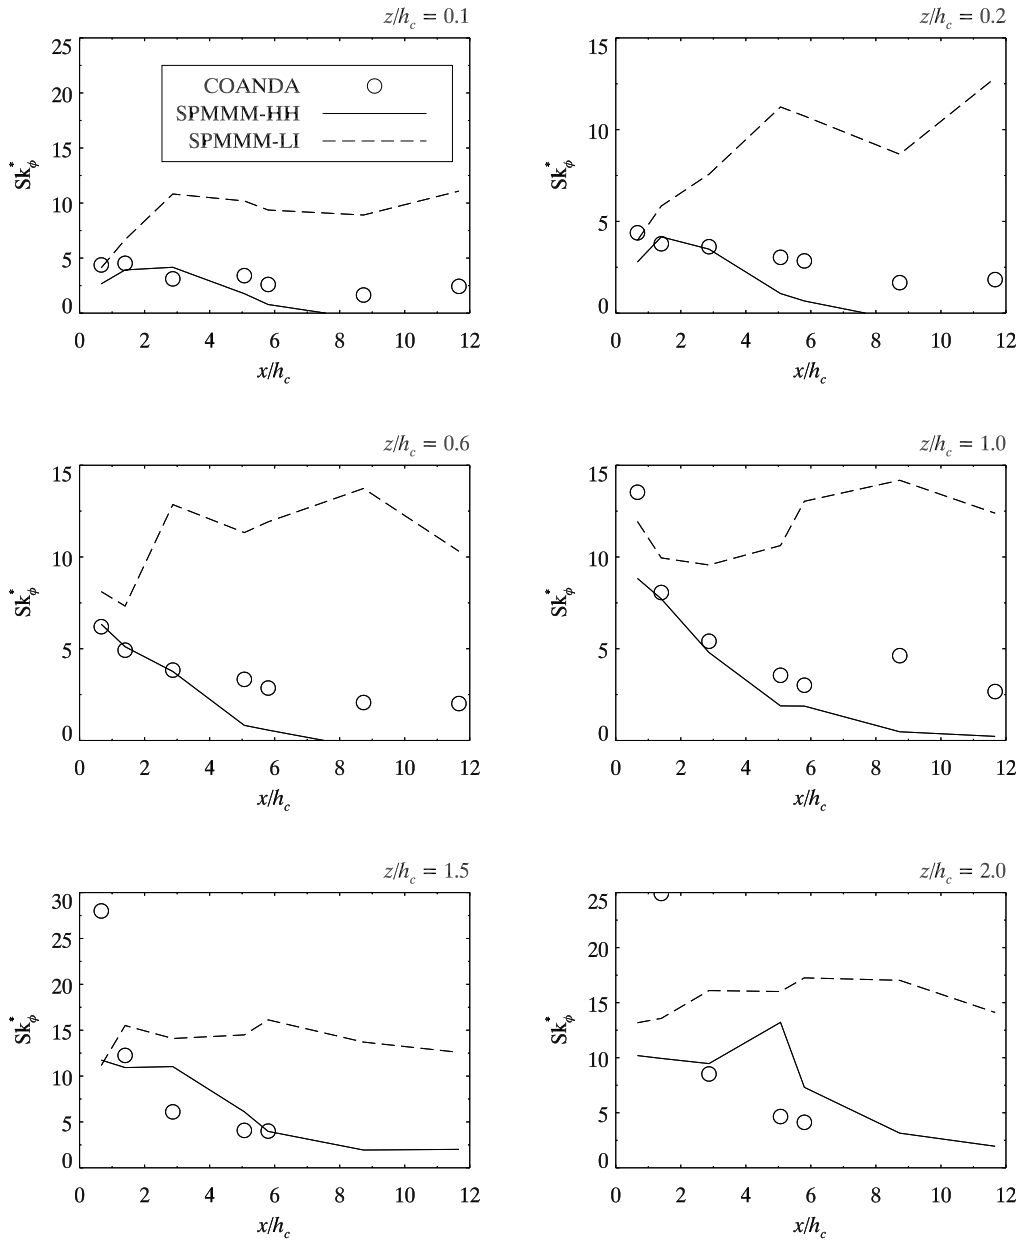


Figure 5.22: Streamwise transects of the skewness of the dimensionless concentration at selected heights on the plume centreline from SPMMM simulations of an elevated point source ($z_s/h_c = 0.2$) in the Tombstone canopy flow.

the performance measures corresponding to Figure 5.22. The SPMMM–HH results generally outperform the SPMMM–LI results. We do note that for $z/h_c \leq 1.0$ the FB is outside of its acceptable range for both simulations, as is the FAC2 at various heights.

z/h_c	flow	FB	NMSE	FAC2	NAE
0.1	HH	0.579	0.634	0.571	0.702
	LI	-0.938	1.468	0.286	0.951
0.2	HH	0.625	0.655	0.429	0.672
	LI	-0.968	1.737	0.286	0.988
0.6	HH	0.461	0.414	0.429	0.492
	LI	-0.999	1.624	0.286	0.999
1.0	HH	0.452	0.331	0.714	0.452
	LI	-0.666	0.755	0.429	0.718
1.5	HH	0.217	0.618	0.800	0.501
	LI	-0.270	0.790	0.200	0.805
2.0	HH	-0.014	0.718	0.250	0.706
	LI	-0.329	0.700	0.500	0.794

Table 5.4: Performance measures for the skewness of the dimensionless concentration on the plume centreline from SPMMM simulations of an elevated point source ($z/h_c = 0.2$) in the Tombstone canopy flow. The performance measures were calculated from the data displayed in Figure 5.22.

Figures 5.23 – 5.25 show the spanwise transects of the skewness of the dimensionless concentration at the same locations as for the mean dimensionless concentration. From Figure 5.23 we see there is comparable performance between the two simulations close to the source. Within the canopy, the agreement with the experimental data is good, except on the plume edges, where the skewness is under-predicted. These *edge effects* are likely not a result of flaw in the model, but an issue of statistical convergence—not enough particles were released to enable the higher-order concentration moments to be com-

puted with good accuracy. The same behavior is observed in measured profiles of skewness and kurtosis due to insufficient sampling. To reliably estimate the skewness or kurtosis on the edge of the plume a *very* large number of particles would have to be released. Above the canopy, the predictions vary from being fair throughout the central region of the plume, and poor on the edges, again due to edge effects. In Figures 5.24 & 5.25 the SPMMM-HH results are seen to produce more accurate predictions than the SPMMM-LI results. Edge effects are once again seen in the results from both simulations.

Statements similar to those made above regarding the skewness can be applied to the kurtosis of the dimensionless concentration on the plume centerline. Streamwise transects at six heights are shown in Figure 5.26, with corresponding performance measures shown in Table 5.5. Again, after predicting similar values of kurtosis close to the source the SPMMM-HH results and the SPMMM-LI results rapidly diverge with the SPMMM-LI results over-predicting the kurtosis. The performance measure of SPMMM-HH are generally within, or slightly outside of, their acceptable range, and better than the SPMMM-LI performance measures.

The spanwise transects of the kurtosis of the dimensionless concentration, as seen in Figures 5.27 – 5.29 (note the logarithmic y -axis), show similar features to the spanwise profiles of skewness shown above: the close to source model predictions (Figure 5.27) are good through the central region of the plume but poor at the edges; farther from the source (Figures 5.28 & 5.29), the SPMMM-HH results show a much better conformance to the experimental data than the SPMMM-LI results through the central region of the plume. Edge effects are again seen in some of the simulation results. The accuracy of the

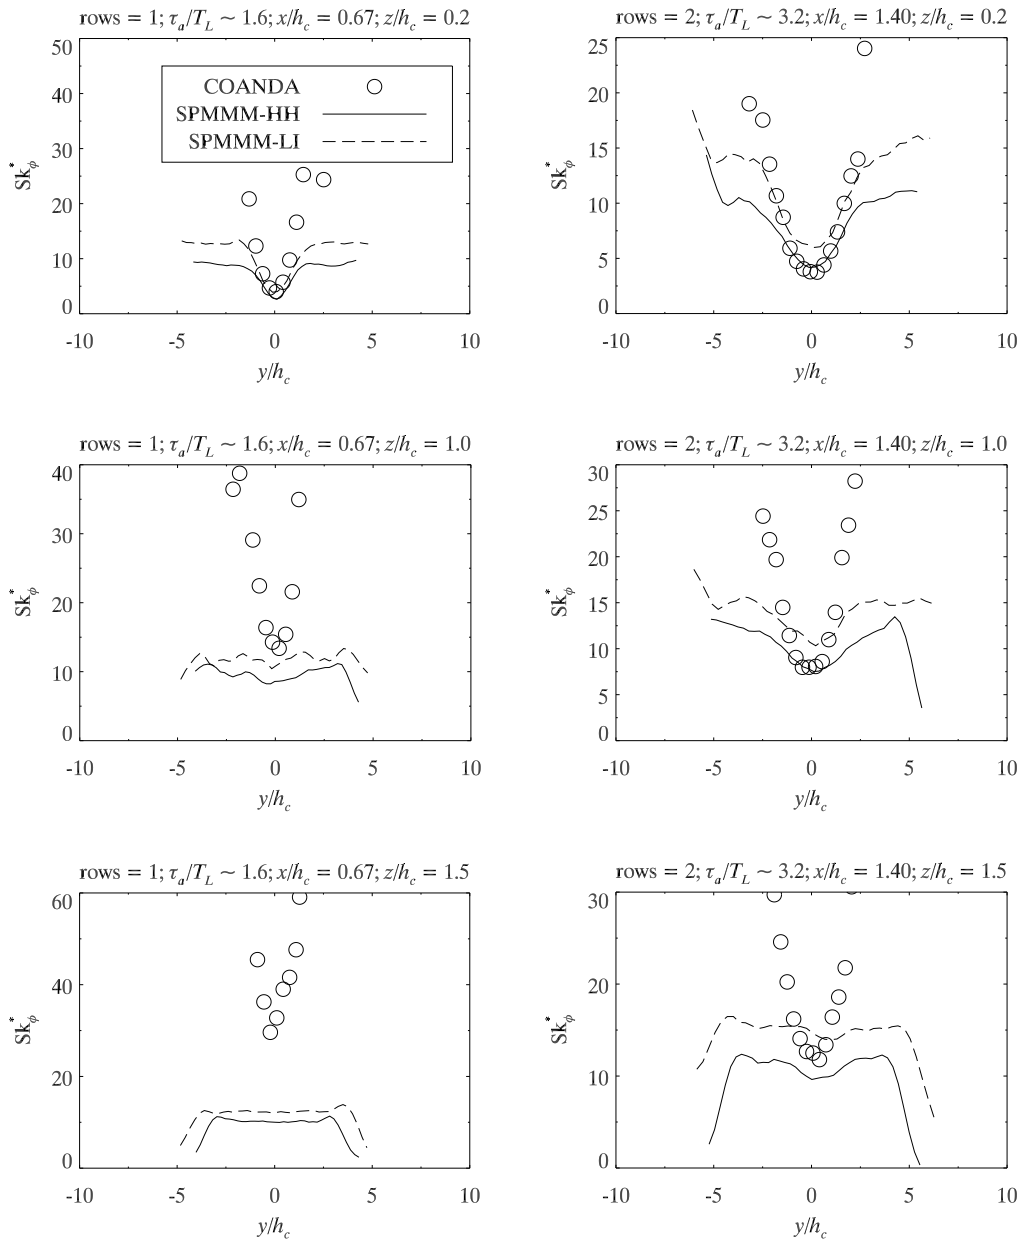


Figure 5.23: Spanwise transects of the skewness of the dimensionless concentration at three heights, one row ($x/h_c = 0.67$; left column) and two rows ($x/h_c = 1.40$; right column) downstream from the source from SPMMM simulations of an elevated point source ($z_s/h_c = 0.2$) in the Tombstone canopy flow.

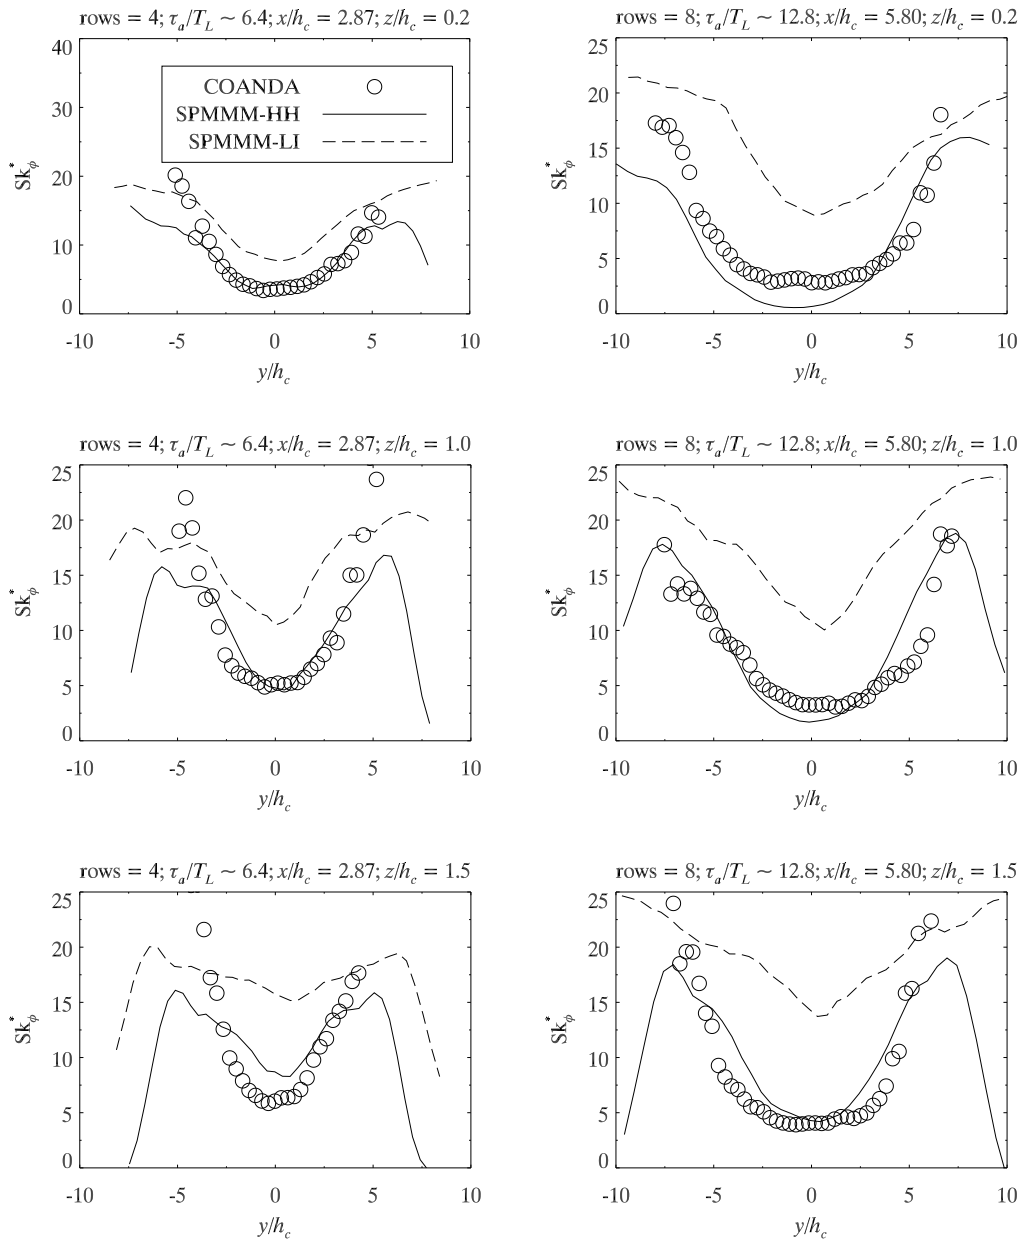


Figure 5.24: Spanwise transects of the skewness of the dimensionless concentration at three heights, four rows ($x/h_c = 2.87$; left column) and eight rows ($x/h_c = 5.80$; right column) downstream from the source from SPMMM simulations of an elevated point source ($z_s/h_c = 0.2$) in the Tombstone canopy flow.

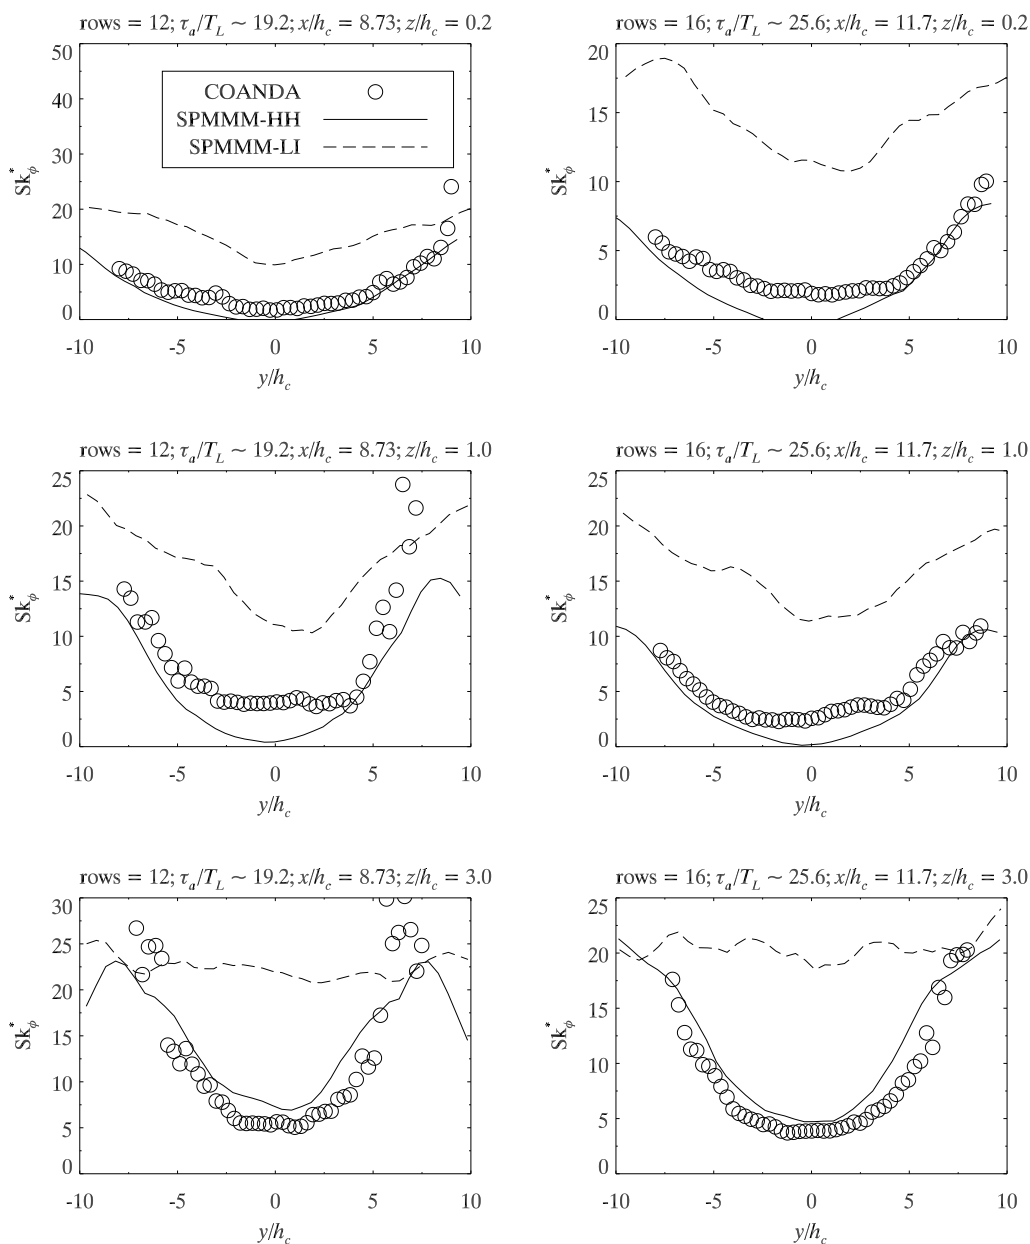


Figure 5.25: Spanwise transects of the skewness of the dimensionless concentration at three heights, twelve rows ($x/h_c = 8.73$; left column) and sixteen rows ($x/h_c = 11.7$; right column) downstream from the source from SPMMM simulations of an elevated point source ($z_s/h_c = 0.2$) in the Tombstone canopy flow.

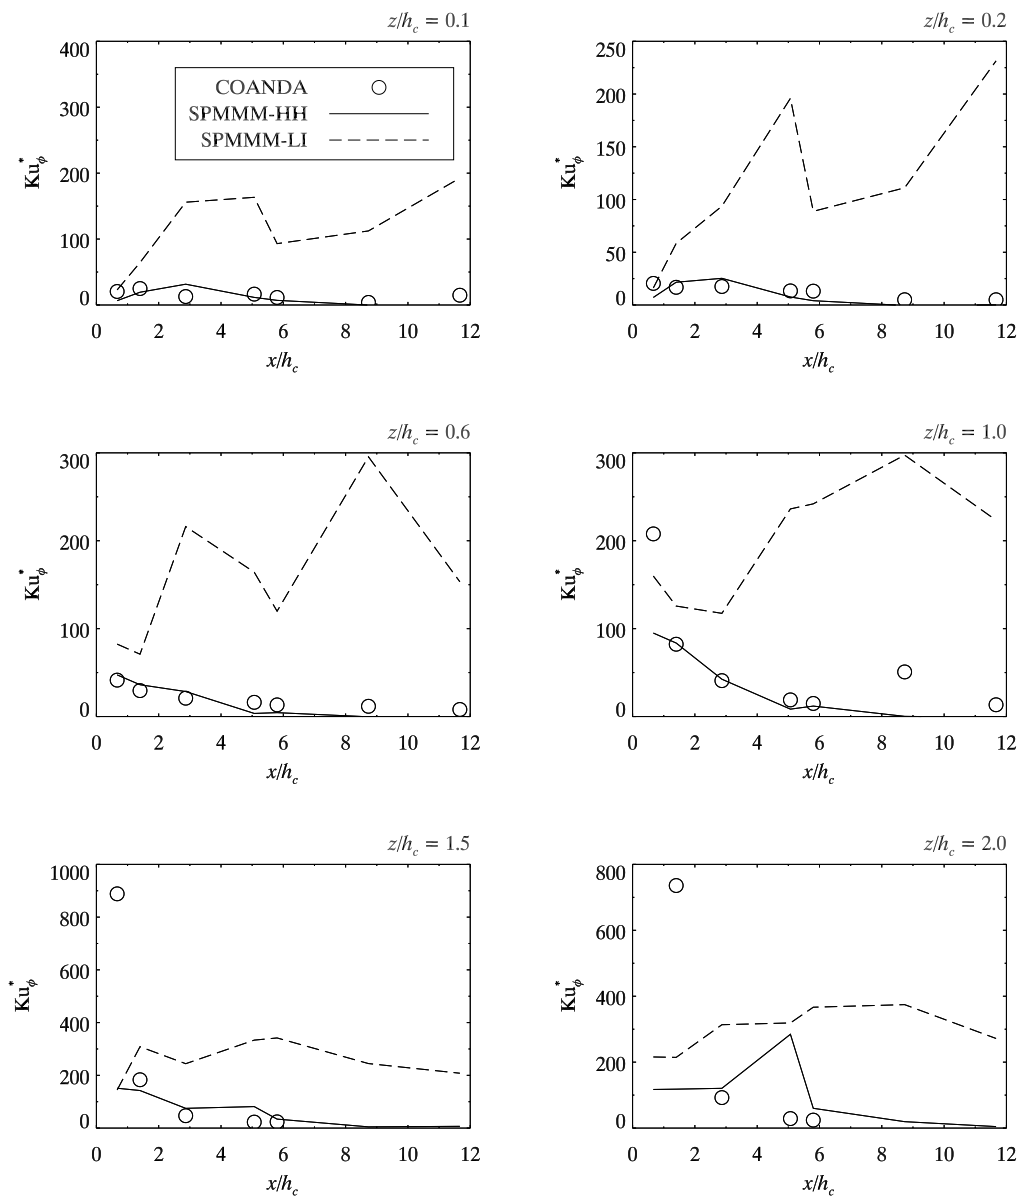


Figure 5.26: Streamwise transects of the kurtosis of the dimensionless concentration at selected heights on the plume centreline from SPMMM simulations of an elevated point source ($z_s/h_c = 0.2$) in the Tombstone canopy flow.

z/h_c	flow	FB	NMSE	FAC2	NAE
0.1	HH	0.329	0.776	0.429	0.745
	LI	-1.539	7.767	0.143	1.539
0.2	HH	0.330	0.504	0.429	0.651
	LI	-1.589	10.53	0.143	1.608
0.6	HH	0.169	0.247	0.429	0.479
	LI	-1.547	7.949	0.143	1.547
1.0	HH	0.562	1.058	0.429	0.580
	LI	-1.063	2.488	0.286	1.168
1.5	HH	0.826	4.875	0.600	1.062
	LI	-0.164	2.521	0.200	1.337
2.0	HH	0.316	3.259	0.250	1.310
	LI	-0.187	1.937	0.250	1.257

Table 5.5: Performance measures for the kurtosis of the dimensionless concentration on the plume centreline from **SPMM** simulations of an elevated point source ($z/h_c = 0.2$) in the Tombstone canopy flow. The performance measures were calculated from the data displayed in Figure 5.26.

