

Hybrid finite-volume/transported PDF method for the simulation of
turbulent reactive flows

by

Venkatramanan Raman

A dissertation submitted to the graduate faculty
in partial fulfillment of the requirements for the degree of
DOCTOR OF PHILOSOPHY

Major: Chemical Engineering

Program of Study Committee:
Rodney O. Fox, Major Professor
Glenn R. Luecke
Francine Battaglia
James C. Hill
Richard Seagrave

Iowa State University
Ames, Iowa
2003

UMI Number: 3105101

Copyright 2003 by
Raman, Venkatramanan

All rights reserved.

INFORMATION TO USERS

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleed-through, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

UMI[®]

UMI Microform 3105101

Copyright 2004 by ProQuest Information and Learning Company.

All rights reserved. This microform edition is protected against unauthorized copying under Title 17, United States Code.

ProQuest Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346

Graduate College
Iowa State University

This is to certify that the doctoral dissertation of
Venkatramanan Raman
has met the dissertation requirements of Iowa State University

Signature was redacted for privacy.

Major Professor

Signature was redacted for privacy.
for the Major Program

Table of Contents

List of Tables	v
List of Figures	vi
Abstract	xiv
1 Introduction	1
Motivation	1
Background	2
Literature Review	4
Objectives	15
Overview	16
2 Basic Theory	17
Introduction	17
Eulerian Approach to Solving NS Equations	18
RANS Equations	19
Turbulence Modeling	21
Compressible Flow Model	24
Probability Density Function Methods	26
Fundamental Statistical Definitions	26
Eulerian PDF Formulation	28
Lagrangian PDF Formulation	30
Scalar Transport using PDF Method	34
Micromixing Models	37
Reaction Computation	41
Hybrid Method	43
Summary	44

3	Numerical Methods	47
	Introduction	47
	Eulerian NS Solver	48
	PDF Scheme	48
	Lagrangian Particle Scheme	50
	Test Case	66
	Summary	69
4	Methane-Air Diffusion Flame Simulation	81
	Introduction	81
	Numerical Procedure	82
	Simulation Conditions	85
	Details of Test Cases	88
	Finite Volume Simulation with Detailed Chemistry	89
	FV-PDF Predictions of the Sandia D Flame	90
	Effect of Model Constants in Turbulence Energy Dissipation Equation	91
	Effect of Mixing Model	93
	Effect of Chemistry Mechanism and Radiation	94
	Conclusions	96
5	Simulation of Gas-phase Chlorination Reactor	123
	Introduction	123
	Reactor Configuration	124
	Simplified Reactor Models	125
	Constant-density PDF Simulation of Premixed-feed Reactor	128
	Simulation Conditions	129
	Simulation Results	129
	Effect of Feed-stream Configuration	132
	Simulation Conditions	132
	Simulation Results	133
	Simulation of Plant-scale Chlorination Reactor	139
	Model Performance under Plant-operating Conditions	141
	Numerical Accuracy of the Time-splitting Scheme	143
	Summary	145

6 Conclusions and Future Work	176
Hybrid Scheme	176
Thermochlorination Reactor	177
Future Directions	180
Bibliography	181

List of Tables

Table 3.1	ISAT performance with variation in error tolerance parameter. . .	64
Table 3.2	ISAT performance with change in the number of particles in the PMSR. Ntree = 8 and Error tolerance = 7×10^{-5}	64
Table 3.3	ISAT performance with change in the number of binary trees. . .	65
Table 5.1	Chemical species.	125
Table 5.2	21 step reaction scheme.	126
Table 5.3	ISAT performance statistics.	128
Table 5.4	PDF simulation statistics. Wall time refers to real time (as if measured by a wall clock) required for the simulation.	129
Table 5.5	Feed stream composition and operating conditions for the plant- scale operation	142

