PDF MODELING OF TURBULENT FLOWS ON UNSTRUCTURED GRIDS

by

József Bakosi A Dissertation Submitted to the Graduate Faculty of George Mason University In Partial Fulfillment of The Requirements for the Degree of Doctor of Philosophy Computational Sciences and Informatics

Committee: 100 4/17/2008 Date:

Dr. Zafer Boybeyi, Dissertation Director

Dr. Nash'at Ahmad, Committee Member

Dr. Pasquale Franzese, Committee Member

Dr. Rainald Löhner, Committee Member

Dr. Dimitrios Papaconstantopoulos, Department Chairperson

Dr. Peter Becker, Associate Dean for Graduate Studies, College of Science

Dr. Vikas Chandhoke, Dean, College of Science

Spring Semester 2008 George Mason University Fairfax, VA PDF modeling of turbulent flows on unstructured grids

A dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy at George Mason University

By

József Bakosi Master of Science University of Miskolc, 1999

Director: Dr. Zafer Boybeyi, Professor Department of Computational and Data Sciences

> Spring Semester 2008 George Mason University Fairfax, VA

Copyright © 2008 by József Bakosi All Rights Reserved

Dedication

This work is dedicated to my family: my parents, my sister and my grandmother

Acknowledgments

I would like to thank my advisor Dr. Zafer Boybeyi, who provided me the opportunity to embark upon this journey. Without his continuous moral and scientific support at every level, this work would not have been possible. Taking on the rather risky move of giving me 100% freedom in research from the beginning certainly deserves my greatest acknowledgments. I am also thoroughly indebted to Dr. Pasquale Franzese for his guidance of this research, for the many lengthy – sometimes philosophical – discussions on turbulence and other topics, for being always available and for his painstaking drive with me through the dungeons of research, scientific publishing and aesthetics. My gratitude extends to Dr. Rainald Löhner from whom I had the opportunity to take classes in CFD and to learn how not to get lost in the details. I found his weekly seminars to be one of the best opportunities to learn critical and down-to-earth thinking. I am also indebted to Dr. Nash'at Ahmad for showing me the example that with a careful balance of school, hard work and family nothing is impossible.

I thank Dr. Thomas Dreeben for the many helpful and insightful discussions on the velocity model and elliptic relaxation. I am also grateful to the reviewers of our journal papers for their valuable comments and suggestions on our initial manuscripts.

I will always remember my professors at home at the Departments of Mathematics, Mechanics, Physics, and Fluid and Heat Engineering at the University of Miskolc, who taught me the foundations of engineering and science: Drs. Bertóti Edgár, Kálovics Ferenc, Kozák Imre, Mészáros József, Páczelt István, Rontó Miklós, Szabó Szilárd, Szeidl György, Takács Csaba and Vince Endre. I further extend my thanks to Prof. Tóth László, who first gave me the opportunity to study abroad and introduced me to research. My admiration towards all of them turned into a passion for science.

I am also thankful to Dr. Sergei Ivanov, my mentor at Fraunhofer in Boston, who encouraged me to do a Ph.D.

I would also like to thank for the continous support of all my friends throughout the years. My friend from college, Fehér Zoltán, for keeping me sane with our long, challenging and fruitful discussions every weekend. Sunil Kumar Appanaboyina for the countless philosophical discussions on culture, life, religion, food and women. Jarek Pietrzykowski who has always been a partner in crime. Fernando Mut for the many discussions on coding, performance and for suggesting the algorithm for the particle search. My housemates Houmam Ali, Andrew Krapf and Chris Lafontaine who made me feel at home right back in college. The wonderful people here at the Laboratory for Atmospheric Hazard Modeling

who have always been helpful and created an atmosphere so that I never felt alone: Dr. Guido Cervone, Laura Clemente, Yasemin Ezber, John Lindeman, Jacek Radzikowski and Priyanka Roy. Special thanks to my international student advisor, shuttle and jazz buddy, Amy Moffitt.

But my greatest thank you goes to my family and friends at home, who never ceased to believe in me, even at times when I myself almost did.

Table of Contents

		Pa	ge
Lis	t of J	$\Gamma_{ m ables}$	iii
Lis	t of F	Figures	ix
Lis	t of S	Symbols	xi
Ab	stract	t	xx
1	Intro	$\operatorname{duction}$	1
2	Gove	erning equations	9
3	Num	erical implementation	23
	3.1	Introduction	23
	3.2	Timestepping procedure	24
	3.3	Solution of the Eulerian equations	24
	3.4	Estimation of Eulerian statistics	27
	3.5	Derivatives of Eulerian statistics	30
	3.6	Estimation of the velocity-conditioned scalar mean	31
	3.7	Particle tracking	33
	3.8	Particle-number control	36
	3.9	No-slip wall-boundary conditions	41
	3.10	Parallel random number generation	44
	3.11	Solution procedure and execution profile	45
4	Chan	nnel flow simulations: results and discussion	51
	4.1	Introduction	51
	4.2	Modeling specifics of channel flow	52
		4.2.1 Modeling the fluid dynamics	53
		4.2.2 Modeling the passive scalar	55
	4.3	Results	58
		4.3.1 Fluid dynamics	58
		4.3.2 Comparison of the IEM and IECM micromixing models	59
		4.3.3 Scalar statistics with the IECM model	65
		4.3.4 Conditional statistics	70
	4.4	The effect of numerical parameters on the results	78

	4.5	Computational cost	. 92
	4.6	Discussion	. 96
5	Stree	t canyon simulations: results and discussion	. 99
	5.1	Introduction	. 99
	5.2	Modeling specifics of the street canyon	. 103
	5.3	Results	. 105
	5.4	Discussion	. 113
6	Cylin	der flow simulations: results and discussion	. 117
	6.1	Introduction	. 117
		6.1.1 A short review of cylinder flow regimes	. 119
		6.1.2 Past experimental and numerical studies	. 121
	6.2	Results	. 124
		6.2.1 Transient and cylinder surface statistics	. 126
		6.2.2 Near wake statistics	. 131
	6.3	Discussion	. 142
7	Sum	nary and discussion	. 146
A	Derivation of the Eulerian PDF transport equation		
В	Computation of the velocity-conditioned scalar mean		
С	Basic particle redistribution algorithm		. 160
D	A more efficient particle redistribution algorithm		
Bil	Bibliography		

List of Tables

Table		Page
3.1	Profile and relative execution times for a timestep	46
3.2	Two designs of a loop to update particles	47
3.3	Comparison of serial and parallel loop performances	49
4.1	Constants to model the joint PDF of velocity and frequency	54
4.2	Model constants for the micromixing timescale	62
4.3	Number of particles required for convergent statistics	91
5.1	Model constants for the full joint PDF	105
5.2	Concentration sampling locations along the street canyon walls	106
6.1	Cylinder surface and recirculation bubble region statistics	129
A.1	Sample space of the Eulerian joint PDF	154
D.1	Particle redistribution benchmarks	165

List of Figures

Figure		Page
3.1	Estimation of Eulerian statistics	29
3.2	Particle tracking with triangular elements; FEM mapping	35
3.3	Particle tracking at concave boundaries	36
3.4	Domain for a testcase to investigate the error of particle redistribution	38
3.5	The number of particles moved by particle redistribution	40
3.6	Results of a simplified testcase with particle redistribution	41
3.7	Evolution of total relative numerical error with particle redistribution \ldots	42
3.8	Comparison of cache misses for different loops	48
3.9	Overall speedups on different types of shared memory machines	50
4.1	Eulerian domain for fully developed turbulent channel flow	56
4.2	Cross-stream profiles of velocity statistics in turbulent channel flow	59
4.3	Comparison of the IEM and IECM models, scalar mean and PDFs $\ . \ . \ .$	63
4.4	Cross-stream profiles of the first four moments of scalar concentration \ldots	66
4.5	Downstream evolutions of scalar mean peak, mean width and r.m.s. width .	68
4.6	PDFs of scalar concentration fluctuations	69
4.7	Velocity fluctuation conditioned on the scalar concentration $\ldots \ldots \ldots$	73
4.8	Mean scalar dissipation conditioned on the concentration $\ . \ . \ . \ .$.	76
4.9	Mean scalar diffusion conditioned on the concentration	77
4.10	<i>Re</i> -dependence of cross-stream velocity statistics in channel flow	79
4.11	The effect of the anisotropic length-scale constant on the Reynolds stresses	81
4.12	The effect of the number of conditioning bins using Fox's projection \ldots .	83
4.13	The effect of the number of conditioning bins with the proposed method $\ . \ .$	85
4.14	The effect of the number of particles using Fox's projection method	87
4.15	The effect of the number of particles with the proposed method	89
4.16	Comparison of computational cost vs. Re	93
5.1	Eulerian grid to compute a street canyon with full resolution $\ldots \ldots \ldots$	106
5.2	Eulerian grid to compute a street canyon with wall-functions $\ . \ . \ . \ .$	107
5.3	Velocty statistics in a street canyon	108

5.4	Turbulent intensities in a street canyon	110
5.5	Mean concentrations along the wall of a street canyon $\ldots \ldots \ldots \ldots \ldots$	111
5.6	Scalar concentration statistics in a street canyon	112
5.7	PDFs of concentration fluctuations in a street canyon	114
6.1	Eulerian mesh for flow around cylinder	125
6.2	Time-evolution of the pressure lift coefficient	127
6.3	Time-evolution of the total and viscous lift coefficients $\ldots \ldots \ldots \ldots$	128
6.4	Time-averaged mean pressure and vorticiy along the cylinder	131
6.5	Streamwise mean velocity behind the cylinder along the centerline	132
6.6	Mean streamwise and cross-stream velocity behind the cylinder \ldots	135
6.7	Cross-stream mean velocity behind the cylinder at $x/D=3.0$	136
6.8	Streamwise and cross-stream Reynolds stress behind the cylinder \ldots .	137
6.9	Turbulent kinetic energy in the recirculation bubble in the cylinder wake	138
6.10	Shear Reynolds stress behind the cylinder	140
6.11	Skewness of the streamwise and cross-stream velocity behind the cylinder $\ .$	142
6.12	Flatness of the streamwise and cross-stream velocity behind the cylinder	143

List of Symbols

Abbreviations

CFL	Courant-Friedrichs-Lewy (Section 3.2)
CNBI	number of conditioning bins, N_c
ср	conditioning pointer of a particle
CFD	computational fluid dynamics
DNS	direct numerical simulation
elp	element indices of particles, "element of particle"
FDF	filtered density function
FEM	finite element method
FV	finite volume
GLM	generalized Langevin model
HP1	Haworth & Pope model 1
HP2	Haworth & Pope model 2
IECM	interaction by exchange with the conditional mean
IEM	interaction by exchange with the mean
IP	isotropization of production
IPMa	isotropization of production model a
IPMb	isotropization of production model b
LES	large eddy simulation
LIPM	Lagrangian isotropization of production model
maxe	the element index with the most number of particles
MAXNPEL	maximum number of particles per element, $N_{p/e}^{\max}$
maxnpel	number of particles in the element with the most number of particles
mine	the element index with the least number of particles

minnpel	number of particles in the element with the least number of particles
nce	number of critical elements
NELEM	number of elements, N_e , nelem
NPAR	total number of particles, N_p
npeli	temporary index array for npels
npels	number of particles per elements, temporary array
npel	number of particles per elements
np	number of particles in bins
parc	particle concentration
parid	indices of particles
psel[12]	linked lists that store the particle indices of each element
PDF	probability density function
PIV	particle image velocimetry
RANS	Reynolds-averaged Navier-Stokes
SLM	simplified Langevin model
UPDF	unsteady PDF
URANS	unsteady Reynolds-averaged Navier-Stokes
vcce	velocity-conditioned scalar concentration in bins

Greek

α	coefficient in particle wall-velocity condition, Equation (2.31)
α	scalar parameter in the simplified particle model, Equation (3.22)
α_i	projection vector in the IECM model, Equation (4.22)
Г	scalar diffusivity
γ_5	constant in the turbulence model, Equation (2.14)
Γ_T	turbulent diffusivity, Equation (4.14)
δ	relative numerical error in kinetic energy, Equation (3.27)
δ_{ij}	Kronecker symbol
ε	dissipation rate of turbulent kinetic energy

η	local coordinate direction via nodes A and C of an Eulerian element, Figure 3.2
η	random variable with a standard uniform distribution, Equation (3.30)
$\theta_{ m sep}$	separation angle
κ	Kármán constant in the law of the wall, Equation (2.34)
ν	kinematic viscosity
$ u_T$	turbulent viscosity, Equation (4.15)
ξ	local coordinate direction via nodes A and B of an Eulerian element, Figure 3.2
ρ	density
$ ho_{lpha}$	normalized velocity-scalar correlation vector, Equation (4.24)
$ ho_{lphaeta}$	normalized velocity-correlation tensor, Equation (4.24)
σ_y^2	mean-square particle displacement in the y direction, $\langle \mathcal{Y}^2 \rangle$, Equation (4.6)
$\sigma_{ m mean}$	width of the mean concentration profile (Section $4.3.3$)
σ_T	turbulent Prandtl number, Equation (4.15)
τ	Kolmogorov timescale (k/ε)
ϕ	scalar concentration
ϕ'	concentration fluctuation $(\psi - \langle \phi \rangle)$
ϕ_0	source strength
$ ilde{\phi}(\psi)$	mean scalar concentration conditional on $U = V$ and $\phi = \psi$, Equation (4.17)
χ	mean scalar dissipation
ψ	sample space variable corresponding to ϕ
ψ_{\max}	local maximum concentration
ψ_{\min}	local minimum concentration
Ω	turbulent frequency
ω	sample space variable of turbulent frequency Ω
ω_I	incident turbulent frequency at a wall, Equation (2.36)
ω_R	reflected turbulent frequency at a wall, Equation (2.36)

Roman

A	label of an element vertex, Figure 3.2
\mathcal{A}	surface area of cylinder
a	loop variable in Equation (4.1)
A_e	area of element e , Figure 3.2
A_i	fluid acceleration (Appendix A)
A_n	average element area around grid node \boldsymbol{n}
A_v	scalar variable in the near-wall turbulence model, Equation (2.13)
В	scalar diffusion, Equation $(A.3)$
b_{ij}	normalized Reynolds stress anisotropy, Equation (2.16)
C	measured concentration in the wind tunnel experiments (Chapter 5)
C_0	coefficient in the Langevin equations (2.10) and (2.28)
C_1	constant in the turbulence model, Equation (2.13)
C_2	constant in the turbulence model, Equation (2.14)
C_3	constant in the turbulent frequency model, Equation (2.20)
C_4	constant in the turbulent frequency model, Equation (2.20)
C_{η}	constant in the near-wall turbulence model, Equation (2.17)
$C_{\rm CFL}$	Courant-Friedrichs-Lewy constant (Section 3.2)
C_{μ}	constant in the $k-\varepsilon$ model
C_{ξ}	wall-normal dependent variable in the near-wall turbulence model, Equation $\left(2.17\right)$
C_L	constant in the near-wall turbulence model, Equation (2.17)
C_L	total lift coefficient
C_P	mean pressure coefficient
C_p	constant for stabilizing the mean pressure, Equation (3.8)
C_s	micromixing model constant
C_T	constant in the turbulent frequency model, Equation (2.20)

 xiv

C_t	micromixing model constant
C_v	constant in the turbulence model, Equation (2.14)
$C_{\omega 1}$	constant in the turbulent frequency model, Equation (2.20)
$C_{\omega 2}$	constant in the turbulent frequency model, Equation (2.20)
C_{Dp}	pressure drag coefficient
C_{Dv}	viscous drag coefficient
C_{Lp}	pressure lift coefficient
C_{Lv}	viscous lift coefficient
C_{Pb}	base pressure coefficient
C_D	total drag coefficient
D	cylinder diameter
d^n	distance of a particle from a wall at timestep n , Equation (3.28)
$dm{r}$	distance of the point r from the source, Equation (5.3)
e_x	unit basis vector in x direction
E	surface roughness parameter in the wall-function, Equation (2.34)
F	force acting on the cylinder
$f(oldsymbol{V},\psi)$	joint PDF of velocity and concentration
f_w	probability of a particle crossing the wall, Equation (3.28)
$f_{U\Omega\phi}(oldsymbol{V},\omega,\psi)$	joint PDF of velocity, turbulent frequency and scalar $(f_{\boldsymbol{Y}}(\boldsymbol{Y}))$
$f_{U\Omega \phi}(oldsymbol{V},\omega \psi)$	joint PDF of \boldsymbol{U} and Ω conditional on $\phi = \psi$
$f_{U\phi oldsymbol{Z}}(oldsymbol{V},\psi,oldsymbol{z})$	multi-point and multi-time joint PDF of \boldsymbol{U},ϕ and \boldsymbol{Z} (Appendix A)
$f_U(V)$	joint PDF of velocity
$f_{\boldsymbol{Y}}(\boldsymbol{Y})$	joint PDF of velocity, turbulent frequency and scalar $(f_{U\Omega\phi}(V,\omega,\psi))$
$f_{Z U\phi}(\boldsymbol{z} \boldsymbol{V},\psi)$	multi-point and multi-time joint PDF of Z conditional on $U = V$ and $\phi = \psi$, Equation (A.8)
$f_{\Omega\phi U}(\omega,\psi V)$	joint PDF of Ω and ϕ conditional on $U = V$, Equation (3.17)
$f_{\phi}(\psi)$	scalar PDF
$g(\psi)$	PDF of a reactive scalar

G_{ij}	coefficient in the GLM, Equations (2.10) and (2.28)
Н	building height (Chapter 5)
h	channel half width (Chapter 4)
h	free stream height (Chapter 5)
H_{ijkl}	fourth-order tensor in the near-wall turbulence model, Equation (2.13)
k	turbulent kinetic energy
k_a	analytical kinetic energy, Equation (3.27)
k_c	computed kinetic energy, Equation (3.27)
k_p	wall-normal component of the turbulent kinetic energy at y_p , Equation (2.33)
L	characteristic lengthscale in the near-wall turbulence model, Equation (2.17)
L	length of line source in the wind tunnel experiments (Chapter 5)
L	length of recirculation bubble (Chapter 6)
L_x	length of domain in x direction (Chapter 5)
n	outward pointing wall-normal
N	number of particles participating in local statistics (Section 3.4)
n	decay exponent in channel flow (Chapter 4)
n	number of nodes of an Eulerian element
$N_{p/e}^{\min}$	required minimum number of particles per element (Section 3.8)
N^j	finite element shape functions (N^A, N^B, N^C) , Equation (3.13)
N_c	number of bins used for velocity conditioning
N_e	number of elements
N_p	number of particles
$N_{p/e}$	number of particles per element initially
Р	particle location
\mathcal{P}	production of turbulent kinetic energy $(-\langle u_i u_j \rangle \partial \langle U_i \rangle / \partial x_j)$
$\langle P \rangle_0$	pressure at infinity
$\left< P \right>_b$	pressure at the back stagnation point of the cylinder
\wp_{ij}	non-local tensor in the elliptic relaxation, Equation (2.13)

Р	pressure
p	pressure fluctuation $(P - \langle P \rangle)$
\hat{Q}_j	nodal value of Q at gridpoint j
\hat{q}_j	value of the function q in node j
q	an unknown function
$Q(oldsymbol{U},\phi)$	random function of \boldsymbol{U} and ϕ
$Q({oldsymbol Y})$	random function with independent variables \boldsymbol{Y}
Q_s	scalar flow rate in the wind tunnel experiments (Chapter 5)
r	location
$r_{ m A}$	vector to node A of an Eulerian element, Figure 3.2
Re	Reynolds number for the street canyon, based on U_0 and H (Chapter 5)
Re_{τ}	friction Reynolds number $(u_{\tau}h/\nu)$
Re_L	turbulence Reynolds number
Re_D	Reynolds number for cylinder, based on U_0 and D (Chapter 6)
r_0	radius of the source
R_L	autocorrelation function of the Lagrangian velocity, Equation (4.5)
r_{\min}	radial location of the minimum mean streamwise velocity in the recirculation bubble (Chapter 6)
St	Strouhal number
$S(\phi)$	chemical source term
S_{ω}	source/sink for the mean turbulent frequency, Equation (2.20)
t	time
$t_{ m m}$	micromixing timescale, Equations (4.3) , (5.3)
T_L	Lagrangian integral timescale, Equation (4.7)
\hat{u}_p	characteristic velocity scale of the turbulence intensity in the vicinity of y_p , Equation (2.33)
$\mathcal{U}_{ ho}$	projected particle velocity, Equation (4.22)
\mathcal{U}_i	Lagrangian particle velocity
\mathcal{U}_i^p	velocity vector of particle p

\mathcal{U}_I	incident streamwise particle velocity at a wall, Equation (2.31)
\mathcal{U}_R	reflected streamwise particle velocity at a wall, Equation (2.31)
$\langle U\rangle_c$	mean velocity at the centerline of the channel
$\left\langle U\right\rangle _{p}$	mean streamwise velocity at y_p , Equation (2.32)
$\overline{\langle u_i u_j \rangle}$	spatial average of $\langle u_i u_j \rangle$ (Section 3.8)
U_0	free stream velocity (Chapter 5 and Chapter 6)
$U_{\rm ref}$	measured free stream velocity at y_{ref} in the wind tunnel experiments (Chapter 5)
u_{τ}	friction velocity
U_c	characteristic velocity, Equation (5.3)
U_e	magnitude of the equilibrium value of the mean velocity at y_p , Equation (2.34)
U_i	Eulerian velocity (U, V, W)
u_i	velocity fluctuation $(V_i - \langle U_i \rangle)$
U_{\min}	minimum mean streamwise velocity in recirculation bubble (Chapter 6)
V	sample space variable corresponding to U
\mathcal{V}_{I}	incident wall-normal particle velocity, Equation (2.30)
\mathcal{V}_R	reflected wall-normal particle velocity, Equation (2.30)
$\langle v^2 \rangle_p$	wall-normal component of the Reynolds stress tensor at y_p , Equation (2.32)
W	scalar-valued Wiener process
W_i'	vector-valued Wiener process (independent of W_i)
W_i	vector-valued Wiener process
\boldsymbol{x}	position
\mathcal{X}_i	Lagrangian particle position
x_{A}	x coordinate of node A, Figure 3.2
x_{CA}	$x_{\rm C} - x_{\rm A}$, Figure 3.2
x_i	position
Y	vector of particle properties / sample space variables $(V_1, V_2, V_3, \omega, \psi)$
\mathcal{Y}	particle position in the y direction

- y^+ distance from a wall in wall units $(u_\tau y/\nu)$
- $y_p \qquad \qquad$ distance of the first grid point from the wall where wall-functions are applied
- y_s cross-stream location of the source (Chapter 4)
- $m{Z}$ vector of multipoint variables in the joint PDF $f_{U\phi Z}(m{V},\psi,m{z})$
- $oldsymbol{z}$ vector of sample space variables corresponding to $oldsymbol{Z}$

Abstract

PDF MODELING OF TURBULENT FLOWS ON UNSTRUCTURED GRIDS József Bakosi, PhD

George Mason University, 2008

Dissertation Director: Dr. Zafer Boybeyi

In probability density function (PDF) methods of turbulent flows, the joint PDF of several flow variables is computed by numerically integrating a system of stochastic differential equations for Lagrangian particles. Because the technique solves a transport equation for the PDF of the velocity and scalars, a mathematically exact treatment of advection, viscous effects and arbitrarily complex chemical reactions is possible; these processes are treated without closure assumptions. A set of algorithms is proposed to provide an efficient solution of the PDF transport equation modeling the joint PDF of turbulent velocity, frequency and concentration of a passive scalar in geometrically complex configurations. An unstructured Eulerian grid is employed to extract Eulerian statistics, to solve for quantities represented at fixed locations of the domain and to track particles. All three aspects regarding the grid make use of the finite element method. Compared to hybrid methods, the current methodology is *stand-alone*, therefore it is consistent both numerically and at the level of turbulence closure without the use of consistency conditions. Since both the turbulent velocity and scalar concentration fields are represented in a stochastic way, the method allows for a direct and close interaction between these fields, which is beneficial in computing accurate scalar statistics.

Boundary conditions implemented along solid bodies are of the free-slip and no-slip type

without the need for ghost elements. Boundary layers at no-slip boundaries are either fully resolved down to the viscous sublayer, explicitly modeling the high anisotropy and inhomogeneity of the low-Reynolds-number wall region without damping or wall-functions or specified via logarithmic wall-functions. As in moment closures and large eddy simulation, these wall-treatments provide the usual trade-off between resolution and computational cost as required by the given application.

Particular attention is focused on modeling the dispersion of passive scalars in inhomogeneous turbulent flows. Two different micromixing models are investigated that incorporate the effect of small scale mixing on the transported scalar: the widely used interaction by exchange with the mean and the interaction by exchange with the conditional mean model. An adaptive algorithm to compute the velocity-conditioned scalar mean is proposed that homogenizes the statistical error over the sample space with no assumption on the shape of the underlying velocity PDF. The development also concentrates on a generally applicable micromixing timescale for complex flow domains.

Several newly developed algorithms are described in detail that facilitate a stable numerical solution in arbitrarily complex flow geometries, including a stabilized mean-pressure projection scheme, the estimation of conditional and unconditional Eulerian statistics and their derivatives from stochastic particle fields employing finite element shapefunctions, particle tracking through unstructured grids, an efficient particle redistribution procedure and techniques related to efficient random number generation.

The algorithm is validated and tested by computing three different turbulent flows: the fully developed turbulent channel flow, a street canyon (or cavity) flow and the turbulent wake behind a circular cylinder at a sub-critical Reynolds number.

The solver has been parallelized and optimized for shared memory and multi-core architectures using the OpenMP standard. Relevant aspects of performance and parallelism on cache-based shared memory machines are discussed and presented in detail. The methodology shows great promise in the simulation of high-Reynolds-number incompressible inert or reactive turbulent flows in realistic configurations.

Chapter 1: Introduction

In engineering industry and atmospheric transport and dispersion modeling there is an increasing use of computational methods to calculate complex turbulent flow fields. Many of these computations depend on the k- ε turbulence model (Jones and Launder, 1972; Bacon *et al.*, 2000), while some are based on second-moment closures (Rotta, 1951; Launder *et al.*, 1975; Hanjalić and Launder, 1972; Speziale *et al.*, 1991). The aim of these statistical methods is to predict the first and second moments of the turbulent velocity field, respectively. In large eddy simulation (LES) the large scale three-dimensional unsteady motions are represented exactly, while the small-scale motions are parameterized. As long as the transport-controlling processes of interest (eg. mass, momentum and heat transfer in shear flows) are resolved, LES predictions can be expected to be insensitive to the details of residual-scale modeling. In applications such as high-Reynolds-number turbulent combustion or near-wall flows, however, where the important rate-controlling processes occur below the resolved scales, the residual-scale models directly influence the model predictions. Since there is no universally 'best' methodology that is applicable for every type of practical flow, it is valuable to develop improvements for the full range of turbulence modeling approaches.

The development of probability density function (PDF) methods is an effort to provide a higher-level statistical description of turbulent flows. The mean velocity and Reynolds stresses are statistics of (and can be obtained from) the PDF of velocity. In PDF methods, a transport equation is solved directly for the PDF of the turbulent velocity field, rather than for its moments as in Reynolds stress closures. Therefore, in principle, a more complete statistical description can be obtained. While for some flows (e.g. homogeneous turbulence) this higher-level description may provide little benefit over second moment closures, in general the fuller description is beneficial in allowing more processes to be treated exactly and in providing more information that can be used in the construction of closure models. Convection, for example, can be exactly represented mathematically in the PDF framework, eliminating the need for a closure assumption (Pope, 2000). Similarly, defining the joint PDF of velocity and species concentrations in a chemically reactive turbulent flow allows for the treatment of chemical reactions without the burden of closure assumptions for the highly nonlinear chemical source terms (Fox, 2003). This latter advantage has been one of the most important incentives for the development of PDF methods, since previous attempts to provide moment closures for this term resulted in errors of several orders of magnitude (Pope, 1990).

The development of PDF methods has mostly been centered on chemically reactive turbulent flows on simple geometries, (e.g. Tang *et al.*, 2000; Xu and Pope, 2000), although applications to more complex configurations (James *et al.*, 2002; Subramaniam and Haworth, 2000) as well as to atmospheric flows (Heinz, 1998; Cassiani *et al.*, 2005b) have also appeared. A large variety of compressible and incompressible laminar flows bounded by bodies of complex geometries have been successfully computed using unstructured grids (Löhner, 2001). The flexibility of these gridding techniques has also been exploited recently in mesoscale atmospheric modeling (Bacon *et al.*, 2000). Significant advances in automatic unstructured grid generation (Löhner, 2000), sophisticated data structures and algorithms, automatic grid refinement and coarsening techniques (Shostko and Löhner, 1995) in recent years have made unstructured grids a common and convenient choice of spatial discretization in computational physics. The success of unstructured grids seems to warrant exploiting their advantages in conjunction with PDF modeling. For reasons to be elaborated on later, in PDF methods the usual choice of representation is the Lagrangian framework with a numerical method employing a large number of Lagrangian particles. A natural way to combine the advantages of existing traditional Eulerian computational fluid dynamics (CFD) codes with PDF methods, therefore, is to develop hybrid methods.

Using structured grids, a hybrid finite-volume (FV)/particle method has been developed by Muradoglu et al. (1999) and Jenny et al. (2001), wherein the mean velocity and pressure fields are supplied by the FV code to the particle code, which in turn computes the Reynolds stress, scalar fluxes and reaction terms. Different types of hybrid algorithms are possible depending on which quantities are computed in the Eulerian and Lagrangian frameworks. For a list of approaches see Muradoglu et al. (1999). Another line of research has been centered on the combination of LES with PDF methods (Givi, 1989; Madnia and Givi, 1993). This approach is based on the definition of the filtered density function (FDF) (Pope, 1990) which is used to provide closure at the residual scale to the filtered LES equations. Depending on the flow variables included in the joint FDF, different variants of the method have been proposed providing a probabilistic treatment at the residual scale for species compositions (Colucci et al., 1998), velocity (Gicquel et al., 2002) and velocity and scalars (Sheikhi et al., 2003). A common feature of these hybrid methods is that certain consistency conditions have to be met, since some fields are computed in both the Eulerian and Lagrangian frameworks. Additionally to the works cited above, further advances on consistency conditions and correction algorithms for hybrid FV/particle codes have been reported by Muradoglu et al. (2001) and Zhang and Haworth (2004), whose authors also extend the hybrid formulation to unstructured grids. Following that line, a hybrid algorithm for unstructured multiblock grids has recently been proposed by Rembold and Jenny (2006). Beside enforcing the consistency of redundantly computed fields, hybrid methods also have to be designed to ensure consistency at the level of the turbulence closure between the two frameworks. For example, the simplified Langevin model (SLM) (Haworth and Pope, 1986) is equivalent to Rotta's model at the Reynolds stress level (Pope, 1994). Thus the use of a $k-\varepsilon$ model in the Eulerian framework and of a SLM PDF model in the Lagrangian framework cannot be consistent (Muradoglu *et al.*, 1999). In the current work, a different approach is taken by representing all turbulent fields by Lagrangian particles and employing the grid (a) to compute only inherently Eulerian quantities (that are only represented in the Eulerian sense), (b) to extract Eulerian statistics and (c) to locate particles throughout the domain. Because the resulting method is not a hybrid one, none of the fields are computed redundantly and the computation can remain fully consistent without the need for correction algorithms. We employ the finite element method (FEM) in all three aspects mentioned above in conjunction with Eulerian grids. The combined application of the FEM and the decoupling of the Eulerian and Lagrangian fields also have important positive consequences regarding particle boundary conditions as compared to the "flux-view" of FV methods.

In the case of turbulent flows around complex geometries the presence of walls requires special treatment, since traditional turbulence models are developed for high Reynolds numbers and need to be modified in the vicinity of walls. This is necessary because the Reynolds number approaches zero at the wall, the highest shear rate occurs near the wall and the impermeability condition on the wall-normal velocity affects the flow up to an integral scale from the wall (Hunt and Graham, 1978). Possible modifications involve damping functions (van Driest, 1956; Lai and So, 1990; Craft and Launder, 1996; Rodi and Mansour, 1993) or wall-functions (Launder and Spalding, 1974; Singhal and Spalding, 1981; Rodi, 1980; Spalding, 1977). In those turbulent flows where a higher level of statistical description is necessary close to walls, adequate representation of the near-wall anisotropy and inhomogeneity is crucial. Durbin (1993) proposed a Reynolds stress closure to address these issues. In his model, the all-important process of pressure redistribution is modeled through an elliptic equation by analogy with the Poisson equation, which governs the pressure in incompressible flows. This represents the non-local effect of the wall on the Reynolds stresses through the fluctuating pressure terms. In an effort to extend PDF methods to wall-bounded turbulent flows, Durbin's elliptic relaxation method has been combined with the generalized Langevin model (Haworth and Pope, 1986) by Dreeben and Pope (1997a, 1998). Wall-function treatment has also been developed for the PDF framework by Dreeben and Pope (1997b), providing the option of the usual trade-off between computational expense and resolution at walls. With minor simplifications these wall-treatmets are closely followed throughout the present study. We compute fully resolved boundary layers with the elliptic relaxation technique and also apply wall-functions in order to investigate their effects on the results and the computational performance.

The dispersion of scalars (e.g. temperature, mass, etc.) in turbulent flows is relevant to a number of scientific phenomena including engineering combustion and atmospheric dispersion of pollutants. Reviews on the subject have been compiled by Shraiman and Siggia (2000) and Warhaft (2000). Several experimental studies have been carried out in order to better understand the behavior of transported scalars in homogeneous isotropic turbulence (Warhaft, 1984; Stapountzis et al., 1986; Sawford, 1995). A literature review of dispersion from a concentrated source in homogeneous but anisotropic turbulent shear flows is given by Karnik and Tayoularis (1989). Inhomogeneous turbulence (e.g. the atmospheric boundary layer or any practical turbulent flow) adds a significant level of complexity to these cases. Extensive measurements of the mean, variance, intermittency, probability density functions and spectra of scalar have been made by Fackrell and Robins (1982) in a turbulent boundary layer. One point statistics in turbulent channel flow have recently been reported by Lavertu and Mydlarski (2005). In urban scale modeling of passive pollutants in the atmosphere, the simplest settings to study turbulent flow and dispersion patterns are street canyons. Different canyon configurations and release scenarios have been studied both experimentally (Hoydysh et al., 1974; Wedding et al., 1977; Rafailids and Schatzmann, 1995; Meroney et al., 1996; Pavageau and Schatzmann, 1999) and numerically (Lee and Park, 1994; Johnson and Hunter, 1995; Baik and Kim, 1999; Huang et al., 2000; Liu and Barth, 2002). A widely studied case in both numerical and experimental fluid dynamics is the flow behind a circular cylinder. Despite its relative simplicity in domain geometry, a myriad of flow behaviors can be explored through a variety of physical circumstances in this flow (Williamson, 1996). Concentrating on the very near wake several aspects of the present PDF model will be explored by computing the turbulent velocity field behind a circular cylinder at a transitional Reynolds number. Direct numerical simulation (DNS) has served as an important counterpart to measurements of turbulence at moderate Reynolds numbers, shedding light on quantities that are difficult to measure (e.g. Lagrangian statistics) and at locations where it is nearly impossible to measure (e.g. close to walls). Turbulent velocity statistics extracted from DNS of channel flow have been reported by Moser *et al.* (1999) and Abe *et al.* (2004), while Vrieling and Nieuwstadt (2003) performed a DNS study of dispersion of plumes from single and double line sources. We will draw several datasets from the above experimental and numerical studies to compare and validate our results pertaining to the channel flow, the street canyon and the wake behind a circular cylinder.

A widely used model to incorporate the effects of small scale mixing on a scalar released in a turbulent flow in the PDF framework is the interaction by exchange with the mean (IEM) model of Villermaux and Devillon (1972) and Dopazo and O'Brien (1974). While this model has the virtue of being simple and efficient, it fails to comply with several physical constraints and desirable properties of an ideal mixing model (Fox, 2003). Although a variety of other mixing models have been proposed to satisfy these properties (Dopazo, 1994), the IEM model remains widely used in practice. Recently, increased attention has been devoted to the interaction by exchange with the conditional mean (IECM) model. Sawford (2004b) has done a comparative study of scalar mixing from line sources in homogeneous turbulence employing both the IEM and IECM models, wherein he demonstrated that the largest differences between the two models occur in the near-field. He also investigated the two models in a double scalar mixing layer (Sawford, 2006) with an emphasis on those conditional statistics that frequently require closure assumptions. Based on the IECM model, PDF micromixing models have been developed for dispersion of passive pollutants in the atmosphere by Luhar and Sawford (2005) and Cassiani *et al.* (2005a,b, 2007a,b). These authors compute scalar statistics in homogeneous turbulence and in neutral, convective and canopy boundary layer by assuming a joint PDF for the turbulent velocity field. However, no previous studies have been conducted on modeling the joint PDF of velocity and a passive scalar from a concentrated source in inhomogeneous flows.

The purpose of this research is to continue to widen the applicability of PDF methods in practical applications, especially to more realistic flow geometries by employing unstructured grids. The current work is a step in that direction, where we combine several models and develop a set of algorithms to compute the joint PDF of the turbulent velocity, characteristic frequency and scalar concentration in complex domains. Complementary to hybrid FV/particle and LES/FDF methods, we provide a different methodology to exploit the advantages of unstructured Eulerian meshes in conjunction with Lagrangian PDF methods. Three flows, a fully developed turbulent channel flow, a street canyon (or cavity) flow and the flow behind a circular cylinder are used to test several aspects of the algorithms.

A series of novel numerical algorithms are proposed to facilitate an efficient solution of the PDF transport equation. A modified pressure projection scheme that has traditionally been used to compute the pressure field in incompressible laminar flows is adapted to the Lagrangian Monte-Carlo solution to compute the mean pressure field in complex domains. Estimation of local Eulerian statistics and their derivatives employing finite element shapefunctions are presented. For the computation of the velocity-conditioned scalar mean required in the IECM model, we propose an adaptive algorithm that makes no assumption on the shape of the underlying velocity PDF and which, using a dynamic procedure, automatically homogenizes the statistical error over the sample space. An efficient particle-tracking procedure for two-dimensional triangles and three-dimensional tetrahedra is presented. Alternatively to particle splitting and merging algorithms, a particle redistribution algorithm is also proposed that ensures the stability of the numerical solution and reduces the need for high number of particles.