SPMMM-HH predictions for the kurtosis is quite good. The magnitude of the statistical noise seen around $-3 \lesssim y/h_c \lesssim 3$ in the SPMMM-LI simulation results in the top four panels of Figure 5.29 is amplified due to the logarithmic y -axis.

For both the simulation utilising horizontally-homogeneous velocity statistics and the simulation utilising locally inhomogeneous velocity statistics it appears that not enough material is escaping the canopy, leading to over-predicted in-canopy mean concentrations and under-predicted above canopy concentrations. The large, positive mean vertical velocity in front of the obstacles, resulting from the flow being deflected up and over the obstacles, helps the SPMMM-LI simulations produce more accurate predictions of the mean and standard deviation of the dimensionless concentration for $x/h_c \lesssim 2.0$, close to the source, but for the most part led to decreased accuracy farther downstream. For the skewness and the kurtosis of the dimensionless concentration, the SPMMM-HH results are more accurate than the SPMMM-LI simulations.

Overall, for point source dispersion in a Tombstone canopy flow, spatial averaging of the velocity statistics resulted in improved performance and accuracy of the SPMMM model predictions. This finding is also noted in earlier work by Wilson and Yee (2000), and may be due in part to the extremely turbulent nature of the Tombstone canopy flow, the tab shaped obstacles are efficient turbulence generators, and the flow field may not be well represented as locally inhomogeneous. Furthermore, the locally inhomogeneous velocity statistics were obtained by interpolating measurements made at a small number of points—there is no guarantee that resulting interpolations are physically valid (and moreover, the gradients of Reynolds stresses are sensitive to

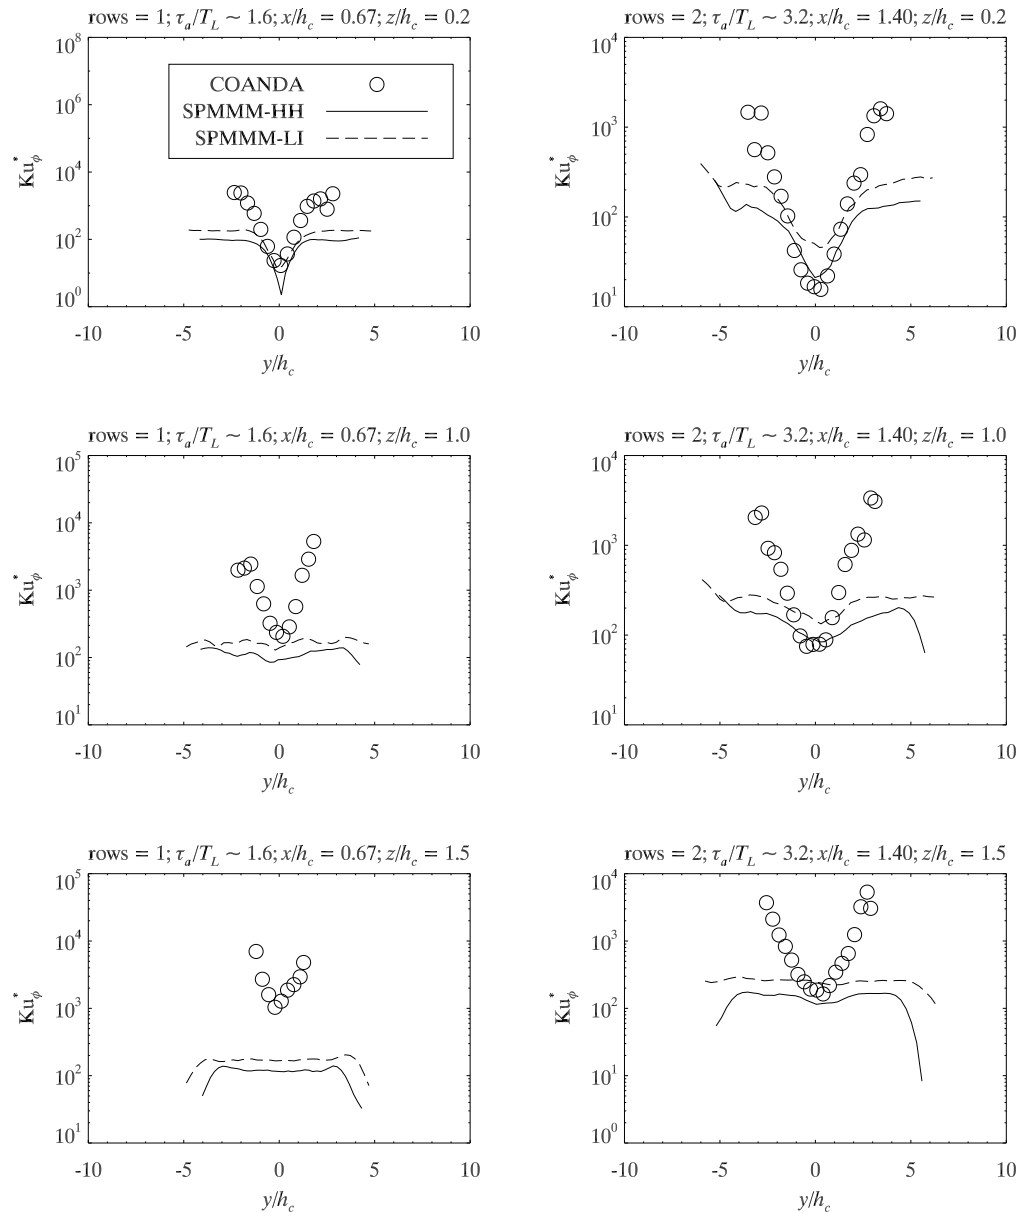


Figure 5.27: Spanwise transects of the kurtosis of the dimensionless concentration at three heights, one row ($x/h_c = 0.67$; left column) and two rows ($x/h_c = 1.40$; right column) downstream from the source from SPMMM simulations of an elevated point source ($z_s/h_c = 0.2$) in the Tombstone canopy flow.

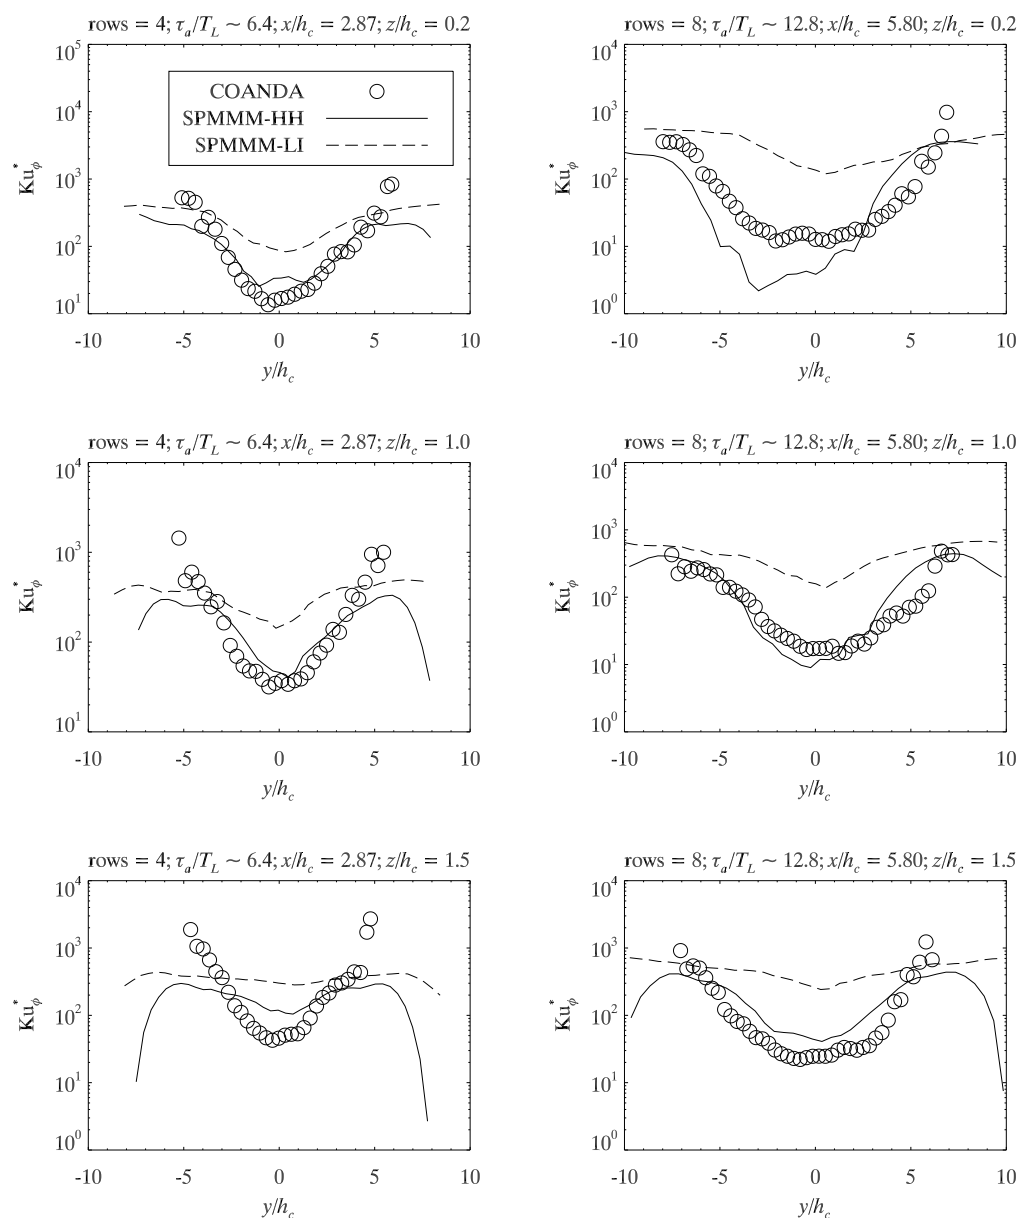


Figure 5.28: Spanwise transects of the kurtosis of the dimensionless concentration at three heights, four rows ($x/h_c = 2.87$; left column) and eight rows ($x/h_c = 5.80$; right column) downstream from the source from SPMMM simulations of an elevated point source ($z_s/h_c = 0.2$) in the Tombstone canopy flow.

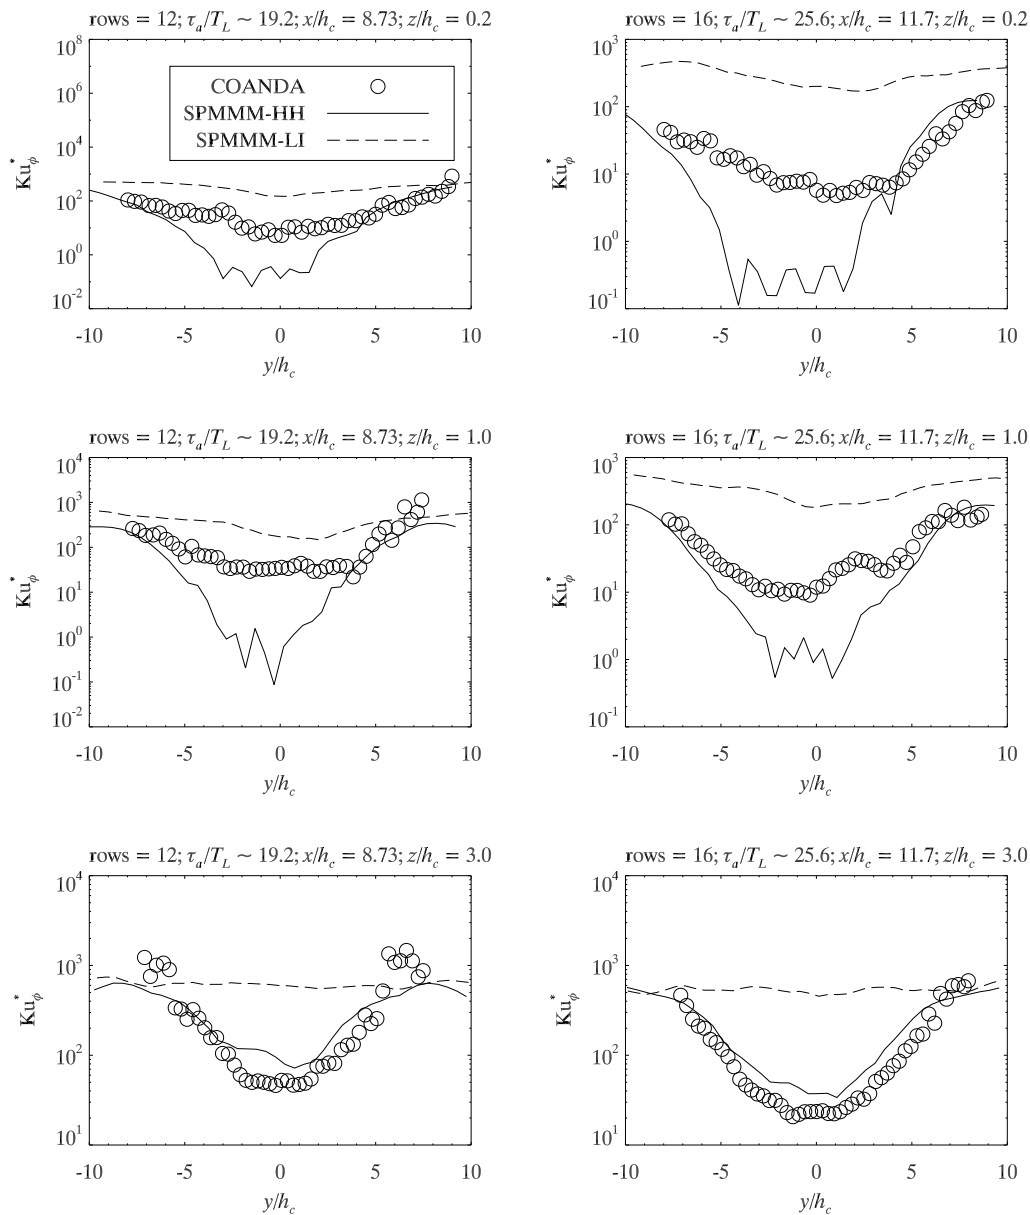


Figure 5.29: Spanwise transects of the kurtosis of the dimensionless concentration at three heights, twelve rows ($x/h_c = 8.73$; left column) and sixteen rows ($x/h_c = 11.7$; right column) downstream from the source from SPMMM simulations of an elevated point source ($z_s/h_c = 0.2$) in the Tombstone canopy flow.

the interpolation, and may strongly influence the predictions of the concentration distribution). These predictions will improve only if these gradients happen to be correct. If not, it could result in drastically poorer predictions in comparison with the simple assumption of horizontal homogeneity, where these gradients in Reynolds stresses disappear—except of course for the vertical gradients. Therefore, it may be better (in terms of model accuracy) to oversimplify a flow field by spatially averaging it than to partially, imperfectly, or poorly represent the inhomogeneities in the field.

The poorer first-order consistency between **MEANS** and **SPMMM** was partially responsible for the poorer fit of the **SPMMM-LI** simulation results to the experimental data. The poorer first-order consistency is related to the presence of rogue velocities and is considered in greater detail in Chapter 6. The fault may also lie in the fact that **SPMMM** does not resolve the obstacles. If the obstacles were represented then the particles would reflect off of them, and perhaps escape the canopy more readily. This mechanism was not investigated in this thesis.

5.4 Chapter Summary

In this chapter **SPMMM** was used to simulate dispersion from continuous, elevated, in-canopy line and point sources in the Tombstone canopy flow. The free parameters of the model were assigned the calibrated values from Chapter 4: $C_0 = 6.0$, $\mu = 0.75$, and $C_r = 0.45$ (the sole exception being one simulation in which $C_0 = 2.0$ and $C_r = 0.18$, for comparison with the Cassiani et al. (2007) simulations).

For the line source dispersion simulations, the necessary velocity statis-

tics to drive the model were extracted from figures in Legg et al. (1986) and interpolated onto the SPMMM grid. The velocity statistics were treated as horizontally-homogeneous. Measurements of the TKE dissipation rate and the Lagrangian integral timescale had been provided by Raupach et al. (1986) and Legg et al. (1986) respectively. With the equation $\varepsilon = 2\sigma_w^2/C_0T_{Lw}$, and by assigning $C_0 = 2.0$, an inferred TKE dissipation rate that matched the experimental profile of the TKE dissipation rate reasonably well was computed. Cassiani et al. (2007) followed this approach, and consequently set the value of the Richardson constant to $C_r \approx C_0/11$, as suggested in Franzese and Cassiani (2007). Given that $C_0 = 2.0$ is on the low end of the recently reported values for the Kolmogorov constant, and that in Chapter 4 the Kolmogorov constant was calibrated to $C_0 = 6.0$, we preferred to not re-tune C_0 and therefore determined the TKE dissipation rate as the sum of the turbulent shear production and the turbulent wake production (the turbulent transport term was neglected as it is not in closed form and would therefore have added another tunable parameter to the model). This approach led to a good match between the TKE dissipation rate used by SPMMM and those reported in Legg et al. (1986), without the need to fix the Kolmogorov constant at a certain value.

The cross-wind line source of the wind-tunnel Tombstone canopy experiments was at a height of $z_s/h_c = 0.85$, with a diameter of $d_s = 0.9$ mm. For the SPMMM simulations the spatial distribution of the source particles was assumed to be Gaussian. It was found that an initial source distribution of $\sigma_0 = 2.0$ produced good agreement with the experimental standard deviation of concentration measurements at $x/h_c = 0.38$. The SPMMM model predictions

of the mean concentration were satisfactory to good above the canopy. Within the canopy, the mean concentration was under-predicted close to the source ($x/h_c = 0.38$ and $x/h_c = 2.78$) and over-predicted farther from the source ($x/h_c = 11.7$). The results for the standard deviation of the concentration were qualitatively good. At $x/h_c = 0.38$ the standard deviation of concentration was under-predicted, as it was above the canopy farther downstream ($x/h_c = 2.78$ and $x/h_c = 11.6$). The standard deviation of the concentration was slightly over-predicted within the canopy at $x/h_c = 2.78$ and $x/h_c = 11.6$. The discrepancies may be related to the velocity statistics, the modelled TKE dissipation rate, or to the fact that SPMMM does not resolve the obstacles that make up the canopy.

To compare SPMMM to the model of Cassiani et al. (2007), a simulation using a TKE dissipation rate inferred from the experimentally measured Lagrangian integral timescale (Legg et al., 1986) was carried out. The agreement was good between the SPMMM results and the Cassiani et al. (2007) results, with differences likely attributable to the implementation of the models: SPMMM utilised single-particle trajectories and a pre-calculated conditional mean concentration field whereas the model of Cassiani et al. (2007) utilised parallel particle trajectories and calculated the conditional mean concentration at each timestep. The vertical profiles of the dimensionless mean concentration from this simulation were more similar to the profiles as simulated by Flesch and Wilson (1992) (whose model was formulated in terms of the Lagrangian integral timescale, as evaluated and communicated directly by Legg et al. (1986)) than the SPMMM simulations that used a TKE dissipation rate determined as the sum of shear production and wake production.

For the point source dispersion simulations in the Tombstone canopy flow, the necessary velocity statistics were interpolated from the experimental water-channel measurements of Hilderman and Chong (2007). Two SPMMM simulations were performed and compared for the point source configuration: one utilising spatially-averaged, horizontally-homogeneous velocity statistics, and the other utilising locally inhomogeneous velocity statistics. The TKE dissipation rate for both simulations was calculated as the sum of the turbulent shear production and the turbulent wake production, as already discussed above.

The 15 mm point source was placed behind an obstacle at a height of $z_s/h_c = 0.2$. It was designed to have a uniform concentration across its diameter and was therefore simulated as a tophat source. Close to the source ($x/h_c \lesssim 2.0$), the use of locally inhomogeneous velocity statistics led to more accurate predictions of the mean and standard deviation of the dimensionless concentration, but also resulted in decreased accuracy farther downstream. For the skewness and the kurtosis of the dimensionless concentration, the SPMMM results with horizontally-homogeneous velocity statistics are more accurate than the SPMMM simulations with locally inhomogeneous velocity statistics. In both simulations it appeared that not enough particles were escaping the canopy, resulting in over-predicted concentrations within the canopy and under-predicted concentrations above the canopy.

From a pragmatic perspective the incorporation of locally inhomogeneous velocity statistics is counterproductive in terms of accuracy, at least in reference to these experiments in a Tombstone canopy flow, as that effort not only complicates the trajectory modelling, but more importantly leads to less accurate results when considering the full downstream fetch and the first four

moments of the concentration field (however, obviously the approach of representing the inhomogeneity of the flow represents “higher fidelity” science, and indeed has been shown to be crucially advantageous in other flows, e.g. Wilson et al. (2009)). If predictions of the concentration statistics are needed only very close to the source then there appears to be some benefit to incorporating the local velocity inhomogeneities.

Causes for this deterioration in performance when utilising locally inhomogeneous velocity statistics may include: exclusion of the large scale inhomogeneities (over-simplification of the extremely turbulent Tombstone canopy flow may in fact be more accurate than a partial representation of the turbulence though the introduction of local inhomogeneity); the local inhomogeneities were interpolated from a small number of measurements and there is no guarantee that resulting interpolations are physically valid—which may result in inaccurate gradients in the Reynolds stresses which could strongly affect the predictions of the concentration field; poorer first-order consistency in the SPMMM-LI simulation, perhaps due to the presence of rogue velocities; the fact that SPMMM does not resolve obstacles which may prevent particles from reflecting their way out of the canopy. While obstacle resolution is not investigated in this thesis the use of a flow field with large scale inhomogeneities is covered in the next chapter, as is an examination of the effects of rogue particle velocities on model consistency.