List of Figures

2.1	ISAT algorithm.	45
2.2	Flowchart of hybrid Eulerian flow-solver/ Transported PDF method. [Jenny et al., 2001].	46
3.1	Flowchart of transported PDF method.	70
3.2	Cell neighborhood for interpolation.	71
3.3	Bilinear mapping used to transform to unit square.	72
3.4	Single cell tracking procedure. Cell normals are indicated by ni	73
3.5	Correction applied to axisymmetric domains.	74
3.6	Shear layer domain used for the test cases. The dotted lines show a rough outline of the shear layer.	75
3.7	Comparison of Finite-Volume/ PDF results of mixture fraction and variance profiles for test case with constant velocity. Dashed line - PDF. Solid line - FV. Symbols - Theoretical value.	76
3.8	Comparison of Finite-Volume/ PDF results of mixture fraction and variance profiles for test case with velocity-shear. Dashed line - PDF. Solid line - FV. Symbols - Theoretical value.	77
3.9	Grid structure for the test case. The individual block orientations have been marked with respective axes.	78
3.10	Enlarged view of the grid. The blocks have been slightly displaced to distinguish block boundaries.	78
3.11	Comparison of Finite-Volume/ PDF results of mixture fraction and variance profiles for the arbitrary test grid. Dashed line - PDF. Solid line - FV. Symbols - Theoretical value.	79
3.12	Comparison of Finite-Volume/ PDF results of mixture fraction and variance profiles for the arbitrary test grid. Dashed line - PDF. Solid line - FV. Symbols - Theoretical value.	80

4.1	Simulation domain used for Sandia D flame.	98
4.2	Comparison of ensemble tracking time per iteration with different particle numbers per cell for varying grid size. (a) 2000 cells and time step of 2.5×10^{-4} s. (b) 2000 cells and time step of 5×10^{-4} s. (c) 8000 cells with time step of 2.5×10^{-4} s.	99
4.3	Axial mean mixture fraction and variance for test case used for optimal time step determination. Solid lines use a time step of 10^{-4} s. Dashed lines: 5×10^{-5} s. Dashed-dotted lines: 2.5×10^{-5} s.	100
4.4	Radial mean mixture fraction and variance at (left) $x/d = 15$ (middle) $x/d = 30$ and (right) $x/d = 45$ for test case used in time step determination. Solid lines: 1×10^{-4} s. Dashed lines 5×10^{-5} s. Dashed-dotted lines: 2.5×10^{-5} s.	101
4.5	(Left) Axial mean temperature and RMS of temperature for two different time steps. (Right) Axial mean species profiles for the two cases. Solid lines: 1×10^{-4} s. Dashed lines 5×10^{-5} s.	102
4.6	Radial mean temperature and RMS of temperature for reacting flow case at axial positions of (left) $x/d = 15$ (middle) $x/d = 30$ and (right) $x/d = 45$ for two different time steps. Solid lines: 1×10^{-4} s. Dashed lines: 5×10^{-5} s.	103
4.7	Centerline profile of axial velocity and RMS axial velocity obtained without a micromixing model in the FV code compared with experimental data. Lines show simulation results and the symbols are experimental data.	104
4.8	Centerline profile of (left) temperature and (right) chemical species obtained without a micromixing model in the FV code compared with experimental data. Lines show simulation results and the symbols are experimental data.	105
4.9	Axial profile of chemical species (lines) and experimental data (symbols) for Case A.	106
4.10	Radial profiles at axial positions of (left) $x/d = 15$, (middle) $x/d = 30$ and (right) $x/d = 45$ for temperature, mixture fraction and chemical species for Case A. Symbols are experimental data and lines are simulation results.	107

4.11	Conditional means at axial locations of (left) $x/d = 15$, (middle) $x/d = 30$ and (right) $x/d = 45$ for selected chemical species for Case A. Symbols are experimental data and lines are simulation results.	108
4.12	Conditional means at axial locations of (left) $x/d = 15$, (middle) $x/d = 30$ and (right) $x/d = 45$ for selected chemical species for Case A. Symbols are experimental data and lines are simulation results.	109
4.13	Mean and RMS axial velocity for Case A (solid line) and Case B (dashed line) compared with experimental data (symbols). . . .	110
4.14	Centerline mean and RMS profile for Case B (solid line) compared with experimental data (symbols).	111
4.15	Radial mean temperature and mean mixture fraction profiles for Case B at axial positions of (left) $x/d = 15$, (middle) $x/d = 30$, and (right) $x/d = 45$. Lines are simulation results and symbols are experimental data.	112
4.16	Conditional means of temperature and selected chemical species for Case B at axial positions of (left) $x/d = 15$, (middle) $x/d = 30$, and (right) $x/d = 45$. Lines are simulation and symbols are experimental data.	113
4.17	(Left) Mean and RMS centerline temperature profile for Case C (solid line) and Case D (dashed line) compared with experimental data (symbols). (Right) Centerline mean and RMS mixture fraction profile for the same cases.	114
4.18	(Left) Mean methane and oxygen profile for case C (solid line) and Case D (dashed line) compared with experimental data (symbols). (Right) Centerline NO and OH profile for the same cases. . . .	115
4.19	Radial profiles of means of temperature and mixture fraction for Case C (solid line) and Case D (dashed line) at axial positions of (left) $x/d = 15$, (middle) $x/d = 30$, and (right) $x/d = 45$. Lines are simulation results and symbols are experimental data. . . .	116
4.20	Conditional means of temperature and selected chemical species for Case C at axial positions of (left) $x/d = 15$, (middle) $x/d = 30$, and (right) $x/d = 45$. Lines are simulation and symbols are experimental data.	117