The solver has been optimized and parallelized for cache-based shared memory and multi-core machines using the OpenMP standard. Accordingly, the discussion on the numerical algorithms highlights several aspects of code design for these high-performance parallel architectures.

The remainder of the dissertation is organized as follows. In Chapter 2 the exact and modeled governing equations are described. Chapter 3 presents details of the solution algorithm with the underlying numerical methods. The method is tested and validated by computing scalar dispersion from concentrated sources in a fully developed turbulent channel flow in Chapter 4 and in a street canyon in Chapter 5 and by computing the velocity field behind a circular cylinder in Chapter 6. Finally, some conclusions are drawn and future directions are discussed in Chapter 7. Several important aspects of the underlying algorithms are detailed in the Appendices.

Chapter 2: Governing equations

The governing system of equations for a passive scalar released in a viscous, Newtonian, incompressible fluid can be derived from Newton's equations of motion (Hirsch, 1988) and is written in the Eulerian framework as

$$\frac{\partial U_i}{\partial x_i} = 0, \qquad (2.1)$$

$$\frac{\partial U_i}{\partial t} + U_j \frac{\partial U_i}{\partial x_j} + \frac{1}{\rho} \frac{\partial P}{\partial x_i} = \nu \nabla^2 U_i, \qquad (2.2)$$

$$\frac{\partial\phi}{\partial t} + U_i \frac{\partial\phi}{\partial x_i} = \Gamma \nabla^2 \phi, \qquad (2.3)$$

where U_i , P, ρ , ν , ϕ and Γ are the Eulerian velocity, pressure, constant density, kinematic viscosity, scalar concentration and scalar diffusivity, respectively. Based on these equations an exact transport equation can be derived for the one-point, one-time Eulerian joint PDF of velocity and concentration $f(\mathbf{V}, \psi; \mathbf{x}, t)$ (Pope, 1985, 2000),

$$\frac{\partial f}{\partial t} + V_i \frac{\partial f}{\partial x_i} = -\frac{\partial}{\partial V_i} \left[\left\langle \nu \nabla^2 U_i - \frac{1}{\rho} \frac{\partial P}{\partial x_i} \middle| \mathbf{V}, \psi \right\rangle f \right] - \frac{\partial}{\partial \psi} \left[\left\langle \Gamma \nabla^2 \phi \middle| \mathbf{V}, \psi \right\rangle f \right], \quad (2.4)$$

where V and ψ denote the sample space variables of the stochastic velocity U(x,t) and concentration $\phi(x,t)$ fields, respectively. (Equation (2.4) is derived in Appendix A.) A remarkable feature of Equation (2.4) is that the effects of convection and viscous diffusion (processes of critical importance in wall-bounded turbulent flows) are in closed form, thus require no modeling assumptions. Other effects, however, require closure assumptions. They are the effects of dissipation of turbulent kinetic energy, pressure redistribution and the small-scale mixing of the transported scalar due to molecular diffusion. The joint PDF $f(\mathbf{V}, \psi; \mathbf{x}, t)$ contains all one-point statistics of the velocity and scalar fields. The price to pay for the increased level of description (compared to traditional moment closures) is that in a general three-dimensional turbulent flow $f(\mathbf{V}, \psi; \mathbf{x}, t)$ is a function of 8 independent variables. This effectively rules out the application of traditional techniques like the finite difference, finite volume or finite element methods for a numerical solution. While in principle this high-dimensional space could be discretized and (after appropriate modeling of the unclosed terms) Equation (2.4) could be solved with the above methods, the preferred choice in the PDF framework is to use a Lagrangian Monte-Carlo formulation. As opposed to the other techniques mentioned, the computational requirements increase only linearly with increasing problem dimension with a Monte-Carlo method. Another advantage of employing a Lagrangian-particle based simulation is that the governing equations may take a significantly simpler form than Equation (2.4).

In a Lagrangian formulation, it is assumed that the motion of fluid particles along their trajectory is well represented by a diffusion process, namely a continuous-time Markov process with continuous sample paths (van Kampen, 2004). Such a process was originally proposed by Langevin (1908) as a stochastic model of a microscopic particle undergoing Brownian motion. Pope (2000) shows that Langevin's equation provides a good model for the velocity of a fluid particle in turbulence. It is important to appreciate that the instantaneous particle velocities, modeled by a Langevin equation, do not represent physical fluid particle velocities individually, rather their combined effect (i.e. their statistics) can model statistics of a turbulent flow. Therefore, the numerical particles can be thought of as an ensemble representation of turbulence, each particle embodying one realization of the flow at a given point in space and time. At a fundamental level, an interesting consequence of this view is that this definition does not require an external (spatial or temporal) filter explicitly,

as the classical Reynolds averaging rules and large eddy simulation filtering do. For example, in unsteady homogeneous or steady inhomogeneous high-Reynolds-number flows, the natural Reynolds-average to define is the spatial and temporal average, respectively. In unsteady *and* inhomogeneous flows however, one is restricted to employ temporal and/or spatial filters leading to the approaches of unsteady Reynolds-averaged Navier-Stokes (URANS) and LES methods, respectively (Pope, 2004). In the PDF framework statistics are defined based on a probability density function. In the current case, for example, the mean velocity and Reynolds stress tensor are obtained from the joint PDF f as

$$\langle U_i \rangle(\boldsymbol{x}, t) \equiv \int_{-\infty}^{\infty} \int_0^{\infty} V_i f(\boldsymbol{V}, \psi; \boldsymbol{x}, t) \mathrm{d}\psi \mathrm{d}\boldsymbol{V},$$
 (2.5)

$$\langle u_i u_j \rangle(\boldsymbol{x}, t) \equiv \int_{-\infty}^{\infty} \int_0^{\infty} (V_i - \langle U_i \rangle) (V_j - \langle U_j \rangle) f(\boldsymbol{V}, \psi; \boldsymbol{x}, t) \mathrm{d}\psi \mathrm{d}\boldsymbol{V},$$
 (2.6)

where the velocity fluctuation is defined as $u_i = V_i - \langle U_i \rangle$. These quantities are welldefined mathematically (van Kampen, 2004; Pope, 2000), independently of the underlying physics, the state of the flow (i.e. homogeneous or inhomogeneous, steady or unsteady), the numerical method and the spatial and temporal discretization. Therefore the promise of a probabilistic view of turbulence (as in PDF methods) at the fundamental level is a more rigorous statistical treatment.

An equivalent model to the Eulerian momentum equation (2.2) in the Lagrangian framework is a system of governing equations for particle position \mathcal{X}_i and velocity \mathcal{U}_i increments (Dreeben and Pope, 1997a)

$$d\mathcal{X}_i = \mathcal{U}_i dt + (2\nu)^{1/2} dW_i, \qquad (2.7)$$

$$d\mathcal{U}_i(t) = -\frac{1}{\rho} \frac{\partial P}{\partial x_i} dt + 2\nu \frac{\partial^2 U_i}{\partial x_j \partial x_j} dt + (2\nu)^{1/2} \frac{\partial U_i}{\partial x_j} dW_j, \qquad (2.8)$$

where the isotropic Wiener process dW_i (Gardiner, 2004) is identical in both equations (numerically, the same exact series of Gaussian random numbers with zero mean and variance dt) and it is understood that the Eulerian fields on the right hand side are evaluated at the particle locations \mathcal{X}_i . Since Equation (2.8) is a diffusion-type stochastic differential equation with a Gaussian white noise (i.e. a Wiener process), it is equivalent to the Fokker-Planck equation that governs the evolution of the probability distribution of the same process (van Kampen, 2004; Dreeben, 1997). Equations (2.7) and (2.8) represent the viscous effects exactly in the Lagrangian framework. Particles governed by these equations are both advected and diffused in physical space. In other words, besides convection the particles diffuse in physical space with coefficient ν , thus they carry momentum as molecules do with identical statistics, as in Brownian motion (Einstein, 1926). After Reynolds decomposition is applied to the velocity and pressure, Equation (2.8) results in

$$d\mathcal{U}_{i}(t) = -\frac{1}{\rho} \frac{\partial \langle P \rangle}{\partial x_{i}} dt + 2\nu \frac{\partial^{2} \langle U_{i} \rangle}{\partial x_{j} \partial x_{j}} dt + (2\nu)^{1/2} \frac{\partial \langle U_{i} \rangle}{\partial x_{j}} dW_{j} - \frac{1}{\rho} \frac{\partial p}{\partial x_{i}} dt + 2\nu \frac{\partial^{2} u_{i}}{\partial x_{j} \partial x_{j}} dt + (2\nu)^{1/2} \frac{\partial u_{i}}{\partial x_{j}} dW_{j},$$
(2.9)

where the last three terms are unclosed. To model these terms, we adopt the generalized Langevin model (GLM) of Haworth and Pope (1986)

$$d\mathcal{U}_{i}(t) = -\frac{1}{\rho} \frac{\partial \langle P \rangle}{\partial x_{i}} dt + 2\nu \frac{\partial^{2} \langle U_{i} \rangle}{\partial x_{j} \partial x_{j}} dt + (2\nu)^{1/2} \frac{\partial \langle U_{i} \rangle}{\partial x_{j}} dW_{j} + G_{ij} \left(\mathcal{U}_{j} - \langle U_{j} \rangle \right) dt + (C_{0}\varepsilon)^{1/2} dW_{i}',$$
(2.10)

where G_{ij} is a second-order tensor function of velocity statistics, C_0 is a positive constant, ε denotes the rate of dissipation of turbulent kinetic energy and dW'_i is another Wiener process. Because of the correspondence between stochastic Lagrangian models and Reynolds stress closures (Pope, 1994), different second order models can be realized with the Langevin equation (2.10), depending on how G_{ij} is specified. An advantage of the GLM family of models is that equation (2.10) ensures realizability as a valid Reynolds stress closure, provided that C_0 is non-negative and that C_0 and G_{ij} are bounded (Pope, 2000). Compared to Reynolds stress closures, the terms in G_{ij} and C_0 represent pressure redistribution and anisotropic dissipation of turbulent kinetic energy. Far from walls, these physical processes can be adequately modeled by appropriate local (algebraic) functions of the velocity statistics. However, such local representation is in contradiction with the large structures interacting with the wall and the viscous wall region (Whizman et al., 1996). The traditionally employed damping or wall-functions, therefore, are of limited validity in an approach aiming at a higher-level statistical description. To address these issues, Durbin (1993) proposed a technique to incorporate the wall-effects on the Reynolds stress tensor in a more natural fashion. In his approach, an elliptic equation is employed to capture the non-locality of the pressure redistribution at the wall, based on the analogy with the Poisson equation which governs the pressure in incompressible flows. The methodology also provides more freedom on controlling the individual components of the Reynolds stress tensor at the wall, such as the suppression of only the wall-normal component representing wall-blocking. Dreeben and Pope (1998) incorporated Durbin's elliptic relaxation technique into the PDF method using the constraint

$$\left(1 + \frac{3}{2}C_0\right)\varepsilon + G_{ij}\langle u_i u_j \rangle = 0, \qquad (2.11)$$

which ensures that the kinetic energy evolves correctly in homogeneous turbulence (Pope, 2000). Introducing the tensor \wp_{ij} to characterize the non-local effects G_{ij} and C_0 are defined as

$$G_{ij} = \frac{\wp_{ij} - \frac{\varepsilon}{2} \delta_{ij}}{k} \quad \text{and} \quad C_0 = \frac{-2\wp_{ij} \langle u_i u_j \rangle}{3k\varepsilon},$$
 (2.12)

where $k = \frac{1}{2} \langle u_i u_i \rangle$ denotes the turbulent kinetic energy. The non-local quantity φ_{ij} is

specified with the following elliptic relaxation equation

$$\wp_{ij} - L^2 \nabla^2 \wp_{ij} = \frac{1 - C_1}{2} k \langle \omega \rangle \delta_{ij} + k H_{ijkl} \frac{\partial \langle U_k \rangle}{\partial x_l}, \qquad (2.13)$$

where the fourth-order tensor H_{ijkl} is given by

$$H_{ijkl} = (C_2 A_v + \frac{1}{3}\gamma_5)\delta_{ik}\delta_{jl} - \frac{1}{3}\gamma_5\delta_{il}\delta_{jk} + \gamma_5b_{ik}\delta_{jl} - \gamma_5b_{il}\delta_{jk}, \qquad (2.14)$$

$$A_v = \min\left[1.0, C_v \frac{\det \langle u_i u_j \rangle}{\left(\frac{2}{3}k\right)^3}\right],\tag{2.15}$$

and

$$b_{ij} = \frac{\langle u_i u_j \rangle}{\langle u_k u_k \rangle} - \frac{1}{3} \delta_{ij} \tag{2.16}$$

is the Reynolds stress anisotropy, $\langle \omega \rangle$ denotes the mean characteristic turbulent frequency and C_1, C_2, γ_5, C_v are model constants. The characteristic lengthscale L is defined by the maximum of the turbulent and Kolmogorov lengthscales

$$L = C_L \max\left[C_{\xi} \frac{k^{3/2}}{\varepsilon}, C_{\eta} \left(\frac{\nu^3}{\varepsilon}\right)^{1/4}\right], \qquad (2.17)$$

with

$$C_{\xi} = 1.0 + 1.3n_i n_i, \tag{2.18}$$

where n_i is the unit wall-normal of the closest wall-element pointing outward of the flow domain, while C_L and C_η are model constants. The definition of C_{ξ} in Equation (2.18) signifies a slight departure from the original model by attributing anisotropic and walldependent behavior to its value. In the case of a channel flow, for example, where the wall is aligned with x, the wall-normal $\mathbf{n} = (0, -1, 0)$. This gives $C_{\xi} = 2.3$ in the computation of \wp_{22} in Equation (2.13) and $C_{\xi} = 1.0$ for all other components. The modification improves the
channel-centerline behavior of the wall-normal Reynolds stress component $\langle v^2 \rangle$ and in turn the cross-stream mixing of the passive scalar. Another departure from the original model is the application of the elliptic term $L^2 \nabla^2 \wp_{ij}$ (as originally proposed by Durbin (1993)) as opposed to $L \nabla^2 (L \wp_{ij})$. This simplification was adopted because no visible improvement has been found by employing the second, numerically more expensive term.

The right hand side of Equation (2.13) can be any local model for pressure redistribution; here we follow Dreeben and Pope (1998) and use the stochastic Lagrangian equivalent of a modified isotropization of production (IP) model proposed by Pope (1994). It is apparent that Equation (2.13) acts like a blending function between the low-Reynolds-number nearwall region and the high-Reynolds-number free turbulence. Close to the wall, the elliptic term on the left hand side brings out the non-local, highly anisotropic behavior of the Reynolds stress tensor, whereas far from the wall the significance of the elliptic term vanishes and the local model on the right hand side is recovered.

The description of the computation of the mean-pressure gradient in Equation (2.10) is deferred to Chapter 3.

The above model needs to be augmented by an equation for a quantity that provides length-, or time-scale information for the turbulence. With traditional moment closures the most common approach is to solve a model equation for the turbulent kinetic energy dissipation rate ε itself as proposed by Hanjalić and Launder (1972). An alternative method is to solve an equation for the mean characteristic turbulent frequency (Wilcox, 1993) $\langle \omega \rangle$ and to define

$$\varepsilon = k \langle \omega \rangle. \tag{2.19}$$

In PDF methods, however, a fully Lagrangian description has been preferred. A Lagrangian stochastic model has been developed for the instantaneous particle frequency ω by van Slooten *et al.* (1998) of which different forms exist, but the simplest formulation can be cast

into

$$d\omega = -C_3 \langle \omega \rangle \left(\omega - \langle \omega \rangle \right) dt - S_\omega \langle \omega \rangle \omega dt + \left(2C_3 C_4 \langle \omega \rangle^2 \omega \right)^{1/2} dW, \qquad (2.20)$$

where S_{ω} is a source/sink term for the mean turbulent frequency

$$S_{\omega} = C_{\omega 2} - C_{\omega 1} \frac{\mathcal{P}}{\varepsilon},\tag{2.21}$$

where $\mathcal{P} = -\langle u_i u_j \rangle \partial \langle U_i \rangle / \partial x_j$ is the production of turbulent kinetic energy, dW is a scalarvalued Wiener-process, while $C_3, C_4, C_{\omega 1}$ and $C_{\omega 2}$ are model constants. Since the no-slip condition would incorrectly force ε to zero at a no-slip wall, Equation (2.19) needs to be modified, thus the dissipation is defined as (Dreeben and Pope, 1998)

$$\varepsilon = \langle \omega \rangle \left(k + \nu C_T^2 \langle \omega \rangle \right), \qquad (2.22)$$

where C_T is also a model constant. A simplification of the original model for the turbulent frequency employed by Dreeben and Pope (1998) is the elimination of the ad-hoc source term involving an additional model constant, since in our case-studies we found no obvious improvements by including it. This completes the model for the joint PDF of velocity and the (now included) characteristic turbulent frequency ω . The specification of particle boundary conditions will be discussed in Chapter 3. The equations to model the joint PDF of velocity and turbulent frequency closely follow the work of Dreeben and Pope (1998). Slight modifications consist of

- the anisotropic definition of lengthscale L in Equations (2.17) and (2.18),
- the application of the elliptic term $L^2 \nabla^2 \wp_{ij}$ instead of $L \nabla^2 (L \wp_{ij})$ in Equation (2.13), and
- the elimination of an ad-hoc source term in Equation (2.21).

Since a passive scalar, by definition, has no effect on the turbulent velocity field, modeling

the pressure redistribution and dissipation have been discussed independently from the scalar, i.e. it has been assumed that in Equation (2.4) the following hold

$$\left\langle \nu \nabla^2 U_i - \frac{1}{\rho} \frac{\partial P}{\partial x_i} \middle| \mathbf{V}, \psi \right\rangle = \left\langle \nu \nabla^2 U_i - \frac{1}{\rho} \frac{\partial P}{\partial x_i} \middle| \mathbf{V} \right\rangle.$$
 (2.23)

However, the opposite, that the micromixing of the scalar can be modeled independently of V, cannot be assumed in general (Pope, 1998). A simple mixing model is the interaction by exchange with the mean (IEM) model (Villermaux and Devillon, 1972; Dopazo and O'Brien, 1974), which models the conditional scalar diffusion in Equation (2.4) independent of the underlying velocity field, i.e. assuming

$$\langle \Gamma \nabla^2 \phi | \mathbf{V}, \psi \rangle \cong \langle \Gamma \nabla^2 \phi | \psi \rangle.$$
 (2.24)

In the Lagrangian framework, the IEM model is written as

$$d\psi = -\frac{1}{t_{\rm m}} \left(\psi - \langle \phi \rangle\right) dt, \qquad (2.25)$$

where $t_{\rm m}$ is a micromixing timescale. It has been pointed out, however, that the assumption that the scalar mixing is independent of the velocity, Equation (2.24), bears no theoretical justification and is at odds with local isotropy of the scalar field (Fox, 1996; Pope, 1998). On the other hand, the interaction by exchange with the conditional mean (IECM) model does take the velocity field into consideration by employing the velocity-conditioned mean instead of the unconditional mean as

$$d\psi = -\frac{1}{t_{\rm m}} \left(\psi - \langle \phi | \boldsymbol{V} \rangle \right) dt.$$
(2.26)

Both the IEM and IECM models represent the physical process of dissipation and reflect the concept of relaxation towards a scalar mean with the characteristic timescale $t_{\rm m}$. The difference is that in the IEM model, all particles that have similar position interact with each other, while in the IECM model only those particles interact that also have similar velocities, e.g. fluid elements that belong to the same eddy.

It can be shown that in the case of homogeneous turbulent mixing with no mean scalar gradient the two models are equivalent since the velocity and scalar fields are uncorrelated (Fox, 1996) and the micromixing timescale $t_{\rm m}$ is proportional to the Kolmogorov timescale $\tau = k/\varepsilon$. In an inhomogeneous case, e.g. a concentrated source, however, there are various stages of the spreading of the plume requiring different characterizations of $t_{\rm m}$. In this case, the formal simplicity of the IEM and IECM models is a drawback, since a single scalar parameter $t_{\rm m}$ has to account for all the correct physics. The timescale should be inhomogeneous and should depend not only on the local turbulence characteristics but also on the source location, type, size, distribution and strength. Because of this complexity, a general flow-independent specification of $t_{\rm m}$ has been elusive. We will define the micromixing timescale for a passive scalar in the following chapters corresponding to the flows modeled.

This completes the model for the joint PDF of turbulent velocity, frequency and scalar. The model is 'complete' in the sense, that the equations are free from flow-dependent specifications (Pope, 2000), thus, in principle, it is generally applicable to any transported passive scalar released into an incompressible, high-Reynolds-number flow.

Defining G_{ij} and C_0 through (2.12) enables the model to adequately capture the nearwall effects in the higher-order statistics when the wall-region has sufficient resolution. In realistic simulations, however, full resolution of high-Reynolds-number boundary layers is not always possible (and may not be necessary), especially on the urban or meso-scale in atmospheric modeling. For such cases a second option is the use of wall-functions instead of the elliptic relaxation to model the near-wall turbulence. Employing wall-functions for no-slip walls provides a trade-off between the accuracy of fully resolved boundary layers and computational speed. The significantly more expensive full resolution is absolutely required in certain cases, such as computing the heat transfer at walls embedded in a flow or detaching boundary layers with high adverse pressure gradients. Conversely, a boundary layer representation by wall-functions is commonly used when the exact details close to walls are not important, and the analysis focuses on the boundary layer effects at farther distances. Wall-functions are widely applied in atmospheric simulations, where full wallresolution is usually prohibitively expensive even at the micro- or urban-scale (Bacon *et al.*, 2000; Lien *et al.*, 2004). It is worth emphasizing that one of the main assumptions used in the development of wall-functions is that the boundary layer remains attached. This is not always the case in simulations of complex flows. However, since wall-functions are the only choice for realistic atmospheric simulations, they are still routinely employed with reasonable success.

To investigate the gain in performance and the effect on the results, we implemented the wall-treatment for complex flow geometries that have been developed for the PDF method by Dreeben and Pope (1997b). Since in this case the viscous effects are not explicitly modeled, the viscous terms do not appear in the particle equations for the position and velocity increments:

$$\mathrm{d}\mathcal{X}_i = \mathcal{U}_i \mathrm{d}t, \tag{2.27}$$

$$d\mathcal{U}_{i}(t) = -\frac{1}{\rho} \frac{\partial \langle P \rangle}{\partial x_{i}} dt + G_{ij} \left(\mathcal{U}_{j} - \langle U_{j} \rangle \right) dt + (C_{0}\varepsilon)^{1/2} dW_{i}^{\prime}.$$
(2.28)

Furthermore, in this case the tensor G_{ij} is defined by the simplified Langevin model (SLM) (Haworth and Pope, 1986) and C_0 is simply a constant:

$$G_{ij} = -\left(\frac{1}{2} + \frac{3}{4}C_0\right) \langle \omega \rangle \delta_{ij} \quad \text{with} \quad C_0 = 3.5.$$
(2.29)

In line with the purpose of wall-functions, boundary conditions have to be imposed on particles that hit the wall so that their combined effect on the statistics at the first gridpoint from the wall will be consistent with the universal logarithmic wall-function in equilibrium flows, i.e. in boundary layers with no significant adverse pressure gradients. The development of boundary conditions based on wall-functions rely on the self-similarity of attached boundary layers close to walls. These conditions are applied usually at the first gridpoint from the wall based on the assumption of constant or linear stress-distribution. This results in the well-known self-similar logarithmic profile for the mean velocity. For the sake of completeness the conditions on particles developed by Dreeben and Pope (1997b) are reported here. The condition for the wall-normal component of the particle velocity reads

$$\mathcal{V}_R = -\mathcal{V}_I,\tag{2.30}$$

where the subscripts R and I denote reflected and incident particle properties, respectively. The reflected streamwise particle velocity is given by

$$\mathcal{U}_R = \mathcal{U}_I + \alpha \mathcal{V}_I, \tag{2.31}$$

where the coefficient α is determined by imposing consistency with the logarithmic law at the distance of the first gridpoint from the wall, y_p :

$$\alpha = \frac{2\hat{u}_p^2 \langle U \rangle_p |\langle U \rangle_p|}{\langle v^2 \rangle_p U_e^2},\tag{2.32}$$

where \hat{u}_p is a characteristic velocity scale of the turbulence intensity in the vicinity of y_p , defined as

$$\hat{u}_p = C_\mu^{1/4} k_p^{1/2}, \tag{2.33}$$

with $C_{\mu} = 0.09$. $\langle U \rangle_p$, $\langle v^2 \rangle_p$ and k_p are, respectively, the mean streamwise velocity, the wall-normal component of the Reynolds stress tensor and the turbulent kinetic energy, all obtained from the particle fields at y_p . In Equation (2.32) U_e is the magnitude of the equilibrium value of the mean velocity at y_p and is specified by the logarithmic law

$$U_e = \frac{u_\tau}{\kappa} \log\left(E\frac{y_p u_\tau}{\nu}\right),\tag{2.34}$$

where $\kappa = 0.41$ is the Kármán constant and the surface roughness parameter E = 8.5 for a smooth wall. The friction velocity u_{τ} is computed from local statistics as

$$u_{\tau} = \sqrt{\hat{u}_{p}^{2} + \gamma_{\tau} \left| \frac{y_{p}}{\rho} \frac{\partial \langle P \rangle}{\partial x} \right|} \quad \text{with} \quad \gamma_{\tau} = \max \left[0; \ \text{sign} \left(\langle uv \rangle \frac{\partial \langle P \rangle}{\partial x} \right) \right]. \tag{2.35}$$

In Equations (2.30-2.35) the streamwise x and wall-normal y coordinate directions are defined according to the local tangential and normal coordinate directions of the particular wall-element in question. In other words, if the wall is not aligned with the flow coordinate system then the vectors \mathcal{U}_i and $\partial \langle P \rangle / \partial x_i$, and the Reynolds stress tensor $\langle u_i u_j \rangle$, need to be appropriately transformed into the wall-element coordinate system before being employed in the above equations. The condition on the turbulent frequency is given by

$$\omega_R = \omega_I \exp\left[\beta \frac{\mathcal{V}_I}{y_p \langle \omega \rangle}\right] \quad \text{with} \quad \beta = -\frac{2}{\frac{1}{2} + \frac{3}{4}C_0 + C_3 + C_{\omega 2} - C_{\omega 1}}.$$
 (2.36)

In summary, the flow is modeled by a large number of Lagrangian particles representing a finite sample of all fluid particles in the domain which can be thought of as different realizations of the underlying stochastic fields. Numerically, each particle has position \mathcal{X}_i and with its velocity \mathcal{U}_i carries its turbulent frequency ω and scalar concentration ψ . For full wall-resolution the particle positions and velocities are advanced according to Equations (2.7) and (2.10) using Equations (2.12-2.18). While in the wall-functions case the positions and velocities are advanced by Equations (2.27) and (2.28) using Equations (2.29-2.36). In both cases, the particle frequencies and scalar concentrations are governed by (2.20) and either (2.25) or (2.26), respectively. The particle equations are discretized and advanced in time by the explicit forward Euler-Maruyama method (Kloeden and Platen, 1999). Even in the case of full wall-resolution using the elliptic relaxation technique, this method was preferred to the more involved exponential scheme that was originally suggested by Dreeben and Pope (1998), since the code is sufficiently stable with the simpler and computationally less expensive Euler-Maruyama method as well.

Chapter 3: Numerical implementation

3.1 Introduction

The numerical solution algorithm is based on the time-dependent particle governing equations: (2.7) and (2.10) in the full wall-resolution case and (2.27) and (2.28) in the wallfunctions case, (2.20) and either (2.25) or (2.26). An adaptive timestepping strategy to advance the system is described in Section 3.2. All Eulerian statistics required in these equations need to be estimated at the particle locations at the given instant in time. This is performed by an unstructured Eulerian grid that discretizes the flow domain, which can be conveniently refined around regions where a higher resolution is necessary. The methods used to compute unconditional statistics, their derivatives and conditional statistics are described in Sections 3.4, 3.5 and 3.6, respectively. The grid is also used to solve the elliptic relaxation equation (2.13) and to solve for the mean pressure required in Equation (2.10). The main characteristics of the solution of these two Eulerian equations together with a projection method to obtain the mean pressure are described in Section 3.3. In order to identify which particles contribute to local statistics, the particles need to be continuously followed as they travel throughout the domain. The particle tracking algorithm that is used for this purpose is described in Section 3.7. In complex configurations, where the spatial resolution can differ significantly from one region to another, an algorithm is necessary to ensure that the number of particles in every computational element is above a certain threshold, so that meaningful statistics can be computed. We present and test an algorithm that accomplishes this task in Section 3.8 and Appendix C and further refine it in Appendix D. The boundary conditions at no-slip walls applied to particles, to the elliptic relaxation equation (2.13) and to the mean pressure are described in Section 3.9. Some aspects of parallel random number generation are described in Section 3.10. An overview of the solution procedure with the execution profile of a timestep is given in Section 3.11.

3.2 Timestepping procedure

To discretize the governing equations in time we apply the explicit forward Euler-Maruyama scheme (Kloeden and Platen, 1999). The size of the timestep is estimated in every step based on the Courant-Friedrichs-Lewy (CFL) (Courant *et al.*, 1928) condition as

$$\Delta t = C_{\text{CFL}} \cdot \min_{n} \frac{\sqrt{A_n}}{||\langle \boldsymbol{U} \rangle_n||_2},\tag{3.1}$$

where A_n is the average element area around gridnode n. According to Equation (3.1) we compute a characteristic timescale for each gridnode by dividing the characteristic edge length (defined by the square-root of the element area) by the length of the mean velocity vector at the given location. Then we choose the smallest characteristic timescale of all gridpoints for the next timestep multiplied by a CFL constant of $C_{CFL} = 0.7$. This ensures that during a single step no information will travel farther than the length of Eulerian elements.

3.3 Solution of the Eulerian equations: mean pressure and elliptic relaxation

In incompressible flows the pressure establishes itself immediately through the pressure-Poisson equation, which is a manifestation of the divergence constraint (2.1) expressing mass conservation. The numerical difficulties arising from the straightforward discretization of this equation in finite difference, finite volume and finite element methods are reviewed by Löhner (2001). Several different methods have been devised to deal with these issues, which stem from the fact that the mass conservation equation decouples from the momentum equation and acts on it only as a constraint, which may result in the decoupling of every second gridpoint thereby numerically destabilizing the solution. Some of these methods are: the use of different functional spaces for the velocity and pressure discretization, artificial viscosities, consistent numerical fluxes, artificial compressibility and pressure projection schemes. For our purposes we adopt the pressure projection approach.

Additionally, in PDF methods due to the stochastic nature of the simulation, the Eulerian statistics and their derivatives are subject to considerable statistical noise. Fox (2003) suggests three different ways of calculating the mean pressure in PDF methods. The first approach is to extract the mean pressure field from a simultaneous consistent Reynolds stress model solved using a standard CFD solver (Correa and Pope, 1992). This approach solves the noise problem although it leads to a redundancy in the velocity model. The second approach attacks the noise problem by computing the so-called 'particle-pressure field' (Delarue and Pope, 1997). This results in a stand-alone transported PDF method and the authors succesfully apply it to compute a compressible turbulent flow. The third approach is the hybrid methodology mentioned in the Introduction, which uses an Eulerian CFD solver to solve for the mean velocity field and a particle-based code to solve for the fluctuating velocity (Muradoglu *et al.*, 1999; Givi, 1989). These methods are made consistent by the careful selection of turbulence models in the Eulerian and Lagrangian frameworks and the use of consistency conditions.

A different approach is proposed here. We adopt a modified version of the pressure projection scheme originally proposed by Chorin (1968) in the finite difference context, which has been widely used in laminar flows. The modification compared to the original projection scheme involves solving for the difference of the pressure between two consecutive timesteps, instead of the pressure field itself. This ensures that at steady state the residuals of the pressure correction vanish (Löhner, 2001). We adopt the scheme in the Lagrangian-Eulerian setting and combine the projection algorithm with the particle equations as follows.

The idea of pressure projection is to first predict the velocity using the current flow variables without taking the divergence constraint into consideration. Then in a second step, the divergence constraint is enforced by solving a pressure-Poisson equation. Finally the velocity is corrected using the new pressure field, resulting in a divergence-free velocity field. Thus, using full wall-resolution and explicit (forward Euler-Maruyama) time-integration of the particle velocity, one complete timestep $(n \rightarrow n + 1)$ is given by:

• Velocity prediction: $\mathcal{U}^n \to \mathcal{U}^*$

$$\mathcal{U}_{i}^{*} = \mathcal{U}_{i}^{n} - \frac{1}{\rho} \frac{\partial \langle P \rangle^{n}}{\partial x_{i}} \Delta t + 2\nu \frac{\partial^{2} \langle U_{i} \rangle^{n}}{\partial x_{j} \partial x_{j}} \Delta t + (2\nu)^{1/2} \frac{\partial \langle U_{i} \rangle^{n}}{\partial x_{j}} \Delta W_{j} + G_{ij} \left(\mathcal{U}_{j}^{n} - \langle U_{j} \rangle^{n} \right) \Delta t + (C_{0}\varepsilon)^{1/2} \Delta W_{i}^{\prime};$$
(3.2)

• Pressure projection: $\langle P \rangle^n \to \langle P \rangle^{n+1}$

$$\nabla \cdot \langle \boldsymbol{U} \rangle^{n+1} = 0, \qquad (3.3)$$

$$\frac{\langle U \rangle^{n+1} - \langle U \rangle^*}{\Delta t} + \frac{1}{\rho} \nabla (\langle P \rangle^{n+1} - \langle P \rangle^n) = 0, \qquad (3.4)$$

which results in

$$\frac{1}{\rho}\nabla^2(\langle P\rangle^{n+1} - \langle P\rangle^n) = \frac{\nabla \cdot \langle U\rangle^*}{\Delta t};$$
(3.5)

• Mean velocity correction: $\langle \boldsymbol{U} \rangle^* \to \langle \boldsymbol{U} \rangle^{n+1}$

$$\langle \boldsymbol{U} \rangle^{n+1} = \langle \boldsymbol{U} \rangle^* - \frac{1}{\rho} \Delta t \nabla (\langle \boldsymbol{P} \rangle^{n+1} - \langle \boldsymbol{P} \rangle^n).$$
(3.6)

Since the velocity field is fully represented by particles, the velocity prediction (3.2) and

correction (3.6) steps are applied to particles. The above procedure ensures that the Poisson equation for the mean pressure is satisfied at all times, thus the joint PDF representing an incompressible flow satisfies realizability, normalization and consistency conditions (Pope, 1985) in every timestep. To stabilize the computation of the mean pressure a small artificial diffusion term is added to the divergence constraint in Equation (3.3)

$$\nabla \cdot \langle \boldsymbol{U} \rangle^{n+1} = C_p \frac{1}{\rho} \nabla^2 \langle \boldsymbol{P} \rangle^n, \qquad (3.7)$$

where C_p is a small constant, e.g. $C_p = 10^{-3}$, which results in the stabilized version of the pressure projection step

$$\frac{1}{\rho}\nabla^2(\langle P\rangle^{n+1} - \langle P\rangle^n) = \frac{1}{\Delta t} \left(\nabla \cdot \langle U\rangle^* - C_p \frac{1}{\rho} \nabla^2 \langle P\rangle^n\right).$$
(3.8)

Both the elliptic relaxation (2.13) and pressure projection (3.8) equations are solved with the finite element method using linear shapefunctions on a grid consisting of triangles (Löhner, 2001). The grid is generated by the general purpose mesh generator, Gmsh, developed by Geuzaine and Remacle (2006). The FEM coefficient matrices are stored in block compressed sparse row format (Saad, 2003). The resulting linear systems are solved by the method of conjugate gradients combined with a Jacobi preconditioner. While the elliptic equation (2.13) for the tensor \wp_{ij} may appear prohibitively memory-hungry and expensive for larger meshes, the equation is well-conditioned and the iterative solution converges in just a few iterations starting from an initial condition using the solution in the previous timestep.