Chapter 6

MUST Canopy Simulations

This chapter represents the culmination of my thesis work on the modelling of concentration fluctuations, in that it identifies a severe problem with the whole approach—namely, an instability of the underlying algorithm for computing particle paths that manifests when the trajectory model is driven by flow statistics that (whether correctly or incorrectly) entail highly irregular variability from node to node. The problem is revealed here in the context of simulations of the Mock Urban Setting Trials (MUST) experiments, which document concentration fluctuations in a point source plume meandering its way through a regular array of volume-occupying obstacles.

The MUST Canopy has been the subject of three experiments: full-scale, atmospheric experiments (Yee and Biltoft, 2004); 1:50 scale wind-tunnel experiments (Gailis and Hill, 2006); and 1:205 scale water-channel experiments (Hilderman and Chong, 2007). The three experimental data sets are thoroughly compared by Yee et al. (2006), who concluded that the water-channel experiments more accurately than the wind-tunnel experiments reproduced quantitatively the detailed plume structure of the full-scale atmospheric experiments (when appropriately scaled). In this thesis we used the experi-

mental data from the water-channel experiments and therefore discuss these experiments more fully. The full-scale and wind-tunnel experiments are briefly described for completeness.

6.1 Experimental and Computational Setup

The full-scale, atmospheric experiments of the Mock Urban Setting Trials were conducted at the US Army Dugway Proving Grounds in the Great Basin Desert of Utah. The canopy consisted of one hundred and twenty shipping containers with dimensions $2.42 \text{ m} \times 12.2 \text{ m} \times 2.52 \text{ m}$ (length, width, and height respectively) that were regularly placed in twelve rows, each containing ten containers. The obstacles were oriented such that the 12.2 m side faced the streamwise component of the mean wind. The average obstacle face to face spacings (i.e., the gap between the obstacles) in the streamwise and spanwise directions were 12.9 m and 7.9 m , respectively. This arrangement of the obstacles resulted in a frontal area index of $\lambda_f = 0.10$ (more than a factor of two lower than the Tombstone Canopy) and a plan area index of $\lambda_p = 0.096$. A plan view of the MUST array can be seen in Figure 6.1. Note that the dimensions shown in this figure correspond to the water-channel experiments (discussed below). Velocity statistics upstream, within, and downstream of the array were collected with 2-D and 3-D sonic anemometers. The majority of the dispersion experiments were undertaken in slightly stable to stable nocturnal conditions, although a small number of the experiments had near-neutral stability.

Pure propylene gas was released from a 0.05 m diameter pipe at five different heights (on different occasions): $z_s/h_c = 0.059, 0.51, 0.71, 1.02,$ and 2.05 .

The flow rate was sufficiently low such that the propylene could be considered a passive tracer, regardless of propylene having a density 1.45 times that of air. Depending on the source height the pipe was either vertically or horizontally oriented. Concentration data were collected using forty-eight digital photo-ionisation detectors, which have a frequency response of 50 Hz and a sensitivity of about 0.01 parts per million by volume of propylene. Due to the high levels of turbulence within the canopy, the resulting concentration measurements do not distinguish between dispersion from a vertically oriented pipe and dispersion from a horizontally oriented pipe. The position of the source varied from being upwind of the array to within the array.

The 1:50 scale wind-tunnel MUST experiments were performed in the 1.5 MW Boundary Layer Wind-Tunnel (measuring 55 m \times 12 m \times 4 m in length, width, and height respectively) at Monash University in Victoria, Australia. The working section of the tunnel is approximately 10–12 metres in length, depending on the type of experiment being conducted. A 1:50 scale version of the MUST Canopy was placed within the tunnel. The length, width, and height of the obstacles were 48 mm \times 244 mm \times 48 mm, respectively. Note that the length and width of the wind-tunnel MUST Canopy obstacles are 1:50 scale, whereas the height of the obstacles is marginally inconsistent, at 1:48 scale. Digital anemometers extracted mean wind speeds while crossed hot-wire anemometers extracted turbulence measurements. Velocity measurements did not extend above $z = 1.00$ m. By this height there was little change in the mean streamwise velocity, and the boundary layer depth was taken to be $\delta \approx 1.00$ m. The ratio of the boundary layer depth to the canopy height was thus $\delta/h_c \approx 20$. The majority of the experiments had neutral stability, but

a few were stably stratified. To produce a stably stratified boundary layer, heating elements in the ceiling of the wind-tunnel were turned on. As the air passed over these elements, it became hotter than the underlying layers, resulting in an inverted stably stratified boundary layer. For the dispersion experiments with stable stratification, the 1:50 scale MUST Canopy and the source were inverted and affixed to the ceiling of the wind-tunnel.

Pure ethylene gas was released from a 10 mm diameter ground-level point source. As the density of ethylene is 0.965 times that of air, and the temperature of the ethylene was the same as the air, the tracer could be considered passive. Concentration data were collected by ten flame ionisation detectors. While qualitatively similar to the full-scale MUST array plume, the wind-tunnel array plume in neutrally stratified flows had lateral and vertical spreads that were 2 to 3 larger than the full-scale array plume, and consequently an under-predicted peak mean concentration. The stably stratified experiments however produced results much more similar to the full-scale field experiments.

The 1:205 scale water-channel MUST experiments were conducted by Coanda Research and Development Corporation in the same facility that had earlier been used for the Tombstone Canopy experiments, discussed in Section 5.1. The length, width, and height of the obstacles were $11.8 \text{ mm} \times 59.4 \text{ mm} \times 12.4 \text{ mm}$, respectively. Figure 6.1 shows the water-channel MUST obstacle array in plan view. Lettered points correspond to velocity statistic data extraction positions. Point F lies in the centre of the array (i.e. $(x_F, y_F) = (0, 0)$). For the experiments used in this thesis, the ground-level point source was at point L. As for the Coanda Tombstone Canopy experiments, laser Doppler velocimetry was used to extract velocity statistics. Velocity data for the vertical




Figure 6.1: The water-channel MUST obstacle array in plan view. The height of the obstacles was 12.4 mm. Lettered points correspond to velocity statistic data extraction positions. The ground-level point source was at point L. Dotted lines correspond to concentration statistic data extraction locations 1, 2, 3, 5, 8, and 11 rows down stream of the source.

direction were not extracted at points A, C, G, I, K, and O. With regards to notation, the line of points A - H, which does not pass through any obstacles, shall be referred to as a *canyon line*, whereas the line of points I - P, which does pass through obstacles, shall be referred to as an *obstacle line*.

From the measured velocity statistics, the friction velocity was evaluated as $u_* = 0.0255 \text{ m s}^{-1}$, the boundary layer depth was determined to be $\delta = 275 \text{ mm}$ and the free-stream velocity to be 0.375 m s^{-1} , giving a Reynolds number (based on δ and $\langle u \rangle_\delta$) of $\text{Re}_\delta = 1.0 \times 10^5$. The ratio of the boundary layer depth to the canopy height was $\delta/h_c = 22.2$, similar to that of the wind-tunnel MUST experiments. The mean streamwise velocity at the height of the canopy was $\langle u \rangle_c \approx 0.165 \text{ m s}^{-1}$, when averaged along the inflow region of the MUST array (i.e., the leading edge points A, B, C, I, J; point K not included because the velocity measurements at this location do not extend to the ground-level). The corresponding Reynolds number (based on $\langle u \rangle_c$ and h_c) is $\text{Re}_c = 2040$. Castro and Robins (1977) suggested the criterion $\text{Re} > 4000$ for acceptable physical modelling of flow around surface mounted cubes, and thus $\text{Re}_c = 2040$ seems to be too low. However, as suggested by MacDonald et al. (1998), for flows around obstacles with a significant non-unity aspect ratio (the aspect ratio of the MUST array obstacles is $59.4 : 12.4 \approx 4.79 : 1$), a more relevant length scale for the interaction of the flow with the obstacles is $h_c^* = A_f^{1/2}$, where A_f is the frontal area of the obstacles (note that for non-normal incidence flow the projected frontal area must be used). With h_c^* and $\langle u \rangle_c$, the Reynolds number for the MUST flow at canopy height is $\text{Re}_c^* = 4470 > 4000$, which is sufficient to ensure Reynolds number independence in the water-channel experiments. We repeat for gravitas that the analysis of Yee et al. (2006) concluded that

the water-channel experiments more accurately reproduced quantitatively the full-scale atmospheric experiments than the wind-tunnel experiments, when appropriately scaled.

Sodium fluorescein dye (which we recall from Chapter 5 is dynamically passive in water) was released from the 2.8 mm inner diameter, vertically oriented tube at a constant rate of $Q = 12 \text{ mL min}^{-1}$. This release rate was low enough to ensure that the dye did not escape the recirculating wake region behind the obstacles during release from the source. Laser induced fluorescence techniques were used to extract concentration data 1, 2, 3, 5, 8, and 11 rows downstream from the source at several heights. The concentration extraction locations are shown as dotted lines in Figure 6.1. The raw concentration data were made dimensionless with equation (5.3) before computing statistics.

To drive the MEANS and SPMM simulations, spanwise periodicity of the measured Coanda velocity statistics was assumed, and the array shown in Figure 6.1 was reduced to include only the area visited by the plume, as shown in Figure 6.2. The dotted lines demarcate the grid used to interpolate the measured velocity data. The shaded point L bin corresponds to point L in Figure 6.1. The interpolation of the velocity statistics was similar to the interpolation used for the locally inhomogeneous Tombstone Canopy velocity statistics (i.e., Section 5.1). The lettered point data shown in Figure 6.2 was interpolated onto a grid consisting of $N_x^{\text{vel}} = 12$ bins in the streamwise direction, $N_y^{\text{vel}} = 25$ bins in the spanwise direction and $N_z^{\text{vel}} = 40$ bins in the vertical direction. To guide the splines to the ground-level, we assumed the velocity statistics remained constant at the $\sim 0.5h_c$ value (the lowest measurement height) for $z < 0.5h_c$. While it may have been more realistic to enforce a

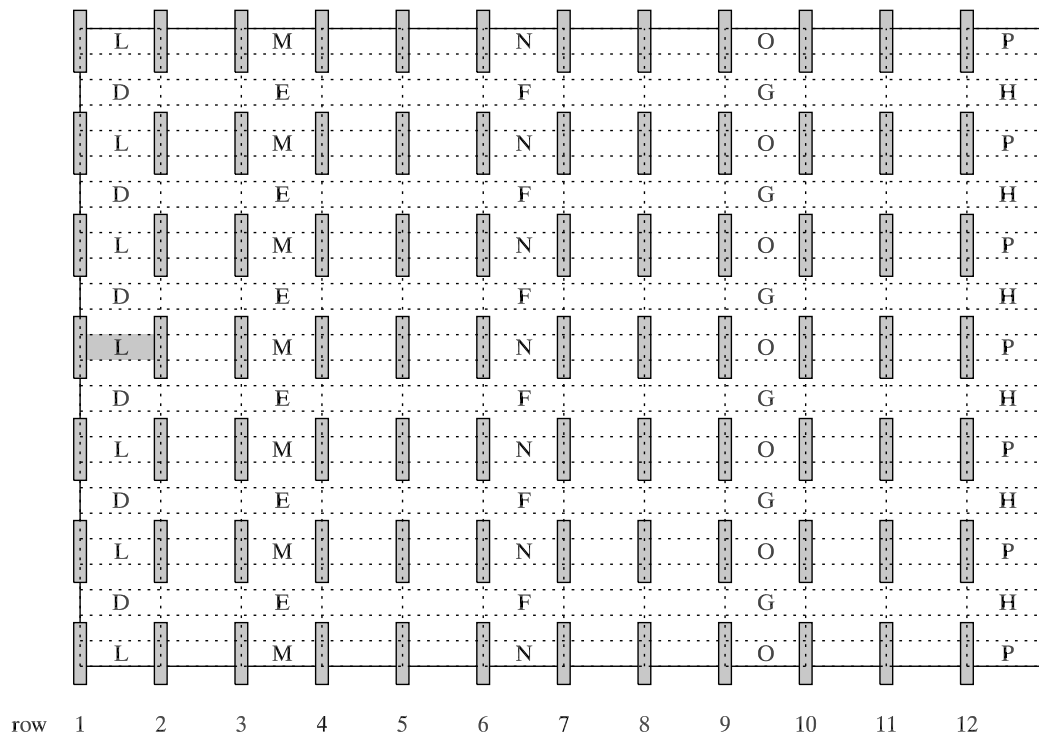


Figure 6.2: The reduced water-channel MUST obstacle array in plan view, as used for the **MEANS** and **SPMMM** simulations. Areas of the array not visited by the plume in the water-channel experiment have been removed. The dotted lines demarcate the grid used to interpolate the measured velocity data. The shaded point L bin corresponds to point L in Figure 6.1.

generic decay of the velocity statistics to the ground-level, the additional detail of the velocity statistics would have been lost in the relatively low resolution of the discretisation grids. There was also no guarantee that the functional form of the decay would have been correct for the MUST array flow, as the lowest measurement height of the velocity statistics was still relatively high at $z/h_c \approx 0.5$.

For each lettered velocity extraction point in Figure 6.2, vertical profiles of the statistics have been produced using linear spline interpolation. Cubic spline interpolation (as used in the Tombstone Canopy simulations) resulted in undesirably large oscillations between the velocity measurement locations near ground-level. This was due to fewer in-canopy velocity statistic measurements to guide the splines in the MUST canopy experiments versus the number of in-canopy measurements in the Tombstone Canopy experiments (e.g., compare with Figure 5.3). The resulting vertical splines can be seen for the mean streamwise velocity at points L and P in Figure 6.3.

The vertical profiles of the velocity statistics from extraction locations D, E, F, G, H, L, M, O, and P were then interpolated onto the $N_x^{\text{vel}} \times N_y^{\text{vel}}$ grid for each vertical level using thin plate splines. Streamwise transects along a line through the source (i.e., along an obstacle line from point L to point P) at four heights are shown in Figure 6.4. Note that heights in this figure (and those that follow) refer to the bin that contains that height. For example, the “ground-level” transect plots those bins on the line from point L to point P that are in contact with the ground, the “ $z/h_c = 1.0$ ” transect plots those bins on the line from point L to point P that contain the height $z/h_c = 1.0$, and so on. The row indices along the x -axis correspond to the row indices




Figure 6.3: Vertical profiles of the mean streamwise velocity at points L and P from the water-channel MUST canopy experiments and the resulting interpolation used to drive MEANS and SPM3M simulations.

in Figures 6.1 and 6.2, as do the letters. Corresponding lines of letters in




Figure 6.4: Streamwise transects of the interpolated mean streamwise velocity along an obstacle line (i.e., from point L to point P) at four heights for the water-channel MUST canopy experiments. These interpolations were used to drive the MEANS and SPMM simulations. The row indices along the x -axis correspond to the rows in Figures 6.1 and 6.2, as do the letters.

Figure 6.2 have identical interpolations. An interesting feature in the transects shown in the above figure is that at the source location (ground-level transect at point L) the mean streamwise velocity is negative (i.e., upstream) with a value of $\langle u \rangle / u_* = -0.28$, showing a recirculating wake region behind the obstacle which shields the source location (see Figure 6.1).

Spanwise transects through the source (i.e., between rows 1 and 2) are shown for four heights in Figure 6.5. The letters along the x -axis correspond to the lettered velocity statistic extraction points in Figures 6.1 and 6.2. As for




Figure 6.5: Spanwise transects of the interpolated mean streamwise velocity at four heights for the water-channel MUST canopy experiments. These interpolations were used to drive the MEANS and SPMM simulations. The letters along the x -axis correspond to the letters in Figures 6.1 and 6.2.

the Tombstone Canopy simulations, the interpolated velocity statistics were allowed to take non-zero values in regions occupied by the canopy obstacles, and particles were permitted to enter the interior volumes of the obstacles. Recall that this was permitted to prevent singularities in the Reynolds stress tensor and to prevent particles from getting stuck (i.e., not moving downstream to x_{\max}) within the simulation domain, since this version of the SPMM model suite is not obstacle resolving.

The TKE dissipation rate was calculated assuming local equilibrium as $\varepsilon = \mathcal{P}_s + \mathcal{P}_w$, as previously discussed at the end of Section 5.1. The mod-




Figure 6.6: Modelled vertical profiles of the TKE dissipation rate at points A and P for the water-channel MUST canopy experiments. The TKE dissipation was calculated assuming local equilibrium ($\varepsilon = \mathcal{P}_s + \mathcal{P}_w$).

elled vertical profiles of the TKE dissipation for points L and P are shown in

Figure 6.6. The other points in the array have similar profiles.

Examples of the other velocity statistics are shown in Figure 6.7, which shows the vertical profiles of the mean velocities and standard deviation of the streamwise velocity, and Figure 6.8, which shows the vertical profiles of the standard deviation of the spanwise and vertical velocities, and the velocity covariances. Both figures show the experimental measurements and the interpolated profiles for two points in the centre of the MUST array: point F (on a canyon line), and point N (on an obstacle line). Above the canopy,

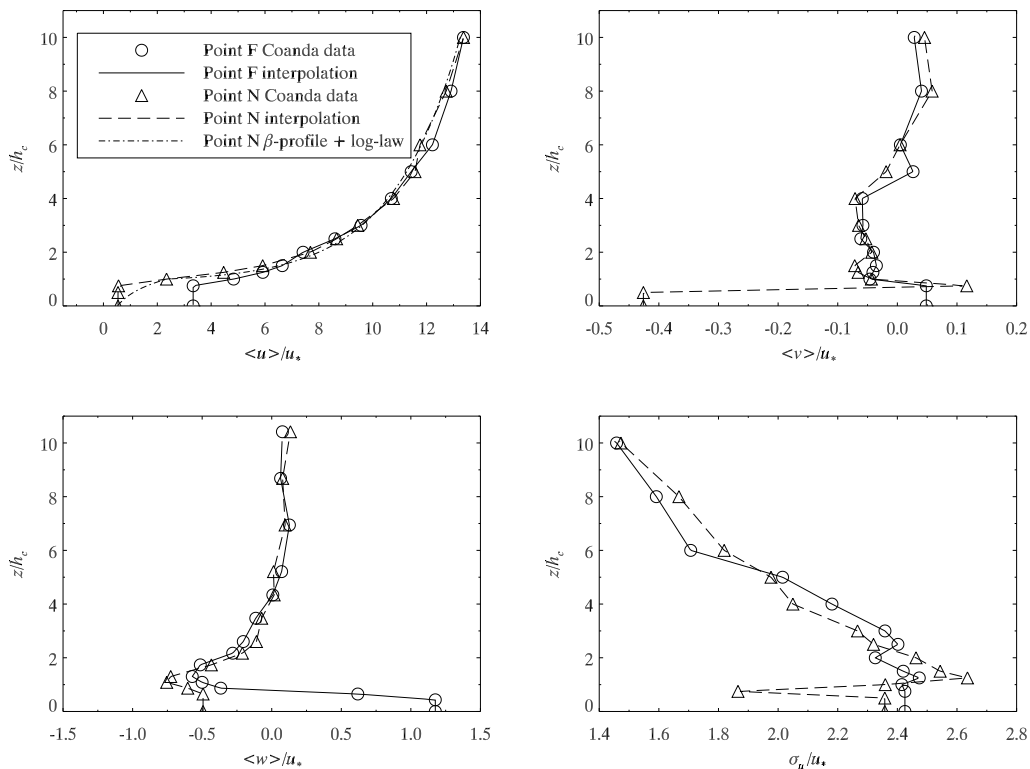


Figure 6.7: Vertical profiles of the mean velocities and the standard deviation of the streamwise velocity for point F (on a canyon line) and point N (on an obstacle line) from the water-channel MUST canopy experiments. The symbols represent the experimental measurements and the lines represent the modelled profiles used to drive MEANS and SPMMM.

the profiles for all of the velocity statistics at both points are similar in shape and magnitude. The covariance $\langle u'v' \rangle$ is an exception to this observation, but its magnitude is very small when compared to the other statistics. Notable

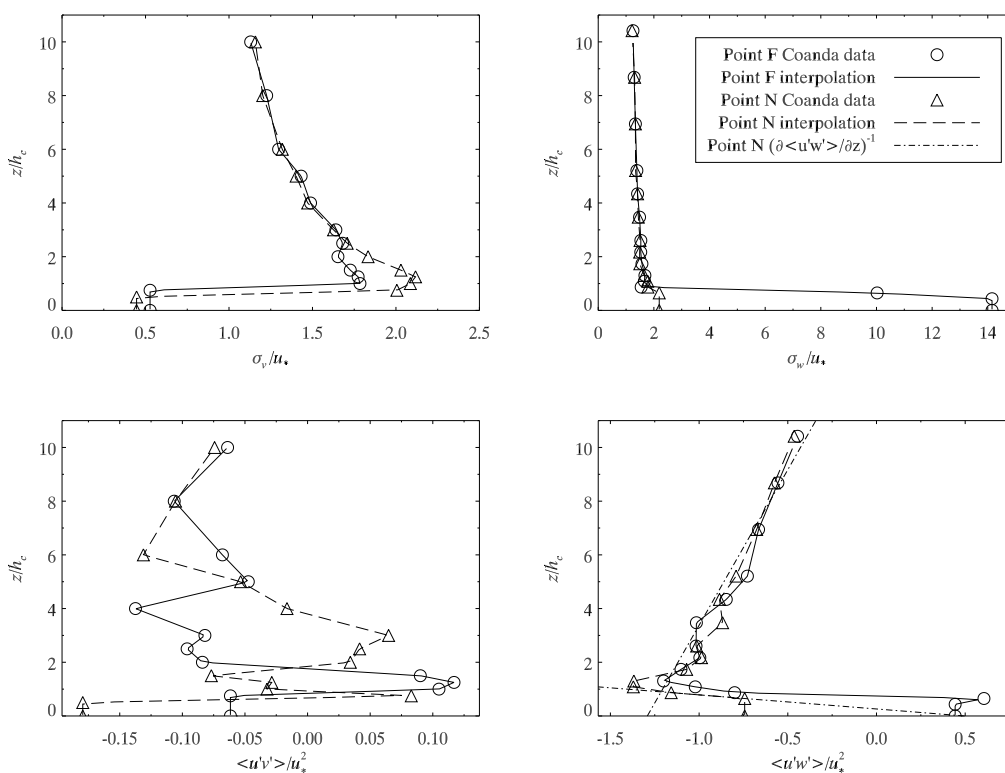


Figure 6.8: Vertical profiles of the standard deviation of the spanwise and vertical velocities and the velocity covariances for point F (canyon line) and point N (obstacle line) from the water-channel MUST canopy experiments. The symbols represent the experimental measurements and the line represent the modelled profiles used to drive **MEANS** and **SPMMM**.

features in the vertical profiles below the canopy height include: the mean streamwise velocity is (not surprisingly) lower for point N, which is behind an obstacle, than for point F; the mean vertical velocity at point F is positive, tending to push material out of the canopy, while it is negative at point N, tending to draw material into the canopy; the magnitude of the standard de-

variation of the vertical velocity at point F (canyon line) is approximately seven times larger than at point N (obstacle line), and its above canopy values. With regards to σ_w , the obstacle line and canyon line point pair N-D is anomalous in that the in-canopy standard deviation of the vertical velocity is larger for the canyon line point than for the obstacle line point (i.e., $\sigma_w^F > \sigma_w^N$ for $z/h_c < 1$), opposite to what is expected (larger standard deviations are expected in the re-circulating wake region behind the obstacles, that is, along the obstacle lines). All other pairs where w velocity data were extracted show the expected behaviour for $z/h_c < 1$: $\sigma_w^L > \sigma_w^D$, $\sigma_w^M > \sigma_w^E$, $\sigma_w^P > \sigma_w^H$ (obstacle line points on the left of the inequality, canyon line points on the right). The anomaly for the point pair N-D serves to highlight the large scale inhomogeneities in the MUST canopy flow.