4.21	Axial profiles of selected chemical species for different mechanisms compared with experimental data. Solid lines are GRI-3.0 mechanism results and dashed lines are ARM results. Symbols are experimental data.	118
4.22	Radial profile of mean temperature and mixture fraction for different mechanisms at axial positions of (left) $x/d = 15$, (middle) $x/d = 30$, and (right) $x/d = 45$. Solid lines are GRI-3.0 mechanism (Case E) results and dashed lines are ARM (Case F) results. Symbols are experimental data.	119
4.23	Conditional means of temperature and selected chemical species for different mechanisms at axial positions of (left) $x/d = 15$, (middle) $x/d = 30$, and (right) $x/d = 45$. Solid lines are GRI-3.0 mechanism (Case E) results and dashed lines are ARM (Case F) results. Symbols are experimental data.	120
4.24	Axial profiles of scalars for Case G that neglects radiation. Solid lines are simulation result and symbols are experimental data.	121
4.25	Conditional means of temperature and selected chemical species for Case G at axial positions of (left) $x/d = 15$, (middle) $x/d = 30$, and (right) $x/d = 45$. Solid lines are simulation results and symbols are experimental data.	122
5.1	Schematic diagram of jet-stirred reactor.	146
5.2	CSTR model - Variation of outlet temperature with inlet temperature at a constant residence time of 7.5 s.	147
5.3	CSTR model - Variation of outlet temperature with residence time at inlet temperature of 365 K.	148
5.4	PMSR model - Effect of micromixing time on final temperature. Residence time is 7.5 s with inlet temperature of 323 K.	149
5.5	Mean velocity (top), turbulent diffusivity (middle), and local mixing time (bottom) profiles obtained from a $k-\omega$ turbulence model with a 15-species reaction mechanism.	150
5.6	Comparison of time-averaged centerline temperature profiles for detailed and reduced chemistry with different micromixing parameters after three residence times.	151

5.7	Temperature (top), chlorine (middle), and chloroform (bottom) profiles with detailed chemistry simulated using a composition PDF method.	152
5.8	Effect of density on flow conditions.	153
5.9	Velocity profile for cold flow condition.	153
5.10	Velocity profile for reacting condition.	154
5.11	Radial temperature profiles for different grid sizes at various axial positions. The solid lines represent a 53×33 grid and the dashed lines represent 53×79 grid. Temperature has been normalized with respect to an arbitrary value of 750 and the maximum axial shift is 0.05 m for this temperature.	155
5.12	Streamlines illustrating recirculation zones for (top) Premixed (middle) non-premixed with chlorine inner jet and (bottom) non-premixed with chlorine outer jet configurations.	156
5.13	Profiles of (a) temperature, (b) chlorine and (c) chloroform for the premixed case. Plot (d) is the Temperature profile with infinitely fast micromixing.	157
5.14	Centerline temperature and product profile for (solid) infinitely fast micromixing case and (dashed) IEM mixing model with pre-mixed inlets.	158
5.15	Profiles of (a) temperature, (b) chlorine and (c) chloroform for the non-premixed case with chlorine as the inner jet. Plot (d) is the Temperature profile with infinitely fast micromixing.	159
5.16	Radial scatter plots at different axial distances of (a) temperature, (b) chlorine and (c) chloroform for the non-premixed case with chlorine as the inner jet. Starting from left, the figures in each column correspond to axial positions of $x/r = 1$, $x/r = 3$, and $x/r = 6.7$, respectively.	160
5.17	Centerline temperature and product profile using (solid) infinitely fast micromixing condition and (dashed) IEM mixing model for non-premixed inlets with chlorine in the inner stream.	161
5.18	Profiles of (a) temperature, (b) chlorine and (c) chloroform for the non-premixed case with chlorine as the outer jet. Plot (d) is the Temperature profile with infinitely fast micromixing.	162