3.4 Estimation of Eulerian statistics

During the numerical solution of the governing equations, Eulerian statistics need to be estimated at different locations of the domain. Since the joint PDF contains information on all one-point statistics of the velocity, frequency and scalar concentration fields, these are readily available through appropriate averages of particle properties. For example, the mean velocity at a specific location in space and time is obtained as the integral over all sample space of the joint PDF $f_{\mathbf{Y}}(\mathbf{Y}; \mathbf{x}, t)$

$$\langle U_i \rangle \equiv \int V_i f_{\boldsymbol{Y}}(\boldsymbol{Y}; \boldsymbol{x}, t) \mathrm{d} \boldsymbol{Y},$$
 (3.9)

where \mathbf{Y} denotes the vector of all sample space variables $\mathbf{Y} = (V_1, V_2, V_3, \omega, \psi)$. For brevity we omit (but assume) the space and time dependence of the statistics. In traditional particle-codes the calculation of statistics is usually performed by kernel estimation using weight-functions (Pope, 2000). In particle-in-cell methods (Grigoryev *et al.*, 2002) an Eulerian mesh covers the computational domain and means are computed in each element or gridpoint. The latter approach is followed here and Equation (3.9) is computed by an ensemble average over all particle velocities in the vicinity of \mathbf{x}

$$\langle U_i \rangle \cong \frac{1}{N} \sum_{p=1}^N \mathcal{U}_i^p,$$
(3.10)

where N is the number of particles participating in the local mean at x and \mathcal{U}_i^p is the velocity vector of particle p. In the first pass an element-based mean is computed considering the particles in a given element, Figure 3.1. In the second pass, these element-based means are transferred to nodes of the grid by calculating the average of the elements surrounding the nodes. Wherever Eulerian statistics are needed at particle locations, like in Equation (2.10), the average of the nodal values are used for all particles residing in a given element. These node-based statistics are also used in the elliptic relaxation (2.13) and pressure projection (3.8) equations. An advantage of this two-pass procedure is that a natural smoothing is inherent in transferring statistics from elements to nodes. Using only nodal statistics to



Figure 3.1: Estimation of Eulerian statistics on unstructured grids. In a first pass, elementbased statistics are computed considering the particles residing in elements. In a second pass, element-based statistics are transferred to nodes by computing the averages of elements surrounding nodes. The nodal averages of each element are then used at particle locations in the Lagrangian governing equations.

update particles also makes the method more robust, since it provides an efficient guard against the unwanted occurrence of empty elements, i.e. elements without any particles. The problem of high statistical error caused by an empty element is mitigated by the other elements surrounding the given node. Linked lists (Löhner, 2001) provide an efficient access of unstructured-grid-based data from memory (e.g. elements surrounding points, points surrounding points, etc.). Once first-order statistics, like the mean velocity, are computed, higher order statistics are calculated by the same procedure. As an example, the Reynolds stress tensor is obtained by

$$\langle u_i u_j \rangle \equiv \int (V_i - \langle U_i \rangle) (V_j - \langle U_j \rangle) f_{\mathbf{Y}}(\mathbf{Y}) \mathrm{d}\mathbf{Y} \cong \frac{1}{N} \sum_{p=1}^N \left(\mathcal{U}_i^p - \langle U_i \rangle \right) \left(\mathcal{U}_j^p - \langle U_j \rangle \right).$$
(3.11)

3.5 Derivatives of Eulerian statistics

From finite element approximation theory, an unkown function $q(\boldsymbol{x})$ given in nodes can be approximated over an element as

$$q(\boldsymbol{x}) = \sum_{j=1}^{n} N^{j}(\boldsymbol{x})\hat{q}_{j}, \qquad (3.12)$$

where n is the number of nodes of the element, \hat{q}_j is the value of the function q in node j and N^j are finite element shapefunctions. For speed and simplicity, we use only a single type of element (triangle) with linear shapefunctions, which are written in the local (ξ, η) coordinate system of the element as (see also Figure 3.2)

$$N^{A} = 1 - \xi - \eta,$$

$$N^{B} = \xi,$$

$$N^{C} = \eta.$$
(3.13)

Employing the approximation in Equation (3.12), the spatial gradient of the expectation of any function $Q(\mathbf{Y}; \mathbf{x}, t)$ can be computed over an element as

$$\frac{\partial Q}{\partial x_i} = \sum_{j=1}^n \frac{\partial N^j}{\partial x_i} \hat{Q}_j, \qquad (3.14)$$

where \hat{Q}_j denotes the nodal value of Q at gridpoint j of the element. The derivatives of the linear shapefunctions in Equation (3.13) in the global (x, y) coordinate system can be derived analytically (Löhner, 2001)

$$\frac{\partial}{\partial x} \begin{bmatrix} N^{\mathrm{A}} \\ N^{\mathrm{B}} \\ N^{\mathrm{C}} \end{bmatrix} = \frac{1}{2A_{e}} \begin{bmatrix} -y_{\mathrm{CA}} + y_{\mathrm{BA}} \\ y_{\mathrm{CA}} \\ -y_{\mathrm{BA}} \end{bmatrix}, \qquad \frac{\partial}{\partial y} \begin{bmatrix} N^{\mathrm{A}} \\ N^{\mathrm{B}} \\ N^{\mathrm{C}} \end{bmatrix} = \frac{1}{2A_{e}} \begin{bmatrix} x_{\mathrm{CA}} - x_{\mathrm{BA}} \\ -x_{\mathrm{CA}} \\ x_{\mathrm{BA}} \end{bmatrix}, \quad (3.15)$$

where A_e is the area of element e. The derivatives are constant functions and are based only on the location of the gridpoints (see also Figure 3.2), e.g. $y_{CA} = y_C - y_A$. If the grid does not change during computation, these derivatives can be pre-computed and stored in advance of timestepping.

Second derivatives are obtained using a two-pass procedure. In the first pass the first derivatives are computed using Equation (3.14) and then transferred to nodes by computing the averages of the elements surrounding nodes. The same procedure is applied to the derivatives in gridpoints in the second pass to obtain second derivatives.

3.6 Estimation of the velocity-conditioned scalar mean

Equation (2.26) requires the estimation of the scalar mean conditioned on the velocity field $\langle \phi | \mathbf{V} \rangle$. In the current case, this is defined as

$$\langle \phi | \mathbf{V} \rangle \equiv \int \psi f_{\Omega \phi | U}(\omega, \psi | \mathbf{V}) \mathrm{d}\omega \mathrm{d}\psi,$$
 (3.16)

where the conditional PDF $f_{\Omega\phi|U}$ is expressed through Bayes' rule using the full PDF $f_{Y}(Y)$ and the marginal PDF of the velocity $f_{U}(V)$ as

$$f_{\Omega\phi|U}(\omega,\psi|V) \equiv \frac{f_Y(Y)}{f_U(V)} \equiv \frac{f_{U\Omega\phi}(V,\omega,\psi)}{f_U(V)}.$$
(3.17)

Mathematically, the conditional mean $\langle \phi | \mathbf{V} \rangle$ defines a mean value for each combination of its conditional variables, i.e. in a three-dimensional flow, in every spatial and temporal location $\langle \phi | \mathbf{V} \rangle$ is a function that associates a scalar value to a vector, $\langle \phi | \mathbf{V} \rangle : \mathbb{R}^3 \to \mathbb{R}$. In practice, this means that the velocity-sample space needs to be discretized (divided into bins) and different mean values have to be computed for each bin using the particles whose velocities fall into the bin. In order to keep the statistical error small this procedure would require a large number of particles in every element. To overcome this difficulty, Fox (1996) proposed a method in which the three-dimensional velocity space is projected onto a onedimensional subspace where the discretization is carried out. This substantially reduces the need for an extensive number of particles. This projection method is exact in homogeneous turbulent shear flows, where the joint velocity PDF is Gaussian. Nevertheless, in more complex situations it can still be incorporated as a modeling assumption.

A more general way of computing the conditional mean is to use three-dimensional binning of the velocity sample space V. In order to homogenize the statistical error over the sample space, the endpoints of the conditioning bins in each direction can be determined so that the distribution of the number of particles falling into the bins is as homogeneous as possible. For a Gaussian velocity PDF this can be accomplished by using statistical tables to define the endpoints (Fox, 1996). If the underlying velocity PDF is not known, however, another strategy is required. Note that there is absolutely no restriction on the distribution of the same (or even the same number) in every dimension and can also vary from element to element. Only some sort of clustering of the particles is needed, i.e. grouping them into subgroups of particles with similar velocities. A simple algorithm that accomplishes this task is as follows.

Without loss of generality, we assume that a sample-space binning of $(2 \times 2 \times 2)$ is desired. In a first step all particles residing in the given element are sorted according to their \mathcal{U} velocity component. Then the first and the second halves of the group are separately sorted according to their \mathcal{V} component. After further dividing both halves into halves again, each quarter is sorted according to the \mathcal{W} component. Finally, halving the quarters once again we compute scalar means for each of these 8 subgroups. Naturally, the binning can be any other structure with higher (even unequal) number of bins if that is desirable, e.g. $(5 \times 5 \times 5)$ or $(4 \times 12 \times 5)$. This procedure defines the bins dynamically based on the criterion that the bin-distribution of the number of particles be as homogeneous as possible. By doing that, it homogenizes the statistical error over the sample space and also ensures that every bin will contain particles. This simple procedure is completely general, independent of the shape and extent of the velocity PDF and dynamically adjusts the bin-distribution to the underlying PDF in every element. It is also robust, since if the number of particles in an element happens to be very low compared to the desired binning, e.g. we only have 5 particles for the 125 bins of a $(5 \times 5 \times 5)$ binning structure, the above sorting & dividing procedure can be stopped at any stage and the subgroups defined up to that stage can already be used to estimate the conditioned means. In other words, if in the above example we require that at least 2 particles should remain in every subgroup we simply stop after the first sort and only use two groups. An algorithm that accomplishes the conditioning step after the particles have been sorted into subgroups is detailed in Appendix B. The statistical error resulting from employing different number of conditioning bins is investigated in more detail in Chapter 4.

3.7 Particle tracking

Particles have to be tracked continuously as they travel throughout the grid in order to identify which element they contribute to when local statistics are computed. A variety of algorithms with different characteristics have been developed to accomplish this task (Grigoryev *et al.*, 2002). Since we use explicit timestepping, the particles will not jump over many elements in a timestep, thus the fastest way to track particles is some sort of known-vicinity algorithm (Löhner, 1995). The two-dimensional particle tracking employed here is as follows. If a particle is not in its old element (where it was in the last timestep), it is searched in the *next best element* of the surrounding elements. The knowledge of the next best element is a feature of the basic interpolation algorithm that is used to decide whether the particle resides in a given element. The interpolation algorithm is based on FEM shapefunctions, which are usually employed for approximating unknowns over elements (as it is used in Section 3.3 to discretize the Eulerian equations and in Section 3.5 to approximate functions and their derivatives) and correspond to a linear mapping between the global and local coordinates of the element, see also Figure 3.2. We use these shapefunctions here for interpolation in two dimensions, but this procedure can also be used in a three-dimensional case with tetrahedra (Löhner, 1995). In the current two-dimensional case, evaluating two of them is sufficient to decide whether the particle is inside of the element. The decision is made by the following condition (see also Figure 3.2)

$$\text{if } \left\{ \begin{array}{l} \left(N^{\mathrm{A}}>0\right) \quad \text{and} \quad \left(N^{\mathrm{C}}>0\right) \quad \text{and} \quad \left(N^{\mathrm{A}}+N^{\mathrm{C}}\right)< A_{e} \end{array} \right\} \qquad (3.18) \\ \text{inside} \\ \end{array}$$

else

outside

where A_e is the total area of the element, while N^A and N^C are the signed half-lengths of the cross-products

$$N^{\mathrm{A}} = \frac{1}{2} |(\boldsymbol{r}_{\mathrm{C}} - \boldsymbol{r}_{\mathrm{B}}) \times (\boldsymbol{r}_{\mathrm{P}} - \boldsymbol{r}_{\mathrm{B}})|, \qquad (3.19)$$

$$N^{\rm C} = \frac{1}{2} |(\boldsymbol{r}_{\rm P} - \boldsymbol{r}_{\rm B}) \times (\boldsymbol{r}_{\rm A} - \boldsymbol{r}_{\rm B})|. \qquad (3.20)$$

Note that these are also the area coordinates of the triangle corresponding to the nodes A and C and also the values of the finite element shapefunctions corresponding to the three nodes, Equations (3.13), evaluated at the particle location **P**. A convenient feature of this procedure is that once the values N^A , N^C and $N^B = A_e - N^A - N^C$ are evaluated, in case the particle is not found in the element, they also give us a hint about the direction of the particle location that is outside of the element. If condition (3.18) is not satisfied, at least one of N^A , N^B and N^C is negative. The next best element is in the direction



Figure 3.2: The decision whether a particle resides in a triangular element is made based on computing cross-products of element-edge vectors and vectors of vertex-particle coordinates. E.g. $N^{\rm A}$ is half of the signed area of the parallelogram spanned by vectors $(\mathbf{r}_{\rm C} - \mathbf{r}_{\rm B})$ and $(\mathbf{r}_{\rm P} - \mathbf{r}_{\rm B})$. Also shown is the local coordinate system (ξ, η) of the triangle after a linear mapping with the finite element shapefunctions in Equations (3.13).

corresponding to the lowest of the three values. Combining this with a data structure (e.g. a linked list (Löhner, 2001)) that stores the element indices surrounding elements, we can easily and efficiently identify which element is most likely to contain the particle or at least which direction to search next. Most of the time, the particles do not jump out of their host elements, but if they do, this procedure finds them in usually 2-3 steps.

The above neighbor-to-neighbor algorithm performs very well in the domain, but it may fail to jump over concave boundaries, resulting in a dead-lock (Löhner, 1995). In order to remedy this problem the following strategy is employed. An element on the boundary has two surrounding elements at most and the ones that would be outside of the domain are tagged in the data structure that stores the three element indices surrounding elements, see also Figure 3.3. If this tagged element is returned as the next best guess, the particle is on the other side of a concave section (or a corner) of the boundary. Since even in this case the particle must be close to its old host element, the particle is searched next in all



Figure 3.3: A particle jumping over a concave corner on the boundary and the next best guess based on its old host element would be through the boundary, outside of the domain. A fall-back procedure finds the new host element of the particle by searching the elements surrounding the nodes (displayed with thicker edges) of its old host element.

elements surrounding the nodes of its old host element. (This is also stored in a linked list for fast access.) This fall-back procedure always finds the particle around a corner, thus a brute-force search is not necessary over all elements.

3.8 Particle-number control

In the setup phase an equal number of particles are uniformly generated into each element with the initial velocities \mathcal{U}_i sampled from a Gaussian distribution with zero mean and variance 2/3, i.e. the initial Reynolds stress tensor is isotropic with unit turbulent kinetic energy, $\langle u_i u_j \rangle = \frac{2}{3} \delta_{ij}$. Initial particle frequencies ω are sampled from a gamma distribution with unit mean and variance 1/4 and the scalar concentration ψ is set to 0.

During the timestepping procedure a sufficient number of particles have to be present in every element at all times to keep the deterministic error due to bias small (Pope, 1995). However, the grid can be refined differently in different regions of the domain, as it is done at walls to resolve the boundary layer or around a concentrated source of a passive scalar to capture the high scalar gradients. Since the particles themselves model real fluid particles, at locations where the grid is refined more particles are necessary for an increased resolution. Therefore it is reasonable to keep the element-distribution of the number of particles as homogeneous as possible. Particle-number control is a delicate procedure in PDF methods, because external modification of the particle locations or properties may result in undesired changes of the local statistics and the joint PDF itself. Nevertheless, particle splitting and merging techniques are routinely applied to keep the particle distribution reasonable and to improve the efficiency and stability of the simulation (Cassiani *et al.*, 2007b). Appendix C describes the algorithm that we developed to keep the number of particles per element above a certain treshold and to guard the simulation against the occurrence of empty elements (i.e. elements without particles).

In what follows, we describe a simple testcase that we use to investigate the error introduced by the particle redisitribution. Note that the traditional way of referring to this procedure is *particle splitting and merging*. Since we do not change the total number of particles throughout the simulation (which is more memory efficient than splitting and merging) we refer to this as *particle redistribution*. To investigate the error, we consider the simplified model equations

$$\mathrm{d}\mathcal{X}_i = \mathcal{U}_i \mathrm{d}t, \tag{3.21}$$

$$d\mathcal{U}_i = -(\mathcal{U}_i - \alpha \langle U_i \rangle) dt + \sqrt{2} dW_i, \qquad (3.22)$$

where α is a scalar parameter and the initial conditions for \mathcal{U}_i are taken to be independent, standardized, normally distributed random variables:

$$\langle U_i \rangle = 0, \qquad \langle u_i u_j \rangle = \delta_{ij}.$$
 (3.23)

Equation (3.22) is characteristic of the Langevin equation (2.10) without viscous effects,



Figure 3.4: A rectangular domain with a stretched grid to test the error introduced by the particle redistribution algorithm using Equations (3.21) and (3.22).

e.g. Equation (2.28), see also Xu and Pope (1997). The mean $\langle U_i \rangle$ of the solution of the stochastic differential equation (3.22) is the solution of the following linear deterministic differential equation (Arnold, 1974)

$$\frac{\mathrm{d}\langle U_i \rangle}{\mathrm{d}t} = -(\langle U_i \rangle - \alpha \langle U_i \rangle), \qquad (3.24)$$

$$\langle U_i \rangle (t=0) = 0.$$
 (3.25)

It can be seen that the trivial solution $\langle U_i \rangle = 0$ satisfies the above deterministic initial value problem. For a nonzero initial condition the solution of Equation (3.22) is stable and reaches steady state if $\alpha < 1$ with $\langle U_i \rangle = 0$ and $\langle u_i u_j \rangle = \delta_{ij}$. For $\alpha > 1$ the equation becomes unstable and the solution grows exponentially, while for $\alpha = 0$ the equation is neutrally stable. For our purposes we use $\alpha = 0.5$. Equations (3.21) and (3.22) are advanced on a rectangular domain with two free-slip walls (from where particles are simply reflected) and a periodic inflow/outflow boundary-pair, see Figure 3.4. The domain is highly stretched on purpose in the y direction. Initially, an equal number of particles are generated into every element, which in the current case results in a spatially inhomogeneous particle distribution. As the timestepping advances the particles naturally tend to evolve into a spatially homogeneous distribution, which may result in empty elements in the highly refined region if the number of particles is too small. This is circumvented by the particle redistribution algorithm. We will test the algorithm by calculating the time-evolutions of the spatial average of the diagonal components of $\langle u_i u_j \rangle$, indicated by $\overline{\langle u_i u_j \rangle}$, using different initial number of particles per element $N_{p/e}$. In order to ensure that the particle redistribution algorithm intervenes on a same level in each case, the ratio

$$\frac{N_{p/e}}{N_{p/e}^{\min}} \propto \text{number of particles moved}$$
(3.26)

is kept constant. In other words, as the initial number of particles $N_{p/e}$ is increased, we increase the required minimum number of particles per element $N_{p/e}^{\min}$ as well, so that the number of particles that will have to be moved is approximately the same, hence the algorithm intervenes at the same level. To verify that this is the case, the number of times the redistribution algorithm is called (the number of particles moved in a timestep) is monitored and plotted in Figure 3.5 for the different cases. Figure 3.6 depicts $\langle u_i u_j \rangle$ for different values of $N_{p/e}$. It can be seen in Figure 3.6 (a) that the algorithm reproduces the analitical solution with a given numerical error. This error, which is always present in the numerical solution of stochastic differential equations, can be decomposed into three different parts: truncation error due to finite-size timesteps, deterministic error (or bias) due to the finite number of particles employed and random (or statistical) error (Pope, 1995). The particle redistribution introduces an additional error which is directly visible by comparing Figures 3.6 (a) and (d). It is also apparent that the bias decreases with increasing number of particles as it can be expected. However, Figures 3.6 (b)-(f) also show that the additional error introduced by the particle redistribution also diminishes as the number of particles



Figure 3.5: The number of particles moved in each timestep by the particle redistribution algorithm for different total number of particles. In the legend the constant $N_{p/e}/N_{p/e}^{\min}$ ratio is displayed.

increase while the intervention of the redistribution, Equation (3.26), is kept at a constant level. This can be seen more directly in Figure 3.7, which depicts the evolution of the total relative numerical error defined as

$$\delta = \frac{k_c - k_a}{k_a},\tag{3.27}$$

where k_c and k_a denote the computed and analytical kinetic energy, respectively. This error incorporates both the usual numerical errors and the additional one due to the particle redistribution algorithm. For comparison, the evolution of the error without particle redistribution is also displayed. Since the total sum of the errors converges to zero, the error introduced by the redistribution algorithm also diminishes and the solution converges to the PDF without redistribution.

We have found that a particle redistribution algorithm of a similar sort (or particle



Figure 3.6: Time-evolutions of the diagonal components of $\langle u_i u_j \rangle$ solving Equations (3.21) and (3.22) employing different number of particles. (a) No redistribution with initial number of particles per element $N_{p/e}=200$; redistribution with (b) $N_{p/e}=50$, (c) $N_{p/e}=100$, (d) $N_{p/e}=200$, (e) $N_{p/e}=400$ and (f) $N_{p/e}=800$, respectively. The ratio $N_{p/e}/N_{p/e}^{\min}=10$ is kept constant for cases (b) to (f). The horizontal line at the ordinate 1 depicts the analitical solution at steady state.

splitting and merging) is essential to provide adequate numerical stability in modeling inhomogeneous flows especially in complex geometries. In addition, it also dramatically reduces the need for high number of particles per elements on stretched grids.

3.9 No-slip wall-boundary conditions

Over any given time-interval a particle undergoing reflected Brownian motion in the vicinity of a wall may strike the wall infinitely many times (Dreeben and Pope, 1998). This means that particles can follow three different trajectories when interacting with walls. The particle either (a) crosses the wall during the timestep and it is behind the wall at the end of the



Figure 3.7: Evolution of the total relative numerical error defined by Equation (3.27) with increasing number of particles. Solid line – with redistribution, dashed line – without redistribution.

timestep or (b) crosses the wall during the timestep but it is not behind the wall at the end of the timestep or (c) does not cross the wall during the timestep. Therefore wall-conditions have to be enforced on particles that follow trajectory (a) and (b). The probability that the particle following trajectory (b) crossed the wall during timestep Δt can be calculated by (Karatzas and Shreve, 1991)

$$f_w = \exp\left(\frac{-d^n d^{n+1}}{\nu \Delta t}\right),\tag{3.28}$$

where d^n denotes the distance of the particle from the wall at timestep n. Thus, particle wall-conditions are applied if

$$d^{n+1} < 0, \quad \text{trajectory (a)}, \tag{3.29}$$

or if

$$d^{n+1} \ge 0$$
 and $\eta < f_w$, trajectory (b), (3.30)

where η is a random variable with a standard uniform distribution. The new particle location is calculated based on perfect reflection from the wall. The particle velocity is set according to the no-slip condition

$$\mathcal{U}_i = 0. \tag{3.31}$$

A boundary condition on the characteristic turbulent frequency ω has to ensure the correct balance of the turbulent kinetic energy at the wall (Dreeben and Pope, 1998) and has to be consistent with the near-wall kinetic energy equation

$$\nu \frac{\partial^2 k}{\partial n^2} + \varepsilon = 0, \qquad (3.32)$$

where n is the outward normal of the wall. Accordingly, the frequency for a particle striking the wall is sampled from a gamma distribution with mean and variance respectively

$$\langle \omega \rangle = \frac{1}{C_T} \frac{\mathrm{d}\sqrt{2k}}{\mathrm{d}y} \quad \text{and} \quad \left\langle \left(\omega - \langle \omega \rangle\right)^2 \right\rangle = C_4 \langle \omega \rangle^2.$$
 (3.33)

For better performance the above particle conditions are only tested and enforced for particles that reside close to walls, i.e. in elements that have at least an edge or a node on a no-slip wall-boundary.

Following Dreeben and Pope (1998), the wall-boundary condition for the elliptic relaxation equation (2.13) is set according to

$$\wp_{ij} = -4.5\varepsilon n_i n_j. \tag{3.34}$$

For the pressure-Poisson equation (3.8), a Neumann-condition is obtained from the Eulerian

mean-momentum equation

$$\frac{\partial \langle U_i \rangle}{\partial t} + \langle U_j \rangle \frac{\partial \langle U_i \rangle}{\partial x_j} + \frac{1}{\rho} \frac{\partial \langle P \rangle}{\partial x_i} = \nu \nabla^2 \langle U_i \rangle - \frac{\partial \langle u_i u_j \rangle}{\partial x_j}, \qquad (3.35)$$

by taking the normal component at a stationary solid wall

$$\frac{1}{\rho} \frac{\partial \langle P \rangle}{\partial x_i} n_i = \nu \frac{\partial^2 \langle U_i \rangle}{\partial x_j \partial x_j} n_i - \frac{\partial \langle u_i u_j \rangle}{\partial x_j} n_i.$$
(3.36)

In the wall-functions case, when the boundary layers along no-lip walls are represented based on the "law of the wall", the advection term in Equation (3.35) is non-zero at y_p , therefore the normal component of this term appears in the Neumann condition for the mean pressure

$$\frac{1}{\rho} \frac{\partial \langle P \rangle}{\partial x_i} n_i = \nu \frac{\partial^2 \langle U_i \rangle}{\partial x_j \partial x_j} n_i - \frac{\partial \langle u_i u_j \rangle}{\partial x_j} n_i - \langle U_j \rangle \frac{\partial \langle U_i \rangle}{\partial x_j} n_i.$$
(3.37)

3.10 Parallel random number generation

The solver has been parallelized and run on different shared memory architectures. Both the initialization and the timestepping require a large number of random numbers with different distributions and characteristics. Two components of the position \mathcal{X}_i and three components of the velocity \mathcal{U}_i are retained for a two-dimensional simulation, therefore the governing equations (2.7), (2.10) and (2.20) altogether require 6 independent Gaussian random numbers for each particle in each timestep. Since these 6 numbers per particle are always needed and are always Gaussian, they can be efficiently stored in a table, which is regenerated in each timestep. Different methods exist to efficiently sample pseudo-random numbers in parallel (Mascagni, 1997). In order to be able to reproduce the simulation results and to avoid surpassing cross-correlations between random number streams, we initialize a single stream and split it into k non-overlapping blocks, where k is the number of parallel threads. Then each of the threads generates from its own corresponding block, avoiding data races with other threads. This can be quite efficient, since a large amount of random numbers are generated at once and each thread accesses only its own portion of the stream. The same block-splitting technique is used to fill another table with uniform random numbers for the boundary condition Equation (3.30). Using this sampling technique, an almost ideal speedup can be achieved when random numbers in tables are regenerated, see also Table 3.1. For those equations in which the number of random numbers is a priori unknown (e.g. sampling a gamma distribution for the wall-condition of Equation (3.33) for particles that struck the wall), a stream is split into k disjoint substreams and the leap-frog technique is used to sample from them in parallel (Entacher et al., 1998). The leapfrog technique could also be used for a priori known number of random numbers, but due to its higher cache-efficiency, block-splitting performs slightly better. (In block-splitting the sampling positions in the streams are much farther from each other and thus the processes are less likely to interfere with each other's caches.) These techniques have been found essential to achieve a good parallel performance for the loop advancing the particles, see also Section 3.11.

3.11 Solution procedure and execution profile

The main stages of one complete timestep in their order of execution are displayed in Table 3.1. Also shown are the percentage of the execution times of each stage relative to a complete timestep and their speedups on a machine with two quad-core processors. The performance data were obtained by running a case that contained approximately 10 million particles and the Eulerian grid consisted of about 20 thousand triangles.

A significant portion of the execution time is spent on advancing the particle-governing equations. This is mostly a loop which can be constructed in two fundamental ways: in an element-based or in a particle-based fashion as displayed in Table 3.2. The main advantage

Table 3.1: Structure and profile of a timestep with relative execution times compared to the time spent on the full timestep and parallel performances of each step on a machine with two quad-core processors. The listing order corresponds to the order of execution. The performance data is characteristic of a case with 10M particles using a grid with 20K triangles, the simulation altogether requiring approximately 1.2GB memory. The processors are two quad-core CPUs (8 cores total), each pair sharing 4MB cache and the CPU-to-memory communication bandwidth.

task	relative	speedup	speedup	speedup	speedup	
	exe-	with	with	with	with	
	cution	2	4	6	8	
	time	CPUs	CPUs	CPUs	CPUs	
• compute the size of the next	0.001~%	not parallelized				
timestep, see Section 3.2						
• solve elliptic relaxation	2.87~%	1.91	4.08	5.76	7.60	
equation (2.13) , see Sec-						
tion 3.3						
• advance particle properties	73.2~%	2.02	4.12	6.16	8.20	
according to Equations (2.7) ,						
(2.10), (2.20) and (2.26)						
\bullet regenerate random number	19.01~%	2.01	3.99	5.79	7.50	
tables, see Section 3.10						
\bullet solve pressure-Poisson equa-	2.0~%	1.86	3.49	4.55	5.02	
tion, see Section 3.3						
\bullet correct mean velocities, see	1.0~%	1.69	1.95	1.94	1.96	
Section 3.3						
• compute Eulerian statistics,	1.6~%	1.22	1.79	1.67	1.77	
see Sections 3.4–3.6						
one complete timestep	99.68~%	1.98	3.95	5.55	7.20	

of the element-based loop is that once the Eulerian statistics are gathered for an element they can be used to update all particles in the element without recomputing them. However, it can be significantly off-balance in parallel, since it is not rare that the number of particles per element can differ by as much as two orders of magnitude at different regions of the domain. Another disadvantage of the element-based loop is that most of the time it accesses Table 3.2: Two fundamental ways of constructing a loop to advance the particle-governing equations (2.7), (2.10), (2.20) and (2.26). (i) – element-based loop, (ii) – particle-based loop.

(i)	
for	all Eulerian elements e gather Eulerian nodal statistics for element e ; compute element-average statistics;
	<pre>for all particles p in element e // update particles in element e</pre>
end	
(ii)	
for	all Eulerian elements e // pre-compute element-average statistics gather Eulerian nodal statistics for element e; compute and store element-average statistics;
end	
for	all particles p // update particles obtain index e of host element for particle p; get element-average Eulerian statistics for element e; advance particle p;
end	

the arrays containing the particle properties, $\mathcal{X}_i, \mathcal{U}_i, \omega, \psi$, in an unordered fashion resulting in increasing cache misses as the timestepping progresses and the particles move throughout the domain, because they get scrambled in memory compared to their spatial locations. Conversely, the big advantages of the particle-based loop are its simplicity and excellent load-balance for parallel execution. The particle-based loop always accesses the arrays containing particle properties consecutively. The effect of the increasing cache misses and the different load-balance on the performance is displayed in Figure 3.8, where the timings of the two loops are compared as the iteration progresses. The element-based loop slows down



Figure 3.8: Performance comparison of the two different loops (displayed in Table 3.2) to advance the particle governing equations (2.7), (2.10), (2.20) and (2.26) for the first 500 timesteps using 8 CPUs. The almost horizontal (red) line represents the particle-based loop, while the curving (black) one is the element-based loop. The problem size is the same as in Table 3.1.

almost fourfold in just 500 timesteps, while the performance degradation of the particlebased loop is negligible. Also, this disparity increases as the number of threads increases, which is shown in Table 3.3, where serial and parallel timings are displayed for both loops with different number of threads. While the element-based loop slightly outperforms the particle-based loop on a single CPU, the high scalability and cache-efficiency of the particlebased loop pays out very well in parallel. In fact its speedup is superlinear, which is due to the fact that as the number of processors increase, more and more data gathered from memory fit into the aggregate cache of the individual CPUs, resulting in faster processing than from central memory.

Cache misses may also be reduced by specifically optimizing for the architecture of shared

Table 3.3: A comparison of serial and parallel performances for a single timestep of the most time-consuming loop, implementing the governing equations to advance particles, Equations (2.7), (2.10), (2.20) and (2.26), using the two different loop-structures displayed in Table 3.2. The data is obtained from the same test simulation as in Table 3.1 using the same hardware. The timings are approximate values after the first 500 timesteps.

	element-based loop		particle-bas	ed loop
number of CPUs	time (ms)	speedup	time (ms)	speedup
1	6909	1.0	8068	1.00
2	4122	1.68	3987	2.03
4	2408	2.87	1943	4.12
6	1979	3.49	1305	6.16
8	1945	3.55	1000	8.20

caches on multi-core CPUs as it has been done in the current case. We have found that this guarantees a good performance on true shared memory machines as well, i.e. on machines whose CPUs do not share their caches and the communication bandwith between the CPU and memory. However, optimizing for non-shared caches and communication bandwidths does not necessarily guarantee optimal performance on multi-core CPUs. These findings clearly show the importance of efficient use of caches. This was also noted with Eulerian CFD codes computing a variety of flows by e.g. Camelli *et al.* (2007).

The parallel performance on higher number of processors is plotted in Figure 3.9. The size of the testproblem is the same as previously in Table 3.1, but the hardware is now a true shared memory machine with separate cache and memory-to-CPU bandwidth for each processor. The code performs reasonably well for this moderate-size problem and the parallel efficiency does not show a sign of leveling out up to the 32 CPUs tested. For comparison, the performance data in Table 3.1 is also shown using mullicore CPUs.

Table 3.1 shows, that the second most time-consuming step in a timestep is the regeneration of the random number tables, which was discussed in Section 3.10. Interestingly, the



Figure 3.9: Overall parallel performance of 100 timesteps taken on two different types of shared memory machines. Solid line and symbols – separate caches and memory-to-CPU bandwidths for each processor, dashed line and open symbols – two quad-core CPUs (8 cores total) each pair sharing a cache and a memory-to-CPU bandwidth. The problem size is the same as in Table 3.1.

solution of the two Eulerian equations, namely the elliptic relaxation equation (2.13) and the pressure-Poisson equation (3.8), only take up about 2-3% of a timestep, respectively. It is worth noting, that the linear system for the elliptic relaxation is nine times larger than that of the pressure-Poisson equation. The former is very well conditioned, while the latter is usually the most time-consuming equation to solve in modeling laminar incompressible flows.
Chapter 4:

Channel flow simulations: results and discussion

4.1 Introduction

In this Chapter, the previously described PDF model is tested in a fully developed, turbulent, long-aspect-ratio channel flow, where a passive scalar is continuously released from concentrated sources. The joint PDF of velocity, characteristic turbulent frequency and concentration of a passive scalar is computed using stochastic equations. The flow is explicitly modeled down to the viscous sublayer by imposing only the no-slip and impermeability condition on particles without the use of damping, or wall-functions. The high-level inhomogeneity and anisotropy of the Reynolds stress tensor at the wall are captured by the elliptic relaxation method. A passive scalar is released from a concentrated source at the channel centerline and in the viscous wall-region. The effect of small-scale mixing on the scalar is mainly modeled by the IECM model. The performance and accuracy of the IECM model compared to the simpler, but more widely used IEM model are evaluated. Several one-point unconditional and conditional statistics are presented in both physical and composition spaces. An emphasis is placed on common approximations of those conditional statistics that require closure assumptions in concentration-only PDF methods, i.e. in methods that assume the underlying turbulent velocity field. The results are compared to the DNS data of Abe et al. (2004) and the experimental data of Lavertu and Mydlarski (2005). The experiments were performed at two different Reynolds numbers ($Re_{\tau} \equiv u_{\tau}h/\nu = 520$ and 1080 based on the friction velocity u_{τ} , the channel half width h, and the kinematic viscosity ν) in a high-aspect-ratio turbulent channel flow, measuring one point statistics of a scalar (temperature) emitted continuously at three different wall-normal source locations from concentrated line sources. Measurements were performed at six different downstream locations between $4.0 \leq x/h \leq 22.0$.

The Chapter is organized as follows. A brief account of the underlying numerical methods with various implementation details specific to this flow are presented in Section 4.2. In Section 4.3, one-point velocity statistics are compared to direct numerical simulation data at $Re_{\tau} = 1080$, and a comparative assessment of the two micromixing models with analytical and experimental data is also given. Detailed statistics of scalar concentration calculated with the IECM micromixing model are presented. Section 4.4 presents a study of the effects of several numerical parameters on the computed results, including the effect of the Reynolds number, the type of velocity conditioning and the number of particles employed. An assessment of the computational cost of the current method is given compared to DNS in Section 4.5. Finally, conclusions pertaining to the channel flow testcase and results are summarized in Section 4.6.

4.2 Modeling specifics of channel flow

The velocity field, in turbulent channel flow after an initial development time, becomes statisticially stationary and homogeneous in the streamwise direction, while it remains inhomogeneous in the wall-normal direction, i.e. the flow becomes statistically one-dimensional. The flow is assumed to be statistically symmetric about the channel centerline. A passive scalar released into this flow is inhomogeneous and three-dimensional. Assuming the channel cross section has a high aspect ratio, we confine our interest to the plane spanned by the wall-normal and streamwise directions, far from the spanwise walls. The computational scheme exploits these features by resolving only one spatial dimension for the velocity statistics and two dimensions for the passive scalar. Although this specialized implementation of the method includes flow-dependent features, it provides good indication of the total computational cost. The description is divided into sections that separately discuss the modeling of the fluid dynamics (Section 4.2.1) and the transported scalar (Section 4.2.2). Both DNS and experimental data are used to validate the results.

4.2.1 Modeling the fluid dynamics

Since the transported scalar is inhomogeneous, both streamwise x and cross-stream y components of the particle positions are retained. A one-dimensional grid is used to compute Eulerian statistics of the velocity and turbulent frequency. An increasing level of refinement is achieved in the vicinity of the wall by obtaining the spacing of the gridpoints from the relation

$$y^{+} = 1 - \cos\left(\frac{\pi}{2}a^{3/4}\right), \qquad 0 \le a < 1,$$
(4.1)

where $y^+ = u_{\tau}y/\nu$ is the distance from the wall non-dimensionalized by the friction velocity u_{τ} and the kinematic viscosity ν and a is a loop-variable that equidistantly divides the interval between 0 and 1 (wall and centerline, respectively) into a desired number of gridpoints. The centerline symmetry of the flow is exploited, thus these statistics are only computed on half of the channel. Using this one-dimensional grid, Eulerian statistics are computed as described in Section 3.4. First and second derivatives of the mean velocity are calculated by first-order accurate finite difference formulas over each element and then transferred to nodes. A constant unit mean streamwise pressure gradient is imposed, which drives the flow and builds up the numerical solution. The cross-stream mean-pressure gradient is obtained by satisfying the cross-stream mean-momentum equation for turbulent channel flow

$$\frac{1}{\rho} \frac{\mathrm{d}\langle P \rangle}{\mathrm{d}y} = -\frac{\mathrm{d}\langle v^2 \rangle}{\mathrm{d}y},\tag{4.2}$$

which implies that the pressure-projection is not necessary for this flow. Since the number of elements does not exceed 100, particle tracking in this one-dimensional case is simply a

C_1	C_2	C_3	C_4	C_T	C_L	C_{η}	C_v	γ_5	$C_{\omega 1}$	$C_{\omega 2}$
1.85	0.63	5.0	0.25	6.0	0.134	72.0	1.4	0.1	0.5	0.73

Table 4.1: Constants for modeling the joint PDF of velocity and frequency.

brute-force check on each element. This is a negligible fraction of the running time, thus there is no need for a more sophisiticated tracking algorithm.