Also shown in Figures 6.7 and 6.8 are interpretations for point N (midway between obstacles on an obstacle line) of the mean streamwise velocity and the $\langle u'w' \rangle$ shear stress in the context of one-dimensional aerodynamics. In the top left panel of Figure 6.7, the β -profile (equation (2.52)) and the displaced log-law (equation (2.53)) are shown to fit the measured streamwise velocity at point N very well. The parameters for these fitted velocity models were: $\beta = 1.5$ and $d = 0.88h_c$ (which is slightly larger than the $d \approx 0.7h_c$ rule of thumb for the displacement height). In the bottom right panel of Figure 6.8, lines of best fit have been fitted to the above and below canopy shear stress profiles. As plotted with the independent variable on the y -axis, the slopes of these lines represent $(\partial \langle u'w' \rangle / \partial z)^{-1}$. In Chapter 2, it was shown that for a horizontally-homogeneous flow the vertical gradient of the shear stress is

balanced principally by the streamwise pressure gradient and the Coriolis force,

$$\frac{\partial \langle u'w' \rangle}{\partial z} = -\frac{1}{\rho} \frac{\partial \langle p \rangle}{\partial x} + f_c \langle v \rangle. \quad (2.32)$$

By assuming horizontal-homogeneity* above the canopy, ignoring Coriolis effects (which are exceedingly small at the scale of the water-channel), and calculating the slope of the above canopy line of best fit, the streamwise pressure gradient above the water-channel MUST Canopy is found to be approximately

$$\frac{\partial \langle p \rangle}{\partial x} \approx -2.58 \times 10^{-3} \text{ Pa m}^{-1}, \quad (6.1)$$

which is comparable to a typical mid-latitude mesoscale atmospheric value of $\sim 10^{-3} \text{ Pa m}^{-1}$ (Parish et al., 2007).

Within the canopy, there is rapid stress attenuation due to drag on the obstacles, and (as also shown in Chapter 2) the dominant terms in the momentum equation are

$$\frac{\partial \langle u'w' \rangle}{\partial z} = -\frac{1}{2} C_D a [\langle u \rangle]^2. \quad (2.49)$$

Since the right-hand side of the above equation is equal to the slope of the below canopy $\langle u'w' \rangle$ line of best-fit, the element area density a (calculated with λ_f) is known, as is the spatially-averaged mean streamwise velocity below the canopy, the effective drag coefficient of the MUST Canopy is $C_D \approx 2.0$. Note that in reality the MUST canopy flow is not one-dimensional and strong pressure gradients across the canopy existed, and so equation (6.1) is severely wrong, and this interpreted drag coefficient is not suggested to be rigorous. A rigorous approach to determining the effective drag coefficient of an array of solid obstacles is discussed in Lien and Yee (2005)

*The $\langle u'w' \rangle$ shear stress was spatially-averaged over the entire MUST Canopy.

Lastly, we note that a suspected outlier at point E in the Coanda measurements of σ_w was excluded from the interpolation, as shown in Figure 6.9. None of the other Coanda velocity statistics measurement locations showed a similar spike in their σ_w profile so high above the canopy.




Figure 6.9: A suspected outlier at point E in the Coanda σ_w measurements. No other Coanda velocity statistic measurement location showed a similar spike in the σ_w profile so high above the canopy. The outlier was excluded from the interpolation.

6.2 Dispersion from a Continuous Point Source

The MUST canopy simulations of dispersion from a continuous point source were severely limited by two factors, which are discussed in turn. First,

within the canopy, the driving velocity statistics had large variations (inhomogeneities) over short length scales (e.g., see the σ_w profile in Figure 6.8), which required that the u , v , and w velocity spaces each be discretised into at *least* 70 bins in order to ensure that the driving velocity PDFs numerically integrated to unity. This number of velocity bins greatly increased the computational resources required by the model, reduced the spatial resolution of the simulations, and resulted in increased statistical error. These problems would be easily solved by using a more powerful computer, and more particles in the simulations. The second (and more severe) limitation was that the steep gradients in the Reynolds stresses near the canopy top (most notably $\partial\sigma_w^2/\partial z$) resulted in the generation of many rogue velocities. As will be demonstrated in Section 6.2.1, the re-initialisation of these rogue velocities breaks the theoretically required consistency between the **MEANS** and **SPMMM** modelled mean concentration.

To demonstrate the second of the above limitations, simulations were run in which $N_\phi = 2 \times 10^7$ particles were used to pre-calculate the unconditional and conditional means with **MEANS**, and $N = 2 \times 10^7$ particles (run as 10 simulations of 2×10^6 particles) were used in the **SPMMM** micromixing simulations. As mentioned above, 70^3 velocity bins ($N_u = N_v = N_w = 70$) had to be used to discretise velocity space, and so the discretisation of the driving velocity statistics in the vertical direction was reduced to 25 bins from 40 in order to reduce the memory required for the simulations below the amount available in the computer. The spatial resolution of the simulation was set to match the discretisation of the driving velocity statistics: $N_x = N_x^{\text{vel}} = 12$, $N_y = N_y^{\text{vel}} = 25$, and $N_z = N_z^{\text{vel}} = 25$. The timestep scale factor was $\mu_t = 0.02$. The

Kolmogorov constant, the micromixing constant, and the Richardson constant were set to the values optimised for the the FR82 experiments in Chapter 4: $C_0 = 6.0$, $\mu = 0.75$, and $C_r = 0.45$, respectively. The initial distribution of particles at the source was assumed to be Gaussian with an initial spread of $\sigma_0 = d_s$.

Figure 6.10 shows the spanwise transects of the mean dimensionless concentration (top left panel), the standard deviation of the dimensionless concentration (top right panel), the skewness of the dimensionless concentration (bottom left panel), and the kurtosis of the dimensionless concentration (bottom right panel), for a position 11 rows downstream of the source (row index 12.5) at $z/h_c = 0.5$. The accuracy of the SPMMM predictions is poor for each statistic as a result of poor first-order consistency between MEANS and SPMMM, an example of which can be seen in Figure 6.11. The figure shows the spanwise transects of the mean concentration 11 rows downstream of the source (row index 12.5), at the top of the canopy $z/h_c = 1.0$. In total, 9.16×10^6 out of 1.44×10^{10} steps had rogue velocities in the MEANS simulation (which computes the underlying conditional mean concentration field for the SPMMM simulations). As a percentage, this represents only 0.064%, a seemingly insignificant amount. However, in the next section it will be shown that it is not the number of rogue velocities that occur that affects the first-order consistency, but rather how often these rogue velocities occur in the same spatial location.

Another discrepancy between SPMMM and the Coanda experimental results shown in Figure 6.10 is the spanwise position of the simulated maximum mean concentration is offset to the left in relation to the experimental results. The probable cause of this is allowing the particles to pass through the interior vol-

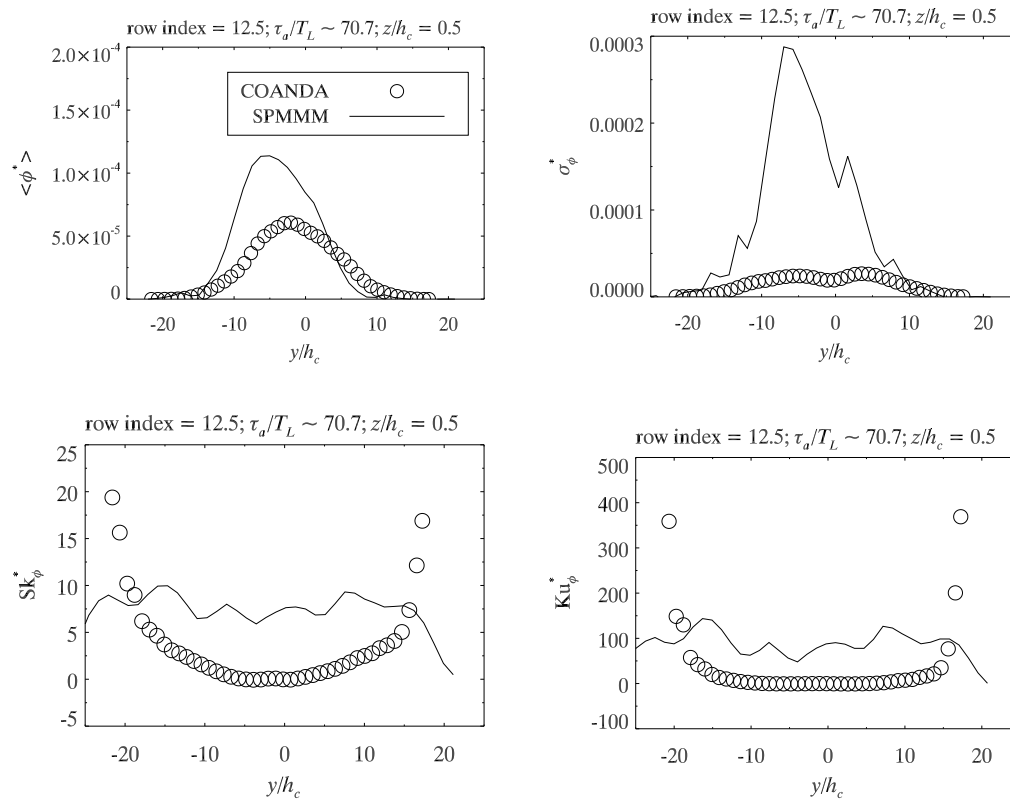


Figure 6.10: Spanwise transects of the first four moments of the dimensionless concentration field from SPMMM simulations of a ground-level point source in the MUST canopy flow. The transects are at a position 11 rows downstream of the source (row index 12.5) and $z/h_c = 0.5$. The spanwise source position is $y/h_c = -3.95$.




Figure 6.11: An example of the first-order inconsistency between the **MEANS** and **SPMMM** simulations for dispersion from a ground-level point source in the MUST canopy flow. The spanwise transects of the mean concentration shown here are for a position 11 rows downstream of the source (row index 12.5) at the top of the canopy.

umes of the obstacles. If the particles were prohibited from passing through the obstacles, they would have to travel around or over them, perhaps reducing this discrepancy. Future versions of SPMMM will resolve the obstacles and reflect particles off of their walls. This may not solve the discrepancy however. Figure 3(d) of Wilson et al. (2009) shows an offset to the right of the spanwise position of the simulated maximum mean concentration in relation to the water-channel MUST canopy experimental results (at a different location than the one used in Figure 6.10). The LS model used by these authors employed a reflection scheme when particles encountered an obstacle, and a computer generated velocity field.

6.2.1 Rogue Velocities Prevent Consistency

To demonstrate the effects that rogue velocities have on the models, two simulations were performed where the particles were restricted to the above canopy region (by setting the reflection height at the bottom boundary of the simulation domain to the canopy height: $z_{\text{rft}} = h_c$). This restriction prevented the particles from encountering the steep in-canopy gradients in the Reynolds stress tensor, thereby almost eliminating the occurrence of rogue velocities. To further simplify the flow, the mean vertical and spanwise velocities were zeroed, along with the $\langle u'v' \rangle$ shear stress ($\langle v \rangle = \langle w \rangle = \langle u'v' \rangle = 0$). Since the velocity statistics varied more slowly above the canopy, twenty velocity bins in each direction ($N_u = N_v = N_w = 20$) was sufficient to ensure that the velocity PDFs numerically integrated to unity. The discretisation of the driving velocity statistics in the vertical direction was returned to the values described in Section 6.1, and the spatial resolution of the simulations was set

to match: $N_x = N_x^{\text{vel}} = 12$, $N_y = N_y^{\text{vel}} = 25$, and $N_z = N_z^{\text{vel}} = 40$. For both simulations, the source height was elevated above the canopy ($z_s/h_c = 1.8$), while keeping the spanwise position at point L ($y_s/h_c = -3.95$) in Figure 6.1 (as in the Coanda experiments). The initial source distribution was set to $\sigma_0 = 7d_s$. These values of z_s and σ_0 were chosen as they provided a reasonable fit to the magnitude of the Coanda measured mean concentration at $z = 2h_c$, two rows downstream of the source (shown below in Figure 6.16), thus permitting some evaluation of the above-canopy performance of SPMMM. The values of all other parameters were the same as for the MUST simulations whose domain encompassed the full flow depth (i.e., down to $z/h_c = 0$) described in Section 6.1.

The first simulation (denoted R300) utilised the velocity statistics as described above and had few rogue trajectories in the MEANS simulations. The second simulation (denoted R301) utilised velocity statistics which did not correct the suspected outlier in σ_w at point E (shown in Figure 6.9), and therefore had produced many more rogue trajectories in the MEANS simulations, probably due to the resulting steeper gradients (i.e., $\partial\sigma_w^2/\partial x_i$) around the outlier. Note that due to the periodicity of the velocity statistics in the spanwise direction, there are multiple outliers in the driving velocity statistics (hereafter called *the outliers*). The number of rogue velocities in the SPMMM simulations was comparable. Table 6.1 summarises the number of particle steps with rogue trajectories for these two simulations. For both simulations, the number of rogue steps as a fraction of the total number of steps would appear to be negligible. In the R300 MEANS simulation, the 41 steps with rogue velocities occurred throughout the simulation domain. In contrast, the majority of the

	MEANS		SPMMM	
	Rogue Steps	Total Steps	Rogue Steps	Total Steps
R300	41	1.72×10^{10}	16649	4.82×10^9
R301	1183	1.70×10^{10}	19612	4.80×10^9

Table 6.1: The number of particle steps with rogue velocities in MUST canopy simulations for two different interpolations of the driving velocity statistics: one ignoring the suspected outlier at the point E Coanda velocity statistics measurement location (R300); and one including it (R301). For both simulations, the particles were prohibited from entering the canopy by setting the reflection height on the bottom boundary to $z_{\text{rft}} = h_c$.

“excess” rogue velocities in the R301 MEANS simulation occurred near the two velocity statistic bins[†] (left and right of the source) containing the outliers (i.e., between row indices 3 and 4; left bin: $-8.89 \leq y/h_c < -6.92$, right bin: $-0.99 \leq y/h_c < 0.99$; $z/h_c = 7.0$), as shown in the y - z scatter plot of rogue velocity positions in Figure 6.12. The shaded boxes are the velocity statistic bins left and right of the source which contain the outliers. There are 911 rogue velocity generation positions plotted in Figure 6.12. Of these, 428 cluster around the left velocity statistic bin, 460 around the right velocity statistic bin, 20 in the space between the two velocity statistic bins, and 3 elsewhere. The repeated occurrence of rogue velocities at these two locations resulted in poor first-order consistency between MEANS and SPMMM in the R301 simulations. The majority of the remaining 272 rogue velocities occurred one row upstream or downstream of the shaded velocity statistic bins shown in Figure 6.12, and at the same spanwise and vertical positions, again suggesting that steeper gradients of the Reynolds stresses caused by the large σ_w^2 value of

[†]Recall the three grids used in MEANS and SPMMM: the velocity statistic grid, used to discretise the driving velocity statistics in physical space ($N_x^{\text{vel}}, N_y^{\text{vel}}, N_z^{\text{vel}}$); the spatial grid, used to discretise the physical space (N_x, N_y, N_z); and the velocity grid, used to discretise the velocity space (N_u, N_v, N_w).





Figure 6.12: The spanwise and vertical positions of the rogue velocities between row indices 3 and 4. The shaded boxes demarcate the spanwise and vertical positions of the velocity statistic bins that contain the outliers. Note how the majority of the rogue velocities occur close to the outliers, particularly on the sides closest to the source ($y_s/h_c = -3.95$).


the outliers in this area of the simulation domain are responsible for the rogue velocity generation.

Figure 6.13 shows the spanwise transects of the mean concentration from the R300 (top panel) and R301 (lower panel) **MEANS** and **SPMMM** simulations. The transects are from row index 3.5 and height $z/h_c = 7.0$, which corresponds to the (x, z) location of the outliers. The spanwise location of the source is $y_s/h_c = -3.95$. The spatial bins containing the outliers span from $-0.99 \leq y/h_c < 0.99$ and $-8.89 \leq y/h_c < -6.92$ (shown as thin shaded rectangles along the y/h_c axis). The small differences between the R300 and R301 **MEANS** transects at this location are due to the differences between the flow fields for the two simulations. Below $z/h_c \approx 5.3$, the flow fields are identical. Above this height, the values of σ_w^2 and its gradients differ between the two flow fields, due to the outliers.

The consistency between **MEANS** and **SPMMM** is good for the R300 simulations, with the small differences in the transects attributable to statistical noise. The consistency in the R301 simulations is poor, with the **SPMMM** simulations overpredicting the mean concentration. Note that the spikes in the **SPMMM** transect in Figure 6.13(b) (where the inconsistencies are worst) are coincident with the edges of the velocity statistics bins that contain the outliers (i.e., the regions of intense rogue velocity generation shown in Figure 6.12). It is important to note that away from regions of rogue velocity generation the first-order consistency between the R301 **MEANS** and **SPMMM** simulations is good to excellent, as shown in the bottom panel of Figure 6.14 at row index 3.5 and height $z/h_c = 4.0$. The R300 results are shown in the top panel for comparison. The statistical noise in the R300 **SPMMM** profiles is less than in Figure 6.13(a) as $z/h_c = 4.0$ is



(a) Results from R300



(b) Results from R301

Figure 6.13: The spanwise transects of the mean concentration from the R300 and R301 MEANS and SPMMM simulations. The transects are from row index 3.5 and height $z/h_c = 7.0$, which corresponds to the (x, z) location of the outliers. The spanwise extents of the spatial bins containing the outliers are shown by the thin shaded rectangles along the y/h_c axis. The spanwise location of the source is $y_s/h_c = -3.95$.

closer to the source than $z/h_c = 7.0$, and thus there are more particles in this region of the simulation domain.

Recall from Chapter 3 that rogue velocities arise due to dynamical and numerical instabilities within the Langevin equations. While a rogue velocity is more probable in regions where there are steep gradients in the Reynolds stresses (such as near the outlier in Figure 6.9), they are not predictable. Rogue velocities affect the first-order consistency by altering the calculation of the conditional mean concentration (equations repeated below for convenience),


$$\langle \phi | \mathbf{u} \rangle = \langle \phi | \mathbf{u} \rangle(x_I, y_J, z_K, u_L, v_M, w_N) = \frac{Q t_r^v}{\mathcal{V} N_\phi^v}, \quad (3.66)$$

where we recall that Q is the source strength, $t_r^v = t_r^v(x_I, y_J, z_K, u_L, v_M, w_N)$ is the conditional residence time, $\mathcal{V} = \mathcal{V}(x_I, y_J, z_K)$ is the volume of the spatial bin, and


$$N_\phi^v = N_\phi^v(x_I, y_J, z_K, u_L, v_M, w_N) = N_\phi f_{\mathbf{u}} \Delta u \Delta v \Delta w, \quad (3.68)$$

is the conditional residence time normalisation constant. The PDF of the driving velocity statistics is $f_{\mathbf{u}}$, and Δu , Δv , and Δw are the streamwise, spanwise, and vertical velocity bin widths, respectively.

When a particle attains a rogue velocity, the three components of its velocity are re-initialised based upon the local velocity statistics. This results in a jump from one velocity bin to another (from the *rogue velocity bin* to the *receiving velocity bin*), resulting in a particle in a velocity bin that it did not travel to in a continuous manner, according the driving velocity PDF. Since the driving velocity statistics are assumed to be Gaussian, there is approximately a 63% chance that the re-initialised velocity is within plus or minus one standard deviation from the mean, which causes velocity bins encompassing



(a) Results from R300




(b) Results from R301

Figure 6.14: The spanwise transects of the mean concentration from the R300 and R301 MEANS and SPMMM simulations. The transects are from row index 3.5 and height $z/h_c = 4.0$. The spanwise extents of the spatial bins containing the outliers are shown by the thin shaded rectangles along the y/h_c axis. The spanwise location of the source is $y_s/h_c = -3.95$.


velocities one standard deviation from the mean to receive the majority of the re-initialised particles. The jump from one velocity bin to another leads to an over-accumulation of conditional residence time t_r^v in the receiving velocity bin, as this residence time that would not have been accumulated if there were no rogue velocities. Complementary to this over-accumulation of residence time in the receiving velocity bin, there is an under-accumulation of residence time in the rogue velocity bin, as the re-initialised particle no longer resides in that bin.

The error in $\langle \phi | \mathbf{u} \rangle$ from the altered values of t_r^v is amplified by the fact that the velocity PDF does not account for the rogue or re-initialised particles. Consequently, the normalisation constants calculated by equation (3.68) are too high for the rogue velocity bins and too low for the re-initialised velocity bins. Note that the unconditional concentration is unaffected by rogue velocities since the unconditional residence times are accumulated for spatial bins only, and are normalised by the number of particles in the simulation (N_ϕ , a constant).

The effects of rogue velocity re-initialisation are very clearly shown in Figure 6.15, which displays the vertical profile of the conditional concentration from the R300 MEANS simulation (top panel), and the R301 MEANS simulation (bottom panel) for velocity bin 4210, which received many re-initialised rogue particles. The spanwise location of these profiles is in the centre of a spatial bin containing an outlier at $y/h_c = 0.0$, where many rogue velocities were generated. The conditioning velocities for velocity bin 4210 were the mean velocities, $\langle \phi | u, v, w \rangle = \langle \phi | u = \langle u \rangle, v = \langle v \rangle, w = \langle w \rangle \rangle$ (numerically, conditioning was accomplished by selecting the velocity bin that contained the mean



(a) Results from R300



(b) Results from R301

Figure 6.15: Comparison of the vertical profiles of $\langle \phi | u, v, w \rangle$ from the R300 and R301 MEANS simulations for a velocity bin that receives re-initialised rogue particles. The concentration was conditioned on the mean velocities, $\langle \phi | u, v, w \rangle = \langle \phi | u = \langle u \rangle, v = \langle v \rangle, w = \langle w \rangle \rangle$. Notice how the location of the large spike centred on $z/h_c \approx 8.0$ corresponds to the location of many rogue velocities in Figure 6.12.

velocities). The R300 profile shows a smooth progression from its maximum value near $z/h_c \approx 1$ to a zero value near $z/h_c \approx 10$. Below $z/h_c \approx 7.0$, the magnitude of the R301 profile is comparable to the R300 profile, but has a serrated appearance. From $z/h_c \gtrsim 7.0$ to $z/h_c \approx 10.0$ there is very large spike in the profile with a magnitude approximately equal to the absolute maximum at $z/h_c = 1.0$. This spike corresponds spatially to a region of rogue velocity generation (compare with Figure 6.12) and is a result of the re-initialisation of rogue velocities through the mechanism described above: excess residence time is being accumulated in the receiving velocity bin which results in t_r^v being too large for this velocity bin; and the velocity PDF does not account for the extra particles (as rogue velocities are caused by dynamical and numerical instabilities and are therefore not predictable), resulting in N_ϕ^v being too low. Note that this large spike is the result of only a few hundred re-initialised rogue velocities, out of some 1.70×10^{10} particle steps.