5.19	Radial scatter plots at different axial distances of (a) temperature, (b) chlorine and (c) chloroform for the non-premixed case with chlorine as the outer jet. Starting from left, the figures in each column correspond to axial positions of $x/r = 1$, $x/r = 3$, and $x/r = 5$, respectively.	163
5.20	Centerline temperature and product profile using (solid) infinitely fast micromixing condition and (dashed) IEM mixing model for non-premixed inlets with chlorine in the outer stream.	164
5.21	Configuration for the plant-scale chlorination reactor with sparger. L_s is the sparger length. R_s is the sparger radius. L and R are the external dimensions of the reactor.	165
5.22	(Top) Velocity-streamline profile and (bottom) mixing time profile for the plant-scale chlorination reactor.	166
5.23	(Top) Temperature, (middle) chlorine, and (bottom) chloroform profiles for the plant-scale reactor.	167
5.24	Axial profile of temperature and selected scalars for the plant-scale chlorination reactor.	168
5.25	Temperature profile and streamline velocity plot of plant scale reactor simulation under stable operating conditions. (Top) Neglecting micromixing, (middle) using micromixing. The top half of the streamline plot corresponds to the infinite-rate mixing case and the bottom half represents the finite-rate mixing case.	169
5.26	Temperature profiles of plant scale reactor simulation under unstable operating conditions. (Top) Neglecting micromixing and (bottom) considering micromixing.	170
5.27	Instantaneous centerline profiles of chlorine mass fraction in ppm near the exit of the reactor. The arrow indicates the direction of increasing simulation time relative the starting solution from a infinite-rate mixing case. Dashed-dotted line - $t = 0.01s$. Dashed line - $t = 3s$. Solid line - $t = 9s$	171
5.28	Simulation conditions for the test of numerical-accuracy. The higher C_ϕ corresponds to infinite-rate mixing and the lower value corresponds to finite-rate mixing. The temperature axis represents the peak temperature found along the centerline in each of the simulations.	172

5.29	Final temperature profiles for the three cases corresponding to the reference points in Fig. 5.28. (top) Infinite-rate mixing lower steady-state, (middle) finite-rate mixing steady-state, and (bottom) Infinite-rate mixing higher steady-state.	173
5.30	Test for grid-independence of finite-rate mixing case. Dashed lines represent the results from the original grid while solid lines show the results from the grid with twice the number of grid points in each direction.	174
5.31	Effect of time step on axial temperature profile. The arrow points to the direction of decreasing time steps. Dashed line - $\Delta t = 10^{-3}$. Dotted line - $\Delta t = 5 \times 10^{-4}$. Solid line - $\Delta t = 2.5 \times 10^{-4}$	175

Abstract

A novel computational scheme is formulated for simulating turbulent reactive flows in complex geometries with detailed chemical kinetics. A Probability Density Function (PDF) based method that handles the scalar transport equation is coupled with an existing Finite-Volume (FV) Reynolds-Averaged Navier-Stokes (RANS) flow solver. The PDF formulation leads to closed chemical source terms and facilitates the use of detailed chemical mechanisms without approximations. The particle-based PDF scheme is modified to handle complex geometries and grid structures. Grid-independent particle evolution schemes that scale linearly with the problem size are implemented in the Monte-Carlo PDF solver. A novel algorithm, *in situ* adaptive tabulation (ISAT) is employed to ensure tractability of complex chemistry involving a multitude of species. Several non-reacting test cases are performed to ascertain the efficiency and accuracy of the method. Simulation results from a turbulent jet-diffusion flame case are compared against experimental data. The effect of micromixing model, turbulence model and the reaction scheme on flame predictions are discussed extensively. Finally, the method is used to analyze the Dow chlorination reactor. Detailed kinetics involving 37 species and 152 reactions as well as a reduced form with 16 species and 21 reactions are used. The effect of inlet configuration on reactor behavior and product distribution is analyzed. Plant-scale reactors exhibit quenching phenomena that cannot be reproduced by conventional simulation methods. The FV-PDF method predicts quenching accurately and provides insight into the dynamics of the reactor near extinction. The accuracy of the fractional time-stepping technique is discussed in the context of apparent multiple steady states observed in a non-premixed feed configuration of the chlorination reactor.