Wall-boundary conditions for the particles are the same as described in Section 3.9, only the situation is simpler here, since the wall is aligned with the coordinate line y = 0. The conditions for the centerline are symmetry conditions, i.e. particles trying to leave the domain through the centerline undergo perfect reflection and the sign of their wallnormal velocity is reversed. Consistently with these particle conditions, boundary conditions are imposed on the Eulerian statistics as well. At the wall, the mean velocity and the Reynolds stress tensor is forced to zero. The mean frequency $\langle \omega \rangle$ is set according to Equation (3.33). At the centerline, the shear Reynolds stress $\langle uv \rangle$ is set to zero. At the wall in the elliptic-relaxation equation (2.13), \wp_{ij} is set according to $\wp_{ij} = -4.5\varepsilon n_i n_j$. In the current case the wall is aligned with y = 0 thus only the wall-normal component is non-zero: $\wp_{22} = -4.5\varepsilon$. At the centerline, symmetry conditions are enforced on \wp_{ij} , i.e. homogeneous Dirichlet-conditions are applied for the off-diagonal components and homogeneous Neumann-conditions for the diagonal components. The initial conditions for the particles are set according to Section 3.8, however the current one-dimensional case enables the use of a sufficient number of particles so that there is no need for particle redistribution. The applied model constants for the joint PDF of velocity and frequency are displayed in Table 4.1.

4.2.2 Modeling the passive scalar

A passive, inert scalar is released from a concentrated source into the modeled fully developed turbulent channel flow, described above. Since the scalar field is inhomogeneous and, in general, not symmetric about the channel centerline, a second, two-dimensional grid is employed to calculate scalar statistics. Employing separate grids for the fluid dynamics and scalar fields enables the grid refinement to be concentrated on different parts of the domain, i.e. the scalar-grid can be refined around the source, while the fluid dynamics-grid is refined at the wall. The two-dimensional mesh is used to calculate Eulerian scalar statistics as described in Section 3.4. Since the scalar statistics are not homogeneous in the streamwise direction, the long rectangular domain is subdivided into several bins (thin vertical stripes, see Figure 4.1) and the following strategy is used to exploit these features. The velocity and turbulent frequency statistics are computed using the one-dimensional grid in which only particles in the first bin participate. The position of these particles are then copied to all downstream bins and (since the fluid dynamics is symmetric about the channel centerline) these particle positions are also mirrored to the upper half of the channel. This means that the particles (as far as positions are concerned) never leave the first bin physically. Since the scalar is passive, only one-way coupling between the two grids is necessary. This is accomplished by using the local velocity statistics computed in the 1d-elements for those 2d-elements that lie the closest to them in the wall-normal coordinate direction. At the wall and centerline boundaries the conditions on the particle properties have already been described in Section 4.2.1. For particles trying to leave the bin through the "inflow/outflow" bin-boundaries a periodic boundary condition is applied, with leaving particles put back on the opposite side. This essentially means that the particle paths remain continuous (as they should), only the code accounts for them as different particles in the computer memory. In order to carry the scalar concentration through bin-boundaries, the particle-scalar ψ



Figure 4.1: The computational domain for the channel flow is subdivided into several bins to exploit the streamwise statistical homogeneity of the turbulent velocity and frequency fields. Particle positions are copied downstream and mirrored to the upper half. Particle scalar concentrations are exchanged through bin-boundaries and the centerline. Note, that the number of particles in the figure does not correspond to the actual number used in the simulation.

is copied downstream (upstream) when the particle tries to leave through the downstream (upstream) bin-boundary. If the particle hits the centerline, its concentration is exchanged with its mirrored pair on the upper half, facilitating a possible non-symmetric behaviour of the scalar. The computation of the velocity-conditioned scalar mean $\langle \phi | \mathbf{V} \rangle$ required in the IECM model (2.26) is carried out with the method described in Section 3.6. The line-source, which in the current two-dimensional case is a point-source, is represented by a circular source with non-dimensional diameter $0.05\nu/u_{\tau}$. The scalar at the source has a constant distribution: particles passing through the source are assigned a constant normalized

unit source strength, i.e. $\psi = \phi_0 = 1$.

The micromixing timescale t_m required to model the viscous diffusion of the scalar is specified based on the following observations. In general, t_m is assumed to be proportional to the timescale of the instantaneous plume (Sawford, 2004b). Once the initial conditions are forgotten, theoretical results (Franzese and Cassiani, 2007) show that the timescale of the instantaneous plume is linear in t in the inertial subrange and is proportional to the turbulence timescale in the far field, when the instantaneous plume grows at the same rate as the mean plume. Based on these considerations the micromixing timescale is computed according to

$$t_{\rm m} = \min\left[C_s \left(\frac{r_0^2}{\varepsilon}\right)^{1/3} + C_t \frac{x}{\langle U \rangle_c}; \ \max\left(\frac{k}{\varepsilon}, C_T \sqrt{\frac{\nu}{\varepsilon}}\right)\right],\tag{4.3}$$

where r_0 denotes the radius of the source, $\langle U \rangle_c$ is the mean velocity at the centerline of the channel, while C_s and C_t are micromixing model constants. This definition reflects the three stages of the spreading of the plume. In the first stage, the timescale of the plume is proportional to that of the source (Batchelor, 1952): accordingly, the first term in the *min* operator represents the effect of the source. In the second stage t_m increases linearly as the scalar is dispersed downstream and the distance x from the source grows (Franzese and Cassiani, 2007). In the final stage, the timescale is capped with the characteristic timescale of the turbulence, which provides an upper limit in the third term of Equation (4.3). Following Durbin (1991) this is defined as the maximum of the turbulent and Kolmogorov timescales: far from the boundaries it becomes k/ε , whereas near a surface, where $k \to 0$, the Kolmogorov timescale provides a lower bound as $C_T(\nu/\varepsilon)^{1/2}$.

4.3 Results

The model has been run for the case of fully developed channel flow at $Re_{\tau} = 1080$ based on the friction velocity u_{τ} and the channel half-width h with a passive scalar released from a concentrated source at the centerline $(y_s/h = 1.0)$ and in the viscous wall region $(y_s/h = 0.067)$. The results are divided into a discussion of the fluid dynamics statistics (4.3.1), a comparison of the two micromixing models (4.3.2) and a presentation of unconditional (4.3.3) and conditional (4.3.4) scalar statistics.

4.3.1 Fluid dynamics

The equations to model the velocity and turbulent frequency have been solved on a 100-cell one dimensional grid with 500 particles per cell. The applied model constants are displayed in Table 4.1. The computed cross-stream profiles of mean streamwise velocity, the non-zero components of the Reynolds stress tensor and the rate of dissipation of turbulent kinetic energy are compared with the DNS data of Abe *et al.* (2004) at $Re_{\tau} = 1020$ in Figure 4.2. Previous PDF modeling studies employing the elliptic relaxation technique (Dreeben and Pope, 1997a, 1998; Wacławczyk *et al.*, 2004) have been conducted up to $Re_{\tau} = 590$. The high-level inhomogeneity and anisotropy in the viscous wall region are well represented by the technique at this higher Reynolds number as well. The purpose of including the parameter C_{ξ} in Equation (2.17) of the wall-normal component of \wp_{ij} is to correct the overprediction of the wall-normal Reynolds stress component $\langle v^2 \rangle$ at the centerline. This facilitates the correct behavior of the mean of the dispersed passive scalar in the center region of the channel (presented in Section 4.3.2).



Figure 4.2: Cross-stream profiles of (a) the mean streamwise velocity, (b) the diagonal components of the Reynolds stress tensor, (c) the shear Reynolds stress and (d) the rate of dissipation of turbulent kinetic energy. Lines – PDF calculation, symbols – DNS data of Abe *et al.* (2004). All quantities are normalized by the friction velocity and the channel half-width. The DNS data is scaled from $Re_{\tau} = 1020$ to 1080.

4.3.2 Comparison of the IEM and IECM micromixing models

An often raised criticism of the IEM model is that there is no physical basis for assuming the molecular mixing to be independent of the velocity field. This assumption gives rise to a spurious (and unphysical) source of scalar flux (Pope, 1998). This behavior of the IEM model has also been demonstrated for line sources in homogeneous grid turbulence (Sawford, 2004b). The situation can be remedied by introducing the velocity-conditioned scalar mean $\langle \phi | \mathbf{V} \rangle$, which leads to the IECM model. Often invoked as a desirable property of micromixing models is that the scalar PDF should tend to a Gaussian for homogeneous turbulent mixing (Pope, 2000; Fox, 2003) (i.e. statistically homogeneous scalar field in homogeneous isotropic turbulence). While mathematically a Gaussian does not satisfy the boundedness property of the advection-diffusion scalar transport equation, it is generally assumed that the limiting form of the PDF can be reasonably approximated by a clipped Gaussian. Also, Chatwin (2002, 2004) argued that in most practical cases, where the flow is inhomogeneous, the scalar PDF is better approximated by non-Gaussian functions, which should ultimately converge to a Dirac delta function about the mean, $\delta(\psi - \langle \phi \rangle)$, where $\langle \phi \rangle$ approaches a positive value in bounded domains and zero in unbounded domains.

In fully developed turbulent channel flow the center region of the channel may be considered approximately homogeneous (Brethouwer and Nieuwstadt, 2001; Vrieling and Nieuwstadt, 2003). Thus for a centerline source, up to a certain downstream distance where the plume still lies completely in the center region, the mean scalar field can be described by Taylor's theory of absolute dispersion (Taylor, 1921). Likewise, numerical simulations are expected to reproduce experimental measurements of grid turbulence. According to the theory, the mean-square particle displacement $\langle \mathcal{Y}^2 \rangle$ is related to the autocorrelation function of the Lagrangian velocity $R_L = \langle v(t)v(t') \rangle / \langle v^2 \rangle$ as

$$\langle \mathcal{Y}^2 \rangle = 2 \langle v^2 \rangle \int_0^t \int_0^{t'} R_L(\xi) \mathrm{d}\xi \mathrm{d}t',$$
(4.4)

where it is assumed that in stationary turbulence R_L depends only on the time difference $\xi = t - t'$. Lagrangian statistics such as $R_L(\xi)$ are difficult to determine experimentally. An analytical expression that is consistent with the theoretically predicted behavior of the Lagrangian spectrum in the inertial subrange is (Arya, 1999)

$$R_L(\xi) = \exp\left(-\frac{|\xi|}{T_L}\right),\tag{4.5}$$

where T_L denotes the Lagrangian integral timescale. Substituting Equation (4.5) into Equation (4.4) the following analytical expression can be obtained for the root-mean-square particle displacement

$$\sigma_y^2 = \langle \mathcal{Y}^2 \rangle = 2 \langle v^2 \rangle T_L^2 \left[\frac{t}{T_L} - 1 + \exp\left(-\frac{t}{T_L}\right) \right].$$
(4.6)

This expression can be used to approximate the spread of the plume that is released at the centerline of the channel. As the Lagrangian timescale we take

$$T_L = \frac{2\langle v^2 \rangle}{C_0 \varepsilon},\tag{4.7}$$

where C_0 is usually taken as the Lagrangian velocity structure function inertial subrange constant (Monin and Yaglom, 1975; Sawford, 2006), which ensures consistency of the Langevin equation (2.10) with the Kolmogorov hypothesis in stationary isotropic turbulence (Pope, 2000). In the current case the value of C_0 is defined by Equation (2.12) and is no longer a constant, but depends on the velocity statistics. For the purpose of the current analytical approximation, however, a constant value (0.8) has been estimated as the spatial average of C_0 computed by Equation (2.12). For the cross-stream Reynolds stress $\langle v^2 \rangle$ and the dissipation rate ε their respective centerline values are employed. In analogy with time tin homogeneous turbulence, we define $t = x/\langle U \rangle_c$, where x is the downstream distance from the source and $\langle U \rangle_c$ is the mean velocity at the centerline. Thus the cross-stream mean

Table 4.2: Model constants of the micromixing timescale $t_{\rm m}$ defined by Equation (4.3) for both the IEM and IECM models.

	source location	C_s	C_t	
centerline	$y_{s}/h = 1.0$	$y^+ = 1080$	0.02	0.7
wall	$y_{s}/h = 0.067$	$y^{+} = 72$	1.5	0.001

scalar profiles predicted by Equation (4.4) are obtained from the Gaussian distribution

$$\langle \phi(y) \rangle = \frac{\phi_0}{\langle U \rangle_c \left(2\pi\sigma_y^2\right)^{1/2}} \exp\left[-\frac{(y-y_s)^2}{2\sigma_y^2}\right],\tag{4.8}$$

where ϕ_0 is the source strength and y_s is the cross-stream location of the source.

After the velocity field converged to a statistically stationary state, a passive scalar is continuously released from a concentrated source. Two release cases have been investigated, where the scalar has been released at the centerline $(y_s/h = 1.0)$ and in the close vicinity of the wall $(y_s/h = 0.067)$. The viscous wall region experiences the most vigorous turbulent activity. The turbulent kinetic energy, its production and its dissipation and the level of anisotropy all experience their peak values in this region, see also Figure 4.2 (b). This suggests a significantly different level of turbulent mixing between the two release cases. Accordingly, the constants that determine the behavior of the micromixing timescales have been selected differently. Both the IEM and IECM models have been investigated with the micromixing timescale defined by Equation (4.3) using the model constants displayed in Table 4.2.

The different behavior of the two models is demonstrated in Figure 4.3, which shows mean concentration profiles for the centerline release computed by both the IEM and IECM models together with the analytical Gaussian solution (4.8) and the experimental data of Lavertu and Mydlarski (2005) for turbulent channel flow. Indeed, the downstream evolution



Figure 4.3: Cross-stream mean concentration profiles normalized by their respective peak values at different downstream locations as computed by the (a) IECM and (b) IEM models for the centerline release. Lines – PDF calculation at solid line, x/h = 4.0, dashed line, x/h = 7.4 and dot-dashed line, x/h = 10.8, hollow symbols – analytical Gaussians using Equation (4.8) at \circ , x/h = 4.0; \triangle , x/h = 7.4 and \Box , x/h = 10.8, filled symbols – experimental data of Lavertu and Mydlarski (2005) at \bullet , x/h = 4.0 and \blacktriangle , x/h = 7.4. Also shown, PDFs of scalar concentration fluctuations at (x/h = 7.4, y/h = 1.0) for the (c) IECM and (d) IEM models. Lines – computation, symbols – experimental data.

of the cross-stream mean concentration profiles computed by the IECM model follows the Gaussians and is expected to deviate far downstream in the vicinity of the walls, where the effect of the walls is no longer negligible. It is also apparent in Figure 4.3 (b) that the IEM model changes the mean concentration, as expected. As discussed by Lavertu and Mydlarski (2005), the measurements of the mean concentration experience the largest uncertainty due to inaccuracies in estimating the free-stream mean. Also, to improve the signal-to-noise ratio far downstream, a thicker wire had to be employed for measurements performed on the second half of the length considered, i.e. x/h > 11.0. These difficulties are probably the main source of the discrepancy between the experimental data and the agreeing analytical and numerical results for the case of the centerline release. Because of these inconsistencies only results for the first half of the measured channel length (x/h < 11.0) are considered in the current study.

The marginal PDF of scalar concentration can be obtained from the joint PDF $f_{Y}(Y)$ by integrating over the velocity and frequency spaces

$$f_{\phi}(\psi) = \int f_{U\Omega\phi}(\mathbf{V},\omega,\psi) \mathrm{d}\mathbf{V} \mathrm{d}\omega.$$
(4.9)

According to experimental data in grid turbulence (Sawford, 2004b) the skewness at the centerline is expected to be negative close to the source and to become positive only farther downstream. At x/h = 7.4, y/h = 1.0 the temperature PDF measured by Lavertu and Mydlarski (2005) suggests positive skewness in accordance with Sawford's (2004b) data. In Figure 4.3 (c) and (d) the normalized PDFs of scalar concentration fluctuations at this location as computed by both models are depicted. As opposed to the IEM model prediction, both the location of the peak and the overall shape of the PDF are captured correctly by the IECM model.

The different behavior of the two micromixing models is apparent in all one point statistics considered, with the IECM model producing a closer agreement to experimental data. The price to pay for the higher accuracy is an additional 30-40% in CPU time as compared to the IEM model. In the remaining section only the IECM model results are considered.

4.3.3 Scalar statistics with the IECM model

Cross-stream distributions of the first four moments of the scalar concentration at different downstream locations are shown in Figure 4.4 for both release scenarios. The results are compared to experimental data where available.

The mean and root-mean-square (r.m.s.) profiles are normalized by their respective peak values. The width of the mean concentration profiles is most affected by the wall-normal Reynolds stress component $\langle v^2 \rangle$ which is responsible for cross-stream mixing. Due to the underprediction of this component by the velocity model throughout most of the inner layer $(y^+ < 800)$ and the uncertainties in the experimental data mentioned in Section 4.3.2, the mean concentration profiles in Figure 4.4 should be considered at most qualitative.

For the wall-release, the r.m.s. profiles display a clear drift of the peaks towards the centerline with increasing distance from the source Figure 4.4 (f). This tendency has also been observed in turbulent boundary layers by Fackrell and Robins (1982) and Raupach and Legg (1983). Since the scalar is statistically symmetric, in the case of the centerline release, no tranverse drift of the r.m.s. profiles is expected, Figure 4.4 (b). Double peaking of the r.m.s. profiles has been observed in homogeneous turbulence by Warhaft (2000) and Karnik and Tavoularis (1989), noting that the profiles are initially double-peaked close to the source, then single-peaked for a short distance and then again double-peaked far downstream. Lavertu and Mydlarski (2005) found no double peaks in their measurements. Corresponding to the channel flow experiments, the PDF simulation exhibits no double-peaking in the r.m.s. profiles. Applying the projection method to compute $\langle \phi | \mathbf{V} \rangle$ as mentioned in Section 3.6 and described in Section 4.4 results in double peaking of the r.m.s. profiles, which is possibly due to a loss of statistical information due to its Gaussian assumption of the velocity PDF.

Skewness profiles are depicted in Figure 4.4 (c) and (g). For both release cases, near the centers of the plumes the skewness is close to zero, indicating that the PDFs of the scalar



Figure 4.4: See next page for caption.

Figure 4.4: Cross-stream distributions of the first four moments of scalar concentration at different downstream locations for (a)–(d) the centerline release $(y_s/h = 1.0)$ and (e)–(h) the wall release $(y_s/h = 0.067)$. Lines – calculations, symbols – experimental data at solid line, •, x/h = 4.0; dashed line, •, x/h = 7.4 and dot-dashed line, •, x/h = 10.8. The horizontal dashed lines for the skewness and kurtosis profiles indicate the Gaussian values of 0 and 3, respectively. Note the logarithmic scale of the kurtosis profiles.

concentration downstream of the sources are approximately symmetric. Towards the edges of the plumes however, the PDFs become very highly positively skewed, with a sudden drop to zero in the skewness outside of the plume. As observed by Lavertu and Mydlarski (2005), the downstream evolutions of the skewness profiles indicate the eventual mixing of the plume, with the high peaks decreasing. In the current simulations the high skewnesspeaks at the edge of the plumes start increasing first to even higher levels (up to about x/h = 10.0) and only then start decreasing. In the case of the wall-release, the negative skewness in the viscous wall region (also apparent in the experimental data) becomes even more pronounced in the buffer layer and in the viscous sublayer, where experimental data is no longer available. The kurtosis values are close to the Gaussian value of 3 at the crossstream location of the sources, but show significant departures towards the edges of the plume.

Figure 4.5 shows downstream evolutions of the peak of the mean and r.m.s. and the width of the mean concentration profiles. In homogeneous isotropic turbulence and homogeneous turbulent shear flow the decay rate of the peak of the mean concentration profiles is reasonably well described by a power law of the form $\langle \phi \rangle_{\text{peak}} \propto x^n$. In the present inhomogeneous flow Lavertu and Mydlarski (2005), based on the experiments, suggest decay exponents of $n \sim -0.7$ and -0.6 for the wall and centerline sources, respectively. These



Figure 4.5: Downstream evolutions of (a), (b) the peak mean scalar concentration, (c), (d) the width of the mean concentration and (e), (f) the peak of the r.m.s. profiles for the centerline and wall releases, respectively. Solid lines – numerical results, symbols – experimental data.



Figure 4.6: Probability density functions of scalar concentration fluctuations at selected downstream locations for the (a) centerline and (b) wall-releases at the cross-stream location of their respective sources (i.e. y/h = 1.0 and y/h = 0.067, respectively). Lines – calculation, symbols – experimental data at solid line, •, x/h = 4.0 and dot-dashed line, •, x/h = 10.8.

evolutions are compared to experimental data in Figure 4.5 (a) and (b). Downstream evolutions of the width of the mean concentration profiles σ_{mean} are plotted in Figure 4.5 (c) and (d) for the two releases. According to the experimental data, these do not exhibit power-law dependence, as is the case in homogeneous flows. Since the simulations are carried out only on the first half of the measured channel length, the three downstream locations are not sufficient to unambiguously decide whether the simulation data exhibits power-law behavior for the peaks and widths of the mean profiles.

The downstream decay of the peak values of the r.m.s. profiles can be well-approximated by a power-law of the form $\langle \phi'^2 \rangle_{\text{peak}}^{1/2} \propto x^n$, similarly to homogeneous shear flow and isotropic grid-generated turbulence, Figure 4.5 (e) and (f). The experiments suggest n = -1 for both releases.

Probability density functions of scalar concentration fluctuations are depicted in Figure 4.6 for both release cases. The cross-stream location of these PDFs are chosen to coincide with that of their respective sources, i.e. y/h = 1.0 for the centerline release and y/h = 0.067 for the wall-release. Two downstream locations are plotted, at the first and at the third location from the sources measured by Lavertu and Mydlarski (2005), at x/h = 4.0 and x/h = 10.8, respectively. While the PDFs for the centerline release are in reasonable agreement with the experiments, some discrepancies are apparent in the wall-release case. A possible reason behind this disparity is the ad-hoc specification of the mixing timescale in Equation (4.3), which is mostly based on theoretical considerations and experimental observations in homogeneous turbulence.

4.3.4 Conditional statistics

The current model solves for the full joint PDF of the turbulent velocity, frequency and scalar concentration. Therefore we can also examine those quantities that require closure assumptions in composition-only PDF methods. These methods are often used in combustion engineering to model complex chemical reactions in a given turbulent flow or in dispersion modeling in the atmospheric boundary layer. In these cases the simplest approach is to assume the shape of the velocity PDF and numerically solve a set of coupled model equations that govern the evolution of the joint PDF of the individual species concentrations in composition space.

The conservation equation for a single reactive scalar is

$$\frac{\partial \phi}{\partial t} + \boldsymbol{U} \cdot \nabla \phi = \Gamma \nabla^2 \phi + S(\phi(\boldsymbol{x}, t)), \qquad (4.10)$$

where $S(\phi)$ is the chemical source term. In high Reynolds number, constant property flow the PDF of a reactive scalar $g(\psi; \boldsymbol{x}, t)$ is governed by (Dopazo, 1994; Pope, 2000)

$$\frac{\partial g}{\partial t} + \langle U_i \rangle \frac{\partial g}{\partial x_i} = \Gamma \nabla^2 g - \frac{\partial}{\partial x_i} \left(g \langle u_i | \psi \rangle \right) - \frac{\partial^2}{\partial \psi^2} \left(g \left\langle \Gamma \frac{\partial \phi}{\partial x_i} \frac{\partial \phi}{\partial x_i} \middle| \psi \right\rangle \right) - \frac{\partial}{\partial \psi} \left[g S(\psi) \right], \quad (4.11)$$

or alternatively

$$\frac{\partial g}{\partial t} + \frac{\partial}{\partial x_i} \Big[g\left(\langle U_i \rangle + \langle u_i | \psi \rangle \right) \Big] = -\frac{\partial}{\partial \psi} \Big\{ g \Big[\langle \Gamma \nabla^2 \phi | \psi \rangle + S(\psi) \Big] \Big\}.$$
(4.12)

An attractive feature of these formulations is that the usually highly nonlinear chemical source term is in closed form. Closure assumptions, however, are necessary for the velocity fluctuations conditional on the scalar concentration $\langle u_i | \psi \rangle$ and the conditional scalar dissipation $\langle 2\Gamma \nabla \phi \cdot \nabla \phi | \psi \rangle$ or the conditional scalar diffusion $\langle \Gamma \nabla^2 \phi | \psi \rangle$. Since for the current case $S(\phi) = 0$, the marginal scalar PDF $f_{\phi}(\psi)$ defined in Equation (4.9) is equal to g, thus in the following we just use f_{ϕ} .

For the convective term Dopazo (1975) applied the linear approximation

$$\langle u_i | \psi \rangle = \frac{\langle u_i \phi' \rangle}{\langle \phi'^2 \rangle} \left(\psi - \langle \phi \rangle \right), \tag{4.13}$$

to compute the centerline evolution of the temperature PDF in a turbulent axisymmetric heated jet. This linear approximation is exact for joint Gaussian velocity and scalar fluctuations. While many experiments (Bezuglov, 1974; Golovanov, 1977; Shcherbina, 1982; Venkataramani and Chevray, 1978; Sreenivasan and Antonia, 1978) confirm the linearity of the conditional mean velocity around the local mean conserved scalar, Kuznetsov and Sabel'nikov (1990) observe that most of the experimental data show departure from this linear relationship when $|\psi - \langle \phi \rangle|$ is large. Experimental data from Sreenivasan and Antonia (1978) and Bilger *et al.* (1991) also show that in inhomogeneous flows the joint PDF of velocity and scalar is not Gaussian, which makes the above linear approximation dubious in a general case. Nevertheless, this linear model is sometimes applied to inhomogeneous scalar fields because of its simplicity. Another commonly employed approximation is to invoke the gradient diffusion hypothesis

$$-f_{\phi}\langle u_i|\psi\rangle = \Gamma_T \frac{\partial f_{\phi}}{\partial x_i},\tag{4.14}$$

where $\Gamma_T(\boldsymbol{x}, t)$ is the turbulent diffusivity. In the current case, we specify the turbulent viscosity ν_T based on the traditional $k - \varepsilon$ closure and relate it to Γ_T with the turbulent Prandtl number σ_T as

$$\Gamma_T = \frac{\nu_T}{\sigma_T} = \frac{C_\mu}{\sigma_T} \frac{k^2}{\varepsilon},\tag{4.15}$$

where $C_{\mu} = 0.09$ is the usual constant in the $k - \varepsilon$ model and we choose $\sigma_T = 0.8$.

In Figure 4.7 (a) the downstream evolution of the cross-stream velocity fluctuation conditioned on the scalar is depicted for the wall-release case. Both locations are at the height of the source, i.e. y/h = 0.067. The concentration axis for both locations is scaled between their respective local minimum and maximum concentration values, ψ_{\min} and ψ_{\max} . Note that the model curves show higher negative velocity for low-concentration particles as the distance from the source increases. This is expected, since particles deep inside the plume can have very low concentrations only if they did not come from the source but traveled very fast from above, so that they did not have much time to exchange concentration with the source material. As the plume spreads, only particles with stronger negative velocity can maintain their low concentration values. Likewise, as the center of the plume moves towards the centerline of the channel, high-concentration particles also need to have stronger negative velocities to escape from exchange during their journey from the plumecenter to our sensors, which is apparent on the right side of the figure. Obviously, the linear approximation (4.13) cannot be expected to capture the non-linearity of the model curves, but except for extremely low and high concentrations it performs reasonably well. On the other hand, the gradient diffusion approximation is capable of capturing most features of the



Figure 4.7: Cross-stream velocity fluctuation conditioned on the scalar concentration for the wall-release $(y_s/h = 0.067)$. Thick lines, IECM model; thin lines, gradient diffusion approximation of Equation (4.14); straight sloping lines, linear approximation of Equation (4.13). (a) downstream evolution at the height of the source y/h = 0.067: solid lines, x/h = 4.0; dot-dashed lines, x/h = 10.8 and (b) cross-stream evolution at x/h = 7.4: solid lines, y/h = 0.067; dot-dashed lines, y/h = 0.67.

IECM model behavior: it successfully reproduces the non-linearity, with some discrepancy at low and high concentrations. It is also apparent that the numerical computation of the derivatives of the PDFs in the gradient diffusion model (4.14) is most sensitive to sampling errors at the concentration extremes due to lower number of particles falling into the concentration bins there.

The cross-stream evolution of the conditioned velocity fluctuation is shown in Figure 4.7 (b). Both sensors are now at the downstream location x/h = 7.4 with increasing distance from the wall at y/h = 0.067 and 0.67. As the sensor moves towards the channel centerline, the detected low-concentration particles need weaker negative velocity to maintain those low concentrations. The sensor locations relative to the plume centerline can be identified by examining the cross-stream velocity of the high concentration particles. The sensors at y/h = 0.067 and 0.67 are below and above the plume centerline, respectively, since high-concentration particles at these locations possess negative and positive crossstream velocities. As is expected, the linear approximation reasonably represents the model behavior for mid-concentrations, while its performance degrades at locations with higher non-Gaussianity, i.e. towards the edge of the plume. The performance of the gradient diffusion model is reasonable, except at the concentrations extremes.

For the IECM micromixing model, the mean dissipation conditioned on the scalar concentration can be computed from (Sawford, 2004a)

$$\left\langle 2\Gamma \frac{\partial \phi}{\partial x_i} \frac{\partial \phi}{\partial x_i} \middle| \psi \right\rangle f_{\phi} = -\frac{2}{t_{\rm m}} \int_0^{\psi} (\psi' - \tilde{\phi}) f_{\phi}(\psi') \mathrm{d}\psi', \qquad (4.16)$$

where

$$\tilde{\phi}(\psi) = \int \langle \phi | \mathbf{V} \rangle f_{\mathbf{U}\Omega|\phi}(\mathbf{V}, \omega | \psi) \mathrm{d}\mathbf{V} \mathrm{d}\omega, \qquad (4.17)$$

in which the scalar-conditioned PDF is defined as $f_{U\Omega|\phi}(\mathbf{V}, \omega|\psi) \equiv f_{U\Omega\phi}(\mathbf{V}, \omega, \psi)/f_{\phi}(\psi)$. The function $\tilde{\phi}(\psi)$ in Equation (4.17) can be obtained by taking the average of $\langle \phi | \mathbf{V} \rangle$ over those particles that reside in the bin centered on ψ . In other words, the concentration values are first conditioned on the velocity field, which is required to advance the particle concentrations according to the IECM model, then are conditioned again by dividing the concentration sample space into bins and computing separate means for each bin. We use a few bins for the velocity conditioning (N_c) and a significantly higher number of bins (200) for the scalar sample space in order to obtain a higher resolution. The integral in Equation (4.16), however, is more problematic. As Sawford (2004a) notes, numerical integration errors that accumulate at extreme concentrations may be amplified when divided by the scalar PDF approaching zero at those locations. Since the integral over all concentrations vanishes, i.e. $\langle 2\Gamma(\nabla\phi)^2 | \psi_{\max} \rangle f_{\phi}(\psi_{\max}) = 0$, for mid-concentrations it can be evaluated either from the left $(\psi_{\min} \rightarrow \psi)$ or from the right $(\psi_{\max} \rightarrow \psi)$. Thus the integration errors at the concentration extremes can be significantly decreased by dividing the domain into two parts, integrating the left side from the left and the right side from the right and merging the two results in the division-point. Due to statistical errors, however, the integral over all concentrations may not vanish. In that case, the nonzero value

$$\int_{\psi_{\min}}^{\psi_{\max}} (\psi - \tilde{\phi}) f_{\phi}(\psi) d\psi$$
(4.18)

can be distributed over the sample space by correcting the integrand with the appropriate fraction of this error in each bin.

The conditional mean dissipation for three different downstream locations is depicted in Figure 4.8 for both release cases. As for the conditional velocity, the abscissas here are also scaled between the local ψ_{\min} and ψ_{\max} . The dissipation is normalized by the mixing timescale $t_{\rm m}$ and the square of the mean scalar peak $\langle \phi \rangle_{\rm peak}^2$ at the corresponding downstream locations. Note that in the case of the wall-release, the dissipation curves are an order of magnitude lower than in the centerline release case. This is mainly a result of the choice of the different micromixing model constants, especially C_t .

In the case of the wall release, the curves exhibit bi-modal shapes at all three downstream locations. This tendency has also been observed by Kailasnath *et al.* (1993) in the wake of a cylinder and by Sawford (2006) in a double-scalar mixing layer and, to a lesser extent, also in homogeneous turbulence (Sawford, 2004a). Sardi *et al.* (1998) suggest that in assumed-PDF methods of turbulent combustion a qualitative representation of the conditional dissipation can be obtained in terms of the inverse PDF. To examine this relationship, the corresponding scalar PDFs are also plotted in Figure 4.8 with the same scaling on the concentration axis as the dissipation curves. It is apparent that these results support this reciprocal connection



Figure 4.8: IECM model predictions for the mean scalar dissipation conditioned on the concentration for (a) the centerline release $(y_s/h = 1.0)$ and (b) the wall-release $(y_s/h = 0.067)$ at different downstream locations: solid line, x/h = 4.0; dashed line, x/h = 7.4 and dot-dashed line, x/h = 10.8. The cross-stream locations are the same as the respective source positions. Note the different scales for the dissipation curves between the different releases. Also shown are the scalar PDFs at the same locations for both releases in (c) and (d), respectively.

except at the extremes: high values of the PDF correspond to low dissipation (and vice versa). This can be observed for both releases, but it is most visible in the wall-release case, where the mid-concentration minimum between the two maxima of the bi-modal dissipation curves correspond to the peaks in the PDFs.

The IECM model (2.26) implies a model for the mean diffusion conditioned on the scalar



Figure 4.9: Mean scalar diffusion conditioned on the concentration as predicted by the IECM and IEM models for (a) the centerline release $(y_s/h = 1.0)$ and (b) the wall-release $(y_s/h = 0.067)$ at different downstream locations. The cross-stream locations are the same as the respective source positions. Solid line, x/h = 4.0; dashed line, x/h = 7.4 and dot-dashed line, x/h = 10.8. The straight lines are the linear predictions of the IEM model of Equation (4.20).

concentration as

$$\langle \Gamma \nabla^2 \phi | \boldsymbol{V}, \psi \rangle = -\frac{1}{t_{\rm m}} (\psi - \langle \phi | \boldsymbol{V} \rangle).$$
 (4.19)

The downstream evolution of the conditional diffusion is depicted in Figure 4.9 for both releases. The concentration axes are scaled as before and the curves are normalized by the scalar variance $\langle \phi'^2 \rangle$, the concentration at the source, ϕ_0 and the mean unconditioned dissipation $\chi = \langle 2\Gamma(\nabla \phi)^2 \rangle$, which is computed by integrating Equation (4.16) over the whole concentration space. Also shown are the predictions according to the IEM model, which is given by the linear relationship (Sawford, 2006)

$$\frac{\left\langle \Gamma \nabla^2 \phi | \psi \right\rangle \left\langle \phi'^2 \right\rangle}{\chi \phi_0} = \frac{1}{2} \left(\frac{\left\langle \phi \right\rangle}{\phi_0} - \psi \right). \tag{4.20}$$

Far downstream as the scalar gets better mixed, the predictions of the IEM and IECM models get closer. This behavior has been observed for other statistics, as well as for other flows such as the double-scalar mixing layer (Sawford, 2006). Kailasnath *et al.* (1993) report experimental data on similar shapes for the conditional diffusion in the turbulent wake of a cylinder.

4.4 The effect of numerical parameters on the results

Previous PDF modeling studies of channel flow in conjunction with elliptic relaxation have been reported at $Re_{\tau} = 395$ (Dreeben and Pope, 1998) and $Re_{\tau} = 590$ (Wacławczyk *et al.*, 2004) based on the friction velocity u_{τ} and the channel half-width h. These works concentrate on model development and employ different methodologies with different model constants and numerical methods, which inevitably result in a different balance of model behavior and numerical errors. To assess the prediction at different Reynolds numbers the current model has been run at $Re_{\tau} = 392$, 642 and 1080 using the model constants displayed in Table 4.1. The velocity statistics for all three cases are depicted in Figure 4.10. The mean velocity is well represented in the viscous sublayer $(y^+ < 5)$ for all three Reynolds numbers. In the buffer layer $(5 < y^+ < 30)$ there is a slight departure from the DNS data as the Reynolds number increases and from $y^+ > 30$, where the log-law should hold, there exists approximate self-similarity, i.e. the universal slopes of the profiles are equally wellrepresented with a slight underprediction far from the wall at higher Reynolds numbers. The viscous wall region $(y^+ < 50)$ contains the highest turbulent activity, where production, dissipation, turbulent kinetic energy and anisotropy reach their peak values. The location of the peaks of the Reynolds stress components are successfully captured by the model at all three Reynolds numbers with their intensity slightly underpredicted. Previous studies using elliptic relaxation in the Reynolds stress framework (i.e. Eulerian RANS models) report excellent agreement for these second-order statistics (Durbin, 1993; Whizman et al.,



Figure 4.10: See next page for caption.

Figure 4.10: Cross-stream velocity statistics for fully developed turbulent channel flow at (first column) $Re_{\tau} = 392$, (middle column) $Re_{\tau} = 642$ and (right column) $Re_{\tau} = 1080$. Lines – PDF calculation, symbols – DNS data of Moser *et al.* (1999), Iwamoto *et al.* (2002) and Abe *et al.* (2004) (scaled from $Re_{\tau} = 1020$), respectively. First two rows – mean streamwise velocity, third row – normal Reynolds stresses, fourth row – shear Reynolds stress and fifth row – rate of dissipation of turbulent kinetic energy. All quantities are normalized by the friction velocity u_{τ} and the channel half-width h.

1996). Wacławczyk *et al.* (2004) also achieve very good agreement with DNS data using a different version of a PDF model than the one applied here. A common characteristic of PDF models is the slight overprediction of the wall-normal Reynolds stress component $\langle v^2 \rangle$ far from the wall. This component is responsible for the cross-stream mixing of a transported scalar released into a flow far from a wall. Therefore in applications where the mean concentration of scalars is important this quantity must be adequately captured. To improve on this situation we introduced a slight modification into the computation of the characteristic lengthscale L in the elliptic relaxation, Equation (2.17) as

$$L = C_L \max\left[C_{\xi} \frac{k^{3/2}}{\varepsilon}, C_{\eta} \left(\frac{\nu^3}{\varepsilon}\right)^{1/4}\right], \qquad (4.21)$$

with $C_{\xi} = 1.0 + 1.3 n_i n_i$, where n_i is the unit wall-normal of the closest wall-element pointing outward of the flow domain. This only affects the diagonal Reynolds stresses which can be seen in Figure 4.11 for the different Reynolds numbers. Decreasing $\langle v^2 \rangle$ at the centerline changes the relative fraction of energy distributed among the diagonal components of the Reynolds stress tensor, consequently the other two components, $\langle u^2 \rangle$ and $\langle w^2 \rangle$, are slightly increased. Obviously, these kind of flow-dependent modifications in the turbulence model are of limited value, since their effects in a general setting may not be easily predictable.