The residence time for each particle was not tracked, but the total residence time for the entire R301 MEANS simulation is known, as is the number of particles used. The average residence time that each particle contributed to the R301 MEANS simulation was ~ 3.65 (in code units). For comparison, the residence time accumulated at the location of the large spike centred at $z/h_c \approx 8.0$ in Figure 6.15(b) was only 33.9 units. While at this height the average residence time per particle is likely to be less than 3.65 due to the large streamwise velocity, the relative size of the 3.65 to 33.9 suggests that rogue re-initialisation could strongly affect the accumulation of conditional residence time.

Attempts to mitigate the effects of the rogue velocities and re-initialisation

were unsuccessful. Altering the number of particles used in the simulations resulted in small differences in the **MEANS** and **SPMMM** transects of mean concentration shown in Figure 6.13, but not beyond the level explainable by statistical noise. Similarly, altering the spatial, velocity or temporal resolution of the models resulted in no appreciable changes to the transects. Simply terminating the trajectory of any particle that attained a rogue velocity was not a solution either. Since the rogue velocities frequently occurred around the same spatial location, the resulting mean concentration field had “holes” in it where the rogue particles were terminated, resulting in poor accuracy when compared to the experimental measurements. A more thorough investigation of the causes of rogue velocities (e.g., Yee and Wilson (2007)) is beyond the scope of this thesis, but is an area of interest for future projects.

The occurrence of rogue velocities and their re-initialisation had little effect on the **SPMMM** simulations if the conditional mean concentration field that a particle samples while it propagates downstream does not have any spikes caused by rogue velocities and re-initialisation. The R300 **MEANS** field in Figure 6.15(a) is an example of such a field, the R301 **MEANS** field in Figure 6.15(b) is not. The underlying conditional mean concentration field is thus responsible for the R300 **MEANS** and **SPMMM** first-order consistency, and the R301 **MEANS** and **SPMMM** first-order inconsistency, even though both **SPMMM** simulations had a similar number of rogue velocities (16649 for R300, 19612 for R301).

6.2.2 Above Canopy Concentration Statistics

Unfortunately, due to the generation of many rogue velocities within the canopy, a simulation of dispersion in the **MUST** canopy flow that extended to

ground-level could not be performed. In this section, the results for an above canopy dispersion simulation are displayed to demonstrate that SPMMM does have the ability to predict concentration fluctuations in an inhomogeneous flow, provided that there are few rogue velocities. The model parameters were assigned the same values as in the R300 and R301 simulations; the flow field differed slightly in that only the mean vertical velocity was zeroed. Without $\langle w \rangle = 0$ it was not possible to match closely to the experimental data the magnitude of the simulated mean concentration at $z/h_c = 2.0$, two rows downstream from the source. These simulations are meant to demonstrate that if SPMMM is driven by velocity statistics that do not produce rogue velocities repeatedly in the same spatial location, then it can, with reasonable accuracy, predict the magnitude of various concentration statistics. No performance measures were calculated for these simulations as they are intended to be illustrative. Lastly, since the spatial resolution of the simulations was relatively low, no smoothing of the data was performed.

Figures 6.16 – 6.19 show the spanwise transects of the first four orders of the dimensionless concentration field at $z/h_c = 2.0$ for six downstream positions: 1, 2, 3, 5, 8, and 11 rows downstream of the source (row indices 2.5, 3.5, 4.5, 6.5, 9.5, and 12.5). The SPMMM data is shown as a line, and the Coanda water-channel data as open circles. An estimation of the characteristic timescale τ_a/T_L (evaluated at $z_s/h_c = 1.8$) is displayed in each panel of the figure. Note that the scales on the ordinate axes may vary between panels. To reduce clutter on the plots, every tenth Coanda experimental data point is plotted in the spanwise transects.

Figure 6.16 shows the spanwise transects of the dimensionless mean con-

centration for a height of $z/h_c = 2.0$ at six downstream positions for the above canopy SPMMM simulations. The most notable feature in all six panels is the spanwise position of the peak mean concentration in the SPMMM simulated transects is offset to the left of the experimental transects. Furthermore, even if the spanwise positions of the maxima were aligned, the lateral spread of the simulated plume is narrower than the water-channel MUST canopy plume. Both of these observations are probably due to the fact that the in-canopy flow has been excluded from the simulation, thus depriving the plume of high turbulence intensity region and the rapid initial plume spread as the plume fills the wake region between the obstacles. However, the obstacle resolving, full flow depth simulations of Wilson et al. (2009) also produced lateral spreads that were narrower than the experimental results. The reduced mean lateral dispersion of the SPMMM plume is seen in all four orders of the concentration statistics. At the farthest two downstream measurement locations, the magnitude of the simulated transects of the mean dimensionless concentration matches the experimental transects with good accuracy, providing ideal locations to compare the higher-order concentration statistics.

Figure 6.17 shows the spanwise transects of the standard deviation of the dimensionless concentration for a height of $z/h_c = 2.0$ at six downstream positions for the above canopy SPMMM simulations. All panels show that SPMMM under-predicts the magnitude of the standard deviation of the dimensionless concentration. It is uncertain how much of the error is due to the altered flow field and how much is intrinsic to the model. However, qualitatively the SPMMM transects resemble the water-channel transects. This is particularly noticeable in the divot in the transect near the transect maximum

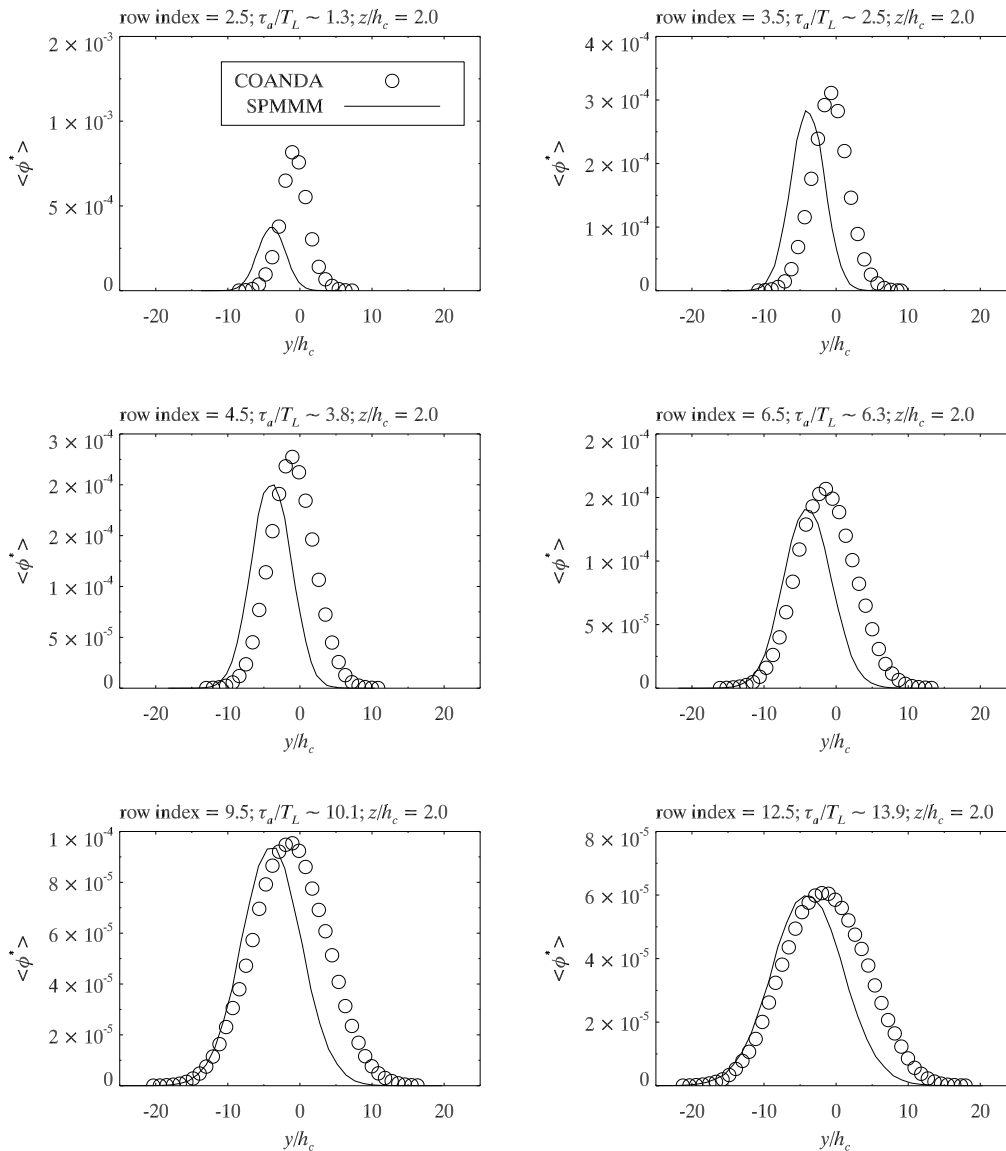


Figure 6.16: Spanwise transects of the dimensionless mean concentration for a height of $z/h_c = 2.0$, at extraction positions 1, 2, 3, 5, 8, and 11 rows downstream from the source (row indices 2.5, 3.5, 4.5, 6.5, 9.5, and 12.5 in Figure 6.1) from SPMMM simulations of an above canopy point source in the MUST canopy flow. Particles were prohibited from entering the canopy in these simulations.

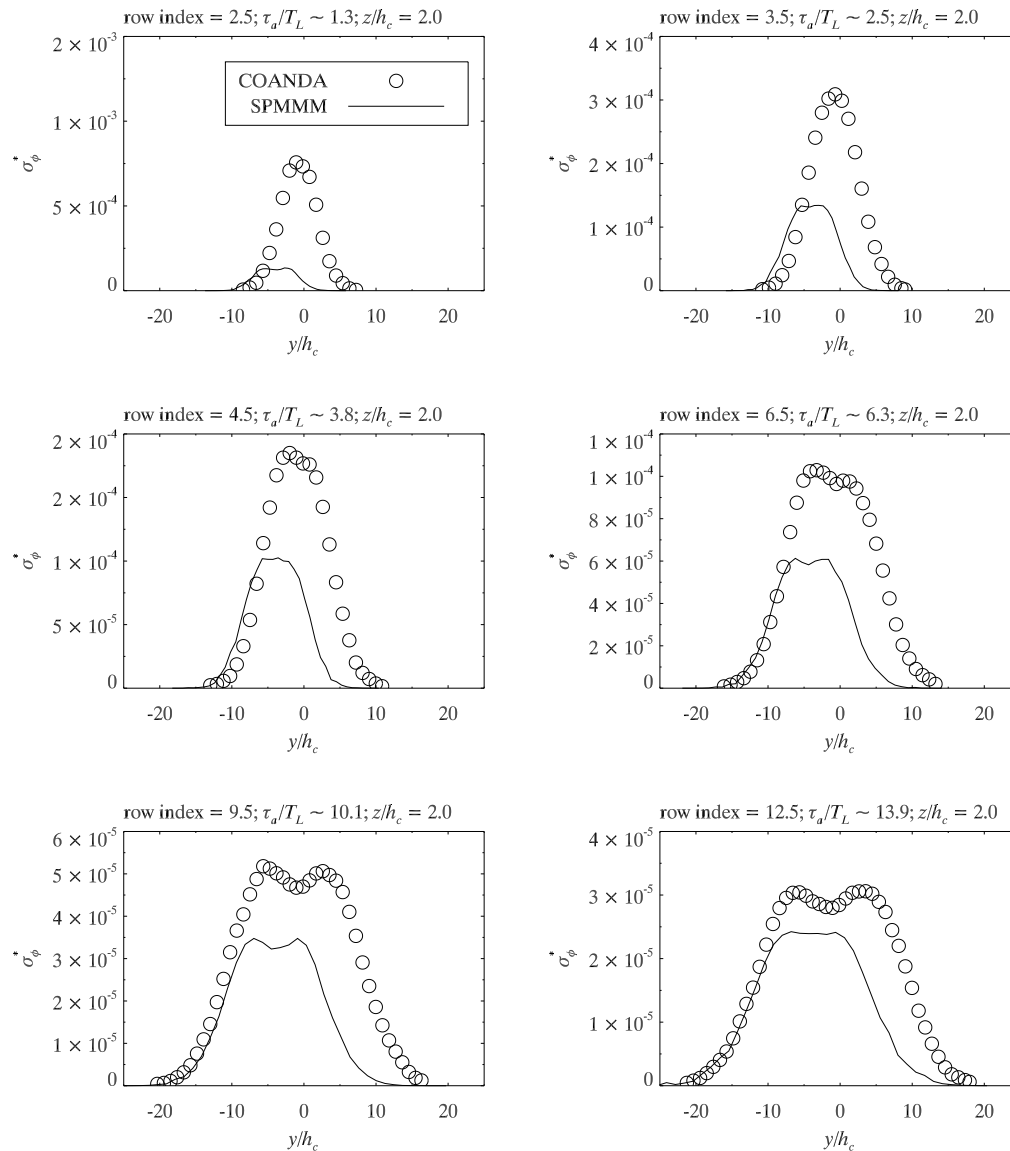


Figure 6.17: Spanwise transects of the standard deviation of the dimensionless concentration for a height of $z/h_c = 2.0$, at extraction positions 1, 2, 3, 5, 8, and 11 rows downstream from the source (row indices 2.5, 3.5, 4.5, 6.5, 9.5, and 12.5 in Figure 6.1) from SPMMM simulations of an above canopy point source in the MUST canopy flow. Particles were prohibited from entering the canopy in these simulations.

at row index 9.5, or eight rows downstream of the source (lower left panel). At row index 12.5, the magnitude of the simulated transect is approaching the magnitude of the experimental transect.

The quantitative accuracy of the SPMMM above canopy simulations improves for the skewness of the dimensionless concentration, as shown by the spanwise transects in Figure 6.18. As for the Tombstone Canopy simulations, edge effects are very clearly seen in the simulation results. Factoring out the shift to the left of the SPMMM transects relative to the experimental ones, and the reduced lateral spread (both of which are probably due to the exclusion of the canopy flow), the qualitative accuracy of the simulation results is also good.

Figure 6.19 shows the spanwise transects of the kurtosis of the dimensionless concentration for a height of $z/h_c = 2.0$ at six downstream positions for the above canopy SPMMM simulations. As for the skewness of the dimensionless concentration, the qualitative and quantitative accuracy of the model predictions is good and edge effects are seen.

6.3 Chapter Summary

A severe limitation of SPMMM has been investigated in this chapter—the generation of rogue velocities by a discrete implementation of Thomson’s LS algorithm and their subsequent re-initialisation breaks the theoretically required first-order consistency between MEANS and SPMMM. This renders SPMMM effectively useless in regions of frequent rogue velocity generation. For the full depth MUST canopy flow, rogue velocities were generated frequently at most locations near the top of the canopy. It has been demonstrated in Section 6.2.1 that when the rogue velocities are distributed throughout the simulation domain

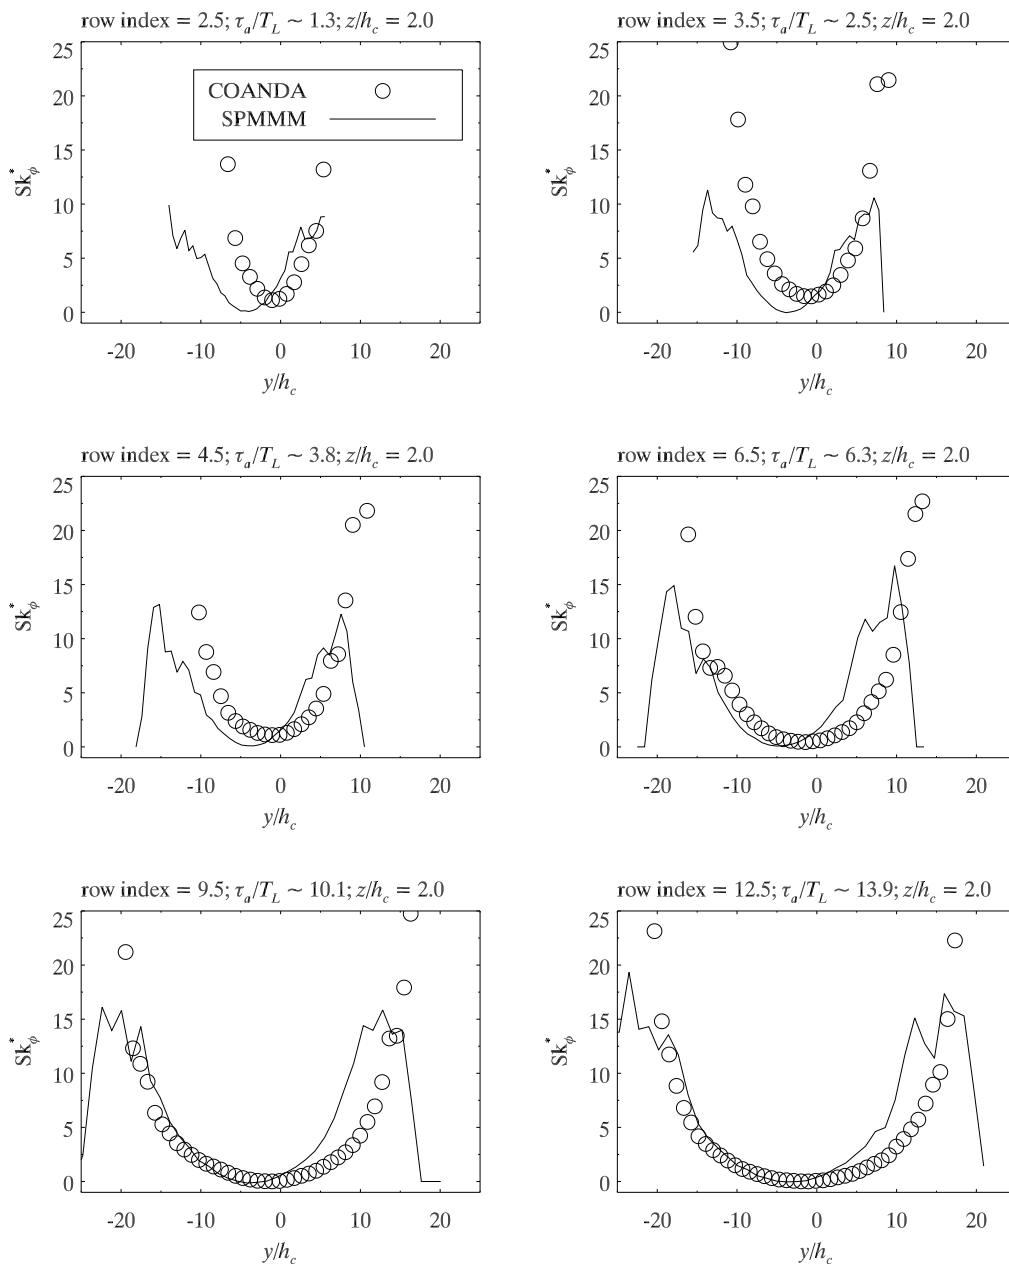


Figure 6.18: Spanwise transects of the skewness of the dimensionless concentration for a height of $z/h_c = 2.0$, at extraction positions 1, 2, 3, 5, 8, and 11 rows downstream from the source (row indices 2.5, 3.5, 4.5, 6.5, 9.5, and 12.5 in Figure 6.1) from SPMM simulations of an above canopy point source in the MUST canopy flow. Particles were prohibited from entering the canopy in these simulations.

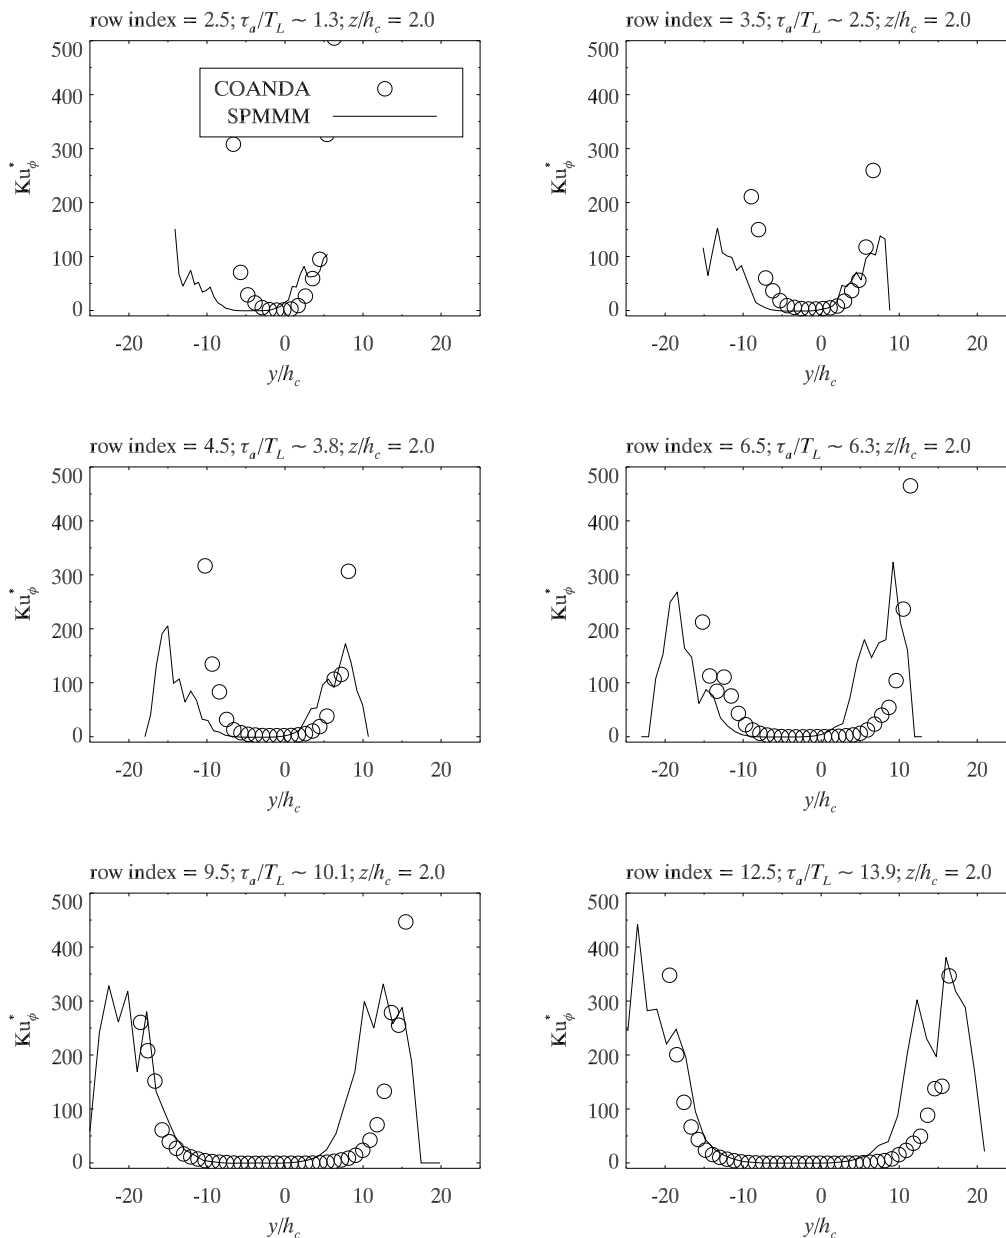


Figure 6.19: Spanwise transects of the kurtosis of the dimensionless concentration for a height of $z/h_c = 2.0$, at extraction positions 1, 2, 3, 5, 8, and 11 rows downstream from the source (row indices 2.5, 3.5, 4.5, 6.5, 9.5, and 12.5 in Figure 6.1) from SPMMM simulations of an above canopy point source in the MUST canopy flow. Particles were prohibited from entering the canopy in these simulations.

the first-order consistency between **MEANS** and **SPMMM** remains good. However, when the rogue velocities are repeatedly generated in a given region, the consistency in that region was found to be poor. Moreover, a *very* small number of rogue velocities in the same region suffice to break the first-order consistency. For example, the R301 **MEANS** simulation showed poor consistency (a factor of ~ 2.3 maximum difference between the **SPMMM** and **MEANS** mean concentrations; Figure 6.13) in a region where the vertical gradient of σ_w^2 was steep and changed sign, even though only 1183 out of 1.70×10^{10} particle steps attained a rogue velocity.