1 Introduction

Motivation

The turbulent flow regime provides distinct advantages for industrial reactor operations. Usually high flow rates are associated with turbulent flow which directly translates to higher throughput of product or an increase in the capacity of the reactor. More significantly, turbulent flow is characterized by enhanced mixing and can be exploited in designing reactors with short contact or residence times. A high degree of mixing can increase the reaction rate which suggests the use of turbulence to control reaction inside the reactor. Turbulence also increases the heat transfer rates and hence operations like wall heating are more effective. Understanding turbulence and the interaction between turbulence and reaction is thus key to a successful optimization and operation of a plant scale reactor.

The industrial approach to understanding turbulence is limited to pilot-scale reactors. Simple models based on conventional chemical engineering principles can be formulated by analyzing the output from the reactor. Though these models work well in predicting product distribution [West et al., 1999], the complex interactions between turbulence and reaction cannot be understood by simplifying either of the processes. In particular, the prediction of transient oscillations and extinction behavior require the flow-field description and a detailed chemistry mechanism. However, the very nature of turbulence does not allow a tractable computational scheme to provide this comprehensive information. In this work, an unconventional solution methodology is employed to remove the restrictions on the size of the chemistry mechanism. In addition, the method is designed to use state-of-the-art turbulent flow solvers available at any particular time thereby obtaining the best possible description of the flow field. In general, a computational framework will be implemented to analyze reactors operating in the turbulent flow regime that may involve complex chemistry mechanisms involving the transport of multiple species.

Background

Turbulent reactive flows have long been considered as one of the most difficult mathematical problems to tackle, principally due to the complex interaction between fluid turbulence and chemical reactions. Recent developments in the numerical methods and the rapid advances in computer architecture have opened up detailed modeling opportunities. Numerical schemes for non-reacting flow are well-advanced and are commonly used for design and other purposes in industry. The RANS method has been used widely for solving such practical flows [Ferziger and Peric, 2002]. Available commercial software can handle almost any geometry or system specification [Fluent Inc., 2001]. In the recent past, Large Eddy Simulation (LES), a new computational technique deemed to be more accurate than the modeled RANS method, has gained popularity. Promising research in adapting LES methods to complex geometries are currently being pursued.

In spite of these tremendous improvements, reactive flows are often considered only as an extension of the basic hydrodynamics. Simulating reactive flows involves solving additional species transport equations that contain chemical source terms. The problem arises from the widely varying time and length scales of the flow and reaction. Reactions occur when reactants are mixed molecularly, that is by molecular diffusion [Fox, 1996]. On the other hand, turbulent energy is contained in length scales much larger than length scales at which molecular diffusion occurs. In all fluid dynamic simulations, only the large scales are resolved and the small scales are modeled. When reactions are involved, the effect of mixing on reaction should also be considered. In Eulerian grid-based simulations, the sub-grid scalar fluctuations are either neglected or modeled by simple closures. Such modeling strategies can lead to large errors in the predictions [Peters, 2000]. In addition, reactions pose problems to numerical stability itself. Reaction kinetics may involve a range of chemical time scales due to different reactions that occur in the domain of interest. An implicit solution scheme then has to take into consideration the smallest time scale to be able to resolve the reaction correctly [Oran and Boris, 2001]. In combustion applications, such time scales can be of the order of 10^{-12} s leading to very high computational cost [Warnatz et al., 1996]. The presence of the such varied time scales will render the underlying differential equations describing flow *stiff*. In other words, slight changes in the state of the system will lead to very large variations in the steady state. There are several solutions to this problem. The simplest is to discard the detailed hydrodynamics of the flow and use zero (Continuously Stirred Tank Reactor - CSTR) or one (Plug Flow Reactor - PFR) dimensional flow models. The chemical kinetics can then be handled easily and accurately. It is quite apparent

that such simplified models cannot account for complex flow patterns that may occur in industrial reactors. In practice, reactor network models or a combination of such models are used to describe the reactor. Non-ideal CSTRs and PFRs are also commonly employed in an attempt to increase the degrees of freedom in parameter space [West et al., 1999].