Figure 4.11: The effect of the modification of the characteristic lengthscale in Equation (2.17) on the diagonal components of the Reynolds stress tensor by employing the additional model constant $C_{\xi} \neq 1$ at (first column) $Re_{\tau} = 392$, (middle column) $Re_{\tau} = 642$ and (right column) $Re_{\tau} = 1080$. Thick lines, $C_{\xi} = 1.0 + 1.3n_in_i$; thin lines, $C_{\xi} = 1.0$; symbols, DNS data as in Figure 4.10.

The only nonzero shear stress component $\langle uv \rangle$ in this flow and the turbulent kinetic energy dissipation rate ε are both in very good agreement with DNS data and even improve as the Reynolds number increases. It is apparent in both Figures 4.10 and 4.11 that the overall prediction of second order statistics improve as the Reynolds number increases. This tendency is expected to continue as the underlying high-Reynolds-number modeling assumptions become better fullfilled.

Into the fully developed flow, a passive scalar has been released from a concentrated source at the channel centerline. A general numerical procedure that can be used to compute the velocity-conditioned scalar mean $\langle \phi | \mathbf{V} \rangle$ in the IECM model has been described

in Section 3.6 and Appendix B. Another method based on the projection of the threedimensional velocity field onto a one-dimensional subspace, where the sample-spatial discretization can be carried out, has been developed and tested in homogeneous turbulence by Fox (1996). In that method, the projected velocity of a particle is found from

$$\mathcal{U}_{\rho} = \alpha_i \mathcal{U}_i, \tag{4.22}$$

where the projection vector α_i is obtained from the following linear relationship

$$\rho_i = \rho_{ij} \alpha_j \tag{4.23}$$

between the normalized velocity-scalar vector and the velocity-correlation tensor (no summation on greek indices)

$$\rho_{\alpha} = \frac{\langle u_{\alpha} \phi' \rangle}{\langle u_{\alpha}^2 \rangle^{1/2} \langle \phi'^2 \rangle^{1/2}}, \qquad \rho_{\alpha\beta} = \frac{\langle u_{\alpha} u_{\beta} \rangle}{\langle u_{\alpha}^2 \rangle^{1/2} \langle u_{\beta}^2 \rangle^{1/2}}, \qquad (4.24)$$

where $\phi' = \psi - \langle \phi \rangle$ denotes the scalar fluctuation. This projection method has been developed (and is exact for) Gaussian velocity PDFs, although it can still be used in inhomogeneous flows with the assumption that the local joint PDF of velocity is not too far from an approximate joint normal distribution. In order to assess the performances and the difference in the predictions, we implemented and compared both methods and tested them with different number of conditioning bins.

To investigate how the choice of the number of conditioning intervals N_c affects the solution with the projection method, several runs have been performed at the highest Reynolds number ($Re_{\tau} = 1080$) with different values for N_c . Some of the unconditional and conditional statistics of the joint PDF are depicted in Figure 4.12. Note that employing $N_c=1$ corresponds to the special case of the IEM model, Equation (2.25). It is apparent that applying only a few intervals already makes a big difference compared to the IEM model



Figure 4.12: Scalar statistics affected by the number of conditioning intervals N_c with computing the velocity-conditioned mean $\langle \phi | \mathbf{V} \rangle$ applying Fox's projection method using Equations (4.22)-(4.24). (a) Cross-stream distribution of the scalar mean at x/h = 4.0, (b) PDF of scalar concentration fluctuations at (x/h = 4.0, y/h = 1.0), (c) mean scalar dissipation conditioned on the concentration at (x/h = 4.0, y/h = 1.0) and (d) mean scalar diffusion conditioned on the concentration at (x/h = 4.0, y/h = 1.0). Dashed line – $N_c=1$ (IEM), dotted line – $N_c=3$, solid line – $N_c=5$, dot-dashed line – $N_c=20$. Symbols on (a) analytical Gaussians according to Taylor (1921) and on (b) experimental data of Lavertu and Mydlarski (2005).

in correcting the prediction of the mean concentration and the PDF of concentration fluctuations also moves towards the experimental data. Increasing N_c may be thought as an approach to increase the resolution of the conditioning (thus better exploiting the advantages of the IECM over the IEM model), however, as Fox (1996) points out, this is of limited value, since the decreasing number of particles per interval increases the statistical error. The current test simulations have been carried out with an initial 500 particles per element and the total number of particles did not change during simulation. Figure 4.12 shows that above $N_c=5$ there is no significant change in the statistics and even at $N_c=20$ the results do not deteriorate. Also displayed in Figure 4.12 are the centerline normalized mean scalar dissipation and diffusion both conditional on the scalar concentration, $\langle \Gamma(\nabla \phi)^2 | \psi \rangle$ and $\langle \Gamma \nabla^2 \phi | \psi \rangle$, respectively. As before, the concentration axes in Figure 4.12 (c) and (d) are scaled between the local minimum and maximum concentration values, ψ_{\min} and ψ_{\max} , in order to zoom in on the interesting part of the concentration space. Using Fox's projection method, the choice of number of conditioning intervals on the velocity space (N_c) has a similar effect on the conditional dissipation and diffusion: they also support the earlier observation that the optimal number of conditioning intervals is at about $N_c=3-5$ to attain convergence.

A different picture reveals itself however, when $\langle \phi | \mathbf{V} \rangle$ is computed with the current method instead of the projection that assumed Gaussianity of the underlying velocity field. The same statistics as shown in Figure 4.12 are plotted in Figure 4.13 for different numbers of conditioning bins, but without employing the projection to compute $\langle \phi | \mathbf{V} \rangle$. The mean profiles do not behave significantly differently, which underlines the earlier observation that employing only a few conditioning bins can already correct the prediction of the mean compared to the IEM model. The PDFs however show significantly higher spikes when compared to their counterparts with projection. The prediction of the conditional dissipation profiles are also different (overall they range about 150% higher) as opposed to that with projection, while the conditional diffusion curves exhibit similar behavior both with and without projection. Figures 4.13 (b-d) also reveal that the currently employed



Figure 4.13: Scalar statistics affected by the number of conditioning intervals when computing the velocity-conditioned mean $\langle \phi | \mathbf{V} \rangle$ with the method described in Section 3.6 and Appendix B. The quantities are the same as in Figure 4.12. Dashed line – $N_c=1$ (IEM), dot-dashed line – $N_c=(3 \times 3 \times 3)$, solid line – $N_c=(5 \times 5 \times 5)$.

finest conditional binning structure of $(5 \times 5 \times 5)$ with an initial 500 particles per element is still not sufficient to achieve convergence for the PDF and these conditional statistics. It is also worth noting, that this is the case for a centerline release and that our sampling location is relatively close to the source and at the centerline, which lies in the "approximately homogeneous" region of the flow.

To examine the effect of the number of particles on the solution, several testruns have been performed with different number of particles employing both methods for computing

 $\langle \phi | V \rangle$. At the Reynolds numbers investigated, $Re_{\tau} = 392, 642$ and 1080, we found the minimum number of particles per elements necessary for a numerically stable solution to be $N_{p/e}=80$, 100 and 150, respectively. Increasing $N_{p/e}$ more than these minimum values would not be necessary to obtain a particle-number-independent velocity PDF, since running the simulation employing up to $N_{p/e}$ =500 resulted in negligible change of the velocity statistics investigated. On the other hand, the scalar statistics exhibit significant differences when different number of particles are employed. Figure 4.14 shows unconditional and conditional statistics of the passive scalar field at $Re_{\tau} = 1080$ using different numbers of particles employing the projection method with $N_c=5$. The cross-stream distribution of the first four moments show that the statistical error due to insufficient number of particles becomes higher towards the edge of the plume, where the joint PDF is most skewed. The discrepancy due to this error is more pronounced in the higher-order statistics. The PDFs of concentration fluctuations and the scalar at the centerline, where the flow can be considered approximately homogeneous, is nearly independent of the number of particles. The prediction of accurate conditional statistics usually requires a large number of particles. This is underlined by the mean conditional dissipation and diffusion in Figures 4.14 (g) and (h) in the center region, which show a slight dependence on $N_{p/e}$. In summary, the velocity statistics are predicted independently of the number of particles. With the projection method to compute $\langle \phi | \mathbf{V} \rangle$, the unconditional scalar statistics (including the PDFs) are predicted approximately independently of the number of particles in the homogeneous center region of the channel, however, the conditional statistics examined there still exhibit a slight particle-number-dependence even with $N_{p/e}=500$. We hypothesize that more complex inhomogeneous and highly skewed flows may require even larger number of particles than the currently employed maximum, 500.

In Figure 4.15 the same scalar statistics as in Figure 4.14 are shown but with $\langle \phi | \mathbf{V} \rangle$ computed with the current method instead of projection for different number of particles


Figure 4.14: See next page for caption.

Figure 4.14: Unconditional and conditional statistics of the passive scalar field affected by the number of particles with $\langle \phi | V \rangle$ computed using the projection method of Equations (4.22)-(4.24) using N_c =5. (a)-(d) Cross-stream distribution of the first four moments at x/h = 4.0, (e) PDF of concentration fluctuations at (x/h = 4.0, y/h = 1.0), (f) PDF of concentration at (x/h = 4.0, y/h = 1.0), (g) mean scalar dissipation conditioned on the concentration at (x/h = 4.0, y/h = 1.0) and (h) mean scalar diffusion conditioned on the concentration at (x/h = 4.0, y/h = 1.0). Dashed line – (initial number of particles per elements) $N_{p/e}$ =150, solid line – $N_{p/e}$ =300 and dot-dashed line – $N_{p/e}$ =500. Symbols on (a) analytical Gaussians according to Taylor (1921), on (b), (c), (e) experimental data of Lavertu and Mydlarski (2005). The horizontal dashed line on (d) indicates the Gaussian kurtosis value of 3.

employing a binning structure of $(5 \times 5 \times 5)$. The technique described in Section 3.6 is robust enough to automatically use less conditioning intervals depending on the number of particles in a given element. Thus, when the simulations were run with $N_{p/e}=150$, 300 and 500, the average number of conditioning bins employed throughout the simulation has been automatically reduced to about 57, 100 and 124, respectively, as compared to the prescribed 125. The scalar mean is predicted equally well as with the projection method showing no sign of dependence on the number of particles, Figure 4.15 (a). Interestingly, the r.m.s. curves do not double-peak if the projection is not used, Figure 4.15 (b) and the width also agrees better with the experimental data. Thus the double-peaks on Figure 4.14 (b) may only be artifacts of the projection. Similarly to using projection, the skewness and kurtosis profiles are predicted with significant particle-number dependence at the edges of the plume. This shows that convergence has not yet been reached with $N_{p/e}=500$ for these higher-order statistics. Also, there is a pronounced flattening at the centerline in the skewness and kurtosis profiles using the projection technique, cf. Figures 4.14 and 4.15 (c-d),



Figure 4.15: See next page for caption.

Figure 4.15: Unconditional and conditional statistics of the passive scalar field affected by the number of particles with $\langle \phi | \mathbf{V} \rangle$ computed with the method described in Section 3.6 and Appendix B using a binning structure of $(5 \times 5 \times 5)$. The legend is the same as in Figure 4.14.

which may also be a side-effect of the projection, since no flattening can be observed in the experimental data. The increasing peaks of the PDFs have already been observed before, when we compared the projection method to the general methodology using different values of N_c . Both Figures 4.15 (e) and (f) show that the PDFs have not converged yet, however, these figures may show the combined effect of increasing both $N_{p/e}$ and N_c , since the conditioning algorithm automatically reduces N_c in case of insufficient number of particles in an Eulerian element. Finally, the conditional dissipation and diffusion curves show a very light dependence on the number of particles applied.

We summarize the findings for the PDF algorithm related to a passive scalar released at the centerline of a fully developed turbulent channel flow as follows:

- the prediction of one-point velocity statistics becomes more accurate with increasing Reynolds number,
- a stable numerical solution and a converged velocity field require about 80–150 particles per element depending on the Reynolds number,
- the prediction of higher-order unconditional scalar statistics and concentration fluctuation PDFs are closer to experimental observations without employing the projection technique to compute $\langle \phi | \mathbf{V} \rangle$,
- conditioned statistics may exhibit a large difference (up to 150%) depending on the application of the projection method, however the lack of experimental data currently prevents us to assess the true error in these quantities,

quantity	particles per element				
velocity statistics, $\langle U_i \rangle$, $\langle u_i u_j \rangle$, k, ε	80–150, slightly increasing				
first two scalar moments, $\langle \phi \rangle$, $\langle \phi'^2 \rangle$ third, fourth and higher-order scalar moments, $\langle \phi'^3 \rangle$, $\langle \phi'^4 \rangle$ scalar concentration PDFs, $f_{\phi}(\phi' \langle \phi'^2 \rangle^{1/2})$, $f_{\phi}(\psi)$ mean conditional scalar dissipation, $\langle \Gamma(\nabla \phi)^2 \psi \rangle$ mean conditional scalar diffusion, $\langle \Gamma \nabla^2 \phi \psi \rangle$	with the Reynolds number 150 500+ 500+ 300 150				

Table 4.3: Minimum number of particles per element required to compute different statistics.

- compared to the simpler IEM model, using the IECM model only with a few conditioning intervals already makes a big difference in correcting the prediction of the scalar mean, both with and without the projection method, for an increase in the overall computational cost of about 30–40%,
- the difference in computational costs of the projection and the current general method used to compute $\langle \phi | \mathbf{V} \rangle$ is negligible,
- with projection, full convergence in the higher-order scalar statistics may require more particles than $N_{p/e}$ =500, while N_c =3–5 was enough to reach convergence in all quantities investigated,
- without projection, full convergence in the higher-order scalar statistics and PDFs may require more particles than $N_{p/e}$ =500, while the binning structure of $(5 \times 5 \times 5)$ was enough to reach convergent unconditional statistics, but this was still not a sufficient conditioning-resolution to achieve convergent concentration PDFs and conditional statistics.

Table 4.3 lists the minimum number of particles per element necessary to accurately compute the one-point statistics investigated in this Chapter.

4.5 Computational cost

Quantitative assessments of the computational cost of PDF methods are sparse in the literature. There is no dedicated study to compare the different stand-alone and hybrid methods side by side or to compare PDF methods to other turbulence modeling techniques. However cost comparisons are useful even if they only provide limited information and are done between different methods at different levels of approximation.

The computational cost of a simulation (the time required to reach convergence with a given accuracy) is largely determined by the resolution requirements, which in the case of a turbulent channel flow mostly amounts to adequately resolving the boundary layer. In an attempt to quantify the increase in cost of the current PDF methodology, several runs have been carried out at different Reynolds numbers between $Re_{\tau} = 100$ and 1080. In all cases only the statistically one-dimensional velocity field has been computed reaching a statistically stationary state, without a scalar release and micromixing. As Re_{τ} is increased, the boundary layer becomes thinner and a finer Eulerian grid is needed to resolve the statistics, which inevitably results in the increase of the number of particles as well. Accordingly, keeping the Courant-number approximately constant, the size of the timestep has to be decreased to achieve the same level of accuracy and stability with increasing Reynolds numbers. This tendency can be examined in Figure 4.16 (a), where the key factors affecting the computational cost vs. Re_{τ} are depicted. These are the smallest element (gridsize), the characteristic flow speed $\langle U \rangle_c / u_\tau$, where $\langle U \rangle_c$ is the mean velocity at the centerline, and the total number of elements N_e or equivalently, the total number of particles N_p . All filled symbols on Figure 4.16 represent the given quantity normalized by the quantity at $Re_{\tau} = 100$. To an approximation, the number of floating-point operations, i.e. the computational cost, is proportional to the number of elements (and the number of particles) and the flow speed and inversely proportional to the gridsize (and the size of the



Figure 4.16: Computational cost of (a) a measured one-dimensional and (b) an extrapolated three-dimensional PDF simulation. Filled symbols and solid lines – PDF calculations, hollow symbols and dashed lines – DNS of channel flow.

timestep). Based on the slope of these three factors on a log-log scale, the approximate slope of the computational cost for the one-dimensional PDF simulation of channel flow can be estimated as

$$\frac{Re_{\tau}^{0.58} \times Re_{\tau}^{0.1}}{Re_{\tau}^{-0.88}} = Re_{\tau}^{1.56}.$$
(4.25)

This approximation based on the three key factors is in reasonable agreement with the measured slope, $Re_{\tau}^{1.72}$, which is based on actual timings. Employing the same arguments, the cost of a three-dimensional PDF simulation may be extrapolated as

$$Re_{\tau}^{1.72} \times Re_{\tau}^{2 \times 0.58} = Re_{\tau}^{2.88}, \tag{4.26}$$

which is displayed in Figure 4.16 (b). For comparison, the slope of the number of required elements for DNS simulations of turbulent channel flow is also displayed, based on the data reported by Abe *et al.* (2004), normalized by N_e at $Re_{\tau} = 180$. This gives the slope of $Re_{\tau}^{2.88}$ which reasonably agrees with $Re_{\tau}^{2.7}$, the prediction of Reynolds (1990) for the total

number of modes required for DNS of channel flow. (For comparison the cost of DNS in homogeneous turbulence grows as $Re_L^{2.25}$ (Pope, 2000) based on the turbulence Reynolds number $Re_L = k^2/(\varepsilon\nu)$). Based on the slope of N_e and Equation (4.25) we approximate the increase in computational cost of DNS for the inhomogeneous channel flow as

$$\frac{Re_{\tau}^{2.88} \times Re_{\tau}^{0.1}}{Re_{\tau}^{-0.88}} = Re_{\tau}^{3.86}.$$
(4.27)

Now we are in a position to quantitatively compare the computational requirements of a three-dimensional PDF to DNS simulations as it is displayed in Figure 4.16 (b). A DNS simulation provides a great wealth of information on the turbulence for a steeply increasing cost at high Reynolds numbers by fully resolving all scales, including dissipation. A statistical technique, such as the current PDF method, approximates certain physical processes, thus it is expected to be less accurate. However, since it does not need to resolve the finest scales, it may be less computationally intensive. Based on Figure 4.16 we observe that a three-dimensional PDF simulation will probably not be as expensive for higher Reynolds numbers as DNS. As depicted in Figure 4.16 (b), the difference in computational cost between DNS and the three-dimensional PDF method is about a decade computing a fully resolved boundary layer at the Reynolds number $Re_{\tau} = 1080$. This means that at this Reynolds number DNS will produce the desired result in 10 times more computing hours than the PDF method. The figure also shows that extrapolating this result to more realistic Reynolds numbers will result in even larger differences in computational costs, DNS being increasingly more expensive than the current PDF method. As an example, resolving the boundary layer at $Re_{\tau} = 10^4$ will take 100 times more CPU time with DNS than with the PDF method.

It is worth noting that the cost in the current case largely amounts to adequately resolving the boundary layer. In general, any method that attempts to fully resolve the boundary layer will have to pay because of the required high resolution, and not necessarily because the method itself is inherently expensive. In a sense, the above assessment is even a bit unfair towards the PDF method since resolving walls is not its main advantage or purpose (although it still performs relatively well in comparison). Also, we estimate that our accounting for the increase in cost due to extrapolating from one to three spatial dimensions, Equation (4.26), is rather conservative, i.e. overpredictive – while the Eulerian statistics are only extracted on a one-dimensional grid, all three components of the particle velocity are already retained, which consitutes as the majority of the computational cost, as it is shown in Section 3.11.

We did not perform comparisons with other methods. A hybrid LES/FDF method for scalars may be expected to have a higher predictive power than the current stand-alone PDF method. However, we do not exclude that a stand-alone PDF method could be less expensive than a hybrid LES/FDF method, since resolution requirements may not have to be as stringent to achieve resolution-independent statistics. The Eulerian LES solution should be filter width and grid independent, which occurs only if a sufficient portion of the turbulent kinetic energy is resolved, i.e. in the case if only the dissipative scales are modeled and the majority of the inertial subrange and energy containing range is resolved. On the other hand, one-point statistical models, like the current stand-alone PDF method, do not need to resolve scales much below the integral scale (Pope, 2000).

Overall, the above assessment of the computational cost certainly cannot be taken in the most general sense as it is based on one simple flow, the fully developed turbulent channel flow, it extrapolates and compares to a method (DNS) that is quite different in both formulation and the results it obtains. Therefore reaching a final conclusion regarding the cost of the methodology is premature. Further assessments based on more flow topologies are needed to provide a better understanding of the computational cost of PDF methods compared to other methods, such as LES and other statistical approaches.

4.6 Discussion

In this Chapter, the previously described stand-alone PDF method has been tested and validated against DNS, analytical and experimental data, computing the dispersion of passive scalars in fully developed turbulent channel flow. The complete PDF-IECM model computes the joint PDF of turbulent velocity, frequency and scalar concentration where the scalar is released from concentrated sources. The flow is represented by a large number of Lagrangian particles and the governing stochastic differential equations have been integrated in time in a Monte-Carlo fashion. The high anisotropy and inhomogeneity at the low-Reynolds-number wall-region have been captured through the elliptic relaxation technique, explicitly modeling the vicinity of the wall down to the viscous sublayer by imposing only the no-slip condition. Durbin (1993) suggested the simple LRR-IP closure of Launder, Reece, and Rodi (1975), originally developed in the Eulerian framework, as a local model used in the elliptic relaxation equation (2.13). Since then, several more sophisticated local Reynolds stress models have been investigated in conjunction with the elliptic relaxation technique (Whizman et al., 1996). In the PDF framework, the Lagrangian modified IP model of Pope (1994) is based on the LRR-IP closure. We introduced an additional model constant C_{ξ} in the definition of the characteristic lengthscale L (2.17) whose curvature determines the behavior of the relaxation and, ultimately, the overall performance of the model in representing the Reynolds stress anisotropy. This resulted in a correction of the original model overprediction of the wall-normal component $\langle v^2 \rangle$ far from the wall, which crucially influences the cross-stream mixing of the transported scalar. However, increasing the constant C_{ξ} adversely affects the level of anisotropy that can be represented by the technique. A more accurate treatment of the Reynolds stresses and scalar mixing should be achieved by a more elaborate second moment closure, such as the nonlinear C-L model of Craft and Launder (1991) or the Lagrangian version of the SSG model of Speziale, Sarkar, and Gatski (1991) suggested by Pope (1994).

An unstructured triangular grid is used to compute Eulerian scalar statistics and to track particles throughout the domain. The main purpose of employing unstructured grids has been to prepare the methodology for more complex flow geometries. A similar particle-in-cell approach has been developed by Muradoglu *et al.* (1999, 2001); Jenny *et al.* (2001); Zhang and Haworth (2004); Rembold and Jenny (2006) and by Ge *et al.* (2007) for the computation of turbulent reactive flows. These approaches combine the advantages of traditional Eulerian computational fluid dynamics (CFD) codes with PDF methods in a hybrid manner. Our aim here is to develop a method that is not a hybrid one, so the consistency between the computed fields can be naturally ensured. The emphasis is placed on generality, employing numerical techniques that assume as little as possible about the shape of the numerically computed joint PDF.

We compared the performance of the IEM and the IECM micromixing models in an inhomogeneous flow with strong viscous effects by modeling both the turbulent velocity field and the scalar mixing. The more sophisticated IECM model provides a closer agreement with experimental data in channel flow for the additional computational expense of 30-40% compared to the IEM model.

Several conditional statistics that often require closure assumptions in PDF models where the velocity field is assumed were extracted and compared to some of their closures. In particular, our conclusions suggest that the scalar-conditioned velocity is well approximated by a linear assumption for mid-concentrations at locations where the velocity PDF is moderately skewed. The gradient diffusion approximation, however, captures most features including the nonlinearity and achieves a closer agreement with the IECM model in slightly more skewed regions of the flow as well. At local concentration extremes and in extremely skewed regions the gradient diffusion approximation markedly departs from the IECM model. The mean scalar dissipation conditioned on the scalar concentration may be well-approximated by the inverse relationship suggested by Sardi *et al.* (1998) in inhomogeneous flows with significant viscous effects as well, except at the concentration extremes. In computing the conditional scalar diffusion, both the IEM and the IECM models produce similar slopes due to the same scalar dissipation rate attained.

The effects of several numerical parameters on the computed results have also been investigated. We found that about a hundred particles per element are enough for a stable numerical solution. However, even 500 particles per element were not enough to obtain particle-number-independent higher-order scalar statistics. Moreover, to obtain accurate higher-order scalar statistics and concentration fluctuation PDFs in inhomogeneous flows, the use of the currently proposed method is advised to compute $\langle \phi | \mathbf{V} \rangle$ as opposed to the projection method assuming Gaussianity.

Chapter 5:

Street canyon simulations: results and discussion

5.1 Introduction

Regulatory bodies, architects and town planners increasingly use computer models in order to assess ventilation and occurrences of hazardous pollutant concentrations in cities. These models are mostly based on the Reynolds-averaged Navier-Stokes (RANS) equations or, more recently, large eddy simulation (LES) techniques. Both of these approaches require a series of modeling assumptions, including most commonly the eddy-viscosity and gradientdiffusion hypotheses. The inherent limitations of these approximations, even in the simplest engineering flows, are well known and detailed for example by Pope (2000). Therefore, there is a clear need to develop higher-order models to overcome these shortcomings. In pollutant dispersion modeling it is also desirable to predict extreme events like peak values or probabilities that concentrations will exceed a certain threshold. In other words, a fuller statistical description of the concentration is required (Chatwin and Sullivan, 1993; Kristensen, 1994; Wilson, 1995; Pavageau and Schatzmann, 1999). These issues have been explored in the unobstructed atmosphere and models capable of predicting these higherorder statistics have also appeared (Franzese, 2003; Cassiani *et al.*, 2005a,b), but more research is necessary to extend these capabilities to cases of built-up areas.

Probability density function (PDF) methods have been developed mainly within the

combustion engineering community as an alternative to moment closure techniques to simulate chemically reactive turbulent flows (Lundgren, 1969; Pope, 1985; Dopazo, 1994). Because many-species chemistry is high-dimensional and highly nonlinear, the biggest challenge in reactive flows is to adequately model the chemical source term. In PDF methods, the closure problem is raised to a statistically higher level by solving for the full PDF of the turbulent flow variables instead of its moments. This has several benefits. Convection, the effect of mean pressure, viscous diffusion and chemical reactions appear in closed form in the PDF transport equation. Therefore these processes are treated mathematically exactly without closure assumptions eliminating the need for gradient-transfer approximations. The effect of fluctuating pressure, dissipation of turbulent kinetic energy and small-scale mixing of scalars still have to be modeled. The rationale is that since the most important physical processes are treated exactly, the errors introduced by modeling assumptions for less important processes amount to a smaller departure from reality. Moreover, the higher level description provides more information which can be used in the construction of closure models.

The PDF transport equation is a high-dimensional scalar equation. Therefore all techniques of solution rely on Monte Carlo methods with Lagrangian particles representing a finite ensemble of fluid particles, because the computational cost of Lagrangian Monte Carlo methods increases only linearly with increasing problem dimensionality, favourably comparing to the more traditional finite difference, finite volume or finite element methods. The numerical development in PDF methods has mainly centered around three distinctive approaches. A common numerical approach is the *standalone Lagrangian* method, where the flow is represented by particles whereas the Eulerian statistics are obtained using kernel estimation (Pope, 2000; Fox, 2003). Another technique is the *hybrid* methodology, which builds on existing Eulerian computational fluid dynamics (CFD) codes based on moment closures (Muradoglu *et al.*, 1999, 2001; Jenny *et al.*, 2001; Rembold and Jenny, 2006). Hybrid methods use particles to solve for certain quantities and provide closures for the Eulerian moment equations using the particle/PDF methodology. A more recent approach is the self-consistent *non-hybrid* method (Bakosi *et al.*, 2007, 2008), which also employs particles to represent the flow, and uses the Eulerian grid only to solve for inherently Eulerian quantities (like the mean pressure) and for efficient particle tracking. Since the latter two approaches extensively employ Eulerian grids, they are particle-in-cell methods (Grigoryev *et al.*, 2002).

After an extensive testing of the methodology in a relatively simple setting, the fully developed turbulent channel flow (Chapter 4), the current Chapter presents an application of the non-hybrid method to a simplified urban-scale case where pollution released from a concentrated line source between idealized buildings is simulated and results are compared to wind-tunnel experiments.

PDF methods in atmospheric modeling have mostly been focused on simulation of passive pollutants, wherein the velocity field (mean and turbulence) is assumed or obtained from experiments (Sawford, 2004b, 2006; Cassiani *et al.*, 2005a,b, 2007a). Instead, the current model directly computes the joint PDF of the turbulent velocity, characteristic turbulent frequency and scalar concentration, thus it extends the use of PDF methods in atmospheric modeling to represent more physics at a higher statistical level. Computing the full joint PDF also has the advantage of providing information on the uncertainty of the simulation on a physically sound basis.

In this Chapter the turbulent boundary layers developing along solid walls are treated in two different ways: either fully resolved or via the application of wall-functions (i.e. the logarithmic "law of the wall"). The full resolution is obtained using Durbin's elliptic relaxation technique (Durbin, 1993), which was incorporated into the PDF methodology by Dreeben and Pope (1997a, 1998). This technique allows for an adequate representation of the near-wall low-Reynolds-number effects, such as the high inhomogeneity and anisotropy of the Reynolds stress tensor and wall-blocking. Wall-conditions for particles based on the logarithmic "law of the wall" in the PDF framework have also been developed by Dreeben and Pope (1997b). These two types of wall-treatments are examined in terms of computational cost / performance trade-off, addressing the question of how important it is to adequately resolve the boundary layers along solid walls in order to obtain reasonable scalar statistics.

At the urban scale the simplest settings to study turbulent flow and dispersion patterns are street canyons. Due to increasing concerns for environmental issues and air quality standards in cities, a wide variety of canyon configurations and release scenarios have been studied both experimentally (Hoydysh *et al.*, 1974; Wedding *et al.*, 1977; Rafailids and Schatzmann, 1995; Meroney *et al.*, 1996; Pavageau and Schatzmann, 1999) and numerically (Lee and Park, 1994; Johnson and Hunter, 1995; Baik and Kim, 1999; Huang *et al.*, 2000; Liu and Barth, 2002). Street canyons have a simple flow geometry, they can be studied in two dimensions and a wealth of experimental and modeling data are available for different street-width to building-height ratios. This makes them ideal candidates for testing a new urban pollution dispersion model. We validate the computed velocity and scalar statistics with the LES simulation results of Liu and Barth (2002) and the wind tunnel measurements of Meroney *et al.* (1996), Pavageau (1996) and Pavageau and Schatzmann (1999). The experiments have been performed in the atmospheric wind tunnel of the University of Hamburg, where the statistics of the pollutant concentration field have been measured in an unusually high number of locations in order to provide fine details inside the street canyon.

The Chapter is organized as follows. In Section 5.2 the specifics of the boundary conditions related to the street canyon are outlined. Several statistics computed using both full wall-resolution and wall-functions are compared to experimental data and large eddy simulation in Section 5.3. Finally, Section 5.4 draws some conclusions and elaborates on possible future directions.

5.2 Modeling specifics of the street canyon

The governing equations for both full wall-resolution and wall-functions cases together with boundary conditions have been described in Chapter 2 and Section 3.9, respectively. In Chapter 3 we also elaborated on several aspects of the numerical techniques that are used to solve the equations. Thus here, only certain specific details that directly relate to the modeling of the street canyon case are described.

The flow geometry can be modeled as statistically two-dimensional if we suppose that the buildings are sufficiently long, like a long street. The particle copying-mirroring strategy used for the channel flow cannot be used here, so the general algorithm is applied. An additional complexity is the computation of the mean pressure in a general way, applying the pressure projection described in Section 3.3. A non-homogeneous Neumann wall-boundary condition for the pressure projection (3.8) has been described in Section 3.9 for both full wall-resolution and wall-functions representations of no-slip walls. The flow is expected to reach a statistically steady state and is driven by a mean-pressure difference between its inflow and outflow. This condition in the free stream (above the buildings) is imposed on the mean pressure as follows.

Assuming that the inflow and outflow are aligned with y, as shown in Figure 5.1, the two-dimensional steady state cross-stream mean-momentum equation holds

$$\frac{1}{\rho}\frac{\partial\langle P\rangle}{\partial y} = -\langle U\rangle\frac{\partial\langle V\rangle}{\partial x} - \langle V\rangle\frac{\partial\langle V\rangle}{\partial y} + \nu\left(\frac{\partial^2\langle V\rangle}{\partial x^2} + \frac{\partial^2\langle V\rangle}{\partial y^2}\right) - \frac{\langle uv\rangle}{\partial x} - \frac{\partial\langle v^2\rangle}{\partial y}.$$
(5.1)

If the inflow and outflow are far enough from the canyon, the flow can be assumed to be an undisturbed turbulent channel flow. Hence we can neglect all terms on the right hand side of Equation (5.1), with the exception of the last term. Thus the inflow and outflow conditions for the mean pressure can be specified according to Equation (4.2). Flow-dependent non-homogeneous Dirichlet conditions have to be imposed in a way that the streamwise gradient

 $\partial \langle P \rangle / \partial x$ is kept at a constant level. This can be achieved by specifying the values of $\langle P \rangle$ at the inflow and outflow based on $\langle P \rangle = -\rho \langle v^2 \rangle$, which will equate their cross-stream derivatives as well. The streamwise gradient $\partial \langle P \rangle / \partial x = \text{const.}$ is applied by shifting up the values of $\langle P \rangle$ at the inflow. Consistently with Equation (3.8) the above condition has to be imposed on the mean-pressure difference in time, $\delta \langle P \rangle = \langle P \rangle^{n+1} - \langle P \rangle^n$. Thus we arrive at the inflow/outflow conditions

$$\delta \langle P \rangle = \begin{cases} -\Delta P \cdot L_x - \rho \langle v^2 \rangle - \langle P \rangle^n, & \text{for inflow points,} \\ \\ -\rho \langle v^2 \rangle - \langle P \rangle^n, & \text{for outflow points,} \end{cases}$$
(5.2)

where $\Delta P < 0$ denotes the imposed constant streamwise mean-pressure gradient over the streamwise length L_x of the domain. This inflow/outflow condition drives the flow and builds up a numerical solution that converges to a statistically stationary state. No conditions are imposed on particles leaving and entering the domain other than periodicity on their streamwise positions. This, in effect, will simulate the "urban roughness" case of Meroney *et al.* (1996), which is a model for a series of street canyons in the streamwise direction. Wall-conditions are imposed on particles that hit wall-elements as described in Section 3.9. On the top of the domain, free-slip conditions are imposed on particles, i.e. perfect reflection on their positions and a sign reversal of their normal velocity component. To model the small-scale mixing of the passive scalar the IECM model is applied with the $(5 \times 5 \times 5)$ binning structure without employing the projection method to compute $\langle \phi | \mathbf{V} \rangle$. To define the micromixing timescale for a scalar released from a concentrated source in a geometrically complex flow domain bounded by no-slip walls, such as a street canyon, we follow Chapter 4 and Bakosi *et al.* (2007, 2008) and specify the inhomogeneous t_m as

$$t_{\rm m}(\boldsymbol{r}) = \min\left[C_s \left(\frac{r_0^2}{\varepsilon}\right)^{1/3} + C_t \frac{d\boldsymbol{r}}{U_c(\boldsymbol{r})}; \ \max\left(\frac{k}{\varepsilon}; \ C_T \sqrt{\frac{\nu}{\varepsilon}}\right)\right],\tag{5.3}$$

Table 5.1: Constants for modeling the joint PDF of velocity, characteristic turbulent frequency and transported passive scalar.

C_1	C_2	C_3	C_4	C_T	C_L	C_{η}	C_v	γ_5	$C_{\omega 1}$	$C_{\omega 2}$	C_s	C_t
1.85	0.63	5.0	0.25	6.0	0.134	72.0	1.4	0.1	0.5	0.73	0.02	0.7

where r_0 denotes the radius of the source, U_c is a characteristic velocity at r which we take as the absolute value of the mean velocity at the given location, d_r is the distance of the point r from the source, while C_s and C_t are model constants. The applied model constants for the micromixing timescale defined by Equation (5.3) are the same as for the centerline-release in channel flow, i.e. $C_s = 0.02$ and $C_t = 0.7$.

The Reynolds number $Re \approx 12000$ based on the maximum free stream velocity U_0 and the building height H. This corresponds to $Re_{\tau} \approx 600$ based on the friction velocity and the free stream height, h = H/2, if the free stream above the buildings is considered as the lower part of an approximate fully developed turbulent channel flow. After the flow has reached a statistically stationary state, time-averaging is used to collect velocity statistics and a continuous scalar is released from a street level line source at the center of the canyon (corresponding to a point-source in two dimensions). The scalar field is also time-averaged after it has reached a stationary state.

5.3 Results

The simulations with the full resolution model have been run with the constants given in Table 5.1, using 300 particles per element. The Eulerian mesh used for this simulation is displayed in Figure 5.1, which shows the considerable refinement along the building walls and tops necessary to solve the boundary layers. In this case, the high anisotropy and inhomogeneity of the Reynolds stress tensor in the vicinity of walls are captured by the



Figure 5.1: Geometry and Eulerian mesh for the computation of turbulent street canyon with full resolution of the wall-boundary layers using elliptic relaxation. The grid is generated by the general purpose mesh generator Gmsh (Geuzaine and Remacle, 2006). The positions labeled by bold numbers indicate the sampling locations for the passive scalar, equivalent with the combined set of measurement tapping holes of Meroney *et al.* (1996), Pavageau (1996) and Pavageau and Schatzmann (1999). In the zoomed area the refinement is depicted, which ensures an adequate resolution of the boundary layer and the vortices forming in the corner.