We believe that rogue velocities break the first-order consistency between **MEANS** and **SPMMM** in two ways: (1) when the rogue velocity is re-initialised it jumps from the rogue velocity bin to the receiving velocity bin, which leads to an under-accumulation of conditional residence time (t_r^v) in the rogue velocity bin, and an over-accumulation in the receiving velocity bin; (2) the conditional residence time normalisation constants (N_ϕ^v : equation (3.68)) are inaccurate for the rogue velocity bin and the receiving velocity bin. Both mechanisms are the result of the rogue velocities not being accounted for by the driving velocity PDF, which results in an inaccurate conditional mean concentration field in the regions of frequent rogue velocity generation and poor consistency between **MEANS** and **SPMMM**. Re-initialising the velocity of a rogue particle based on the local velocity statistics seems reasonable, however, as was shown in this chapter, the re-initialisation strongly alters the conditional mean concentration field, rendering the single-particle LS-IECM approach ineffective for predicting the higher order moments of the concentration field. It has been demonstrated by performing above canopy simulations (where few rogue velocities occurred)

that SPMM can qualitatively reproduce features of the first four moments of the concentration field for dispersion in the MUST canopy flow, suggesting that if rogue velocities can be eliminated then a viable approach to modelling concentration fluctuations can be found within the single-particle LS-IECM framework.

Chapter 7

Summary and Conclusions

7.1 Summary

The dispersion of a non-reactive, passive tracer from a variety of source configurations in three different flows with neutral stability has been simulated using a single-particle trajectory LS-IECM micromixing model called **SPMMM**, which can (in theory) predict all moments of the concentration field due to a source of rather general specification (point, line, area). In practice however, the availability of computational resources and time limit the predictive powers of **SPMMM**. This is an ambitious type of model for it can accommodate whatever estimate of the environmental winds one is able to supply—in fully three dimensions—and is attuned to the turbulence statistics as embodied in the velocity variances and the TKE dissipation rate. The model monitors and updates the position, velocity, and concentration of a particle. The conditional mean concentration field used by **SPMMM** was pre-calculated by a program called **MEANS**. There are (at a minimum) three free parameters in the model: the Kolmogorov constant C_0 , the initial source width σ_0 , and the micromixing timescale t_m . The parametrisation of the micromixing timescale

used in this thesis was proposed by Cassiani, Franzese and Giostra (2005a) and contains two dimensionless free parameters: the micromixing constant μ (a fitting free parameter) and the Richardson constant C_r . In total, there are four free parameters in **SPMMM**. Of these, only σ_0 and μ are truly free, as C_0 has been constrained to an acceptable range of 2–7, and it has been suggested that $C_r \approx C_0/11$ (Franzese and Cassiani, 2007).

A similar model using simultaneously computed particle trajectories and one-dimensional velocity statistics has recently been applied to dispersion within the neutral boundary layer (Cassiani, Franzese and Giostra, 2005a), within the convective boundary layer (Cassiani, Franzese and Giostra, 2005b) and within a canopy layer (Cassiani, Radicchi and Giostra, 2005; Cassiani et al., 2007). For convenience, we hereafter refer to the Cassiani et al. model as **SLS-IECM** (for Simultaneous LS-IECM), which is a name assigned by us, not the authors of the model. While the use of simultaneous trajectories allows **SLS-IECM** to incorporate chemical reaction, it is much more difficult to parallelise since the particles are interactive in that at each time step the conditional mean concentration must be calculated based on the particles that occupy a sub-domain of the simulation domain. Moreover, parallelising **SLS-IECM** will have a large computational communication overhead as neighbouring particles may be on different processors. At each time step, the processors will have to pause, share particle data, then continue, reducing the parallel performance of **SLS-IECM**. In contrast, the single-particle trajectory framework of **SPMMM** allows for trivial parallelisation and a direct increase in performance; the time required to run **SPMMM** on N_p computer processors is $\sim 1/N_p$ the time required to run **SPMMM** on one computer processor. The only time the processors must

communicate is when sharing the plume extent data, at the beginning of the simulation. While the single-particle framework of **SPMMM** precludes the possibility of simulating chemically reactive species, given that there are many non-reactive species that are interesting to study (including all the ones simulated by SLS-IECM to date), the computational performance benefits supported the use of single-particle trajectories in **SPMMM**.

Comparison of **SPMMM** and SLS-IECM has been performed for two different dispersion experiments: (1) point source dispersion in the (Fackrell and Robins, 1982) neutral wall shear layer flow; (2) line source dispersion in the Tombstone canopy flow (Legg et al., 1986). The results of the comparisons are discussed below.

SPMMM was evaluated by simulating dispersion from ground-level and elevated ($z_s = 0.19\delta$) point sources in the Fackrell and Robins (1982) neutral wall shear flow. The free parameters were calibrated to $C_0 = 6.0$, $\sigma_0 = 0.8d_s$, $\mu = 0.75$, and $C_r = 0.45$, using the streamwise transect of the concentration fluctuations on the plume centreline for a plume resulting from continuous tracer release from an elevated point source with a diameter of 9 mm. Note that σ_0 was re-calibrated for each flow, whereas the other parameters kept the values listed above. These values provided a good overall agreement between the experimental data and the model predictions when considering the first two moments of the concentration field. The agreement of the modelled first-order statistics and the experimental ones was excellent (Figures 4.11, 4.12, 4.13). The **SPMMM** first-order concentration statistics were also found to be consistent with the **MEANS** results (Figures 4.7, 4.11, 4.12, 4.13), as theoretically required of the IECM model.

For the streamwise transects of the concentration fluctuation intensity, **SPMMM** produced profiles that qualitatively looked the same but showed some discrepancy with respect to the experimental data (e.g., Figure 4.19). For the 3 mm source, the initial rise ($x/\delta \sim 0.5$) in the fluctuation intensity was approximately twenty percent too low. For the 9 mm source the agreement was good at all locations (probably because this was the data the model parameters were calibrated to). For the 15 mm, 25 mm, and 35 mm sources the initial rise was well captured, but in the range $0.80 \lesssim x/\delta \lesssim 2.0$ the modelled profiles were ten to fifteen percent too low. The modelled vertical profiles of concentration variance (Figure 4.20) matched the experimental profiles well near to the source, but became poorer farther downstream. Above the height of the maximum concentration variance, **SPMMM** over-predicted the concentration variance, a feature which generally matched the **SLS-IECM** results. Below, **SPMMM** under-predicted the concentration variance, and performed poorer than **SLS-IECM**. As both **SPMMM** and **SLS-IECM** used the same parametrisation of the micromixing timescale, the difference in performance between the two models is believed to be related to their implementation: **SPMMM** utilised single-particle trajectories and a pre-calculated conditional mean concentration field whereas **SLS-IECM** utilised parallel particle trajectories and calculated the conditional mean concentration at each timestep.

In terms of concentration variance, **SPMMM** performed better in the simulations of dispersion from a 15 mm ground-level source, showing excellent agreement with the experimental data for both the mean concentration (Figure 4.21) and variance of the concentration (Figure 4.22). Recall that for wall shear layer flow, the **TKE** production and dissipation rates increase towards the

ground. There is also increased stretching, twisting, and folding of the material lines due to increased velocity shear near ground-level. These effects result in increased mixing and dissipation of the concentration fluctuations near the ground. Evidently, the parametrisation of the micromixing timescale used in this thesis captures reasonably well these effects.

The numerical error in the **SPMMM** simulation results behaved as expected. The statistical error in mean concentration and concentration variance went approximately as $\text{RMS}(\Sigma) \propto N^{-1/2}$ (Figure 4.23). The bias error was found to be approximately zero (Figure 4.24), and the discretisation error on the plume centreline was found to be inversely proportional to the number of bins on the z -axis, $|S| \propto N_z^{-1}$ (Figure 4.25).

Simulations of dispersion from an elevated, in-canopy, cross-wind, continuous line source ($z_s = 0.85h_c$) and from an elevated, in-canopy continuous point source ($z_s = 0.2h_c$) have been performed for the Tombstone canopy flow. The velocity statistics for the line source simulations were polynomial interpolations of the velocity measurements from the Raupach et al. (1986) wind-tunnel experiments (and the errata: Raupach et al. (1987)). For the point-source simulations, the velocity statistics were taken from the water-channel experiments of Hilderman and Chong (2007). Two sets of simulations have been conducted for the point source configuration: one using horizontally-homogeneous flow statistics; the other using locally inhomogeneous flow statistics (where some of the small scale inhomogeneities between the canopy obstacles were represented while neglecting the larger scale inhomogeneities).

Overall the performance of **SPMMM** was acceptable. For the line source simulations with the TKE dissipation rate calculated assuming local equilibrium,

the fidelity of the modelled vertical profile of mean concentration to the experimental data of Legg et al. (1986) was satisfactory to good above the canopy (Figure 5.7). Within the canopy, the mean concentration was under-predicted close to the source and over-predicted farther from the source. The in-canopy results were poorer than predictions made by previous investigators (Flesch and Wilson, 1992; Cassiani et al., 2007). However, these discrepancies relative to the results from previous investigators were largely eliminated by adopting (as had the earlier authors referred to above) the TKE dissipation rate profile inferred from the profile of the Lagrangian integral timescale (Figure 5.11). The results for the standard deviation of the concentration were qualitatively and quantitatively good, although somewhat under-predicted near the source (Figures 5.8 and 5.12). This was likely due the greatly simplified horizontally-homogeneous velocity statistics that drove the models.

At most measurement locations, the point source simulations utilising horizontally-homogeneous velocity statistics performed better than the point source simulations that utilised locally inhomogeneous velocity statistics (Figures 5.14 - 5.29). The exception was when predicting the mean concentration and standard deviation close to the source ($x/h_c \lesssim 2.0$), when the use of locally inhomogeneous velocity statistics led to more accurate predictions. In both simulations it appeared that not enough particles were escaping the canopy, resulting in over-predicted mean concentrations within the canopy and under-predicted mean concentrations above the canopy. Particles may have escaped the canopy more readily if they were prohibited from entering the interior volumes of the obstacles and instead reflected off of them. This will be investigated in future versions of **SPMMM**. For the skewness and the kurtosis of the

concentration, the **SPMMM** simulations utilising horizontally-homogeneous velocity statistics performed better than the **SPMMM** simulations utilising locally inhomogeneous velocity statistics.

The generally poorer performance of **SPMMM** utilising locally inhomogeneous velocity statistics in the Tombstone canopy flow simulations may have been caused by: (1) there being no guarantee that the process of interpolating a relatively small number of velocity measurements onto a grid to drive the model resulted in a physically valid and accurate description of the actual flow in the water-channel experiments; (2) the exclusion of large scale inhomogeneities; (3) the presence of many rogue velocities; and (4) the fact that **SPMMM** does not resolve obstacles.

The **MUST** canopy simulations revealed two limitations of **SPMMM**. To numerically integrate the driving velocity PDF such that the integrals were unity at all spatial locations, the u , v , and w velocity spaces each had to be discretised into at least 70 bins, increasing the memory requirements of the model. This technical problem is easily overcome by using a more powerful computer. However, if more computational power is not available, then the spatial resolution of simulations in highly disturbed flows will have to be reduced.

The second and more severe limitation is that recurring generation of rogue velocities and their subsequent re-initialisation at a spatial location effectively renders **SPMMM** useless by breaking the theoretically required first-order consistency between **MEANS** and **SPMMM** as seen in Figure 6.11. The steep gradients of the Reynolds stresses near the top of the **MUST** Canopy resulted in approximately 0.064% of all particle steps attaining a rogue velocity. Even though this is a very small percentage, it strongly affected the accuracy of the **SPMMM** results

(Figure 6.10).

To investigate this phenomenon further, simulations have been performed in which the particles were limited to the above canopy region of the flow by setting the reflection height to the canopy height (the R300 run; Table 6.1). This reduced the incidence of rogue velocities to just 41 steps out of 1.72×10^{10} . These rogue velocities occurred randomly throughout the simulation domain and the resulting first-order consistency between **MEANS** and **SPMMM** was very good, with the small differences attributable to statistical noise (Figure 6.13(a)). A second run (R301) was performed in a flow that was identical to the R300 except for a steep vertical gradient of σ_w^2 caused by a suspected outlier in the Coanda measured velocity statistics (Figure 6.9; recall that this outlier was excluded from the interpolation of the experimental velocity statistics for all other MUST Canopy simulations). An additional 1143 rogue velocities were generated in the R301 run, most of which were near to the outliers (Figure 6.12; recall that there were multiple outliers in the driving velocity field due to the periodicity of the velocity statistics in the spanwise direction).

The spatial locations of the rogue velocities and the largest inconsistencies between **MEANS** and **SPMMM** were coincident (Figure 6.13(b)), suggesting a relationship between the two phenomena. Furthermore, the R300 run and the R301 run both showed excellent first-order consistency between **MEANS** and **SPMMM** away from regions of rogue velocity generation (Figure 6.14). Rogue velocities arise through dynamical and numerical instabilities within the Langevin equations, and they are not represented by the velocity PDF, which is used to calculate the normalisation constant for the conditional residence time in each

velocity bin (equation (3.68)). Since rogue velocities are not accounted for by the velocity PDF, the normalisation constants may be incorrect for those velocity bins that have many occurrences of rogue velocities, and those bins which receive the re-initialised rogue velocities. The process of re-initialisation may lead to an over-accumulation of residence time in the velocity bins that receive the re-initialised rogue particles, and an under-accumulation of residence time in the velocity bins where the rogue velocities occurred, since in both cases particles are travelling through velocity bins that they would not have otherwise reached at that time step if the rogue velocity had not occurred. The result was vastly different conditional mean concentrations in the R300 and R301 simulations (Figure 6.15). More research into these phenomena is needed.

By tuning SPMMM such that the simulated spanwise transects of the dimensionless mean concentration roughly matched the Coanda experimental transects for the MUST canopy flow (Figure 6.16), a qualitative assessment of the ability of SPMMM to simulate higher-order concentration fluctuations in an inhomogeneous flow has been performed. For the standard deviation, the skewness, and the kurtosis of the dimensionless mean concentration (Figures 6.17 – 6.19), the SPMMM simulations reproduced the general shape, and for the skewness and kurtosis, the magnitude of the experimental spanwise transects reasonably well. For all four moments, the simulated transects were shifted to the left in relation to the experimental transects, and the lateral spread was too small. The exclusion of the canopy flow is probably partly the cause of both of these discrepancies. Without the canopy flow, the plume is deprived of the high turbulence intensity region and the rapid initial plume spread as the plume

fills the wake region between the obstacles. Another probable cause is allowing the particles to pass through the interior volumes of the obstacles, instead of treating the canopy obstacles as solid barriers that force the particles to reflect off of them, or go around or over them. Future versions of SPMMM will incorporate solid canopy obstacles.

7.2 Conclusions

The use of the interaction by exchange with the conditional mean micromixing model under a single-particle Lagrangian stochastic framework shows promise for simulating concentration fluctuations in plumes dispersing in canopy flows. For horizontally-homogeneous flows, reasonably accurate predictions of the higher-order moments of the concentration field can be realised at low computational cost (e.g., 2 GB of RAM, 3 GHz processor, and a few hours of simulation time). Given the non-interactive nature of the particles in this style of modelling, the simulation times of the model scale directly with the number of processors used when using parallel processors, allowing very quick predictions to be made if necessary (provided a flow field is available to drive the model). This direct performance increase is the principal advantage to the single-particle trajectory implementation of the IECM micromixing model over the simultaneous particle trajectory implementation. Both approaches produced comparable results for dispersion in horizontally-homogeneous flows.

More work is needed to allow the use of the single-particle LS-IECM model with high turbulence intensity, inhomogeneous flows (such as the MUST canopy flow). In particular, more research into the exact cause and prevention of rogue velocities is needed. Yee and Wilson (2007) investigated the cause of

rogue velocities in strongly inhomogeneous flows, and suggested a more stable method for integrating stochastic differential equations (SDEs). If Step 3 of this integration scheme is omitted, then no rogue velocities will result in the integration of the SDEs, but at a loss of the well-mixed condition for the model. Given this, would it be useful to apply this integration scheme (omitting Step 3) in conjunction with **MEANS** and **SPMMM**, which guarantees the removal of rogue velocities, albeit whilst sacrificing the well-mixed condition? However, if the well-mixed condition leads to a dynamically unstable model, should we blindly accept this model even if one can't use it for useful predictions in complex flows? While this area of research is related to the LS model, the IECM model is rendered ineffective when driven by a flow field that produces rogue velocities repetitively in the same spatial location. As shown by the above **MUST** Canopy simulations, **SPMMM** appears to be able to qualitatively predict concentration fluctuations for plumes dispersion in low turbulence intensity flows, such as the one above the **MUST** canopy.

To move towards higher fidelity science, future versions of **SPMMM** will resolve the obstacles that compose the canopy and reflect particles off of them. To reduce the number of velocity bins required to ensure that the numerical integration of the velocity PDF is unity at all spatial locations, a direct counting approach to determine the conditional residence time normalisations constants may be employed in future versions of **SPMMM**. The use of the analytical velocity PDF to calculate N_ϕ^v does reduce the memory requirements of **SPMMM** by 50% when simulating dispersion in weakly disturbed flows (e.g. the FR82 flow, the horizontally-homogeneous Tombstone canopy flow, the above canopy **MUST** flow) but resulted in very high memory requirements when simulating

dispersion in highly disturbed flows (e.g., full depth MUST flow). While the parametrisation of the micromixing timescale used in this thesis was shown to be robust since the calibrated values of $\mu = 0.75$ and $C_r = 0.45$ produced reasonable results in three different flows, it is always possible to improve upon a model. Alterations to, or an alternative parametrisation of the micromixing timescale is a natural area for future research. Also, more accurate interpolations of the experimental flow fields used to drive **MEANS** and **SPMMM** could perhaps be obtained by imposing mass consistency and non-divergence on the interpolated velocity field.

Where this thesis has expanded the boundaries of micromixing modelling is by using **SPMMM** to simulate: (1) point source dispersion in the Tombstone Canopy flow using a horizontally-homogeneous description of the flow; (2) point source dispersion in the Tombstone Canopy flow using a locally-inhomogeneous (three-dimensional) description of the flow; and (3) point source dispersion in the MUST canopy flow using a three-dimensional inhomogeneous description of the flow. More broadly, this thesis is further evidence that providing an increasingly detailed description of a disturbed flow field as the basis for computation of trajectories (and the concentration field they imply) need not necessarily lead to a refined description of dispersion. Even more basically, it is evident from the work reported here that existing well-mixed first-order Lagrangian stochastic models do not properly handle highly disturbed flow, so that further progress is very much needed in this (by now, fairly mature) subject.

Lastly, while this thesis dealt largely with the theoretical side of modelling concentration fluctuations, let us consider a potential real-world application

for SPMMM. Imagine a little town on the prairie. One kilometre to the west of this town there is a petrochemical plant that produces large quantities of ethylene, which will eventually be used to produce ethylene vinyl acetate, a component of virtually all modern footwear. The material safety data sheet for ethylene notes that it is flammable in air at concentrations of approximately 3%–36% by volume (provided that there is a source of ignition), and that it is chemically stable (i.e., low reactivity in the absence of catalysts).

Now further imagine that on a winter's night, the atmosphere is neutrally stratified and the winds are blowing steadily across the prairie from the west*, directly past the petrochemical plant towards the town. Without warning, safety systems that monitor pressure fail and pipes rupture, ejecting a constant stream of ethylene† into the atmosphere. Should an evacuation of the town be ordered? Or will the ethylene disperse below the lower limit of flammability before it reaches the town? The answer to the former question depends upon the answer to the latter question, which can be addressed by SPMMM.

Figure 7.1 shows contour plots of the ground-level ($z = 0 - 3$ m) percentage by volume of ethylene for this accident. The top panel displays contours of the mean concentration of the ethylene plume. From this panel the town appears to be safe from any explosion hazard from the ethylene plume, as it has dispersed below the lower flammability limit before reaching the town. Based on this information, it appears that an evacuation will not be necessary. However, within the plume are concentration fluctuations, some of which may elevate the concentration of the plume to within the range of flammability of

*Given the winds and the terrain, the flow is well approximated by a wall shear layer, much like the FR82 flow encountered in Chapter 4.

†Recall from Chapter 5 that ethylene is a passive tracer in air.

ethylene. The bottom panel displays contours of the mean concentration plus one standard deviation of concentration ($\langle\phi\rangle + \sigma_\phi$) for the plume. Note that when the concentration fluctuations are included, ethylene of sufficiently high concentration to combust reaches into the town, necessitating an evacuation.

This simple example simulation highlights that an increased knowledge of the underlying structure of a concentration field can assist emergency planners and officials in making more informed decisions, thereby reducing risk. Fortunately, in this example, the town's people were evacuated (being sure to extinguish all sources of ignition) while the hazardous materials response team at the petrochemical plant dealt with the leak, and no one was injured.



(a) Contours of $\langle \phi \rangle$ (b) Contours of $\langle \phi \rangle + \sigma_\phi$

Figure 7.1: Contours of the percentage by volume of ethylene for a simulated accident at a petrochemical plant. The top panel shows the mean concentration of ethylene; the bottom panel shows the mean concentration plus one standard deviation of concentration, highlighting the risk to the town.

Bibliography

- Borgas, M. S. and Sawford, B. L.: 1994, A family of stochastic models for two-particle dispersion in isotropic homogeneous stationary turbulence, *Journal of Fluid Mechanics* **279**, 69–99.
- Brown, G. L. and Roshko, A.: 1974, On density effects and large structure in turbulent mixing layers, *Journal of Fluid Mechanics* **64**, 775–816.
- Brunet, Y., Finnigan, J. J. and Raupach, M. R.: 1994, A wind tunnel study of air flow in waving wheat: Single-point velocity statistics, *Boundary-Layer Meteorology* **70**, 95–132.
- Cassiani, M. A., Franzese, P. A. and Giostra, U. A.: 2005a, A PDF micromixing model of dispersion for atmospheric flow. Part I: development of model, application to homogeneous turbulence and to a neutral boundary layer , *Atmospheric Environment* **39**, 1457–1469.
- Cassiani, M., Franzese, P. and Giostra, U.: 2005b, A PDF micromixing model of dispersion for atmospheric flow. Part II: application to convective boundary layer, *Atmospheric Environment* **39**, 1471–1479.