The other alternative is to simplify reaction kinetics. In general very fast reactions can be approximated by steady state or equilibrium solutions. It is also possible to obtain quasi-steady state solutions where the fast and slow reactions are identified. Based on the effective reaction rates, the species are then assumed to be in a steady state with respect to these equations. Similarly, it is possible to couple the effects of the minor species and form global mechanisms with a reduced set of chemical species that involve reaction rates representative of the complete mechanism [Sung et al., 1998; Jones and Khaki, 1998]. Such kinetics usually limit the time-scale variations and also reduce the number of partial differential equations to be solved, since only those species in the reduced mechanism need to be tracked. Inclusion of detailed flow modeling provides a more physical basis for the process. As is true with all models, such reduced chemistry mechanisms are not universal and the reduction constraints may not be valid in the entire computational domain. It should be recognized that these modeling techniques provide an understanding of the physics in a qualitative sense.

Unrelated to these grid-based techniques lies another approach - the Probability Density Function (PDF) method - that tries to model the statistical properties of the flow. These schemes solve for the mean and the first/second-order moments of various flow properties like velocity, chemical species, and enthalpy. The formulation of the PDF scheme ensures that the chemical source term is closed and the chemical kinetics are handled directly [Pope, 1985]. Also the pressure fluctuations and the mean pressure gradient appear closed. Unlike the RANS methods, it also allows a molecular mixing model to be used to describe mixing of the scalars at unresolved scales. However, numerical implementations of this method have been restricted to Cartesian rectangular grids and the stability of the scheme, especially the computation of the derivatives from a statistically fluctuating field, limits the use only to some test flows [Subramaniam and Haworth, 2000]. In spite of these limitations, the unparalleled ability to handle detailed kinetics and to model molecular diffusion are key attractions. From this discussion, it can be seen that there are two classes of methods. On the one hand, the Eulerian grid-based methods are well suited for solving the momentum and turbulence equation. On the other hand, the PDF methods can handle scalar transport effectively. Hence, a

turbulent reactive flow solver can be formulated based on a hybrid scheme that uses a Eulerian method for the flow solution and a PDF scheme for scalar transport equations [Correa and Pope, 1992]. The hybrid scheme has been in existence for over a decade but the applications have been restricted primarily due to the computational expense of the PDF method [Wouters, 1998]. The exact implementation and the definition of the scope of each of these methods will be explained later. The hybrid method introduces modularity in the handling of the underlying equations. Now each of the different schemes, like the flow solver, the PDF formulation or the mixing model can be replaced without affecting the overall structure. Hence the effect of the each of these models on the solution can be determined independently and also as a function of the other models.

The motivation for this particular work arises from the chlorination reactor operated by The Dow Chemical Company. The plant-scale reactor produces chlorinated hydrocarbons from primary chlorinated derivatives (e.g., methyl chloride) and pure chlorine. Usually, a premixed feed configuration is used with a stable product yield. However, operational uncertainties combined with composition variations in the inlet streams lead to partially segregated feed. Problems ranging from reactor quenching to soot formation are encountered. Several simplified reactor models were tried but with limited success. It was deemed that a complete CFD simulation is required to identify the problems and more importantly, to develop a solution strategy to handle design and scale-up of similar reactors. In subsequent chapters, a comprehensive simulation platform will be created for this purpose and used for the chlorination reactor. The hybrid method explained before will form the backbone of this methodology. The evolution of such hybrid methods and the reactive flow methods, in general, will be discussed next.

Literature Review

The aerospace industry provided the initial thrust for research in computational fluid dynamics. Successful design and process engineering resulting from CFD led to its adaptation in other engineering areas as well. Chemical engineering is among the latest fields to embrace the cost advantages and sheer technical superiority of the simulations based design strategies [Fox, 1996]. In general, CFD deals with laminar and turbulent flows, both reacting and non-reacting. However, the high mixing rates and enhanced reactions in turbulent flow makes it the choice of regime to operate plant-scale reactors. Even before reaction is considered, the solution to the Navier-Stokes equation for a non-reacting flow is itself a hard numerical problem. This involves resolving all the

time and length scales and is computationally intractable for inhomogeneous flows with detailed geometry. Several simulations in low Reynolds number flows have yielded useful statistics for building models [Kim et al., 1987; Rogers and Moser, 1994]. Similarly DNS simulations of passive scalar mixing have been used for developing mixing models to be used in PDF methods [Eswaran and Pope, 1988; Juneja and Pope, 1996]. DNS of reacting scalars is even more expensive and has been limited mostly to single-step Arrhenius kinetics, that is representative of reaction in combustion applications [Boersma, 1999]. Similar to non-reacting simulations, these datasets can be used to model quantities like the sub-grid scalar dissipation rate that can then be used in other CFD solution methods. As explained by Pope [2000], a large portion of the length scales that are resolved lie in the dissipative range and only a small fraction of the length scales lie in the energy-containing and inertial range. Artificial modifications have been used to decrease computational expense like low-wavenumber forcing or sparse-mode methods. However, these methods do not directly solve the Navier-Stokes equation and contain models that may not represent the true physics. One other such modification has become a realizable tool for turbulence simulation.