Table 5.2: Concentration sampling locations at building walls and tops according to the experimental measurement holes of Meroney *et al.* (1996), Pavageau and Schatzmann (1999) and Pavageau (1996). See also Figure 5.1.

#	1	2	3	4	5	6	7	8	9	10	11	12	13	14
x	0.5	1.0	1.5	2.0	2.0	2.0	2.0	2.0	2.0	2.5	2.0	2.0	4.0	4.0
y	2.0	2.0	2.0	2.0	1.93	1.5	1.33	1.0	0.67	0.5	0.33	0.17	0.17	0.33
#	15	16	17	18	19	20	21	22	23	24				
x	4.0	4.0	4.0	4.0	4.0	4.0	4.0	4.5	5.0	5.5				
y	0.5	0.67	1.0	1.33	1.5	1.93	2.0	2.0	2.0	2.0				



Figure 5.2: Geometry and Eulerian mesh for the computation of turbulent street canyon with wall-functions at $Re \approx 12000$. The domain is stripped at no-slip walls so that it does not include the close vicinity of the wall at $y^+ < 30$. The positions for sampling the scalar concentrations are the same as in Figure 5.1.

elliptic relaxation technique, using Equation (2.13).

The simulations using wall-functions were performed on the Eulerian mesh displayed in Figure 5.2, also using 300 particles per element. We implemented the particle-boundary conditions for arbitrary geometry described in Chapter 2. Note that the first gridpoint where the boundary conditions based on wall-functions are to be applied should not be closer to the wall than $y^+ = u_{\tau}y/\nu = 30$, where y^+ is the non-dimensional distance from the wall in wall-units, but sufficiently close to the wall to still be in the inertial sublayer (Dreeben and Pope, 1997b). Accordingly, the grid in Figure 5.2 only contains the domain stripped from the wall-region at $y^+ < 30$.

Turbulence and scalar statistics are obtained entirely from the particles that represent both the flow itself and the scalar concentration field. The Eulerian meshes displayed in



Figure 5.3: Velocity vectors (first row) and iso-contours of turbulent kinetic energy (second row) of the fully developed turbulent street canyon at $Re \approx 12000$ based on the maximum free stream velocity U_0 and the building height H. Left – full resolution with elliptic relaxation, right – coarse simulation with wall-functions.

Figure 5.1 for the full resolution and in Figure 5.2 for the wall-functions cases are used to extract the statistics, to track the particles throughout the domain and to solve the Eulerian equations: Equation (2.13) and the mean-pressure-Poisson equation (3.8) in the fully resolved case and only the latter in the wall-functions case.

In Figure 5.3, the mean velocity vectorfield and the iso-contours of the turbulent kinetic energy are displayed for both fully resolved and wall-functions simulations. It is apparent that the full resolution captures even the smaller counterrotating eddies at the internal corners of the canyon, while the coarse grid-resolution with wall-functions only captures the overall flow-pattern characteristic of the flow, such as the big steadily rotating eddy inside the canyon. The turbulent kinetic energy field is captured in a similar manner. Both methods reproduce the highest turbulence activity at the building height above the canyon, with a maximum at the windward building corner. The full resolution simulation shows a more detailed spatial distribution of energy, whereas the coarse resolution of the wall-functions simulation still allows to capture the overall pattern.

In Figure 5.4, two of the normalized turbulent intensities, $\langle u^2 \rangle^{1/2} / U_0$ and $\langle w^2 \rangle^{1/2} / U_0$, are displayed for both simulation cases and compared with the large eddy simulation results of Liu and Barth (2002). In the LES simulations the filtered momentum equations are solved by the Galerkin finite element method using brick three-dimensional elements, while the residual stresses are modeled by the Smagorinsky closure.

The full resolution simulation shows a very good agreement with the LES. The contour plots of $\langle u^2 \rangle^{1/2}/U_0$ correctly display two local maxima, at the windward external and at the leeward internal corners. The contour plots of $\langle w^2 \rangle^{1/2}/U_0$ show distributed high values at the building level above the canyon, along the windward internal corner and wall, and at the street level downstream of the source. By contrast, the wall-functions contour plots are in general less detailed, failing to reproduce the internal maximum of $\langle u^2 \rangle^{1/2}/U_0$, and showing a more uniform representation of $\langle w^2 \rangle^{1/2}/U_0$.

Several wind tunnel measurements have been carried out for this configuration, measuring concentration statistics above the buildings, at the walls and inside the canyon, for a scalar continuously released from a street level line source at the center of the canyon (Meroney *et al.*, 1996; Pavageau and Schatzmann, 1999; Pavageau, 1996). To examine the concentration values along the building walls and tops, we sampled the computed mean concentration field at the locations depicted in Figure 5.1 and listed in Table 5.2.

The excellent agreement of the results using both full resolution and wall-functions with a number of experiments is shown in Figure 5.5. The concentration peak is precisely captured at the internal leeward corner and the model accurately reproduces the pattern



Figure 5.4: Dimensionless turbulent intensities $\langle u^2 \rangle^{1/2}/U_0$ (first column) and $\langle w^2 \rangle^{1/2}/U_0$ (second column) computed using full wall-resolution (first row) and using wall-functions (second row) at $Re \approx 12000$ compared with the LES results (third row) of Liu and Barth (2002).

of concentration along both walls including the higher values along the leeward wall.

In Figure 5.6, the first two statistical moments of the concentration inside the canyon are compared with experimental data and LES. The agreement with observations indicates that both the fluid dynamics and the micromixing components of the model provide a good representation of the real field. This is shown in the figures where one can observe the



Figure 5.5: Distribution of mean concentrations at the boundary of the street canyon. The experimental data are in terms of the ratio $CU_{ref}HL/Q_s$, where C is the actual measured mean concentration (ppm), U_{ref} is the free-stream mean velocity (m/s) taken at the reference height $y_{ref} \approx 11H$ and Q_s/L is the line source strength (m²/s) in which Q_s denotes the scalar flow rate and L is the source length. The calculation results are scaled to the concentration range of the experiments. References for experimental data: \triangle Meroney *et al.* (1996); \diamond , \forall , Pavageau and Schatzmann (1999); \Box Pavageau (1996). See also Figure 5.1 and Table 5.2 for the measurement locations.

effects of the two driving mechanisms of transport of concentration by the large eddy inside the canyon as well as diffusion by the turbulent eddies.

Because the one-point one-time joint PDF contains all higher statistics and correlations of the velocity and scalar fields resulting from a close, low-level interaction between the two



Figure 5.6: See next page for caption.

Figure 5.6: Comparison of the spatial distribution of the normalized mean $CU_{\rm ref}HL/Q_{\rm s}$ (left column) and variance $\langle c^2 \rangle (U_{\rm ref}HL/Q_{\rm s})^2$ (right column) of the scalar released at the center of the street level. The normalization and the scaling of the calculated results are the same as in Figure 5.5. First row – PDF calculations with full wall resolution, second row – PDF calculations with wall-functions, third row – experimental data of Pavageau and Schatzmann (1999) and fourth row – LES calculations of Liu and Barth (2002).

fields, a great wealth of statistical information is available for atmospheric transport and dispersion calculations. As an example, the time-averaged PDFs of scalar concentration fluctuations are depicted in Figure 5.7 at selected locations of the domain for the full resolution case. While near the source (Figure 5.7 left) the PDF is slightly skewed, but not far from a Gaussian, the distribution of fluctuations can become very complex especially due to intermittency effects, as shown by the multi-modal PDF in Figure 5.7 right.

The performance gain obtained by applying wall-functions as opposed to full resolution was about two orders of magnitude already at this moderate Reynolds number. The gain for higher Reynolds numbers is expected to increase more than linearly.

5.4 Discussion

In this Chapter the PDF method described in the previous chapters was tested by computing the dispersion of a passive pollutant released from a point source. The Eulerian unstructured grid, consisting of triangular element type, is used to estimate Eulerian statistics, to track particles throughout the domain and to solve for inherently Eulerian quantities. The boundary layers developing close to solid walls are fully captured with an elliptic relaxation technique, but can also be represented by wall-functions, which use a coarser grid resolution and require significantly less particles, resulting in substantial savings in computational



Figure 5.7: Probability density functions of scalar concentration fluctuations (left) at x = 3, y = 0.2 and (right) at x = 3, y = 2 using full resolution at walls.

cost. We found that the one-point statistics of the joint PDF of velocity and scalar are wellcaptured by the wall-functions approximation. In view of its affordable computational load and reasonable accuracy, this approximation appears to hold a realistic potential for application of the PDF method in atmospheric simulations, where the natural extension of the work is the implementation of the model in three spatial dimensions.

In hybrid PDF models developed for complex chemically reacting flows, numerical treatments for boundary conditions have been included for symmetric, inflow, outflow and freeslip walls employing the ghost-cell approach common in finite volume methods (Rembold and Jenny, 2006). The representation of no-slip boundaries adds a significant challenge to the above cases. This is partly due to the increased computational expense because of the higher Eulerian grid resolution required if the boundary layers are to be fully resolved. In addition, there is an increased complexity in specifying the no-slip particle conditions for both fully resolved and wall-functions representations. We presented an implementation of both approaches to treat no-slip boundaries with unstructured grids in conjunction with the finite element method. This obviates further complications with ghost-cells. In the case of full wall-resolution we employed the Lagrangian equivalent of a modified isotropization of production (IP) model as originally suggested by Dreeben and Pope (1998). The elliptic relaxation technique, however, allows for the application of any turbulence model developed for high-Reynolds-number turbulence (Durbin, 1993; Whizman *et al.*, 1996). The standard test case for developing near-wall models is the fully developed turbulent channel flow. In this case, we explored the simpler Rotta (1951) model, which is the Eulerian equivalent of the simplified Langevin model (SLM) in the Lagrangian framework (Pope, 1994). This is simply achieved by eliminating the term involving the fourth-order tensor H_{ijkl} from the right hand side of Equation (2.13). While the SLM makes no attempt to represent the effect of rapid pressure (in fact it is strictly correct only in decaying homogeneous turbulence), it is widely applied due to its is simplicity and robustness. Our experience showed a slight degradation of the computed velocity statistics (as compared to direct numerical simulation) using SLM for the case of channel flow. Since we experienced no significant increase in computational expense or decrease in numerical stability, we kept the original IP model.

Similarly, in the case of wall-functions, several choices are available regarding the employed turbulence model. The methodology developed by Dreeben and Pope (1997b) uses the SLM, but it is general enough to include other more complex closures, such as the Haworth & Pope models (HP1 and HP2) (Haworth and Pope, 1986, 1987), the different variants of the IP models (IPMa, IPMb, LIPM) (Pope, 1994) or the Lagrangian version of the SSG model of Speziale, Sarkar, and Gatski (1991). All these closures can be collected under the umbrella of the generalized Langevin model, by specifying its constants as described by Pope (1994). These models have been all developed for high-Reynolds-number turbulence and need to be modified in the vicinity of no-slip walls. Including them in the wall-function formulation is possible by specifying the reflected particle frequency at the wall as $\omega_R = \omega_I \exp(-2\mathcal{V}_I \langle \omega v \rangle_p / \langle \omega v^2 \rangle_p)$ instead of Equation (2.36). This involves the additional computation of the statistics $\langle \omega v \rangle$ and $\langle \omega v^2 \rangle$ at y_p , which does not increase the computational cost significantly, but may result in a numerically less stable condition since the originally constant parameter β which appears using the SLM has been changed to a variable that fluctuates during simulation. We implemented and tested all the above turbulence models using the wall-functions technique. Without any modification of the model constants we found the IPMa and SLM to be the most stable, providing very similar results. Thus we kept the original (and simplest) SLM along with Equation (2.36).

The most widely employed closure to model the small scale mixing of the passive scalar in the Lagrangian framework is the interaction by exchange with the mean (IEM) model (Villermaux and Devillon, 1972; Dopazo and O'Brien, 1974). This simple and efficient model, however, fails to comply with several physical constraints and desirable properties of an ideal mixing model (Fox, 2003). The interaction by exchange with the conditional mean (IECM) model overcomes some of the difficulties inherent in the IEM model. In this Chapter we justify the sole use of the IECM model by its being more physical and more accurate, but we acknowledge that it markedly increases the computational cost.

Chapter 6: Cylinder flow simulations: results and discussion

6.1 Introduction

As a third validation testcase we simulate the turbulent flow in the wake of a circular cylinder. This classical example has been widely studied both experimentally and numerically, therefore a large amount of data have been accumulated about its flow dynamics. Although the domain geometry is relatively simple, the flow exhibits a variety of vastly different behaviors depending on the Reynolds number, ranging from a steady laminar state through unsteady but periodic laminar vortex shedding to transitional and fully developed turbulence. We select the Reynolds number $Re_D = 3900$ (based on the cylinder diameter and the free stream velocity), mainly because it corresponds to a transitional flow in the near wake behind the cylinder. Secondly, this Reynolds number has also been studied extensively with both LES and DNS, thus a quantitative comparison of several flow statistics computed by other methods is also possible. From the modeling viewpoint this Reynolds number is a challenging tasks to undertake. At this Reynolds number the separating boundary layers along the cylinder surface are fully laminar. Transition to turbulence occurs in the very near wake due to shear layer instabilities, which is followed by a region dominated by vortex shedding dynamics where the wake becomes fully turbulent and the coherent structures gradually give place to fully developed turbulence. Since these features require a solver to perform relatively well in all laminar, transitional and turbulent regions of the flow, this case appears to be a good candidate to identify the limitations of the current method. Another reason to compute this flow is to further evaluate the current PDF methodology using unstructured grids and no-slip walls with curvature in complex geometries.

One of the key components simulating this flow is the adequate resolution of the separating boundary layers which decisively determines the flow behavior downstream and crucially influences the accuracy of the numerical solution. Accordingly, LES studies with sufficient wall resolution have been successful in predicting both cylinder surface and downstream wake statistics relatively accurately. On the other hand, RANS models, due to their inherent high-Reynolds-number assumption, have usually failed to predict both the wake and the mean integrated statistics along the cylinder surface, such as the drag (even with adequate wall resoution). Employing wall-functions at the cylinder surface may also be problematic, since wall-functions are built on the fundamental assumption that the boundary layer is turbulent and remains attached. Neither of these assumptions are correct along the cylinder surface at a sub-critical Reynolds number. The separating laminar boundary layers along a curved geometry provides a tough testcase for the elliptic relaxation technique as well. Although this type of wall-treatment can be tought of as a set of sophisticated blending functions for near-wall turbulence, its fundamental assumptions are less restrictive compared to wall-functions. It also represents all components of the Reynolds stress tensor at the wall instead of relying on the turbulent viscosity hypothesis. Although this technique has originally been developed for turbulent boundary layers, it seems compelling to investigate its performance modeling a separating laminar boundary layer transitioning to turbulence.

Another complication is that the flow is highly unsteady and the turbulence is mechanically generated in the domain by the obstacle. In such situations an adequately resolved LES/DNS may perform well both far and in the vicinity of walls since it solves both the small wall-generated vortices and the large eddies far from walls. In other situations, however, where a given level of turbulence is required to be present but the turbulence-generating obstacles are not required to be part of the domain, other means are necessary to provide the right level of fluctuations in LES which are not always obvious. A RANS model has no problems handling these latter situations since it represents the turbulent kinetic energy explicitly in its formulation. On the other hand, solving turbulence which is generated within the domain may be a difficult task for both RANS and URANS models, due to their above mentioned limitations close to walls.

Unsteady PDF methods have been developed based on the LES methodology defining the filtered density function (FDF) which is used to provide closure for the filtered equations. The development has resulted in the hybrid FV/particle methods. Following the same logical sequence that led to unsteady RANS based on RANS simulations, it seems relevant to investigate an unsteady PDF (UPDF) methodology based on steady PDF methods. Here we will apply the current model, developed and tested for steady flows, for a transient flow in the same way as RANS models are applied to obtain time-dependent statistics resulting in URANS.

In Section 6.1.1 a short review of the circular cylinder flow regimes are given. This is followed by an overview of the literature regarding experimental and numerical studies investigating the cylinder near wake at sub-critical Reynolds numbers, Section 6.1.2. In Section 6.2 several computed velocity statistics are examined and compared to LES, DNS and experimental data where available. Finally, Section 6.3 sums up the findings regarding this testcase.

6.1.1 A short review of cylinder flow regimes

Reviews on the physics of the cylinder flow have been compiled by Morkovin (1964); Berger and Wille (1972); Norberg (1987) and more recently by Williamson (1996). Only a short overview is given in the following.

The single relevant parameter of the flow over a circular cylinder is the Reynolds number, defined here as $Re_D = U_0 D/\nu$, where U_0 is the free stream velocity, D is the cylinder diameter and ν denotes the kinematic viscosity.

At Re_D lower than approximately 40, the flow is laminar and steady. The boundary layer separates at $Re_D \approx 3-5$ resulting in two symmetric counter-rotating vortices behind the cylinder. This recirculating region grows linearly with Reynolds number and the velocity profiles at the end of the recirculating region exhibit self-similarity.

At Reynolds numbers higher than 40 the vortices become unstable which initiates periodic vortex shedding resulting in a Kármán vortex street. The non-dimensionalized frequency of the separating vortices is the Strouhal number ($St = nD/U_0$) which is used to characterize the unsteadyness of the flow related to the periodic vortex street. For up to about $Re_D = 150$ the flow remains laminar and the Strouhal number increases with Reynolds number, then reaches a plateau of ~ 0.21. Transition to three-dimensionality starts at $Re_D = 180 - 260$ due to the appearing streamwise vortices in the wake.

At the sub-critical Reynolds number range, between 300 and 2×10^5 , the separating boundary layers are still fully laminar along the cylinder surface and transition into turbulence occurs in the near wake due to shear layer instabilities. At the lowest Reynolds numbers in this range the flow becomes fully turbulent only about 40-50 diameters downstream, where the periodic vortices have been completely diffused. As the Reynolds number increases this transition moves closer to the cylinder. At the highest Reynolds numbers in this range the transition in the shear layers occurs very close to the separation points.

In the critical Reynolds number range, between 2×10^5 and 3.5×10^6 , two significant changes occur that crucially influence the drag on the cylinder. At $Re_D \approx 3.6 \times 10^5$ the drag coefficient drops abruptly (from 1.2 to 0.3) due to a sudden increase in the base pressure behind the cylinder. The separating laminar boundary layer along the cylinder surface transitions to turbulence and reattaches then finally separates again. The separation point moves towards the downstream side of the cylinder and the width of the wake decreases to less than 1 cylinder diameter. In the range 5×10^5 and 3.5×10^6 the base pressure decreases which increases the drag from 0.3 to 0.7 which remains at this value up to about $Re_D = 10^7$. In the post-critical regime, above 3×10^6 , the boundary layer transitions to turbulence before separating and the Strouhal number stays approximately constant at 0.27.

6.1.2 Past experimental and numerical studies

Ma et al. (2000) divide the cylinder wake into three regions at sub-critical Reynolds numbers: the near wake up to about 10 diameters downstream, the intermediate wake up to fifty diameters and the far or self-preserving wake beyond that (Matsumura and Antonia, 1993). There are relatively few experiments available in the near wake due to difficulties and special arrangements required in order to obtain accurate data, as in the experiments of Cantwell and Coles (1983) who provided measurements up to x/D = 8 for the Reynolds number $Re_D = 140\,000$. Employing particle image velocimetry (PIV) Lourenco & Shih (1993, see Beaudan and Moin 1994) have obtained data on the first two moments of the velocity field in the recirculation region at $Re_D = 3900$. Ong and Wallace (1996) reported data on the first four moments of the velocity and its spectra based on hot-wire measurements conducted with an X-array probe between $3 \le x/D \le 10$ at the same Reynolds number. Both the cylinder surface and near wake statistics are particularly sensitive to experimental disturbances, such as acoustic noise levels, cylinder vibrations, surface roughness and other geometric parameters in this Reynolds number range (Norberg, 1987). This is exemplified by the different lengths of the recirculation bubbles obtained by the experiments of Lourenco and Shih (1993), Ong and Wallace (1996) and Govardhan & Williamson (2000, see Ma et al. 2000, Figure 1). The possible causes of the discrepancy among the experimental datasets are discussed in more detail by Noca *et al.* (1998).

A summary of the literature regarding numerical simulations of the cylinder flow up to the middle of the last decade at different Reynolds numbers is given by Beaudan and Moin (1994). Their overall conclusion is that two-dimensional Navier-Stokes simulations at transitional Reynolds numbers (between 150 and 300) are capable of predicting Strouhal numbers and drag coefficients, but become unreliable in the sub-critical regime. Although the flow geometry is nominally two-dimensional, three-dimensional effects at these higher Reynolds numbers become non-negligible. Steady RANS simulations employing the $k - \varepsilon$ model predict inaccurate mean velocity and Reynolds stress distributions in the near wake and produce mixed results for the integrated statistics over the cylinder surface (Beaudan and Moin, 1994). This is perhaps little surprise, since in this flow the eddy-viscosity is anisotropic and negative in regions where history and transport effects dominate over production of Reynolds stresses, indicating the inadequacy of the turbulent-viscosity hypothesis for this flow (Franke et al., 1989). Underresolved DNS improve on RANS simulations by better capturing the drag coefficients up to the critical Reynolds number 10^6 , which is attributed to better resolving the three-dimensionality of the flow (Beaudan and Moin, 1994). From this viewpoint it will be interesting to see how the current PDF model performs: although the spanwise components of the particle positions are not retained, the velocity field is three-dimensional in the sense of fluctuations. In other words, while all three components of the particle velocities are retained to represent spanwise fluctuations, mean spanwise motions due to streamwise and cross-stream vorticity are not represented and it is assumed that $\langle W \rangle = 0$.

A systematic LES study at $Re_D = 3900$ has been undertaken by Beaudan and Moin (1994) whose main objective was to evaluate the performance of the dynamic residualstress model (Germano *et al.*, 1991) in a flow where RANS simulations have been known to have difficulties. They performed simulations without closure, with the fixed-coefficient Smagorinsky-model and with the dynamic model. Both two and three-dimensional cases have been computed to assess the importance of representing three-dimensional effects. Another goal was to evaluate the performance of higher order upwind schemes for the advection terms, including fifth and seventh order finite difference approximations. The
work further demonstrates the necessity of three-dimensional calculations for this flow, documenting consistent improvements in all quantities examined when compared to twodimensional simulations with no residual-stress model (other than the numerical diffusion inherent in upwind schemes). Regarding the spatial discretization, they conclude that even higher order upwind schemes are not suitable for LES due to their numerical diffusion which may be comparable to the subfilter-scale diffusion. Following this line of work Mittal and Moin (1996) used central differencing in order to better control the numerical diffusion, while Kravchenko and Moin (2000) employed a high order B-spline-based finite element method obtaining Reynolds stress distributions in closer agreement with their respective experimental profiles. The above series of LES studies show that the computed statistics in the near wake may be significantly influenced by the choice of the discretization scheme for the advection term. However, the choice of different models for the unresolved stress is clearly less important. Another concern in LES, just like in RANS, is that the use of eddyviscosity-based models for the unresolved scales in non-equilibrium flows are questionable (Liu and Liu, 1997).

Preliminary results on DNS of the cylinder flow at sub-critical Reynolds numbers have been reported by Tomboulides *et al.* (1993) and Henderson and Karniadakis (1995), but full resolution of the near wake has become possible only recently. Ma and Karniadakis (1997) have performed direct simulations based on hierarchical spectral methods employing unstructured grids. The study compared DNS and the LES results of Beaudan and Moin (1994) and Mittal and Moin (1996). This work was followed by more detailed numerical studies by Ma *et al.* (2000) and more recently by Dong *et al.* (2006), who combined experimental imaging (PIV) and DNS performing both experiments and numerical simulations at $Re_D = 3900$ and 10 000 in order to investigate the near wake focusing on the onset of shear-layer instabilities and Reynolds number effects. We will compare results from the current PDF simulations with many of the experimental and numerical datasets mentioned above.

6.2 Results

Several PDF simulations have been carried out to model the unsteady flow around a circular cylinder at $Re_D = 3900$ employing the grid displayed in Figure 6.1. The refinement in the vicinity of the cylinder amounts to 156 elements along the circumference with an average size of $4.5 \times 10^{-3}D$ in the radial and 0.02D in the circumferential direction. This is slightly coarser than the coarsest case in the LES simulations of Kravchenko and Moin (2000) and corresponds to about half the resolution of the LES study of Mahesh et al. (2004). The total number of elements is approximately 50K triangles. The number of particles per element initially is set to 50 and the CFL number is 0.8 (kept at this constant level) using forward Euler-Maruyama timestepping with the adaptive technique described in Section 3.2. Once the boundary layers start separating from the cylinder surface, these parameters result in a particle redistribution of approximately 200–300 particles each timestep, requiring a minimum of 5 particles in each element. This extent of redistribution (only $\sim 0.01\%$ of all the particles redistributed in each timestep) is sufficient enough to have a non-negligible negative effect on the overall performance of the code, so a more efficient particle redistribution procedure has been developed, which by itself is 200 times faster than the basic one described in Appendix C and results in an overall speedup of 15 times for the whole code. The details of this new algorithm are described in Appendix D.

The initial conditions are as follows: the particle velocity is assigned a joint Gaussian distribution with a low-level turbulent kinetic energy, $k/U_0^2 = 0.01$ homogeneously on the whole domain, while the particle frequencies are sampled from a gamma distribution with unit mean and variance 1/4. Free-slip conditions are imposed on the cross-stream bound-aries, i.e. a particle hitting the wall is simply reflected with opposite cross-stream velocity



Figure 6.1: Eulerian mesh for computing the near wake of the cylinder flow at $Re_D = 3900$ based on the cylinder diameter D and free stream velocity U_0 . The refinement along the wall amounts to 156 wall elements with an average size of $4.5 \times 10^{-3}D$ in the radial and 0.02D in the circumferential direction.

 \mathcal{V} . Particles leaving at the outflow are relocated at the inflow leaving their cross-stream position \mathcal{Y} intact and setting their velocity to $\mathcal{U}_i = (U_0, 0, 0)$, which corresponds to a joint delta distribution, i.e. incoming laminar flow with streamwise velocity U_0 . At the cylinder surface, no-slip conditions are imposed on particles as described in Section 3.9. The no-slip wall-conditions are enforced on the extracted velocity statistics as well. A homogeneous Dirichlet condition is imposed on the mean pressure at the outflow and homogeneous Neumann conditions on every other outer boundary. At the cylinder the Neumann condition (3.36) is enforced for the mean pressure. The boundary conditions for the elliptic relaxation tensor are $\varphi_{ij} = -4.5\varepsilon n_i n_j$ at the cylinder wall and homogeneous Neumann conditions along all other boundaries. The applied model constants are the same as before and displayed in Table 4.1.

The current unsteady PDF simulations have been carried out in a similar fashion as an unsteady RANS simulation. In URANS the model equations developed for computing the time-averaged statistics for an inhomogeneous flow are solved in a time-accurate manner, sampling the solution at certain timesteps. This can be thought of as filtering in time with the filter width defined as the time between two consecutive timesteps. Similarly, in the current UPDF simulations we take the equations originally developed for steady flows and solve them with a time-accurate numerical algorithm and sample results at specified timesteps.

In the following, we examine flow statistics regarding the transient nature of the flow as well as integrated quantities along the cylinder surface (Section 6.2.1) and time-averaged fields in the near wake (Section 6.2.2).

6.2.1 Transient and cylinder surface statistics

A common parameter used to examine the cylinder flow is the Strouhal number which is defined as the non-dimensional form of the vortex shedding frequency, n, as

$$St = \frac{nD}{U_0}.$$
(6.1)

There are many quantities from which the Strouhal number can be extracted from a simulation, the time evolution of the cross-stream component of the force acting on the body, i.e. the lift, being the most common one. We evaluate the force F on the cylinder surface \mathcal{A} by

$$F_{i} = \int_{\mathcal{A}} \left(-\langle P \rangle \delta_{ij} + \rho \nu \frac{\partial \langle U_{i} \rangle}{\partial x_{j}} \right) n_{j} \mathrm{d}\mathcal{A}, \tag{6.2}$$

where n_j is the wall-normal. The drag and lift can be obtained by taking the streamwise (i = 1) and cross-stream (i = 2) components of \mathbf{F} . The drag and lift coefficients, C_D and C_L , are the non-dimensional components of the force $\mathbf{F} = F_x \mathbf{e}_x + F_y \mathbf{e}_y$ and can be



Figure 6.2: Time-evolution of the pressure lift coefficient. The black solid line denotes the computed instantaneous value in every timestep, while the red line is its 100-point running average.

decomposed into pressure and viscous parts:

$$C_D = \frac{F_x}{(1/2)\rho U_0^2 D^2} = C_{Dp} + C_{Dv}, \qquad (6.3)$$

$$C_L = \frac{F_y}{(1/2)\rho U_0^2 D^2} = C_{Lp} + C_{Lv}.$$
(6.4)

The time-evolution of the pressure lift coefficient C_{Lp} is plotted in Figure 6.2. Since Monte-Carlo PDF simulations are stochastic by nature, there is a considerable statistical noise in all quantities computed, especially in the ones based on the mean pressure. Nevertheless, applying a moving time-average with a window of 100 timesteps the Strouhal number of ~ 0.2 can be easily extracted. Note that the value of St is not sensitive to the size of the window. Table 6.1 shows how this value compares to past experiments and numerical



Figure 6.3: Time-evolution of the total and viscous lift coefficients at $Re_D = 3900$. (a) threedimensional LES of Beaudan and Moin (1994) using a dynamic model for the unresolved scales, (b) current PDF simulations. The red lines on (b) indicate that they have been obtained using 100-point running averages from instantaneous data similar to the one in Figure 6.2. Solid lines – total lift, dashed lines – viscous lift.

simulations. In Figure 6.3 the evolution of the total and viscous lift coefficients are compared to the LES results of Beaudan and Moin (1994) also performed at $Re_D = 3900$. The PDF simulation successfully reproduces the irregularity of the vortex shedding at this Reynolds number, which is also apparent in the three-dimensional LES and has also been observed in experiments, such as the oil-flow visualizations of Schewe (1986) at the critical Reynolds number 2.64×10^5 just before the drag crisis occurs. Also shown in Table 6.1 are the total drag C_D and base pressure coefficients

$$C_{Pb} = \frac{\langle P \rangle_b - P_0}{(1/2)\rho U_0^2},$$
(6.5)

where $\langle P \rangle_b$ and P_0 are the pressures at the back stagnation point and at infinity, respectively, which are also quantities frequently examined in cylinder flow studies. In general, both of these quantities, calculated by the PDF method, are in good agreement with the LES, DNS

Table 6.1: Cylinder surface and recirculation bubble region statistics. The large eddy simulation data employing upwind, central difference and B-spline schemes for the advection term are provided by Beaudan and Moin (1994); Mittal and Moin (1996) and Kravchenko and Moin (2000), respectively, all performed in three dimensions at $Re_D = 3900$. The DNS data corresponds to the high resolution Case I of Ma *et al.* (2000) and Dong *et al.* (2006) also at $Re_D = 3900$. References for the experimental data: St - Ong and Wallace (1996); C_{Pb} , $C_D - \text{Norberg}$ (1987) at $Re_D = 4020$; $\theta_{\text{sep}} - \text{Son}$ and Hanratty (1969) at $Re_D = 5000$; L/D, $\langle U \rangle_{\min}/U_0$, $r_{\min}/D - \text{PIV}$ of Dong *et al.* (2006) at $Re_D = 4000$. The last column labeled by "PDF" denotes the current PDF simulation.

	Upwind	Central	B-spline	DNS	Expt.	PDF
	LES	LES	LES			
Strouhal number, St	0.203	0.207	0.21	0.203	0.21	0.2
Base pressure coefficient, C_{Pb}	-0.95	-0.93	-0.94	-0.96	-0.99	-0.79
Total drag coefficient, C_D	1.0	1.0	1.04	0.981	0.98	1.04
Separation angle, θ_{sep}	85.8°	86.9°	88.0°	89.0°	86.0°	89.3°
Length of recirculation bub-	1.36	1.4	1.35	1.12	1.47	1.53
ble, L/D						
Minimum streamwise velocity	-0.32	-0.35	-0.37	-0.291	-0.252	-0.34
in bubble, $\langle U \rangle_{\rm min}/U_0$						
Location of $\langle U \rangle_{\min}$ in bubble,	0.88			0.88	1.01	1.1
r_{\min}/D						

and the experimental data at a slightly higher Reynolds number with the base pressure slightly overpredicted. Beaudan and Moin (1994) report substantially lower base pressure (-2.16) and consequently higher drag (1.74) from a two-dimensional LES simulation with no subfilter-scale model. They found large discrepancies in other quantities as well, such as the amplitude and regularity of the lift, the skin-friction coefficient and the complete absence of an attached recirculation bubble. Since their three-dimensional simulations are in close agreement with experiments, they conclude that three-dimensional effects strongly influence the near-wake at this Reynolds number and that modeling the three-dimensionality of the flow is therefore essential. The current PDF simulations are two-dimensional in the sense of mean motions. While all three components of the particle velocities are retained, their positions are only allowed in the x - y plane and the spanwise position is not represented. In other words, the turbulent fluctuations are modeled as three-dimensional, while the mean flow is two-dimensional. The generally close agreement of the current and subsequent PDF results with both three-dimensional simulations and experimental data suggests that retaining the three-dimensional fluctuations is essential.

Next we examine time-averaged statistics along the cylinder surface. These and subsequent time-averaged quantities have been collected after the quasi-periodic vortex shedding has been started, in the time-range of $60 \le tU_0/D \le 200$, which amounts to approximately 28 vortex shedding cycles. The mean pressure coefficient

$$C_P = \frac{\langle P \rangle - P_0}{(1/2)\rho U_0^2}$$
(6.6)

is plotted in Figure 6.4 (a) along with three-dimensional LES, DNS and experimental data. The overall agreement is very good, except for a slightly higher mean pressure at the front stagnation point, which is most likely due to the closeness of the inflow boundary to the cylinder – only 3.5 diameters upstream, Figure 6.1. The spanwise component of the mean vorticity, computed as

$$\Omega_z = \frac{1}{2} \left(\frac{\partial \langle V \rangle}{\partial x} - \frac{\partial \langle U \rangle}{\partial y} \right), \tag{6.7}$$

is plotted in Figure 6.4 (b). The PDF simulation accurately predicts the location of the boundary layer separation (where the vorticity becomes zero) indicated by the close agreement of the vorticity distribution with DNS data and by the correct separation angle of $\theta_{sep} = 89.3^{\circ}$, see also Table 6.1.



Figure 6.4: Time-averaged mean (a) pressure coefficient and (b) spanwise vorticity distributions along the cylinder wall. Red solid lines – PDF simulation, dashed lines – DNS of Ma *et al.* (2000), blue dot-dashed lines – LES of Kravchenko and Moin (2000), symbols – experimental data of (a) Norberg (1987) at $Re_D = 3000$ and (b) Son and Hanratty (1969) at $Re_D = 5000$.

6.2.2 Near wake statistics

Capturing the correct point of separation is crucial in predicting the correct statistics in the recirculation bubble as well, where transition to turbulence occurs as the thin shear layers become unstable. Predicting the transition has been a challenging task in both experiments and numerical simulations designed to obtain flow statistics at sub-critical Reynolds numbers. Measurements are found to be very sensitive to experimental disturbances as well as to other details, such as the cylinder aspect ratio (i.e. the spanwise length of the domain compared to the cylinder diameter). Some of the influencing factors in simulations are the type of the spatial and temporal discretization schemes which are directly related to the extent of numerical dissipation, the type and amount of subfilter-scale diffusion in LES and the spanwise size of the domain in both LES and DNS (Beaudan and Moin, 1994; Ma et al., 2000). Because of these difficulties, substantially different experimental datasets are



Figure 6.5: Streamwise mean velocity behind the cylinder along the centerline at $Re_D = 3900$. Red solid line – PDF simulation, blue dot-dashed line – LES of Kravchenko and Moin (2000), symbols – experiments of \Box , Lourenco & Shih (1993, see Beaudan and Moin 1994), \circ , Ong and Wallace (1996) and \times , Govardhan & Williamson (2000, see Ma *et al.* 2000).

available regarding the recirculation bubble. This is exemplified by Figure 6.5, where the mean streamwise velocity from the current PDF simulation is plotted together with several experimental datasets and the LES of Kravchenko and Moin (2000) employing a higher order B-spline-based finite element method and a dynamic subfilter model.

Due to the aforementioned difficulties, major differences have also been found in the cross-stream shape of the mean streamwise velocity in the bubble. Moin and co-workers (Beaudan and Moin, 1994; Mittal and Moin, 1996; Kravchenko and Moin, 2000) consistently obtained profiles closer to a *U-shape* from their LES simulations, in disagreement with the experimental data of Lourenco and Shih (1993) and the LES of Frohlich *et al.*

(1998) generating V-shape profiles. Kravchenko and Moin (2000) discuss in length the possible sources of the differences. They point out that immediately behind the cylinder, at x/D = 0.58, both the experiments and numerical simulations predict a U-shape profile, which evolves into a V-shape farther downstream. The shape of the mean velocity profiles is directly related to the level of fluctuations and therefore the transition in the shear layers. Smaller fluctuations result in U-shape, while larger fluctuations results in a more mixed and diffused V-shape profile for the mean velocity, see Figures 22 and 23 of Kravchenko and Moin (2000). Also, the length of the laminar shear layers is larger for U-shape and shorter for V-shape profiles, indicating that the onset of instability (the transition to turbulence) occurs farther and closer to the cylinder, respectively. In a direct simulation the precise point where the level of fluctuations becomes large enough to initiate the instability of the shear layers is influenced by many factors including the inherent numerical and the additional subfilter diffusion. Accordingly, simulations performed on coarser grids tend to produce V-shape profiles while finer grids result in U-shape profiles, see Figures 24 and 25 of Kravchenko and Moin (2000). Increasing the value of the Smagorinsky-constant also results in more diffusion, however its effect is more pronounced on the fluctuations, resulting in a U-shape profile for the mean velocity after the fluctuations have been attenuated. This major influence of the subfilter-scale diffusion in LES has been shown by the systematic DNS and LES study of Ma et al. (2000), see Figures 7 and 8 therein, who also found the aspect ratio (i.e. the spanwise length of the domain) to be a decisive factor affecting the shape of the mean velocity profiles. From high resolution DNS simulations, they find two distinct converged states, arriving at either U or V-shapes, depending on the spanwise size of the domain employed, see their Figures 5 and 6. The narrower domain corresponds to the size used by Moin *et al.* and converges to U-shape, while the twice as wider domain produces a V-shape profile in close agreement with the experiments of Lourenco and Shih (1993).