- Cassiani, M., Radicchi, A. and Albertson, J. D.: 2007, Modelling of concentration fluctuations in canopy turbulence, *Boundary Layer Meteorology* **122**, 655–681.
- Cassiani, M., Radicchi, A. and Giostra, U.: 2005, Probability density function modelling of concentration in and above a canopy layer, *Agricultural and Forest Meteorology* **133**, 153–165.
- Castro, I. P. and Robins, A. G.: 1977, The flow around a surface-mounted cube in uniform and turbulent streams, *Journal of Fluid Mechanics* **79**, 307–335.
- Chang, J., C. and Hanna, S. R.: 2004, Air quality performance evaluation, *Meteorology and Atmospheric Physics* **87**, 167–196.
- Coppin, P. A., Raupach, M. R. and Legg, B. J.: 1986, Experiments on Scalar Dispersion Within a Model Plant Canopy. Part 2: An Elevated Plane Source, *Boundary Layer Meteorology* **35**, 167–191.
- Davidson, P., A.: 2004, *Turbulence*, Oxford University Press, 657pp.
- Dopazo, C. and O'Brien, E. E.: 1974, Functional formulation of nonisothermal turbulent reactive flows, *Physics of Fluids* **17**, 1968–1975.
- Dopazo, C., Valiño, L. and Fueyo, N.: 1997, Statistical Description of the Turbulent Mixing of Scalar Fields, *International Journal of Modern Physics B* **11**, 2975–3014.
- Durbin, P. A.: 1980, A stochastic model of two-particle dispersion and concentration fluctuations in homogeneous turbulence, *Journal of Fluid Mechanics* **100**, 279–302.

- Fackrell, J. E.: 1980, A flame ionisation detector for measuring fluctuating concentration, *Journal of Physics E Scientific Instruments* **13**, 888–893.
- Fackrell, J. E. and Robins, A. G.: 1982, Concentration fluctuations and fluxes in plumes from point sources in a turbulent boundary layer, *Journal of Fluid Mechanics* **117**, 1–26.
- Finnigan, J. K., Shaw, R. H. and Patton, E. G.: 2009, Turbulence structure above a vegetation canopy, *Journal of Fluid Mechanics* **637**, 387–424.
- Flesch, T. K. and Wilson, J. D.: 1992, A two-dimensional trajectory-simulation model for non-Gaussian, inhomogeneous turbulence within plant canopies, *Boundary-Layer Meteorology* **61**, 349–374.
- Fox, R. O.: 1996, On velocity-conditioned scalar mixing in homogeneous turbulence, *Physics of Fluids* **8**, 2678–2691.
- Fox, R., O.: 2003, *Computational Models for Turbulent Reacting Flows*, Cambridge University Press, 418pp.
- Franzese, P.: 2003, Lagrangian Stochastic modelling of a fluctuating plume in the convective boundary layer, *Atmospheric Environment* **37**, 1691–1701.
- Franzese, P. and Cassiani, M.: 2007, A statistical theory of turbulent relative dispersion, *Journal of Fluid Mechanics* **571**, 391–417.
- Gailis, R. M. and Hill, A.: 2006, A Wind-Tunnel Simulation of Plume Dispersion Within a Large Array of Obstacles, *Boundary-Layer Meteorology* **119**, 289–338.

- Gailis, R. M., Hill, A., Yee, E. and Hilderman, T.: 2006, Extension of a fluctuating plume model of tracer dispersion to a sheared boundary layer and to a large array of obstacles, *Boundary Layer Meteorology* **In Press**.
- Gardiner, C. W.: 1983, *Handbook of Stochastic Methods for Physics Chemistry and the Natural Sciences*, Springer, 442 pp.
- Garratt, J. R.: 1992, *The atmospheric boundary layer*, Cambridge University Press, 316 pp.
- Gifford, F. A.: 1959, Statistical properties of a fluctuating plume dispersion model, *Advances in Geophysics* **6**, 117–137.
- Hanna, S. R.: 1989, Confidence limits for air quality model evaluations as estimated by bootstrap and jackknife resampling methods, *Atmospheric Environment* **23**, 1385–1398.
- Hanna, S. R., Chang, J. C. and Strimaitis, D. G.: 1993, Hazardous gas model evaluation with field observations, *Atmospheric Environment* **27**, 2265–2285.
- Hilderman, T. and Chong, R.: 2007, *A laboratory study of momentum and passive scalar transport and diffusion within and above a model urban canopy – final report, DRDC Document No. DRDC Suffield CR 2008-025. Technical Report*, Defence R&D Canada – Suffield, Ralston, Alberta, 78 pp.
- Hsieh, K. J., Lien, F. S. and Yee, E.: 2007, Numerical modeling of passive scalar dispersion in an urban canopy layer. , *Journal of Wind Engineering and Industrial Aerodynamics* **95**, 1611–1636.
- Kaimal, J. C. and Finnigan, J. J.: 1994, *Atmospheric Boundary Layer Flows, Their Structure and Measurement*, Oxford University Press, 289 pp.

- Kolmogorov, A. N.: 1941, The local structure of turbulence in incompressible viscous fluid for very large Reynolds numbers (Russian), *Dokl. Akad. Nauk. SSSR* **30**, 299–303.
- Kolmogorov, A. N.: 1962, A refinement of previous hypotheses concerning the local structure of turbulence in a viscous incompressible fluid at high Reynolds number, *Journal of Fluid Mechanics* **13**, 82–85.
- Kolmogorov, A. N.: 1991, The local structure of turbulence in incompressible viscous fluid for very large Reynolds numbers, *Royal Society of London Proceedings Series A* **434**, 9–13.
- Launder, B. E.: 1978, Heat and Mass Transfer, *Topics in Applied Physics Vol. 12, Turbulence*, Editor: P. Bradshaw, pp. 231–287.
- Legg, B. J., Raupach, M. R. and Coppin, P. A.: 1986, Experiments on Scalar Dispersion Within a Model Plant Canopy. Part 3: An Elevated Line Source, *Boundary Layer Meteorology* **35**, 277–302.
- Lien, F.-S. and Yee, E.: 2005, Numerical modelling of the turbulent flow developing within and over a 3-d building array, part iii: a distributed drag force approach, its implementation and application, *Boundary-Layer Meteorology* **114**, 287–313.
- Luhar, A. K., Hibberd, M. F. and Borgas, M. S.: 2000, A skewed meandering plume model for concentration statistics in the convective boundary layer, *Atmospheric Environment* **34**, 3599–3616.

- Luhar, A. K. and Sawford, B. L.: 2005, Micromixing modelling of concentration fluctuations in inhomogeneous turbulence in the convective boundary layer, *Boundary-Layer Meteorology* **114**, 1–30.
- Lumley, J. L. and Yaglom, A. M.: 2001, A Century of Turbulence, *Flow, Turbulence and Combustion* **66**, 241–286.
- MacDonald, R. W., Griffiths, R. F. and Hall, D. J.: 1998, An improved method for estimation of surface roughness of obstacle arrays, *Atmospheric Environment* **32**, 1857–1864.
- Monin, A. S. and Yaglom, A. M.: 1975, *Statistical Fluid Mechanics II*, MIT Press, 874 pp.
- Obukhov, A. M.: 1941a, Spectral energy distribution in a turbulent flow (Russian), *Dokl. Akad. Nauk. SSSR* **32**, 22–24.
- Obukhov, A. M.: 1941b, Spectral energy distribution in a turbulent flow (Russian), *Izv. Akad. Nauk. SSSR* **4-5**, 453–466.
- Parish, T. R., Burkhardt, D. and Rodi, A. R.: 2007, Determination of the Horizontal Pressure Gradient Force Using Global Positioning System on board an Instrumented Aircraft, *Journal of atmospheric and oceanic technology* **24**, 521–528.
- Poggi, D., Katul, G. G. and Cassiani, M.: 2008, On the anomolous behaviour of the Lagrangian structure function similarity constant inside dense canopies, *Atmospheric Environment* **42**, 4212–4231.
- Pope, S. B.: 1985, PDF methods for turbulent reactive flows, *Progress in Energy and Combustion Science* **11**, 119–192.

- Pope, S. B.: 1998, The vanishing effect of molecular diffusivity on turbulent dispersion: implications for turbulent mixing and the scalar flux, *Journal of Fluid Mechanics* **359**, 299–312.
- Pope, S. B.: 2000, *Turbulent Flows*, Cambridge University Press: Cambridge UK, 806 pp.
- Raupach, M. R., Coppin, P. A. and Legg, B. J.: 1986, Experiments on Scalar Dispersion Within a Model Plant Canopy. Part 1: The Turbulence Structure, *Boundary Layer Meteorology* **35**, 21–52.
- Raupach, M. R., Coppin, P. A. and Legg, B. J.: 1987, Erratum: Experiments on Scalar Dispersion Within a Model Plant Canopy. Part 1: The Turbulence Structure, *Boundary Layer Meteorology* **39**, 423–434.
- Reynolds, O.: 1894, On the dynamical theory of incompressible viscous flows and the determination of the criterion, *Philosophical Transaction of the Royal Society of London Series A* **186**, 123–161.
- Richardson, L. F.: 1922, *Weather Prediction by Numerical Process*, Cambridge University Press.
- Richardson, L. F.: 1926, Atmospheric Diffusion shown on a Distance-Neighbour Graph, *Proceedings of the Royal Society of London Series A* **110**, 709–737.
- Rodean, H. C.: 1996, *Stochastic Lagrangian Models of Turbulent Diffusion*, American Meteorology Society, 84 pp.
- Sawford, B. L.: 2004a, Conditional scalar mixing statistics in homogeneous isotropic turbulence, *New Journal of Physics* **6**, 1–30.

- Sawford, B. L.: 2004b, Micro-Mixing Modelling of Scalar Fluctuations for Plumes in Homogeneous Turbulence, *Flow, Turbulence, and Combustion* **72**, 133–160.
- Sreenivasan, K. R.: 1995, On the universality of the Kolmogorov constant, *Physics of Fluids* **7**, 2778–2784.
- Stull, R. B.: 1988, *An introduction to boundary layer meteorology*, Atmospheric Sciences Library, Dordrecht: Kluwer, 1988, 666 pp.
- Sykes, R. I., Lewellen, W. S. and Parker, S. F.: 1984, A turbulent-transport model for concentration fluctuations and fluxes, *Journal of Fluid Mechanics* **139**, 193–218.
- Taylor, G. I.: 1921, Diffusion by continuous movements, *Proceedings of the London Mathematical Society* **20**, 196–212.
- Thomson, D. J.: 1987, Criteria for the selection of stochastic models of particle trajectories in turbulent flows, *Journal of Fluid Mechanics* **180**, 529–556.
- Thomson, D. J.: 1990, A stochastic model for the motion of particle pairs in isotropic high-Reynolds-number turbulence, and its application to the problem of concentration variance, *Journal of Fluid Mechanics* **210**, 113–153.
- Thomson, D. J.: 1996, The second-order moment structure of dispersing plumes and puffs, *Journal of Fluid Mechanics* **320**, 305–329.
- Van Kampen, N. G.: 1981, *Stochastic Processes in Physics and Chemistry*, North-Holland Publishing Company, 419 pp.

- Wang, B.-C., Yee, E. and Lien, F.-S.: 2009, Numerical study of dispersing pollutant clouds in a built-up environment, *International Journal of Heat and Fluid Flow* **30**, 3–19.
- Wilson, J. D., Finnigan, J. J. and Raupach, M. R.: 1998, A first-order closure for disturbed plant-canopy flows, and its application to winds in a canopy on a ridge, *Quarterly Journal of the Royal Meteorological Society* **124**, 705–732.
- Wilson, J. D. and Sawford, B. L.: 1996, Review of Lagrangian stochastic models for trajectories in the turbulent atmosphere, *Boundary-Layer Meteorology* **78**, 191–210.
- Wilson, J. D. and Yee, E.: 2000, Wind transport in an idealised urban canopy, *Preprints, 3rd Symposium on the Urban Environment*, American Meteorological Society, pp. 40–41.
- Wilson, J. D., Yee, E., Ek, N. and d’Amours, R.: 2009, Lagrangian simulation of wind transport in the urban environment, *Quarterly Journal of the Royal Meteorological Society* **135**, 1586–1602.
- Wyngaard, J. C. and Coté, O. R.: 1971, The Budgets of Turbulent Kinetic Energy and Temperature Variance in the Atmospheric Surface Layer., *Journal of Atmospheric Sciences* **28**, 190–201.
- Xu, J. and Pope, S. B.: 1999, Assessment of numerical accuracy of PDF/monte carlo methods of turbulent reacting flows, *Journal of Computational Physics* **152**, 192–230.

- Yee, E. and Biltoft, C. A.: 2004, Concentration Fluctuation Measurements in a Plume Dispersing Through a Regular Array of Obstacles, *Boundary-Layer Meteorology* **111**, 363–415.
- Yee, E., Chan, R., Kosteniuk, P. R., Chandler, G. M., Biltoft, C. A. and Bowers, J. F.: 1994, Incorporation of internal fluctuations in a meandering plume model of concentration fluctuations , *Boundary-Layer Meteorology* **67**, 11–39.
- Yee, E., Gailis, R. M., Hill, A., Hilderman, T. and Kiel, D.: 2006, Comparison of wind-tunnel and water-channel simulations of plume dispersion through a large array of obstacles with a scaled field experiment, *Boundary Layer Meteorology* **In Press**.
- Yee, E., Wang, B.-C. and Lien, F.-S.: 2009, Probabilistic Model for Concentration Fluctuations in Compact-Source Plumes in an Urban Environment, *Boundary-Layer Meteorology* **130**, 169–208.
- Yee, E. and Wilson, D. J.: 2000, A comparison of the detailed structure in dispersing tracer plumes measured in grid-generated turbulence with a meandering plume incorporating internal fluctuations., *Boundary-Layer Meteorology* **94**, 253–296.
- Yee, E. and Wilson, J. D.: 2007, Instability in Lagrangian stochastic trajectory models, and a method for its cure, *Boundary-Layer Meteorology* **122**, 243–261.

Yue, W., Parlange, M. B., Meneveau, C., Zhu, W., van Hout, R. and Katz, J.:
2007, Large-eddy simulation of plant canopy flows using plant-scale representation, *Boundary-Layer Meteorology* **124**, 183–203.

Appendix A

Mathematical and Statistical Tools

In this appendix various mathematical and statistical tools are discussed. Further information can be found in Monin and Yaglom (1975); Gardiner (1983); Stull (1988); Pope (2000), for example.

A.1 Eulerian and Lagrangian Coordinates

There are two systems of coordinates that are frequently used to describe the motion of a fluid - the Lagrangian and Eulerian systems. The Eulerian coordinate system measures meteorological quantities at a fixed point in space \mathbf{x} . The vast majority of meteorological data is Eulerian with the fixed point in space being the location of the instrument. In this work Eulerian position and velocity are denoted by lower case symbols, \mathbf{x} and \mathbf{u} respectively.

The Lagrangian coordinate system follows a *fluid particle* as it moves through a fluid. A fluid particle is a point that, by definition, moves at the local fluid velocity. Lagrangian position and velocity are denoted by upper case symbols, \mathbf{X} and \mathbf{U} respectively. Lagrangian measurements are made with respect to a reference position \mathbf{X}_0 at a fixed time t_0 . This reference position

is referred to as the *Lagrangian coordinate* or the *material coordinate*. The position of the fluid particle is defined by the position at the reference time $\mathbf{X}(t_0, \mathbf{X}_0) = \mathbf{X}_0$. By definition Lagrangian measurements are made with respect to the reference position and the notation is usually simplified by leaving it out of the argument. Thus the position of the Lagrangian particle can be defined by

$$\mathbf{X}(t_0) = \mathbf{X}_0, \quad (\text{A.1})$$

and

$$\frac{\partial}{\partial t} \mathbf{X}(t) = \mathbf{U}(t) = \mathbf{u}(\mathbf{X}(t), t). \quad (\text{A.2})$$

The rate of change of the Lagrangian velocity can be found by taking the partial derivative with respect to time of equation (A.2),

$$\begin{aligned} \frac{\partial}{\partial t} \mathbf{U}(t) &= \frac{\partial}{\partial t} \mathbf{u}(\mathbf{X}(t), t) \\ &= \left(\frac{\partial}{\partial t} \mathbf{u}(\mathbf{x}, t) \right)_{\mathbf{x}=\mathbf{X}(t)} + \frac{\partial}{\partial t} X_i(t) \left(\frac{\partial}{\partial x_i} \mathbf{u}(\mathbf{x}, t) \right)_{\mathbf{x}=\mathbf{X}(t)} \\ &= \left(\frac{\partial}{\partial t} \mathbf{u}(\mathbf{x}, t) + u_i(\mathbf{x}, t) \frac{\partial}{\partial x_i} \mathbf{u}(\mathbf{x}, t) \right)_{\mathbf{x}=\mathbf{X}(t)} \\ &= \left(\frac{d}{dt} \mathbf{u}(\mathbf{x}, t) \right)_{\mathbf{x}=\mathbf{X}(t)}, \end{aligned} \quad (\text{A.3})$$

where we made use of the chain rule. The Lagrangian or *material derivative* is thus defined by

$$\frac{d}{dt} \equiv \frac{\partial}{\partial t} + u_i \frac{\partial}{\partial x_i} = \frac{\partial}{\partial t} + \mathbf{u} \cdot \nabla. \quad (\text{A.4})$$

The material derivative is so named because it tracks the variation for a particular material particle, that is, the time derivative of a property ϕ of a material particle, called a material property. Concentration is an example of a material property.

A.2 Tensor Notation

In fluid dynamics it is common to encounter both vector and tensor notation. Tensor notation is used to compactly represent equations. It is therefore a worthwhile exercise to review some of features of the notation. There are many rules to tensor notation but the following three will help the beginner.

Rule 1: There is an implied summation over the values 1, 2, 3 whenever two identical indices appear in one term. For example, consider the expression $x_i y_i$. Since there are two identical indices in it summation is implied and the expression expands to $x_1 y_1 + x_2 y_2 + x_3 y_3$.

Rule 2: Whenever there is a free index in a term then that index appears unsummed in all terms in that equation. Therefore the equation with a free index effectively represents three equations, one for each value of the free index. The equation $x_i = y_i$ has a free index in it. Thus this equation actually represents three equations: $x_1 = y_1$, $x_2 = y_2$, and $x_3 = y_3$.

Rule 3: An index cannot appear more than twice in any term. For example ϵ_{ijkk} is a valid expression while ϵ_{ijii} is not.

Terms such as x_{ij} represent nine terms and are second order tensors. A vector, x_i , is a first order tensor and a scalar is a zeroth order tensor. In this thesis **bold** terms represent vectors (i.e., $\mathbf{x} = x_i$). The *Kronecker delta*

$$\delta_{ij} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases}, \quad (\text{A.5})$$

and the *alternating unit tensor*

$$\epsilon_{ijk} = \begin{cases} 1 & \text{if } ijk = 123, 231, 312 \\ -1 & \text{if } ijk = 321, 213, 132, \\ 0 & \text{if two or more indices are alike} \end{cases} \quad (\text{A.6})$$

are useful in the representation of terms with direction dependence.

Mathematical operators have tensor representations as well. The below definitions are for Cartesian coordinates and assume that the necessary derivatives exist. The *gradient* of a scalar field $f = f(x_1, x_2, x_3)$ is

$$\nabla f = \left(\frac{\partial f}{\partial x_1}, \frac{\partial f}{\partial x_2}, \frac{\partial f}{\partial x_3} \right) = \frac{\partial f}{\partial x_i}. \quad (\text{A.7})$$

The *divergence* of a vector field $\mathbf{F} = (F_1, F_2, F_3)$ (where $F_1 = F_1(x_1, x_2, x_3)$, etc.) is

$$\nabla \cdot \mathbf{F} = \frac{\partial F_1}{\partial x_1} + \frac{\partial F_2}{\partial x_2} + \frac{\partial F_3}{\partial x_3} = \frac{\partial F_i}{\partial x_i}. \quad (\text{A.8})$$

The *scalar Laplacian* of the scalar field f is

$$\nabla^2 f = \nabla \cdot (\nabla f) = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2} + \frac{\partial^2 f}{\partial z^2} = \frac{\partial}{\partial x_i} \left(\frac{\partial f}{\partial x_i} \right) = \frac{\partial^2 f}{\partial x_i \partial x_i}. \quad (\text{A.9})$$

Similarly, the *vector Laplacian* of the vector field \mathbf{F} is

$$\nabla^2 \mathbf{F} = \nabla(\nabla \cdot \mathbf{F}) - \nabla \times (\nabla \times \mathbf{F}) = \frac{\partial^2 \mathbf{F}_i}{\partial x_j \partial x_j}. \quad (\text{A.10})$$

Example

Let us convert the constant density Navier-Stokes (NS) equation for a divergence free velocity field from vector to tensor notation and expand it. In Lagrangian vector form the NS equation is

$$\frac{d\mathbf{u}}{dt} = -\frac{1}{\rho} \nabla p + \nu \nabla^2 \mathbf{u}, \quad (\text{A.11})$$

where ρ is the fluid density, p is the modified pressure (which includes gravitational terms), and ν is the kinematic viscosity of the fluid. The Coriolis term has been suppressed for simplicity. Making use of tensor notation, the Lagrangian form of the NS equation is

$$\frac{du_i}{dt} = -\frac{1}{\rho} \frac{\partial p}{\partial x_i} + \nu \frac{\partial^2 u_i}{\partial x_j \partial x_j}. \quad (\text{A.12})$$

Since the index i is unsummed in the term to the left of the equal sign it represents three equations. It also appears in every other term of the equation in accordance with Rule 2. Upon expanding we get

$$\begin{aligned} \frac{du_1}{dt} &= -\frac{1}{\rho} \frac{\partial p}{\partial x_1} + \nu \frac{\partial^2 u_1}{\partial x_j \partial x_j}, \\ \frac{du_2}{dt} &= -\frac{1}{\rho} \frac{\partial p}{\partial x_2} + \nu \frac{\partial^2 u_2}{\partial x_j \partial x_j}, \\ \frac{du_3}{dt} &= -\frac{1}{\rho} \frac{\partial p}{\partial x_3} + \nu \frac{\partial^2 u_3}{\partial x_j \partial x_j}. \end{aligned} \quad (\text{A.13})$$

The index j in the last term on the right-hand side of these equations appears twice, implying summation in accordance with Rule 1. Applying this convention leaves

$$\begin{aligned} \frac{du_1}{dt} &= -\frac{1}{\rho} \frac{\partial p}{\partial x_1} + \nu \left(\frac{\partial^2 u_1}{\partial x_1 \partial x_1} + \frac{\partial^2 u_1}{\partial x_2 \partial x_2} + \frac{\partial^2 u_1}{\partial x_3 \partial x_3} \right), \\ \frac{du_2}{dt} &= -\frac{1}{\rho} \frac{\partial p}{\partial x_2} + \nu \left(\frac{\partial^2 u_2}{\partial x_1 \partial x_1} + \frac{\partial^2 u_2}{\partial x_2 \partial x_2} + \frac{\partial^2 u_2}{\partial x_3 \partial x_3} \right), \\ \frac{du_3}{dt} &= -\frac{1}{\rho} \frac{\partial p}{\partial x_3} + \nu \left(\frac{\partial^2 u_3}{\partial x_1 \partial x_1} + \frac{\partial^2 u_3}{\partial x_2 \partial x_2} + \frac{\partial^2 u_3}{\partial x_3 \partial x_3} \right). \end{aligned} \quad (\text{A.14})$$

Applying the conventions $\mathbf{x} = (x_1, x_2, x_3) = (x, y, z)$ and $\mathbf{u} = (u_1, u_2, u_3) = (u, v, w)$ we are left with

$$\begin{aligned}\frac{du}{dt} &= -\frac{1}{\rho} \frac{\partial p}{\partial x} + \nu \left(\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} + \frac{\partial^2 u}{\partial z^2} \right), \\ \frac{dv}{dt} &= -\frac{1}{\rho} \frac{\partial p}{\partial y} + \nu \left(\frac{\partial^2 v}{\partial x^2} + \frac{\partial^2 v}{\partial y^2} + \frac{\partial^2 v}{\partial z^2} \right), \\ \frac{dw}{dt} &= -\frac{1}{\rho} \frac{\partial p}{\partial z} + \nu \left(\frac{\partial^2 w}{\partial x^2} + \frac{\partial^2 w}{\partial y^2} + \frac{\partial^2 w}{\partial z^2} \right).\end{aligned}\tag{A.15}$$

Thus the tensor equation

$$\frac{du_i}{dt} = -\frac{1}{\rho} \frac{\partial p}{\partial x_i} + \nu \frac{\partial^2 u_i}{\partial x_j \partial x_j}.$$

compactly represents the group of equations in (A.15).