The series of studies mentioned above makes it clear that the ability of large eddy simulation to capture the precise point of instability heavily depends on the correct balance of phsyical, numerical and subfilter-scale diffusion. The same issue is present in PDFtype methods as well, with either subfilter diffusion (in LES/FDF) or modeled turbulent diffusion (in UPDF), therefore we can expect similar difficulties in these methods as well. Although modifying the model constants may improve certain predictions, it is always of limited value and we did not explore it. More importantly, grid-, and particle-number independence should be established.

Cross-stream distributions of the mean streamwise velocity $\langle U \rangle$ obtained from the current PDF simulation is plotted at different downstream locations in Figure 6.6 (a) and (c) along with DNS and experimental data. We see that the PDF simulation correctly predicts the V-shape of the streamwise velocity in the bubble with the minimum at the centerline slightly underpredicted towards the end of the bubble indicating a strong mean backflow there. Farther downstream, where the turbulence is dominated by vortex dynamics, the prediction is also very reasonable. Beaudan and Moin (1994) also examine the errors in the experiments of Lourenco & Shih based on the expected symmetries and anti-symmetries in the mean velocity and Reynolds stress components. They find that the errors in the mean streamwise velocities $\langle U \rangle$ are at 5% of the maximum local velocity past 1 diameter downstream, while cross-stream velocities $\langle V \rangle$ exhibit errors comparable to their actual values, i.e. close to 100%, in the first 3.5 diameters which increases to 200-300% farther downstream. Ong and Wallace (1996) report experimental uncertainties of 2% for their mean velocities.

The time-averaged cross-stream velocities $\langle V \rangle$ produced by the PDF simulation are displayed in Figure 6.6 (b) and (d). Up to the streamwise length examined, x/D = 10, the anti-symmetric shape of the profiles are correctly captured with their magnitude gradually diminishing as the flow gets better mixed downstream. Immediately behind the cylinder,



Figure 6.6: Mean (a), (c) streamwise and (b), (d) cross-stream velocity at different downstream locations behind the cylinder at $Re_D = 3900$. Red solid lines – PDF simulation, black dashed lines – DNS of Ma *et al.* (2000), symbols – experiments of \Box , Lourenco and Shih (1993) and \circ , Ong and Wallace (1996).

at x/D = 1.06 and 1.54, the profiles resemble both the DNS and the experimental data but with less pronounced extrema. The prediction of $\langle V \rangle$ improves farther downstream in the bubble, x/D = 2.02, when compared to DNS data and stays close to the experiments of Ong and Wallace (1996) for $x/D \geq 3$. It is worth noting that the two experimental datasets we use to compare the results do not match each other as shown in Figure 6.7. The maximum magnitude of the velocity and the spread of the wake are different with asymmetries evident



Figure 6.7: Time-averaged cross-stream mean velocity behind the cylinder at x/D = 3.0 at $Re_D = 3900$. Red solid line – PDF simulation, black dashed line – DNS of Ma *et al.* (2000), symbols – experiments of \Box , Lourenco and Shih (1993) and \circ , Ong and Wallace (1996).

in both datasets. Both the DNS data and the PDF simulation follow the experiments of Ong and Wallace (1996) more closely.

The streamwise and cross-stream components of the Reynolds stress tensor, $\langle u^2 \rangle$ and $\langle v^2 \rangle$, are displayed in Figure 6.8 at several downstream locations. The error analysis of Beaudan & Moin suggests 20% error in the Reynolds stress components for the experimental data of Lourenco and Shih (1993) between $1.0 \leq x/D \leq 2.5$ and slightly increasing farther downstream. Ong and Wallace (1996) provide the experimental error in these quantities as 2% for all their measured length, $x/D \geq 3$. The streamwise Reynolds stress component $\langle u^2 \rangle$ in the recirculation bubble is in excellent agreement with the experiments and DNS data, Figure 6.8 (a). The locations and the extent of the double peaks and their local minimum at the centerline are all predicted accurately, in close agreement with the experiments. Farther



Figure 6.8: Streamwise (a), (c) and cross-stream (b), (d) Reynolds stress components at different downstream locations behind the cylinder at $Re_D = 3900$. Red solid lines – PDF simulation, black dashed lines – DNS of Ma *et al.* (2000), symbols – experiments of \Box , Lourenco and Shih (1993) and \circ , Ong and Wallace (1996).

downstream, Figure 6.8 (c), as the double peaks diminish in magnitude and gradually give place to single peaks indicating a more mixed state, the PDF predictions slightly diverge from the DNS data, underpredicting the level of fluctuations at $x/D \ge 7$.

Figures 6.8 (b) and (d) show that the cross-stream fluctuations $\langle v^2 \rangle$ are severely underpredicted throughout the whole length of the wake examined. This may be due to the presence of an excessive level of turbulent and/or numerical diffusion originating from the



Figure 6.9: Time-averaged cross-stream profiles of turbulent kinetic energy in the recirculation bubble in the wake of a cylinder at x/D = 1.0 and x/D = 2.0 at $Re_D = 3900$. Red solid lines – PDF simulation, blue dot-dashed lines – LES of Beaudan and Moin (1994).

turbulence model and a coarse grid (or overdiffusive spatial discretization), respectively. Ma et al. (2000) show that increasing the subfilter-scale diffusion in LES can have a disastrous effect on the second moments, especially in the recirculation bubble where the transition to turbulence occurs, where the turbulent kinetic energy reaches its highest levels. We can examine the turbulent kinetic energy, $k = (\langle u^2 \rangle + \langle v^2 \rangle + \langle w^2 \rangle)/2$, in the bubble in Figure 6.9. Close to the cylinder, at x/D = 1, the level of energy is in reasonable agreement with the dynamic LES simulations of Beaudan & Moin. In the LES simulations the energy grows almost threefold by the end of the bubble, x/D = 2, while this growth in the PDF simulations is almost negligible or dissipated. We suspect that a high level of local diffusion (modeled and/or numerical) attenuates both cross-stream $\langle v^2 \rangle$ and spanwise $\langle w^2 \rangle$ fluctuations, which only dissipates further downstream. Another possible explanation for the discrepancy in the $\langle v^2 \rangle$ and $\langle w^2 \rangle$ components is the lack of representation of the mean motions in the third, spanwise dimension. Other than these factors, we also suspected another possible source of numerical dissipation, namely the way the local Eulerian statistics are computed, as described in Section 3.4. Computing ensemble averages in elements, then transferring them into gridpoints, and finally at the particle positions, employing nodal averages in elements, also have a smoothing/diffusive effect. Therefore another algorithm has been implemented in which the element-based ensemble averages are directly used in updating particles, without the intermediate step of transferring statistics to and from nodes. Originally, the main reason for the more complex two-step procedure was to mitigate the dire effects of elements without particles, however, this is not strictly necessary if one applies a particle redistribution procedure. A series of numerical tests with both algorithms, however, resulted in no significant change in the fields (i.e. it was not less diffusive). Because the simpler algorithm was not measurably more efficient than its current counterpart, we kept the two-step procedure. Further investigations are necessary to pinpoint the exact cause of the discrepancy between our simulations and the agreeing experimental and DNS data. These may include higher spatial refinement in the bubble, higher order (less diffusive) timestepping scheme and simulation in fully three-dimensional space.

Figure 6.10 displays the shear Reynolds stress component $\langle uv \rangle$ at different downstream locations in the wake. We see that the agreement with DNS and experimental data is quite good from the end of the recirculation bubble, x/D = 2.02. Immediately behind the cylinder, x/D = 1.06 and x/D = 1.54, the predictions follow the DNS and experiments in shape, but the peaks of the profiles are diffused. Also, the anti-symmetric double peaks at x/D = 1.06, apparent in both the DNS and the experiments, are only recognizable by the slightly flattening shear stress profile at the centerline. We suspect that a simulation with a more refined Eulerian grid would help in predicting the very near wake shear stress even more accurately.



Figure 6.10: Time-averaged shear Reynolds stress at different downstream locations behind the cylinder at $Re_D = 3900$. Red solid lines – PDF simulation, black dashed lines – DNS of Ma *et al.* (2000), symbols – experiments of \Box , Lourenco and Shih (1993) and \circ , Ong and Wallace (1996).

We now turn our attention to higher order velocity statistics. In general, third and fourth order moments of the velocity field are rarely investigated in the literature regarding this flow. The most widely applied turbulence modeling techniques, $k - \varepsilon$ and Reynolds stress models, would not be economical and little is known about the reliance of the turbulent viscosity hypothesis at these higher level of closures. However, higher order moments of passive tracers (or the full concentration PDF) would be valuable in atmospheric pollution modeling, where there is a need to predict extreme events and probabilities in concentration fields. Lagrangian dispersion models are capable of providing this information and are routinely used to compute scalar fields, usually in conjunction with traditional CFD-type (URANS or LES) models that provide them the mean and fluctuating velocity. In these applications the micromixing of the scalar (which determines the scalar PDF) is commonly modeled without taking the velocity field into account, which is not justified theoretically (Pope, 1998) and may result in erroneous predictions as we demonstrated in Chapter 4 for a very simple flow. PDF methods naturally account for the close interaction between the turbulent velocity and scalar fields, therefore potentially more accurate predictions can be achieved. In principle, direct simulations (LES and DNS) could also provide this higher level of velocity information for Lagrangian micromixing models, however, a more natural approach is to jointly model the PDF of velocity and concentrations, which requires an accurate representation of the higher order velocity statistics. Therefore we now examine the prediction of the skewness and flatness of the velocity field behind the cylinder.

Cross-stream profiles of the skewness of the streamwise and cross-stream velocity components are displayed in Figure 6.11 at several downstream locations. Ong and Wallace (1996) provide experimental data for the skewness and flatness for $x/D \ge 3.0$. The profiles are plotted for the usual locations, both inside and outside the bubble. In general, all skewness predictions are in good agreement with the experimental data with the streamwise component slightly underpredicted, especially at the edge of the wake close to the end of the recirculation bubble, x/D = 4.0, Figure 6.11 (c). However, this discrepancy gradually diminshes farther downstream. The prediction of flatness is similary good as shown in Figure 6.12. Although the flatnesses of both streamwise and cross-stream velocity are overpredicted immediately after the bubble in the highly skewed regions of the wake edges, the predictions greatly improve further downstream. It is worth noting here, that both skewness and flatness profiles are normalized by the velocity fluctuations which become numerically very small towards the edges of the wake, see Figure 6.8, thus the actual third and fourth moments, e.g. $\langle u^3 \rangle / U_0^3$ and $\langle u^4 \rangle / U_0^4$, would be better candidates for examining the accuracy of these higher order statistics. The good agreement of the skewness and flatness profiles indicates the model accurately captures the shape of the velocity PDF.



Figure 6.11: Skewness of the streamwise (a), (c) and cross-stream (b), (d) velocity components at different downstream locations behind the cylinder at $Re_D = 3900$. Red solid lines – PDF simulation, symbols – experiments of Ong and Wallace (1996). The horizontal dashed lines at each location indicate the Gaussian skewness value of 0.

6.3 Discussion

The laminar-to-turbulent transitional flow in the near wake of a circular cylinder at the subcritical Reynolds number of 3900 has been computed with a PDF method. The method has been applied for the first time to compute a flow with a complex geometry bounded by no-slip walls with significant curvature. The elliptic relaxation technique that is used to



Figure 6.12: Flatness of the streamwise (a), (c) and cross-stream (b), (d) velocity components at different downstream locations behind the cylinder at $Re_D = 3900$. Red solid lines – PDF simulation, symbols – experiments of Ong and Wallace (1996). The horizontal dashed lines at each location indicate the Gaussian flatness value of 3.

represent all components of the Reynolds stress tensor in the low-Reynolds-number wallregion has also been applied for the first time for highly curved boundaries with significant adverse pressure gradient resulting in boundary layer separation. Although mean spanwise motions are not represented in the current case thus only a two-dimensional Eulerian grid is employed to extract statistics, all three dimensions of the fluctuating velocity are retained.

Transient and time-averaged statistics of the joint PDF of the three velocity components have been compared to LES, DNS and experimental data. The predictions show significant improvement compared to past pure Eulerian RANS simulations. The quality and accuracy of the PDF results are comparable to three-dimensional LES and DNS predictions for all the quantities examined. One exception is the cross-stream Reynolds stress component $\langle v^2 \rangle$, which is noticeably underpredicted. This can be due to an overdiffusive numerical scheme (the lack of resolution of the Eulerian grid and/or the low-order accuracy of the temporal discretization), the lack of representation of the spanwise mean motions or a too large modeled turbulent diffusion. The advantages of the method can be summarized as follows:

- higher level statistical description of the stochastic fields than traditional RANS-type closures,
- a close interaction between the stochastic velocity and scalar fields,
- mathematically exact representation of advection, viscous diffusion, the effect of mean pressure and complex chemical reactions; these physical processes are treated without closure assumptions,
- the ability to follow highly distorted material surfaces accurately,
- the possibility for a relatively straightforward inclusion of history-dependent constitutive relations,
- excellent parallel performance.

A natural next step building on this work is to include the dispersion of passive scalars and develop a universal micromixing timescale that can be used in complex geometries in conjunction with the IEM/IECM models. Following this line a further step could be the inclusion of chemical reactions, in which the biggest advantage of the whole methodology lies, since it could be used to simulate chemically reactive turbulent flows surrounded by realistic no-slip walls in complex geometries, without the burden of the closure of the chemical source terms. Another obvious and straightforward way to expand on this work is to add the third dimension to the particle position and implement three-dimensional tetrahedra as the Eulerian grid, resulting in a fully three-dimensional code.

Further improvements to the code can be realized by employing edge-based data structures and solution techniques (Barth, 1991; Mavriplis, 1991; Peraire *et al.*, 1992; Luo *et al.*, 1993) to build the finite element coefficient matrices and the right hand sides for the two Eulerian equations, the elliptic relaxation and mean pressure projection. Although currently the solution of these equations takes only a small fraction of the running time, see Section 3.11, this will most likely change in three dimensions. In this case the solution will be more efficient with an edge-based solver because of the reduction of indirect addressing compared to the redundant element-based solution.

Currently, the simple Jacobi preconditioner is used to improve the convergence of the conjugate gradients solver for the mean pressure. More sophisticated preconditioners could also be explored to reduce the number of iterations, which will also significantly increase in three dimensions.

Porting the code to 3D will also result in excessive memory requirements if all nine components of the elliptic relaxation tensor are to be stored in every gridpoint as it is done throughout this study. More efficient elliptic relaxation could be achieved by employing different derivatives of the elliptic relaxation technique that store only a scalar variable instead of all 9 components (e.g. Wacławczyk *et al.*, 2004) and/or storing and solving \wp_{ij} only in the vicinity of walls, where the lengthscale *L*, Equation (2.17), has significant curvature.

Chapter 7: Summary and discussion

This work has presented a series of numerical methods to compute the one-point one-time joint PDF of turbulent velocity, characteristic frequency and scalar concentrations in high-Reynolds-number incompressible turbulent flows with complex geometries. Following the terminology coined by Muradoglu *et al.* (1999), we call the current methodology *non-hybrid* since an Eulerian CFD solver is not used in conjunction with the particle code to solve the PDF equations, i.e. the method is stand-alone. The method does belong to the familiy of particle-in-cell methods, where the Eulerian grid is used solely for: (i) estimating Eulerian statistics; (ii) tracking particles in the domain; and (iii) solving for quantities that are only represented in the Eulerian sense (i.e. mean pressure and elliptic relaxation). Compared to hybrid models, the current non-hybrid method assures that none of the fields are computed redundantly, therefore the simulation is kept consistent both numerically and at the level of turbulence closure without the need to enforce consistency conditions.

Adequate wall-treatment on the higher order statistics of the velocity field is achieved with an elliptic relaxation technique without damping or wall-functions, i.e. the boundary layers at solid (no-slip) walls are fully resolved. On the other hand, for application areas where full resolution of the turbulent boundary layers is not an option, we provide a treatment consistent with wall-functions that are commonly used in moment closures. The validation examples demonstrate the applicability of the algorithm in two-dimensional flows. A natural future direction along these lines is the extension to three spatial dimensions. This should be straightforward, since all the numerical methods are general enough and the extension only pertains to the Eulerian grid (e.g. tetrahedra instead of triangles) and an additional equation for the particle position, since all three components of particle velocities are already represented.

A significant challenge in stand-alone transported PDF methods is the accurate and stable computation of the mean pressure. This is mainly due to the following reasons: the mean velocity and Reynolds stresses have to be estimated from a noisy particle field and the pressure-Poisson equation requires their first and second derivatives, respectively, which are even noisier. We described a method to compute the mean pressure in conjunction with particle/PDF methods that only requires first derivatives of the mean velocity, which is based on a pressure-projection technique that has already been widely used and tested in laminar flows.

The two Eulerian equations needed by the algorithm are both solved on unstructured Eulerian grids with the finite element method. The last couple of decades have seen great strides in automatic unstructured grid generation, grid refinement and coarsening techniques and the development of highly sophisticated grid-based data structures that minimize cache misses. Using the algorithm presented in this work all this knowledge pertaining to unstructured meshes can be utilized in conjunction with the PDF equations and complex flow geometries. Employing finite elements together with particle/PDF methods also has the advantage of greatly simplifying boundary conditions for particles – no ghost elements are required as in finite volume methods. Furthermore, finite element approximation functions are not only used for particle tracking but also provide an elegant way of estimating derivatives of statistics from particle fields.

We also described a general algorithm that can be used to calculate the velocityconditioned scalar mean for the IECM micromixing model. The procedure homogenizes the statistical error over the sample space for arbitrary velocity PDFs by dynamically adjusting the number of bins and their distribution.

A particle-redistribution algorithm has also been described that provides stability by ensuring that no Eulerian elements remain without particles at any time during the simulation. This task has traditionally been accomplished via particle splitting and merging techniques. However, computationally it is more efficient not to introduce or eliminate particles during timestepping, so that the arrays storing physical properties can keep their original size and can remain consecutively accessible with minimal cache misses. Both particle splitting and merging algorithms and the current redistribution procedure do change the local PDF, and this is certainly an undesired effect. We are not aware of any algorithm in the literature which accomplishes particle-number control without altering the underlying numerically computed joint PDF. In particle splitting and merging algorithms mass may be conserved, but fulfillment of all mass, momentum and kinetic energy conservation is in general not possible in single events, only statistically (Rembold and Jenny, 2006). Our method is no exception. We presented an error analysi employing a simplified set of particle equations on a homogeneous example. We believe, that further tests are certainly necessary to investigate the error introduced by the redistribution algorithm in inhomogeneous flows. Also it is worth pointing out that there is no clear or established benchmark to investigate the effects of redistribution algorithms, as the effects may be space-dependent, and their importance is relative to the application.

We also proposed a general form for the micromixing timescale that can be used in a flow-, and geometry-independent manner for modeling the effect of small-scale mixing on a transported passive scalar released from a concentrated source. Although the computed concentration results compare well with analytical and experimental data for the testcases, this is to be considered under heavy development regarding both its mathematical expression and modeling constants.

The solver has been parallelized with the OpenMP standard, which easily allows the exploitation of multi-core workstations mainly used for production codes. Our performance study has shown a good parallel speedup up to 32 CPUs tested on shared memory machines using single-, dual-, and quad-core CPUs. We also ported the code to Intel's Cluster OpenMP technology, which allows an OpenMP program to run on a beowulf-type cluster of networked workstations requiring a minor programing effort compared to an MPI-based implementation. However, we found that the algorithm with its current design is not suitable for Cluster OpenMP.

Three testcases of increasing complexity have been presented to demonstrate the applicability of the algorithm. The resulting fields show a good agreement when compared to DNS and experimental data where available. In the future, more tests with cases of different complexity will definitely need to be carried out. This is especially true for the micromixing timescale, which has to be tested in different flows and for different source scenarios to assess the validity of its form and its modeling constants.

The hybrid methods that combine existing Eulerian CFD solvers with the PDF methodology are based on RANS and LES methods. Both of these lines of development concentrate on the modeling of chemical reactions which appear in mathematically closed form in the PDF framework. The Eulerian governing equations consist of the fully compressible equations for conservation of mass, momentum and energy. This system is augmented by a set of stochastic equations for Lagrangian particles that represent species' concentrations and may also provide turbulence closure depending on how the fluctuating velocity field is represented. Furthermore, the mean pressure is obtained from an equation of state. In these hybrid methods, since the preferred way of representing flow variables is via Favreaveraging, the density must also appear explicitly at the Lagrangian particle level. Since the mean continuity equation is also required to solve the mean Eulerian governing equations, the density is represented redundantly. Consequently, consistency must be ensured explicitly. The currently proposed non-hybrid method is stand-alone and represents mass consistency without redundancy, only by ensuring that the mean velocity field is divergencefree at all times by solving a Poisson equation for the mean pressure. Thus particles (and particle number-density) do not represent mass (or real fluid density) and no additional mass-consistency condition is required. Therefore the problem of a high degree of variation in particle-number density between different regions of the flow field only amounts to different statistical errors, without the additional computational errors introduced by inconsistency. This is possible precisely, because we solve the fully incompressible equations, where the density is constant. This advantage, however, may certainly become a disadvantage in turbulent chemistry even in an otherwise incompressible flow if the stochastic density variations due to chemical reactions have to be represented accurately. Therefore the limitations of the current methodology regarding its applicability in conjunction with chemical reactions remain to be seen.

Regarding computational costs, Pope (2000) places PDF methods somewhere between Reynolds stress closures and large eddy simulation. We showed that he most expensive parts of the current non-hybrid method are the advancement of particle properties and random number generation, which together account for more than 90% of the computational cost. Both hybrid and non-hybrid methods need to advance and track particles, generate random numbers, estimate Eulerian statistics, solve for the mean pressure and ensure sufficient number of particles everywhere on the flow domain. Therefore we expect the computational costs of these components to be comparable for the two methodologies. In addition to the above, hybrid methods need to enforce consistency conditions and solve the Eulerian system of governing equations as well. This very approximate analysis suggests that there is no clear reason to think that the total cost of the two methods will be very different. However, it would be valuable to perform thorough side-by-side comparisons among the different standalone (fully Lagrangian), hybrid RANS, hybrid LES and non-hybrid methods in order to have a better picture on their relative computational costs. Additionally to these approaches, it would be interesting to explore a method that solves the fully incompressible equations, just like the current non-hybrid method, but in the Eulerian framework, which represents only the scalar concentrations by Lagrangian particles. We expect the computational cost of such a method to be significantly less than the current method, since the velocity field would not be represented by particles. This would be beneficial in situations where higher order statistics of the velocity are not required and a close interaction between the stochastic velocity and scalar concentration fields is not important. The value of such a method may be limited in applications of turbulent chemistry, but the higher level statistical description of the scalar fields could be advantageous in atmospheric pollution modeling.

On the other hand, the close interaction between the stochastic velocity and scalar concentration fields (that both hybrid methods for velocity and scalars and the current nonhybrid method can provide) is important as a research tool to shed light on micromixing mechanisms; to provide information and data on higher order statistics; and in applications where accurate modeling of the micromixing of scalars is required.

Appendix A: Derivation of the Eulerian PDF transport equation

In PDF methods of turbulent flows, the velocity and transported scalar fields are considered as time dependent, multivariate random fields (Pope, 2000). In other words, as opposed to a deterministic theory, the three components of the velocity and the scalar concentration are represented by a joint probability distribution function containing the full one-point, onetime statistics of the velocity and the scalar. It should be emphasized, however, that this one-point, one-time description does not contain information regarding other points in space and time, therefore – as it will be shown – does not provide a complete description of the random velocity and scalar fields. As a consequence, the one-point, one-time PDF contains no information about the length-scale or frequency of the fluctuations, thus appropriate models are necessary to supply this missing information in the form of models. In an incompressible turbulent flow containing scalar tracers, the state of the fluid at any location is described by the instantaneous Eulerian velocity U(x, t), pressure P(x, t) and the species' mass concentrations $\phi(x, t)$. In the following, the PDF transport equation (2.4) is derived starting from the system of Eulerian governing equations (2.1-2.3).

Let $f(\mathbf{V}, \psi; \mathbf{x}, t)$ denote the one-point, one-time Eulerian joint PDF of the random velocity $\mathbf{U}(\mathbf{x}, t)$ and scalar $\phi(\mathbf{x}, t)$, where the three dimensional Euclidean space (V_1, V_2, V_3) is the sample space of the random velocity vector $\mathbf{U} = (U_1, U_2, U_3)$ and ψ is the sample space variable of the random scalar concentration ϕ . Table A.1 summarizes the random flow variables and their sample spaces in the joint PDF $f(\mathbf{V}, \psi; \mathbf{x}, t)$. This can also be viewed as an eight-dimensional scalar-valued function having a unique value at each location of the eight-dimensional Euclidean state-space

$$f(V_1, V_2, V_3, \psi; x_1, x_2, x_3, t) \colon \mathbb{R}^8 \to \mathbb{R},$$
 (A.1)

thus for every different set of eight independent variables $(V_1, V_2, V_3, \psi; x_1, x_2, x_3, t)$ the function f corresponds to a single scalar. An other way to look at this, is to have 3+1 scalar functions $(\mathbb{R} \to \mathbb{R})$ at each point in space and time. This increased dimensionality is characteristic of PDF methods, since every single random scalar variable is represented with its probability density distribution, so instead of a single scalar, a scalar function, its probability distribution, is taken into account.

The transport equation for the joint PDF $f(\mathbf{V}, \psi; \mathbf{x}, t)$ can be derived from the conservation equations (2.2-2.3), which are rewritten here in a more convenient form:

$$\frac{\mathrm{D}U_i}{\mathrm{D}t} = A_i, \quad \text{where} \quad A_i = \nu \nabla^2 U_i - \frac{1}{\rho} \frac{\partial P}{\partial x_i}, \tag{A.2}$$

$$\frac{D\phi}{Dt} = B$$
, where $B = \Gamma \nabla^2 \phi$, (A.3)

where the substantial derivative is denoted by

$$\frac{\mathrm{D}}{\mathrm{D}t} \equiv \frac{\partial}{\partial t} + U_j \frac{\partial}{\partial x_j}.\tag{A.4}$$

There are several ways of deriving PDF transport equations. A useful method involving delta functions is described by Pope (2000), but here a different approach is followed, which has been used by Pope (1985) and also by Fox (2003). The method is based on equating two independent expressions for $\langle DQ/Dt \rangle$, where $Q(U, \phi)$ is "almost" any function¹. The first expression for $\langle DQ/Dt \rangle$ is obtained by employing the definition of the substantial derivative (A.4) and the mathematical expectation of a random function

 $^{{}^{1}}Q(\boldsymbol{U},\phi)$ is an arbitrary function, however, it has necessary properties so that its statistics, like Equations (2.5) and (2.6), are not divergent.

Table A.1: Random flow variables and their corresponding sample spaces in the joint PDF $f(\mathbf{V}, \psi; \mathbf{x}, t)$.

Quantity	Random variable	Sample space	
velocity	$\boldsymbol{U} = (U_1, U_2, U_3)$	$oldsymbol{V}=(V_1,V_2,V_3)$	
scalar	ϕ	ψ	

$$\langle Q(\boldsymbol{U},\phi)\rangle = \int Q(\boldsymbol{V},\psi)f(\boldsymbol{V},\psi)\mathrm{d}\boldsymbol{V}\mathrm{d}\psi$$
:

$$\left\langle \frac{\mathrm{D}Q}{\mathrm{D}t} \right\rangle = \left\langle \frac{\partial Q(\boldsymbol{U},\phi)}{\partial t} \right\rangle + \left\langle U_i \frac{\partial Q(\boldsymbol{U},\phi)}{\partial x_i} \right\rangle$$

$$= \frac{\partial}{\partial t} \int Q(\boldsymbol{V},\psi) f(\boldsymbol{V},\psi;\boldsymbol{x},t) \mathrm{d}\boldsymbol{V} \mathrm{d}\psi + \frac{\partial}{\partial x_i} \int V_i Q(\boldsymbol{V},\psi) f(\boldsymbol{V},\psi;\boldsymbol{x},t) \mathrm{d}\boldsymbol{V} \mathrm{d}\psi \quad (A.5)$$

$$= \int Q(\boldsymbol{V},\psi) \left\{ \frac{\partial f}{\partial t} + V_i \frac{\partial f}{\partial x_i} \right\} \mathrm{d}\boldsymbol{V} \mathrm{d}\psi.$$

The second expression can be deduced by relating changes in Q to changes in U and ϕ as

$$\frac{\mathrm{D}Q}{\mathrm{D}t} = \frac{\partial Q}{\partial U_i} \frac{\mathrm{D}U_i}{\mathrm{D}t} + \frac{\partial Q}{\partial \phi} \frac{\mathrm{D}\phi}{\mathrm{D}t},\tag{A.6}$$

where the material derivatives can be replaced by A_i and B from Equations (A.2) and (A.3) and we also take take the expectation as

$$\left\langle \frac{\mathrm{D}Q}{\mathrm{D}t} \right\rangle = \left\langle \frac{\partial Q}{\partial U_i} A_i \right\rangle + \left\langle \frac{\partial Q}{\partial \phi} B \right\rangle. \tag{A.7}$$

Note, that in general, A_i and B depend on multi-point information of the random fields Uand ϕ , for example, they depend on the velocity and scalar gradients and Laplacians. Since these quantities are not contained in the one-point, one-time PDF $f(V, \psi; x, t)$, let these additional unknowns be collected into the vector Z(x, t). Furthermore, let the joint PDF of U, ϕ and Z be $f_{U\phi Z}(V, \psi, z; x, t)$. According to Bayes' rule (van Kampen, 2004), this can be written as the product of a conditional and a marginal PDF as

$$f_{\boldsymbol{U}\phi\boldsymbol{Z}}(\boldsymbol{V},\psi,\boldsymbol{z}) = f_{\boldsymbol{Z}|\boldsymbol{U}\phi}(\boldsymbol{z}|\boldsymbol{V},\psi)f(\boldsymbol{V},\psi).$$
(A.8)

Thus the first term on the right hand side of Equation (A.7) can be rewritten as

$$\left\langle \frac{\partial Q}{\partial U_i} A_i \right\rangle = \int \frac{\partial Q(\boldsymbol{V}, \psi)}{\partial V_i} A_i(\boldsymbol{V}, \psi, \boldsymbol{z}) f_{U\phi \boldsymbol{z}}(\boldsymbol{V}, \psi, \boldsymbol{z}) \mathrm{d} \boldsymbol{V} \mathrm{d} \psi \mathrm{d} \boldsymbol{z}$$

$$= \int \frac{\partial Q(\boldsymbol{V}, \psi)}{\partial V_i} \langle A_i | \boldsymbol{V}, \psi \rangle f(\boldsymbol{V}, \psi) \mathrm{d} \boldsymbol{V} \mathrm{d} \psi, \qquad (A.9)$$

where the conditional expectation of the acceleration A_i is

$$\langle A_i | \boldsymbol{V}, \psi \rangle = \int A_i(\boldsymbol{V}, \psi, \boldsymbol{z}) f_{\boldsymbol{z} | \boldsymbol{U} \phi}(\boldsymbol{z} | \boldsymbol{V}, \psi) \mathrm{d} \boldsymbol{z}.$$
 (A.10)

Note, that $\langle A_i | \mathbf{V}, \psi \rangle$ is a function only of \mathbf{V} and ψ (additionally to the implicit dependence on \mathbf{x} and t) since all the unknowns \mathbf{Z} have been integrated out. Integration by parts yields

$$\left\langle \frac{\partial Q}{\partial U_{i}} A_{i} \right\rangle = \int \frac{\partial}{\partial V_{i}} \Big[Q(\boldsymbol{V}, \psi) \langle A_{i} | \boldsymbol{V}, \psi \rangle f(\boldsymbol{V}, \psi) \Big] d\boldsymbol{V} d\psi - \int Q(\boldsymbol{V}, \psi) \frac{\partial}{\partial V_{i}} \Big[\langle A_{i} | \boldsymbol{V}, \psi \rangle f(\boldsymbol{V}, \psi) \Big] d\boldsymbol{V} d\psi,$$
(A.11)

where the first term on the right hand side vanishes provided that Q is monotonic as |V|tends to infinity and the expectation $\langle A_i Q \rangle$ exists.² A similar procedure can be followed for the term containing B in Equation (A.7) to obtain

$$\left\langle \frac{\partial Q}{\partial \phi} B \right\rangle = -\int Q(\mathbf{V}, \psi) \frac{\partial}{\partial \psi} \Big[\langle B | \mathbf{V}, \psi \rangle f(\mathbf{V}, \psi) \Big] \mathrm{d}\mathbf{V} \mathrm{d}\psi.$$
(A.12)

²These conditions are given in practically all flow conditions that one can encounter, see also (Pope, 1985).

Substituting Equations (A.11) and (A.12) into Equation (A.7) the second expression for $\langle DQ/Dt \rangle$ can be obtained:

$$\left\langle \frac{\mathrm{D}Q}{\mathrm{D}t} \right\rangle = -\int Q(\mathbf{V},\psi) \left\{ \frac{\partial}{\partial V_i} \left[\langle A_i | \mathbf{V},\psi \rangle f(\mathbf{V},\psi) \right] + \frac{\partial}{\partial \psi} \left[\langle B | \mathbf{V},\psi \rangle f(\mathbf{V},\psi) \right] \right\} \mathrm{d}\mathbf{V} \mathrm{d}\psi.$$
(A.13)

Substracting the second expression for $\langle DQ/Dt \rangle$ (A.13) from the first expression (A.5) we obtain

$$\int Q(\mathbf{V},\psi) \left\{ \frac{\partial f}{\partial t} + V_i \frac{\partial f}{\partial x_i} + \frac{\partial}{\partial V_i} \left[\langle A_i | \mathbf{V}, \psi \rangle f \right] + \frac{\partial}{\partial \psi} \left[\langle B | \mathbf{V}, \psi \rangle f \right] \right\} \mathrm{d}\mathbf{V} \mathrm{d}\psi = 0.$$
(A.14)

Since this equation holds for "almost" any function $Q(\mathbf{V}, \psi)$, the term in the brackets must sum to zero, thus we obtain the transport equation for the velocity-scalar joint PDF $f(\mathbf{V}, \psi; \mathbf{x}, t)$

$$\frac{\partial f}{\partial t} + V_i \frac{\partial f}{\partial x_i} = -\frac{\partial}{\partial V_i} \Big[\langle A_i | \boldsymbol{V}, \psi \rangle f \Big] - \frac{\partial}{\partial \psi} \Big[\langle B | \boldsymbol{V}, \psi \rangle f \Big].$$
(A.15)

As it can be seen, in physical space the joint PDF f evolves due to the velocity field V, in velocity space due to the conditional acceleration $\langle A_i | V, \psi \rangle$ and in concentration space due to the conditional diffusion term $\langle B | V, \psi \rangle$. Now, we substitute the right hand sides of the momentum and scalar conservation equations from (A.2) and (A.3) and arrive at Equation (2.4)

$$\frac{\partial f}{\partial t} + V_i \frac{\partial f}{\partial x_i} = -\frac{\partial}{\partial V_i} \left[\left\langle \nu \nabla^2 U_i - \frac{1}{\rho} \frac{\partial P}{\partial x_i} \middle| \mathbf{V}, \psi \right\rangle f \right] - \frac{\partial}{\partial \psi} \left[\left\langle \Gamma \nabla^2 \phi \middle| \mathbf{V}, \psi \right\rangle f \right].$$
(A.16)

After decomposing the pressure P into its mean $\langle P \rangle$ and fluctuating part p, another useful

form of this equation is derived by Pope (2000)

$$\frac{\partial f}{\partial t} + V_i \frac{\partial f}{\partial x_i} = \nu \frac{\partial^2 f}{\partial x_i \partial x_i} + \frac{1}{\rho} \frac{\partial \langle P \rangle}{\partial x_i} \frac{\partial f}{\partial V_i} - \frac{\partial^2}{\partial V_i \partial V_j} \left[f \left\langle \nu \frac{\partial U_i}{\partial x_k} \frac{\partial U_j}{\partial x_k} \middle| \mathbf{V}, \psi \right\rangle \right]
+ \frac{\partial}{\partial V_i} \left[f \left\langle \frac{1}{\rho} \frac{\partial p}{\partial x_i} \middle| \mathbf{V}, \psi \right\rangle \right] - \frac{\partial}{\partial \psi} \left[f \left\langle \Gamma \nabla^2 \phi \middle| \mathbf{V}, \psi \right\rangle \right],$$
(A.17)

where it is apparent that convection, the effect of mean pressure and viscous diffusion appear mathematically exactly. On the other hand, the physical processes of dissipation of turbulent kinetic energy, pressure redistribution and the small-scale mixing of the scalar, denoted by the three conditional expectations which require modelling assumptions, can also be recognized explicitly.

Appendix B: Computation of the velocity-conditioned scalar mean

In Section 3.6 a numerical strategy to estimate the velocity-conditioned scalar mean $\langle \phi | \mathbf{V} \rangle$ required in Equation (2.26) is detailed. An algorithm that accomplishes the conditioning step after the particles have been sorted in element **e** into subgroups may be written as follows. Let $\text{CNBI}(=N_c)$, $\text{NELEM}(=N_e)$, $\text{NPAR}(=N_p)$ and $\text{MAXNPEL}(=N_{p/e}^{\max})$ denote the number of conditioning bins, the total number of elements of the Eulerian grid, the total number of particles and the maximum number of particles per elements, respectively. Furthermore, let the arrays np[CNBI], vcce[NELEM*CNBI], npel[NELEM], parid[MAXNPEL] and parc[NPAR] represent the number of particles in bins, the velocity-conditioned scalar concentration in bins of each element, the actual number of particles in each element, the indices of the particles residing in element **e** and the particle concentrations, respectively. (Note the use of C-style indexing, i.e. the array indices start from 0. Comments are initiated by "//" and typeset in gray.)

sort parid[0:MAXNPEL-1] according to the sorting & dividing procedure;

initialize np[0:CNBI-1] = vcce[0:CNBI-1] = n = 0;

for all particles in element e

- // compute bin index
- i = CNBI*n/npel[e];

// increase number of particles in bin i

np[i] = np[i] + 1;
// add particle concentration to bin i

vcce[e*CNBI+i] = vcce[e*CBI+i] + parc[parid[n]];

// store conditioning pointer for particle

cp[parid[n]] = i;

// increase number of particles considered

n = n + 1;

end

// finish computing conditional mean in bin i

for all bin i

```
vcce[e*CNBI+i] = vcce[e*CNBI+i]/np[i];
```

end

After this algorithm, the array cp[NPAR] will contain conditioning pointers for each particle relative to their host element, so that the velocity-conditioned scalar mean $\langle \phi | \mathbf{V} \rangle$ for particle p in element e can be obtained as vcce[e*CNBI+cp[p]].

Appendix C: Basic particle redistribution algorithm

In Section 3.8 the need for a particle redistribution algorithm is emphasized. What follows is such an algorithm that we employ in order to keep the number of particles/element above a certain treshold.

do {

find the elements (mine, maxe) containing the
 smallest and largest number of particles (minnpel, maxnpel);
if { (minnpel < MINNPEL) and (minnpel ≠ maxnpel) }
 move a particle from element maxe to mine;
regenerate array npel and linked lists psel1, psel2;
} while { (minnpel < MINNPEL) and (minnpel ≠ maxnpel) };</pre>

The loop stops if the required minimum number of particles/element MINNPEL(= $N_{p/e}^{\min}$) is reached or the element-distribution of particles becomes homogeneous over the elements. Any particle may be moved from element maxe to mine as long as the local statistics are not altered. In principle, this can be achieved if the properties ($\mathcal{U}_i, \omega, \psi$) of the newly arriving particle in element mine are sampled from the local joint PDF. A quick way of doing this is to initialize the particle properties by copying a randomly chosen particle already residing in element mine. Since the joint PDF is represented by a finite number of particles, taking out a particle from element maxe and putting it into mine will alter the local statistics in both elements, even if the new properties are copied from a neighbor. Since maxe contained the largest number of particles on the whole domain, we are less concerned about the effect of a single leaving particle since the local PDF is well represented there. However, the effect of the newly introduced particle in element mine, where the joint PDF was already poorly represented, is of higher importance. Thus in Section 3.8 we investigate the error introduced by the above particle redistribution using a simplified governing equation.

Array npel stores the number of particles in each element, while the linked lists psel1 and psel2 stores the particle indices in each element. These arrays are regenerated after each particle moved, since finding the elements with the smallest and largest number of particles requires npel, while moving a particle requires a particle index from the old and the new host element.

Appendix D: A more efficient particle redistribution algorithm

Appendix C introduces the basic idea of the algorithm that is used to ensure enough particles in every element at all times. That algorithm is simple and robust, however it is not very efficient because of the brute-force nature of finding the elements with the smallest and largest number of particles. It executes these searches before each particle is moved, which, in principle, may not be required. Also, after it moves each particle, it regenerates the linked lists that store the indices of particles in each element (**psel1** and **psel2**) and the array that stores the number of particles per elements (**npe1**). Again, although this is not strictly required after moving each particle, the simple organization of that **while** loop requires it.

A much more efficient way of performing the above redistribution is as follows. First, in a temporary array we sort the indices of the elements into the order of increasing number of particles. Now we have all the elements that contain the smallest and the largest number of particles clustered in the bottom and top part of that array. Then the redistribution step consists of a nested loop over only the critical elements (that have less particles than MINNPEL) with an inner loop that iterates over the number of missing particles in the given element. The body of the loop is the same as before, i.e. removing a particle from an element that has many particles and adding it into one that does not have enough, copying the particle properties from a random particle that already resides there. In pseudo/C code this whole procedure is written as follows. (The array elp[NPAR] stores the element index of each particle. Array indexing starts from 0. The structure of the linked lists psel1 and psel2 follows Löhner (2000). Comments are initiated by "//" and typeset in gray.) // make a copy of array npel and its indices to npels and npeli

// also count the number of critical elements (the ones that have less particles than MINNPEL)

sort temporary array npels and drag along the indices npeli;

// now we have the elements with the least number of particles at the bottom of
// array npels and the elements with the most number of particles at the top

// redistribute particles from the top to the bottom

// loop over critical elements from the bottom up until we reach MINNPEL

for (e = 0; e < nce; e++)

{

// loop over the number of missing particles in each critical element

```
for ( p = 0; p < MINNPEL-npels[e]; p++ )</pre>
```

{

// get element index from the top where particle will be moved from

i = npeli[nelem-e-1];

// get a particle index from the top (this one will be moved to the bottom)

pi = psel1[psel2[i]+1+p]; // starting from the first one, get next

// get element index in the bottom where particle will be moved to

j = npeli[e];

// get a particle index in the bottom (whose properties will be used to // initialize the newly arriving particle)

if (npels[e] > 0) // if there is at least one particle

// starting from the first get the next one, restart if exhausted

copy particle properties from particle pj to pi;

```
elp[pi] = j; // store particle's new element number
}
npel[i] = npel[i] - p; // take out p particles from top element
npel[j] = npel[j] + p; // put p particles into bottom element
```

regenerate array npel and linked lists psel1, psel2;

}

The above procedure removes particles from elements at the top of array npels and adds them into elements at the bottom, initializing the newly arriving particle properties with

Table D.1: Timings for the two particle redistribution algorithms, described in the current chapter and in Appendix C. The case is the computation of the cylinder flow detailed in Chapter 6, employing an Eulerian mesh of approximately 50K triangles, 2.5 million particles with requiring a minimum of 5 particles per elements at all times (MINNPEL=5), which results in a redistribution of 200-300 particles moved in every timestep. The data are relevant to a single timestep using 8 processor cores.

algorithm	number of parti-	clock time for redis-	total time of a	%
	cles redistributed	tribution only, ms	whole timestep, ms	
basic	261	16881.6	18371	91.8
improved	270	85.8	1251	6.8
speedup		196.8	14.7	

one of those already in the critical element. In essence, this is the same as in Appendix C, but now we only operate on the elements that contain the smallest and largest number of particles. The big advantages are that now the brute-force searches are completely eliminated, we only access data which we have to modify (further reducing a large number of cache misses) and the array npel and linked lists psel1 and psel2 have to be regenerated only once and not for all particles moved. Additionally, the parallelization of the new algorithm is simpler. The brute-force searches required at least one synchronization point (when a new minimum or maximum was found and had to be updated), while parallelization of the new algorithm is trivial and requires no synchronization at all.

Simple tests indicate that this algorithm in itself is about 200 times faster than the one described in Appendix C using only about 2.5 million particles with 50K Eulerian elements. Table D.1 shows some timings comparing the two different algorithms computing the cylinder case, where continuous redistribution of about 200-300 particles per timestep is required after the vortex shedding has been initiated. We see that the old algorithm accounts for more than 90% of the total running time, while the new one a mere 6.8%, resulting in

an overall speedup of the code of almost 15 times. The improvement is expected to be even more significant with larger cases, more complex flows and more processors. Bibliography

Bibliography

Abe, H., Kawamura, H., Matsuto, Y., 2004. Surface heat-flux fluctuations in a turbulent channel flow up to $Re_{\tau} = 1020$ with Pr = 0.025 and 0.71. Int. J. Heat Fluid Fl. 25 (3), 404–419.

Arnold, L., 1974. Stochastic Differential Equations: Theory and Applications. Krieger Publishing Company.

Arya, P. S., 1999. Air pollution meteorology and dispersion. Oxford University Press.

Bacon, D. P., Ahmad, N. N., Boybeyi, Z., Dunn, T. J., Hall, M. S., Lee, P. C. S., Sarma, R. A., Turner, M. D., III., K. T. W., Young, S. H., Zack, J. W., 2000. A dynamically adapting weather and dispersion model: The Operational Multiscale Environment Model with Grid Adaptivity (OMEGA). Mon. Weather Rev. 128, 2044–2076.

Baik, J.-J., Kim, J.-J., 1999. A numerical study of flow and pollutation dispersion characteristic in urban street canyons. J. Appl. Meteorol. 38, 1576–1589.

Bakosi, J., Franzese, P., Boybeyi, Z., 2007. Probability density function modeling of scalar mixing from concentrated sources in turbulent channel flow. Phys. Fluids 19 (11), 115106. URL http://link.aip.org/link/?PHF/19/115106/1

Bakosi, J., Franzese, P., Boybeyi, Z., 2008. A non-hybrid method for the PDF equations of turbulent flows on unstructured grids. Accepted for publication in J. Comput. Phys.

Barth, T., 1991. A 3-D Upwind Euler Solver for Unstructured Meshes. AIAA-91-1548-CP.

Batchelor, G. K., 1952. Diffusion in a field of homogeneous turbulence II. The relative motion of particles. P. Camb. Philos. Soc. 48 (2), 345–362.

Beaudan, P., Moin, P., 1994. Numerical experiments on the flow past a circular cylinder at subcritical Reynolds number. Tech. Rep. Report TF-62, Department of Mechanical Engineering, Stanford University.

Berger, E., Wille, R., 1972. Periodic flow phenomena. Annu. Rev. Fluid Mech. 4, 313–340.

Bezuglov, V. A., 1974. Measurement of velocity of concentration transport. Tech. Rep. Pt. 1, Dolgoprudnyi, Moscow Physical-Technical Institute.

Bilger, R. W., V., L. R. S. L., Krishnamoorthy, 1991. Reaction in a scalar mixing layer. J. Fluid Mech. 233, 211.

Brethouwer, G., Nieuwstadt, F. T. M., 2001. DNS of mixing and reaction of two species in a turbulent channel flow: a validation of the conditional moment closure. Flow Turbul. Combust. 66, 209–239.

Camelli, F., Löhner, R., Mestreau, E. L., 2007. Timings of an unstructured-grid CFD code on common hardware platforms and compilers. AIAA-07-1107.

Cantwell, B., Coles, D., 1983. An experimental study of entrainment and transport in the turbulent near wake of a circular cylinder. J. Fluid Mech. 136, 321–374.

Cassiani, M., Franzese, P., Giostra, U., 2005a. A PDF micromixing model of dispersion for atmospheric flow. Part I: development of the model, application to homogeneous turbulence and to neutral boundary layer. Atmos. Environ. 39 (8), 1457–1469.

Cassiani, M., Franzese, P., Giostra, U., 2005b. A PDF micromixing model of dispersion for atmospheric flow. Part II: application to convective boundary layer. Atmos. Environ. 39 (8), 1471–1479.

Cassiani, M., Radicchi, A., Albertson, J. D., 2007a. Modelling of concentration fluctuations in canopy turbulence. Bound.-Lay. Meteorol. 122 (3), 655–681.

Cassiani, M., Radicchi, A., Albertson, J. D., Giostra, U., 2007b. An efficient algorithm for scalar PDF modelling in incompressible turbulent flow; numerical analysis with evaluation of IEM and IECM micro-mixing models. J. Comput. Phys. 223, 519–550.

Chatwin, P. C., 2002. Some remarks on modelling the PDF of the concentration of a dispersing scalar in turbulence. Eur. J. Appl. Math. 13, 95–108.

Chatwin, P. C., 2004. Singular PDFs of a dispersing scalar in turbulence. Flow Turbul. Combust. 72, 273–285.

Chatwin, P. C., Sullivan, P. J., 1993. The structure and magnitude of concentration fluctuations. Bound.-Lay. Meteorol. 62, 269–280.

Chorin, A. J., 1968. Numerical solution of the Navier-Stokes equations. Math. Comput. 22, 745–762.

Colucci, P. J., Jaberi, F. A., Givi, P., Pope, S. B., 1998. Filtered density function for large eddy simulation of turbulent reacting flows. Phys. Fluids 10 (2), 499–515. URL http://link.aip.org/link/?PHF/10/499/1

Correa, S. M., Pope, S. B., 1992. Comparison of a Monte Carlo PDF/finite-volume mean flow model with bluff-body raman data. In: Twenty-fourth Symposium (International) on Combustion. Pittsburgh, PA: The Combustion Institute, pp. 279–285.

Courant, R., Friedrichs, K., Lewy, H., 1928. Über die partiellen differenzengleichungen der mathematischen physik. Mathematische Annalen 100 (1), 32–74.

Craft, T. J., Launder, B. E., 1991. Computation of impinging flows using second moment closures. In: Proc. 8th symposium on Turbulent Shear Flows, T. U. Munich, Germany. pp. 5–8.

Craft, T. J., Launder, B. E., 1996. Reynolds stress closure designed for complex geometries. Int. J. Heat Fluid Fl. 17, 245–254.

Delarue, B. J., Pope, S. B., 1997. Application of PDF methods to compressible turbulent flows. Phys. Fluids 9 (9), 2704–2715.

URL http://link.aip.org/link/?PHF/9/2704/1

Dong, S., Karniadakis, G. E., Ekmekci, A., Rockwell, D., 2006. A combined direct numerical simulation–particle image velocimetry study of the turbulent near wake. J. Fluid Mech. 569, 185–207.

Dopazo, C., 1975. Probability density function approach for a turbulent axisymmetric heated jet. Centerline evolution. Phys. Fluids 18 (4), 397–404. URL http://link.aip.org/link/?PFL/18/397/1

Dopazo, C., 1994. Recent developments in pdf methods. In: Libby, P. A. (Ed.), Turbulent reactive flows. Academic, New York.

Dopazo, C., O'Brien, E. E., 1974. An approach to the autoignition of a turbulent mixture. Acta Astronaut. 1, 1239–1266.

Dreeben, T. D., 1997. PDF modeling of near-wall turbulent flows. Ph.D. thesis, Cornell University.

Dreeben, T. D., Pope, S. B., 1997a. Probability density function and Reynolds-stress modeling of near-wall turbulent flows. Phys. Fluids 9 (1), 154–163. URL http://link.aip.org/link/?PHF/9/154/1

Dreeben, T. D., Pope, S. B., 1997b. Wall-function treatment in pdf methods for turbulent flows. Phys. Fluids 9 (9), 2692–2703.

URL http://link.aip.org/link/?PHF/9/2692/1

Dreeben, T. D., Pope, S. B., 1998. Probability density function/Monte Carlo simulation of near-wall turbulent flows. J. Fluid Mech. 357, 141–166.

Durbin, P. A., 1991. Near-wall turbulence modeling without damping functions. Theor. Comp. Fluid Dyn. 3, 1–13.

Durbin, P. A., 1993. A Reynolds stress model for near-wall turbulence. J. Fluid Mech. 249, 465–498.

Einstein, A., 1926. Investigations on the theory of the brownian movement. Methuen, London, translation by A. D. Cowper.

Entacher, K., Uhl, A., Wegenkittl, S., 1998. Linear congruential generators for parallel Monte-Carlo: the Leap-Frog case. Monte Carlo Methods and Applications 4 (1), 1–16. URL http://citeseer.ist.psu.edu/entacher97linear.html

Fackrell, J. E., Robins, A. G., 1982. Concentration fluctuations and fluxes in plumes from point sources in a turbulent boundary layer. J. Fluid Mech. 117, 1–26.

Fox, R. O., 1996. On velocity-conditioned scalar mixing in homogeneous turbulence. Phys. Fluids 8 (10), 2678–2691.

URL http://link.aip.org/link/?PHF/8/2678/1

Fox, R. O., 2003. Computational models for turbulent reacting flows. Cambridge University Press.

Franke, R., Rodi, W., Schönung, B., 1989. Analysis of experimental vortex shedding data with respect to turbulence modeling. In: 7th Symposium on Turbulent Shear Flows, Stanford University.

Franzese, P., 2003. Lagrangian stochastic modeling of a fluctuating plume in the convective boundary layer. Atmos. Environ. 37, 1691–1701.

Franzese, P., Cassiani, M., 2007. A statistical theory of turbulent relative dispersion. J. Fluid Mech. 571, 391–417.

Frohlich, J., Rodi, W., Kessler, P., Parpais, S., Bertoglio, J., Laurence, D., 1998. Large eddy simulation of flow around circular cylinders on structured and unstructured grids. In: Hirschel, E. H. (Ed.), Notes on Numerical Fluid Mechanics. Vieweg.

Gardiner, C. W., 2004. Handbook of stochastic methods for physics, chemistry and the natural sciences, 3rd Edition. Springer-Verlag, Berlin Heidelberg New York.

Ge, H.-W., Zhu, M.-M., Chen, Y.-L., Gutheil, E., 2007. Hybrid unstready RANS and PDF method for turbulent non-reactive and reactive flows. Flow Turbul. Combust. 78, 91–109.

Germano, M., Piomelli, U., Moin, P., Cabot, W. H., 1991. A dynamic subgrid-scale eddy viscosity model. Phys. Fluids 3 (7), 1760–1765.

URL http://link.aip.org/link/?PFA/3/1760/1

Geuzaine, C., Remacle, J., 2006. Gmsh: a three-dimensional finite element mesh generator with built-in pre- and post-processing facilities. URL http://www.geuz.org/gmsh

Gicquel, L. Y. M., Givi, P., Jaberi, F. A., Pope, S. B., 2002. Velocity filtered density function for large eddy simulation of turbulent flows. Phys. Fluids 14 (3), 1196–1213. URL http://link.aip.org/link/?PHF/14/1196/1

Givi, P., 1989. Model free simulations of turbulent reactive flows. Prog. Energ. Combust. 15.

Golovanov, Y. V., 1977. Experimental investigation of statistical characteristics of turbulent fluctuations of contaminant concentration in a submerged jet. Ph.D. thesis, Dolgoprudnyi, Moscow Physical-Technical Institute.

Grigoryev, Y. N., Vshivkov, V. A., Fedoruk, M. P., 2002. Numerical "particle-in-cell" methods: theory and applications. Utrecht, Boston.

Hanjalić, K., Launder, B. E., 1972. A Reynolds stress model of turbulence and its application to thin shear flows. J. Fluid Mech. 52, 609–638.

Haworth, D. C., Pope, S. B., 1986. A generalized Langevin model for turbulent flows. Phys. Fluids 29 (2), 387–405.

URL http://link.aip.org/link/?PFL/29/387/1

Haworth, D. C., Pope, S. B., 1987. A pdf modeling study of self-similar turbulent free shear flows. Phys. Fluids 30 (4), 1026–1044.

URL http://link.aip.org/link/?PFL/30/1026/1

Heinz, S., 1998. Connections between lagrangian stochastic models and the closure theory of turbulence for stratified flows. Int. J. Heat Fluid Fl. 19, 193–200.

Henderson, R., Karniadakis, G. E., 1995. Unstructured spectral element methods for simulation of turbulent flows. J. Comput. Phys. 122, 191–217.

Hirsch, C., 1988. Numerical computation of internal and external flows. Vol. 1. The MIT Press.

Hoydysh, W. G., Griffiths, R. A., Ogawa, Y., 1974. A scale model study of the dispersion of pollution in street canyons. In: 67th Annual Meeting of the Air Pollution Control Association, Denver, Colorado, APCA Paper No. 74-157.

Huang, H., Akutsu, Y., Arai, M., Tamura, M., 2000. A two-dimensional air quality model in an urban street canyon: Elevation and sensitivity analysis. Atmos. Environ. 34, 689–698.

Hunt, J. C. R., Graham, J. M. R., 1978. Free-stream turbulence near plane boundaries. J. Fluid Mech. 84, 209–235.

Iwamoto, K., Suzuki, Y., Kasagi, N., 2002. Reynolds number effect on wall turbulence: toward effective feedback control. Int. J. Heat Fluid Fl. 23, 678–689.

James, S., Anand, M. S., Pope, S. B., 2002. The lagrangian pdf transport method for simulations of gas turbine combustor flows. AIAA Paper (2002-4017).

Jenny, P., Pope, S. B., Muradoglu, M., Caughey, D. A., 2001. A hybrid algorithm for the joint PDF equation of turbulent reactive flows. J. Comput. Phys. 166, 218–252.

Johnson, G. T., Hunter, L. J., 1995. A numerical study of dispersion of passive scalar in city canyons. Bound.-Lay. Meteorol. 75, 235–262.

Jones, W. P., Launder, B. E., 1972. The prediction of laminarization with a two-equation model of turbulence. Int. J. Heat Mass Tran. 15 (2), 301–314.

Kailasnath, P., Sreenivasan, K. R., Saylor, J. R., 1993. Conditional scalar dissipation rates in turbulent wakes, jets, and boundary layers. Phys. Fluids 5 (12), 3207–3215. URL http://link.aip.org/link/?PFA/5/3207/1

Karatzas, I., Shreve, S. E., 1991. Brownian motion and stochastic calculus, 2nd Edition. Springer.

Karnik, U., Tavoularis, S., 1989. Measurements of heat diffusion from a continuous line source in a uniformly sheared turbulent flow. J. Fluid Mech. 202, 233–261.

Kloeden, P. E., Platen, E., 1999. Numerical Solution of Stochastic Differential Equations. Springer, Berlin.

Kravchenko, A. G., Moin, P., 2000. Numerical studies of flow over a circular cylinder at $Re_d = 3900$. Phys. Fluids 12 (2), 403–417.

URL http://link.aip.org/link/?PHF/12/403/1

Kristensen, L., 1994. Recurrence of extreme concnetrations. In: S.-V. Grining, M. M. Millán (Eds.), Air Pollution Modelling and Its Application. X. Plenum Press, New York.

Kuznetsov, V. R., Sabel'nikov, V. A., 1990. Turbulence and combustion, English Edition. Hemisphere, New York.

Lai, Y. G., So, R. M. C., 1990. On near wall turbulent flow modelling. J. Fluid Mech. 221, 641–673.

Langevin, P., 1908. Sur la théorie du mouvement Brownien. Comptes Rendus Acad. Sci., Paris 146, 530–533.

Launder, B. E., Reece, G. J., Rodi, W., 1975. Progress in the development of a Reynoldsstress turbulent closure. J. Fluid Mech. 68, 537.

Launder, B. E., Spalding, D. B., 1974. The numerical computation of turbulent flows. Comput. Meth. Appl. M. 3, 269–289.

Lavertu, R. A., Mydlarski, L., 2005. Scalar mixing from a concentrated source in turbulent channel flow. J. Fluid Mech. 528, 135–172.

Lee, I. Y., Park, H. M., 1994. Parameterization of the pollutant transport and dispersion in urban street canyons. Atmos. Environ. 28, 2343–2349. Lien, F. S., Yee, E., Cheng, Y., 2004. Simulation of mean flow and turbulence over a 2d building array using high-resolution CFD and a distributed drag force approach. J. Wind Eng. Ind. Aerod. 92 (2), 117–158.

Liu, C., Liu, Z., 1997. Advances in DNS/LES. In: Proc. first AFOSR Intl. Conf. on DNS/LES. Greyden Press.

Liu, C.-H., Barth, M. C., 2002. Large-eddy simulation of flow and scalar transport in a modeled street canyon. J. Appl. Meteorol. 41 (6), 660–673.

Löhner, R., 1995. Robust, vectorized search algorithms for interpolation on unstructured grids. J. Comput. Phys. 118, 380–387.

Löhner, R., 2000. Advances in unstructured grid generation. In: ECCOMAS 2000.

Löhner, R., 2001. Applied CFD techniques: an introduction based on finite element methods. John Wiley & Sons. U.K.

Lourenco, L. M., Shih, C., 1993. Characterisitics of the plane turbulent near wake of a circular cylinder. A particle image velocimetry study. Tech. Rep. See the Beaudan and Moin report titled *Numerical experiments on the flow past a circular cylinder at subcritical Reynolds number*, Department of Mechanical Engineering, Stanford University.

Luhar, A. K., Sawford, B. L., 2005. Micromixing modelling of mean and fluctuating scalar fields in the convective boundary layer. Atmos. Environ. 39, 6673–6685.

Lundgren, T. S., 1969. Model equation for nonhomogeneous turbulence. Phys. Fluids 12 (3), 485–497.

URL http://link.aip.org/link/?PFL/12/485/1

Luo, H., Baum, J. D., Löhner, R., Cabello, J., 1993. Adaptive Edge-Based Finite Element Schemes for the Euler and Navier-Stokes Equations. AIAA-93-0336.

Ma, X., Karamanos, G.-S., Karniadakis, G. E., 2000. Dynamics and low-dimensionality of a turbulent near wake. J. Fluid Mech. 410, 29–65.

Ma, X., Karniadakis, G. E., 1997. The spectrum of the turbulent near wake: A comparison of DNS and LES. In: Advances in DNS and LES, Proc. first AFOSR Intl. Conf. on DNS/LES. Greyden Press, pp. 129–136.

Madnia, C. K., Givi, P., 1993. Direct numerical simulation and large eddy simulation of reacting homogeneous turbulence. In: B. Galperin, S. A. Orszag (Eds.), Large Eddy Simulations of Complex Engineering and Geophysical Flows. Cambridge University Press, pp. 315–346.

Mahesh, K., Constantinescu, G., Moin, P., 2004. A numerical method for large-eddy simulation in complex geometries. J. Comput. Phys. 197, 215–240.

Mascagni, M., 1997. Some methods of parallel pseudorandom number generation. In: Schreiber, R., Heath, M., Ranade, A. (Eds.), Proceedings of the IMA Workshop on Algorithms for Parallel Processing, Springer-Verlag.

URL http://citeseer.ist.psu.edu/mascagni97some.html

Matsumura, M., Antonia, R., 1993. Momentum and heat transport in the turbulent intermediate wake of a circular cylinder. J. Fluid Mech. 250, 277–296.

Mavriplis, D., 1991. Three-Dimensional Unstructured Multigrid for the Euler Equations. AIAA-91-1549-CP.

Meroney, R. N., Pavageau, M., Rafailidis, S., Schatzmann, M., 1996. Study of line source characteristics for 2-d physical modelling of pollutant dispersion in street canyons. J. Wind Eng. Ind. Aerod. 62 (1), 37–56.

Mittal, R., Moin, P., 1996. Large-eddy simulation of flow past a circular cylinder. In: APS Bulletin, vol. 41 (9), 49th DFD Meeting.

Monin, A. S., Yaglom, A. M., 1975. Statistical fluid mechanics. Vol. 2. The MIT Press.

Morkovin, M. V., 1964. Flow around a circular cylinder. A kaleidoscope of challenging fluid phenomena. In: ASME Symp. on Fully Separated Flows, Philadelphia, Pa. pp. 102–119.

Moser, R. D., Kim, J., Mansour, N. N., 1999. Direct numerical simulation of turbulent channel flow up to $Re_{\tau} = 590$. Phys. Fluids 11 (4), 943–945. URL http://link.aip.org/link/?PHF/11/943/1

Muradoglu, M., Jenny, P., Pope, S. B., Caughey, D. A., 1999. A consistent hybrid finite-volume/particle method for the PDF equations of turbulent reactive flows. J. Comput. Phys. 154, 342–371.

Muradoglu, M., Pope, S. B., Caughey, D. A., 2001. The hybrid method for the PDF equations of turbulent reactive flows: consistency conditions and correction algorithms. J. Comput. Phys. 172, 841–878.

Noca, F., Park, H., Gharib, M., 1998. Vortex formation length of a circular cylinder (300 < Re < 4000) using DPIV. In: Proc. on Bluff Body Wakes and Vortex-Induced Vibration, ASME Fluids Engineering Division, Washington, D.C.

Norberg, C., 1987. Effects of Reynolds number and a low-intensity free-stream turbulence on the flow around a circular cylinder. Tech. Rep. Publication No. 87/2, Department of Applied Thermodynamics and Fluid Mechanics, Chalmer University of Technology, Gothenburg, Sweden.

Ong, L., Wallace, J., 1996. The velocity field of the turbulent very near wake of a circular cylinder. Exp. Fluids 20, 441–453.

Pavageau, M., 1996. Concentration fluctuations in urban street canyons – groundwork for future studies. Tech. rep., Meteorological Institute of the University of Hamburg.

Pavageau, M., Schatzmann, M., 1999. Wind tunnel measurements of concentration fluctuations in an urban street canyon. Atmos. Environ. 33, 3961–3971.

Peraire, J., Peiro, J., Morgan, K., 1992. A Three-Dimensional Finite Element Multigrid Solver for the Euler Equations. AIAA-92-0449.

Pope, S. B., 1985. PDF methods for turbulent reactive flows. Prog. Energ. Combust. 11, 119–192.

Pope, S. B., 1990. Computations of turbulent combustion: progress and challenges. Proc. Combust. Inst. 23, 591–612.

Pope, S. B., 1994. On the relationship between stochastic Lagrangian models of turbulence and second-moment closures. Phys. Fluids 6 (2), 973–985. URL http://link.aip.org/link/?PHF/6/973/1

Pope, S. B., 1995. Particle method for turbulent flows: Integration of stochastic model equations. J. Comput. Phys. 117, 332–349.

Pope, S. B., 1998. The vanishing effect of molecular diffusivity on turbulent dispersion: implications for turbulent mixing and the scalar flux. J. Fluid Mech. 359, 299–312.

Pope, S. B., 2000. Turbulent flows. Cambridge University Press, Cambridge.

Pope, S. B., 2004. Ten questions concerning the large-eddy simulation of turbulent flows. New J. Phys. 6 (35), 1–24.

Rafailids, S., Schatzmann, M., 1995. Physical modelling of car exhaust disperision in urban street canyons. In: Proc. 21st Int. Meeting on Air Pollution Modelling and Its Applications, Baltimore, Nov. 6-10.

Raupach, M. R., Legg, B. J., 1983. Turbulent dispersion from an elevated line source: measurements of wind-concentration moments and budgets. J. Fluid Mech. 136, 111–137.

Rembold, B., Jenny, P., 2006. A multiblock joint pdf finite-volume hybrid algorithm for the computation of turbulent flows in complex geometries. J. Comput. Phys. 220, 59–87.

Reynolds, W. C., 1990. The potential and limitations of direct and large eddy simulations. In: Lumley, J. L. (Ed.), *Whither Turbulence? Turbulence at the crossroads*. Springer-Verlag, Berlin, pp. 313–343.

Rodi, W., 1980. Turbulence Models and Their Application in Hydraulics – A State-of-theart Review. Intl. Assoc. for Hydraulics Research, Delft, The Netherlands.

Rodi, W., Mansour, N. N., 1993. Low Reynolds number $k - \varepsilon$ modelling with the aid of direct numerical simulation data. J. Fluid Mech. 250, 509–529.

Rotta, J. C., 1951. Statistiche theorie nichthomogener turbulenz. Z. Phys. 129, 547.

Saad, Y., 2003. Iterative methods for sparse linear systems, 2nd Edition. Society for Indurstrial and Applied Mathematics, Philadelphia.

Sardi, K., Taylor, A. M. K. P., Whitelaw, J. H., 1998. Conditional scalar dissipation statistics in a turbulent counterflow. J. Fluid Mech. 361, 1–24.

Sawford, B. L., 1995. A simple representation of a developing contaminant concentration field. J. Fluid Mech. 289, 141–157.

Sawford, B. L., 2004a. Conditional scalar mixing statistics in homogeneous isotropic turbulence. New J. Phys. 6 (55), 1–30.

Sawford, B. L., 2004b. Micro-mixing modelling of scalar fluctuations for plumes in homogeneous turbulence. Flow Turbul. Combust. 72, 133–160.

Sawford, B. L., 2006. Lagrangian modeling of scalar statistics in a double scalar mixing layer. Phys. Fluids 18 (8), 085108.

URL http://link.aip.org/link/?PHF/18/085108/1

Schewe, G., 1986. Sensitivity of transition phenomena to small perturbations in the flow round a circular cylinder. J. Fluid Mech. 172, 33–46.

Shcherbina, Y. A., 1982. Statistical characteristics of turbulent transport. Tech. rep., Dolgoprudnyi, Moscow Physical-Technical Institute.

Sheikhi, M. R. H., Drozda, T. G., Givi, P., Pope, S. B., 2003. Velocity-scalar filtered density function for large eddy simulation of turbulent flows. Phys. Fluids 15 (8), 2321–2337. URL http://link.aip.org/link/?PHF/15/2321/1

Shostko, A., Löhner, R., 1995. Parallel 3-D H-refinement. AIAA-95-1662-CP.

Shraiman, B. I., Siggia, E. D., 2000. Scalar turbulence. Nature 405, 639–646.

Singhal, A. K., Spalding, D. B., 1981. Predictions of two-dimensional boundary layers with the aid of the $k - \varepsilon$ model of turbulence. Comput. Meth. Appl. M. 25, 365–383.

Son, J., Hanratty, T. J., 1969. Velocity gradients at the wall for flow around a cylinder at Reynolds numbers 5×10^3 to 10^5 . J. Fluid Mech. 35, 353–368.

Spalding, D. B., 1977. GENMIX – A General Computer Program for Two-Dimensional Parabolic Phenomena. Intl. Assoc. for Hydraulics Research.

Speziale, C. G., Sarkar, S., Gatski, T. B., 1991. Modelling the pressure-strain correlation of turbulence: an invariant dynamical systems approach. J. Fluid Mech. 227, 245–272.

Sreenivasan, K. R., Antonia, R. A., 1978. Joint probability densities and quadrant contributions in a heated turbulent round jet. AIAA J. 16, 867–868.

Stapountzis, H., Sawford, B. L., Hunt, J. C. R., Britter, R. E., 1986. Structure of the temperature field downwind of a line source in grid turbulence. J. Fluid Mech. 165, 401–424.

Subramaniam, S., Haworth, D. C., 2000. A probability density function method for turbulent mixing and combustion on three-dimensional unstructured deforming meshes. International Journal of Engine Research 1 (2), 171–190.

URL http://journals.pepublishing.com/content/yt4q223365kn071w/

Tang, Q., Xu, J., Pope, S. B., 2000. Probability density function calculations of local extinction and NO production in piloted-jet turbulent methane/air flames. P. Combust. Inst. 28, 133–192.

Taylor, G. I., 1921. Diffusion by continuous movements. P. Lond. Math. Soc. 20, 196–211.

Tomboulides, A., Orszag, S., Karniadakis, G. E., 1993. Direct and large-eddy simulations of axisymmetric wakes. AIAA Paper, AIAA–1993–0546.

van Driest, E. R., 1956. On the turbulent flow near a wall. J. Aeronaut. Sci. 23 (11), 1007–1011.

van Kampen, N. G., 2004. Stochastic processes in physics and chemistry, 2nd Edition. North Holland, Elsevier B.V., Amsterdam, The Netherlands.

van Slooten, P. R., Jayesh, Pope, S. B., 1998. Advances in PDF modeling for inhomogeneous turbulent flows. Phys. Fluids 10 (1), 246–265.

URL http://link.aip.org/link/?PHF/10/246/1

Venkataramani, K. S., Chevray, R., 1978. Statistical features of heat transfer in grid generated turbulence: constant gradient case. J. Fluid Mech. 86, 513.

Villermaux, J., Devillon, J. C., 1972. Représentation de la coalescence et de la redispersion des domaines de ségrégation dans un fluide par un modèle d'interaction phénoménologique. In: Proceedings of the 2nd Intl. Symp. on Chemical Reaction Engineering. Elsevier, New York, pp. 1–13.

Vrieling, A. J., Nieuwstadt, F. T. M., 2003. Turbulent dispersion from nearby point sources – interference of the concentration statistics. Atmos. Environ. 37, 4493–4506.

Wacławczyk, M., Pozorski, J., Minier, J.-P., 2004. Probability density function computation of turbulent flows with a new near-wall model. Phys. Fluids 16 (5), 1410–1422. URL http://link.aip.org/link/?PHF/16/1410/1

Warhaft, Z., 1984. The interference of thermal fields from line sources in grid turbulence. J. Fluid Mech. 144, 363–387.

Warhaft, Z., 2000. Passive scalars in turbulent flows. Annu. Rev. Fluid Mech. 32, 203–240.

Wedding, J. B., Lambert, D. J., Cermak, J. E., 1977. A wind tunnel study of gaseous pollutants in city street canyons. Air Pollut. Control Assoc. J. 27, 557–566.

Whizman, V., Laurence, D., Kanniche, M., Durbin, P. A., Demuren, A., 1996. Modeling near-wall effects in second-moment closures by elliptic relaxation. Int. J. Heat Fluid Fl. 17, 255–266.

Wilcox, D. C., 1993. Turbulence modeling for CFD. DCW Industries, La Cañada, CA.

Williamson, C. H. K., 1996. Vortex dynamics in the cylinder wake. Annu. Rev. Fluid Mech. 28, 477–539.

Wilson, D. J., 1995. Concentration fluctuations and averaging time in vapor clouds. Center for Chemical Process Safety. American Institute of Chemical Engineers, New York.

Xu, J., Pope, S. B., 1997. Analysis of numerical errors in solving particle langevin equations. Tech. Rep. FDA 97–07, Cornell University.

Xu, J., Pope, S. B., 2000. PDF calculations of turbulent nonpremixed flames with local extinction. Combust. Flame 123, 281–307.

Zhang, Y. Z., Haworth, D. C., 2004. A general mass consistency algorithm for hybrid particle/finite-volume PDF methods. J. Comput. Phys. 194, 156–193.

Curriculum Vitae

József Bakosi was born on September 24, 1975, in Nyíregyháza, Hungary and completed his Masters in Mechanical Engineering at the University of Miskolc in 1999. After working for a couple of years in Hungary as a research and teaching assistant, then later as a development engineer, he did a full-time internship in precision manufacturing at the Center for Manufacturing Innovation of Fraunhofer USA Inc. in Boston, MA, from 2001 to 2003. He joined the Comprehensive Atmospheric Modeling Program of the School of Computational Sciences at GMU in September, 2003.