A.3 Probability Density Functions

The material in this section is adapted from Pope (2000). Let $\mathbf{v} = (v_1, v_2, v_3)$ be the sample space variable (i.e., a dummy variable) corresponding to $\mathbf{u} = (u_1, u_2, u_3)$, the Eulerian velocity. A single component of \mathbf{u} (say u_1) is a *random variable* while taken together (u_1, u_2, u_3) are *joint random variables* (or any two of them; i.e., u_1 and $u_3 \dots$). Both \mathbf{u} and \mathbf{v} are functions of space and time, $\mathbf{u}(\mathbf{x}, t)$ and $\mathbf{v}(\mathbf{x}, t)$. Suppose we wish to know the probability that the velocity is less than some value v_i^B . Let this set be denoted by $B \equiv \{u_i < v_i^B\}$. The probability of such an event occurring is thus $Prob(B) = Prob\{u_i < v_i^B\}$ and is bounded such that $0 \leq Prob(B) \leq 1$. An impossible event has a probability of 0 and a certain event has a probability of 1.

The *cumulative distribution function* (CDF) is defined as

$$\begin{aligned}F_{\mathbf{u}}(\mathbf{v}) &\equiv Prob\{u_i \leq v_i\} \\ &= Prob\{u_1 \leq v_1, u_2 \leq v_2, u_3 \leq v_3\},\end{aligned}\tag{A.16}$$

and can be used to determine the probability of an event occurring. In terms of the CDF, the probability of event B occurring is $Prob(B) = F_{\mathbf{u}}(v_i^B)$. The CDF is a non-decreasing, non-dimensional function with three properties. Since $u_i > -\infty$ is certain

$$F_{\mathbf{u}}(-\infty) = 0, \quad (\text{A.17})$$

since $u_i > \infty$ is impossible

$$F_{\mathbf{u}}(\infty) = 1, \quad (\text{A.18})$$

and for $v_i^A > v_i^B$

$$F_{\mathbf{u}}(v_i^B) \geq F_{\mathbf{u}}(v_i^A) \quad (\text{A.19})$$

Note, here the short hand $F_{\mathbf{u}}(\infty) = \lim_{v_i \rightarrow \infty} F_{\mathbf{u}}(v_i)$ is used. Also, more generally, if discrete components are allowed in $F_{\mathbf{u}}(v_i)$, another required condition would be

$$\lim_{b \rightarrow v_i^+} F_{\mathbf{u}}(b) = F_{\mathbf{u}}(v_i), \quad (\text{A.20})$$

that is, at any point v_i the limit from the right must be equal to $F_{\mathbf{u}}(v_i)$.

The *probability density function* (PDF) is defined as the derivative of the CDF

$$f_{\mathbf{u}}(\mathbf{v}) = f_{\mathbf{u}}(v_1, v_2, v_3) \equiv \frac{\partial^3}{\partial v_1 \partial v_2 \partial v_3} F_{\mathbf{u}}(v_1, v_2, v_3), \quad (\text{A.21})$$

provided that the CDF is continuous and the derivative exists. The PDF is the probability per unit length in the sample space. It has units of inverse u .

The four most important properties of the PDF are

$$f_{\mathbf{u}}(\mathbf{v}) \geq 0, \quad (\text{A.22})$$

$$\int_{\mathbf{v}} f_{\mathbf{u}}(\mathbf{v}) d\mathbf{v} = 1, \quad (\text{A.23})$$

$$f_{\mathbf{u}}(-\infty) = f_{\mathbf{u}}(\infty) = 0, \quad (\text{A.24})$$

$$Prob\{\mathbf{v}_a < \mathbf{u} \leq \mathbf{v}_b\} = F_{\mathbf{u}}(\mathbf{v}_b) - F_{\mathbf{u}}(\mathbf{v}_a) = \int_{\mathbf{v}_a}^{\mathbf{v}_b} f_{\mathbf{u}}(\mathbf{v}) d\mathbf{v}. \quad (\text{A.25})$$

Let ψ be the sample space variable corresponding to the scalar concentration ϕ and let \mathbf{u} and ϕ be joint random variables. The *joint* PDF for \mathbf{u} and ϕ is

$$f_{\mathbf{u}\phi} = f_{\mathbf{u}\phi}(\mathbf{v}, \psi) \equiv \frac{\partial^4}{\partial v_1 \partial v_2 \partial v_3 \partial \psi} F(v_1, v_2, v_3, \psi). \quad (\text{A.26})$$

denoted $f_{\mathbf{u}\phi}$ or $f_{\phi\mathbf{u}}$ (note that equation (A.21) is also a joint PDF for the joint random variables u_1, u_2, u_3). The value of the ϕ may depend on the value of \mathbf{u} in certain circumstances. In this situation the *conditional* PDF of ϕ conditioned on $\mathbf{u} = \mathbf{v}$ is

$$f_{\phi|\mathbf{u}} = f_{\phi|\mathbf{u}}(\phi|\mathbf{v}) \equiv \frac{f_{\mathbf{u}\phi}(\mathbf{v}, \psi)}{f_{\mathbf{u}}(\mathbf{v})}. \quad (\text{A.27})$$

With respect to notation, “ $|\mathbf{v}$ ” is short form for “ $|\mathbf{u} = \mathbf{v}$ ” and is read “conditional on $\mathbf{u} = \mathbf{v}$ ”, or “given $\mathbf{u} = \mathbf{v}$ ” or simply “given \mathbf{v} ”. Dependence on position and time are included in the CDF and the PDF via \mathbf{u} and ϕ . For example, we could write

$$f_{\mathbf{u}\phi}[\mathbf{v}(\mathbf{x}, t), \psi(\mathbf{x}, t); \mathbf{x}, t], \quad (\text{A.28})$$

to represent the joint PDF of velocity and the scalar but we will generally prefer the more compact notation $f_{\mathbf{u}\phi}(\mathbf{v}, \psi)$ or $f_{\mathbf{u}\phi}$. Equation (A.28) is a function with respect to items to the right of the semi-colon and a density with respect to items to the left of the semi-colon.

The joint random variables u_1 , u_2 , u_3 , and ϕ are said to be *independent* if knowledge of one provides no information in regard to the others. For independent variables, conditioning has no effect and thus the conditional PDF is the same as the *marginal* PDF (f_{u_1} for example). For independent joint random variables the joint PDF decomposes to the product of the marginal PDFs

$$f_{\mathbf{u}}(v_1, v_2, v_3) = f_{u_1}(v_1)f_{u_2}(v_2)f_{u_3}(v_3), \quad (\text{A.29})$$

and

$$f_{\mathbf{u}\phi}(v_1, v_2, v_3, \psi) = f_{u_1}(v_1)f_{u_2}(v_2)f_{u_3}(v_3)f_{\phi}(\psi). \quad (\text{A.30})$$

Using this fact and substituting the above two equations into the definition of the conditional PDF (equation (A.27)) we have

$$\begin{aligned} f_{\phi|\mathbf{u}}(\phi|\mathbf{v}) &\equiv \frac{f_{\mathbf{u}\phi}(\mathbf{v}, \psi)}{f_{\mathbf{u}}(\mathbf{v})} \\ &= \frac{f_{u_1}(v_1)f_{u_2}(v_2)f_{u_3}(v_3)f_{\phi}(\psi)}{f_{u_1}(v_1)f_{u_2}(v_2)f_{u_3}(v_3)} \\ &= f_{\phi}(\psi), \end{aligned} \quad (\text{A.31})$$

which is the marginal PDF of the scalar, demonstrating that conditioning has no effect on independent joint random variables.

A.3.1 Example: The Gaussian or Normal Distribution

The *Gaussian* or *normal* distribution is used extensively throughout this thesis. If w is a normally distributed random variable with mean $\langle w \rangle$ and variance σ_w^2 (described in Section A.4.1 and Section A.4.2 below) then the CDF of w is

$$F_w(w) = \frac{1}{2} \left[1 + \operatorname{erf} \left(\frac{w - \langle w \rangle}{\sqrt{2}\sigma_w} \right) \right], \quad (\text{A.32})$$

and the PDF is

$$f_w(w) = \frac{1}{\sqrt{2\pi}\sigma_w} \exp\left[-\frac{(w - \langle w \rangle)^2}{2\sigma_w^2}\right]. \quad (\text{A.33})$$

The *error function* is defined by

$$\text{erf}(w) \equiv \frac{2}{\sqrt{\pi}} \int_0^w \exp(-\hat{w}^2) d\hat{w}, \quad (\text{A.34})$$

where \hat{w} is a dummy variable used for integration. If w is a *standardised* Gaussian random variable then $\langle w \rangle = 0$ and $\sigma_w^2 = 1$. Figure (A.1) shows the CDF and PDF of the standardised Gaussian distribution in the top and bottom panels respectively. The extreme negative and positive ends of the PDF are commonly referred to as the *tails* of the PDF.

A.4 Statistical Tools

In this section we use the notation of bracketed suffixes to exclude a term from the summation convention of Rule 1 in Section A.2. For example

$$u_{(i)}u_{(i)} \neq u_1u_1 + u_2u_2 + u_3u_3. \quad (\text{A.35})$$

Instead, the bracketed suffixes act as a free index (like Rule 2 in Section A.2) and

$$u_{(i)}u_{(i)} = u_1u_1; \quad u_2u_2; \quad u_3u_3. \quad (\text{A.36})$$

A.4.1 The Unconditional and Conditional Means

With the PDF we can now define some statistical quantities. The definitions are made in terms of variables that will be encountered throughout the thesis but generalisation is straightforward. The (unconditional) *mean*, also




Figure A.1: The CDF and the PDF of the standardised Gaussian or normal distribution, as given by equations (A.32) and (A.33) respectively, with $\langle w \rangle = 0$ and $\sigma_w^2 = 1$. The top panel shows the CDF and the bottom panel shows the PDF. The extreme negative and positive ends of the PDF are commonly referred to as the tails of the PDF.

known as the *expectation value*, is defined as

$$\langle u_i \rangle \equiv \int_{\mathbf{v}} v_i f_{\mathbf{u}}(\mathbf{v}) d\mathbf{v} = \iiint_{-\infty}^{\infty} v_i f_{\mathbf{u}}(v_1, v_2, v_3) dv_1 dv_2 dv_3, \quad (\text{A.37})$$

which, using the definition of the *marginal PDF*, for example,

$$f_{u_3}(v_3) \equiv \iint_{-\infty}^{\infty} f_{\mathbf{u}}(v_1, v_2, v_3) dv_1 dv_2, \quad (\text{A.38})$$

simplifies to

$$\langle u_i \rangle = \int_{-\infty}^{\infty} v_{(i)} f_{u_{(i)}}(v_{(i)}) dv_{(i)}. \quad (\text{A.39})$$

Note the short hand notation

$$\int_{\mathbf{v}} (\) d\mathbf{v} \equiv \iiint_{-\infty}^{\infty} (\) dv_1 dv_2 dv_3, \quad (\text{A.40})$$

introduced to express integration over the entire velocity sample space. If $Q(\mathbf{u}) = Q(u_1, u_2, u_3)$ is a function of \mathbf{u} , the mean is

$$\langle Q(\mathbf{u}) \rangle = \langle Q(u_1, u_2, u_3) \rangle \equiv \int_{\mathbf{v}} Q(\mathbf{v}) f_{\mathbf{u}}(\mathbf{v}) d\mathbf{v}, \quad (\text{A.41})$$

which only exists if the integral converges absolutely. The mean behaves as a linear operator and therefore has some useful properties such as associativity and commutability. If a and b are scalars and $Q(u_i)$ and $R(u_i)$ are functions of a random variable u_i we have

$$\langle [aQ(u_i) + bR(u_i)] \rangle = a\langle Q(u_i) \rangle + b\langle R(u_i) \rangle. \quad (\text{A.42})$$

The mean also commutes with differentiation

$$\left\langle \frac{dQ(u_i)}{du_i} \right\rangle = \frac{d\langle Q(u_i) \rangle}{du_i}, \quad (\text{A.43})$$

and since $\langle u_i \rangle$ is not a random variable

$$\langle \langle u_i \rangle \rangle = \langle u_i \rangle. \quad (\text{A.44})$$

These properties will be used in Section (2.3) to derive the mean flow equations.

The mean of the scalar conditioned on the velocity, or the *conditional mean*, is defined as

$$\langle \phi | \mathbf{v} \rangle \equiv \int_{-\infty}^{\infty} \psi f_{\phi | \mathbf{u}}(\psi | \mathbf{v}) d\psi. \quad (\text{A.45})$$

To recover the unconditional mean concentration simply take the mean over the velocity space,

$$\begin{aligned} \langle \langle \phi | \mathbf{v} \rangle \rangle &= \iint_{-\infty}^{\infty} \psi f_{\phi | \mathbf{u}} f_{\mathbf{u}} d\psi d\mathbf{v} \\ &= \iint_{-\infty}^{\infty} \psi f_{\mathbf{u}\phi} d\psi d\mathbf{v} \\ &= \langle \phi \rangle, \end{aligned} \quad (\text{A.46})$$

where we made use of equation (A.27) to go from a conditional PDF to a joint PDF.

A.4.2 Fluctuations and Moments

With the definition of the mean, the *fluctuation* of u_i can be defined as

$$u'_i \equiv u_i - \langle u_i \rangle. \quad (\text{A.47})$$

It may be interesting to know how the data is spread about its mean value.

The *variance*, a measure of dispersion or spread about the mean, is defined as

$$\sigma_{u_i}^2 = \text{var}(u_i) = \langle u_i'^2 \rangle \equiv \int_{-\infty}^{\infty} (v_{(i)} - \langle u_{(i)} \rangle)^2 f_{u_{(i)}}(v_{(i)}) dv_{(i)}. \quad (\text{A.48})$$

Note the relationship between the fluctuations and the variance. The *standard deviation* is simply the square root of the variance

$$\sigma_{u_i} = \text{stdev}(u_i) = \sqrt{\text{var}(u_i)} = \sqrt{\sigma_{u_i}^2} = \sqrt{\langle u_i'^2 \rangle}. \quad (\text{A.49})$$

The standard deviation is also known as the root-mean-square (RMS) of fluctuation. A distribution with a standard deviation of zero would contain only elements whose values are identical to the mean. As the standard deviation of the distribution gets larger, the values of the elements spread farther and farther from the mean. For a set of normally distributed random variables (such as shown in Figure A.1) approximately 68% of the elements of the set are found within one standard deviation of the mean, approximately 95% are within two standard deviations from the mean, and over 99% are within three standard deviations from the mean.

With the mean and the standard deviation a random variable can be *standardised* by subtracting the mean and dividing by the standard deviation,

$$u_i^{\text{standard}} \equiv \frac{u_i - \langle u_i \rangle}{\sigma_{u_i}}. \quad (\text{A.50})$$

Standardised random variables have the convenient property of having a mean of zero and a variance of unity.

Equation (A.48) can be generalised to the n^{th} power (note that n is an exponent, not an index),

$$\mu_n \equiv \langle u_i^n \rangle = \int_{-\infty}^{\infty} (v_{(i)} - \langle u_{(i)} \rangle)^n f_{u_{(i)}}(v_{(i)}) dv_{(i)}, \quad (\text{A.51})$$

where μ_n is the n^{th} *moment* of u_i about the mean (also called the n^{th} central moment). The variance is thus the second-order moment about the mean. Moments too can be standardised,

$$\mu_n^{\text{standard}} = \frac{\langle u_i^n \rangle}{\sigma_u^n} = \frac{\mu_n}{\sigma_u^n}. \quad (\text{A.52})$$

The third-order standardised moment μ_3^{standard} is called the *skewness* (denoted by Sk), which is a measure of the symmetry of a random distribution

about its point of maximum probability density (the maximum of a PDF is called the *mode*). A distribution with a skewness of zero is symmetric about its mode. If the skewness is negative then more elements of the distribution have values lower than would be expected for a distribution with zero skewness. This results in a longer tail to the left of the mode, as shown in Figure A.2(a), which shows a negatively skewed PDF and one with zero skewness. Distributions with positive skewness have a longer tail to the right of the mode, indicating that more elements of the distribution have values higher than would be expected for a distribution with zero skewness.




Figure A.2: Examples of skewness in probability density functions. The zero skewness PDF in this figure is the Gaussian PDF.

The fourth-order standardised moment μ_4^{standard} is called the *kurtosis* (denoted by Ku), which is a measure of the flatness or peakedness of a random distribution. The term kurtosis is often exchanged freely with the term *excess kurtosis*, which is the amount of kurtosis relative to the Gaussian distribution ($\text{Ku}_{\text{Gauss}} = 3$), and is simply calculated as $\mu_4^{\text{standard}} - 3$. Figure A.3(a) displays three PDFs with differing amounts of kurtotic excess (the Gaussian

distribution is represented by the solid line). Figure A.3(b) displays the same three PDFs but with a logarithmic y -axis to more clearly display the behaviour of the tails. If a distribution: has zero excess kurtosis then it is said to be *mesokurtic*; has negative excess kurtosis then it is said to be *platykurtic* and has a lower, wider peak about its mean and “thinner” tails than the Gaussian distribution; has positive excess kurtosis then it is said to be *leptokurtic* and has a sharper peak about its mean and “fatter” tails than the Gaussian distribution. From the plots in Figure A.3 it is seen that as the kurtosis increases

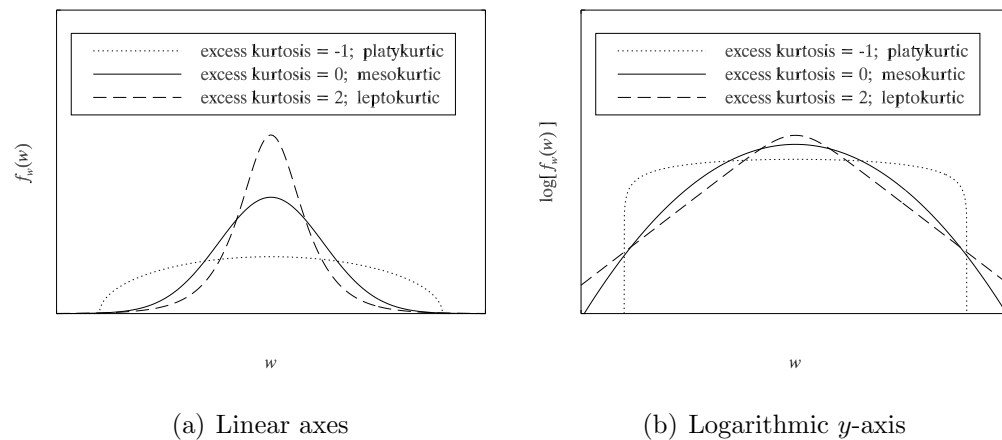


Figure A.3: Examples of excess kurtosis in probability density functions. The Gaussian distribution is represented by the solid line. The domain of each plot is the same.

the data becomes more clustered about the mean, indicating lower variability in the data set. At the same time, however, the tails get fatter, which corresponds to an increased likelihood of very rare events occurring. Therefore, as the kurtosis increases, more of the variance in a distribution is due to the occurrence of these very rare events.

Very rarely are the definitions for the mean, variance, skewness, and kurtosis shown above used to calculate these quantities. Consider the set of N

concentration measurements $\{\phi_I\}$ where $I \in [1, N]$. Estimates for the theoretical statistical quantities defined above can be easily computed from the set $\{\phi_I\}$. The *sample* statistical quantities are denoted by a circumflex (e.g., $\langle \hat{\phi} \rangle$)

The sample mean is calculated as

$$\langle \hat{\phi} \rangle = \frac{1}{N} \sum_{I=1}^N \phi_I, \quad (\text{A.53})$$

the sample variance as

$$\hat{\sigma}_\phi^2 = \frac{1}{N-1} \sum_{I=1}^N (\phi_I - \langle \hat{\phi} \rangle)^2, \quad (\text{A.54})$$

the sample skewness as

$$\hat{\text{Sk}}_\phi = \frac{1}{N} \sum_{I=1}^N \left(\frac{\phi_I - \langle \hat{\phi} \rangle}{\sigma_\phi} \right)^3, \quad (\text{A.55})$$

and the sample excess kurtosis (which is referred to simply as the kurtosis in the main body of this thesis) as

$$\hat{\text{Ku}}_\phi = \frac{1}{N} \sum_{I=1}^N \left(\frac{\phi_I - \langle \hat{\phi} \rangle}{\sigma_\phi} \right)^4 - 3. \quad (\text{A.56})$$

While the above moments focus on only a single random variable the *covariance* is a measure of how one joint random variable changes with respect to another joint random variable. Take for example the x and z components of the velocity \mathbf{u} . The covariance of u_1 and u_3 is defined as

$$\text{cov}(u_1, u_3) = \langle u'_1 u'_3 \rangle \equiv \int_{\mathbf{v}} (v_1 - \langle u_1 \rangle)(v_3 - \langle u_3 \rangle) f_{\mathbf{u}}(\mathbf{v}) d\mathbf{v}. \quad (\text{A.57})$$

The *correlation coefficient* is defined by

$$\rho_{13} \equiv \frac{\langle u'_1 u'_3 \rangle}{[\langle u'^2_1 \rangle \langle u'^2_3 \rangle]^{\frac{1}{2}}} = \frac{\langle u'_1 u'_3 \rangle}{\sigma_{u_1} \sigma_{u_3}}. \quad (\text{A.58})$$

If $\rho_{13} = 0$ then u_1 and u_3 are uncorrelated. If $\rho_{13} = 1$ then u_1 and u_3 are perfectly correlated. If $\rho_{13} = -1$ then u_1 and u_3 are perfectly negatively correlated. Equation (A.57) can be generalised as

$$C_{ij} = \text{cov}(u_i, u_j) = \langle u'_i u'_j \rangle \equiv \int_{\mathbf{v}} (v_i - \langle u_i \rangle)(v_j - \langle u_j \rangle) f_{\mathbf{u}}(\mathbf{v}) d\mathbf{v}, \quad (\text{A.59})$$

where C_{ij} is the covariance matrix. As for the variance, note the covariance is also related to the fluctuations. If $\mathbf{u}' = (u'_1, u'_2, u'_3)$ is joint-normally distributed, then its joint PDF by definition is

$$f_{\mathbf{u}}(\mathbf{u}') = \frac{[\det(C^{-1})]^{1/2}}{(2\pi)^{3/2}} \exp\left(-\frac{1}{2} u'_i C_{ij}^{-1} u'_j\right), \quad (\text{A.60})$$

where $C = C_{ij}$ is the covariance matrix and C^{-1} is its inverse. This equation represents a three-dimensional Gaussian distribution.