Large Eddy Simulations (LES) use the fact that most of the energy is contained in the large scales of motion [Smagorinsky, 1963; Lilly, 1967]. LES methods resolve these large scales and model the dissipative small scales using a sub-grid scale model. A filtering operation is carried out to decompose the turbulent flow field into resolved and residual terms. The resolved variables are solved using the momentum equations. The residual terms are closed using an eddy-diffusivity model. The closure constant termed as the Smagorinsky constant can be determined using a dynamic model [Germano, 1991; Lilly, 1992; Moin et al., 1991]. The filtering operation is well defined only on Cartesian grids and extensions to irregular and unstructured grids are an active area of research [Akselvoll and Moin, 1996; Haworth and Jansen, 2000]. Direct solution of the scalar transport equations for reacting flows using LES requires modeling of the sub-grid residual terms for the scalar fluctuations [Reveillon and Vervisch, 1998]. For reacting flows with large density changes, the solution methods for LES are based mainly on models derived from RANS methods. Hence the discussion is postponed until the RANS methods are explained.

In the recent past, solution schemes to RANS equations have emerged as an important CFD tool. In this technique, time-averaged equations are solved using a closure for the Reynolds stress terms. Essentially, the method resolves the energy containing integral scales [Peters, 2000]. Modeling approaches use either the turbulent viscosity

hypothesis [Boussinesq, 1877] or a second order closure for the Reynolds stress terms. In the former case, the definition of the turbulent-viscosity closes the model. This requires a characteristic local velocity and length scale. All turbulence models are aimed at providing this closure. In general, the turbulent viscosity is expressed as a function of the local kinetic energy and the dissipation. Hence the problem is transformed into modeling these two quantities. The intrinsic assumption is that the turbulence is isotropic. The simplest model is to specify a uniform turbulent viscosity. The mixing-length model [Baldwin and Lomax, 1978] uses a characteristic length scale but incorporates the local strain rate for the time scale. When using the turbulent kinetic energy and dissipation, several two-equation models exist that solve for the profile of these quantities in the entire domain. The transport equations for these models use a production and dissipation term dependent on the local flow fields and need to be solved simultaneously with the momentum equations. The two equation turbulence models like $k-\epsilon$ [Rodi and Mansour, 1993] and the $k-\omega$ [Wilcox, 1994] methods belong to this class. The assumption of local isotropy is not valid in general and may lead to serious deterioration in the simulation predictions. Reynolds stress closures [Launder et al., 1975] have different variations that in general model all the six components of the stress tensor. In addition, the dissipation or the turbulent frequency is modeled to provide a time scale. This eliminates the need for the turbulent-viscosity hypothesis but increases the computational expense. Algebraic stress closures try to minimize the equations solved by providing algebraic relations for some of the terms and have been used with limited success.

RANS methods have been used for a wide variety of industrial flows both in laminar and turbulent regimes. Several commercial codes exist that can handle complex geometries, unsteady transient flows and flows with simple reactions. Introduction of detailed chemistry or mechanisms with widely varying chemical time scales lead to instabilities in the solver [Oran and Boris, 2001]. Also, the computational cost of handling large chemical mechanisms and the corresponding number of partial differential equations make detailed chemistry calculations prohibitive. Several alternative techniques exist. For non-premixed flows, the use of flamelet equations [Williams, 1975] are common. A conserved scalar, namely the mixture fraction and its variance are solved explicitly using the species transport equation. The flamelet equations define the propagation of the flame normal to the flame surface which is tracked along the stoichiometric mixture fraction iso-contour. For a detailed description of the model implementation refer to Peters [2000]. Based on a modeled scalar dissipation rate (using the variance and a local time scale), a flamelet table is constructed *a priori* and then used as a look-up

table during the simulation. Numerous variations like the unsteady flamelet model and the Lagrangian flamelet model [Pitsch, 2003] have been used. The conditional moment closure (CMC) method introduced independently by [Klimenko, 1995] and [Bilger, 1993] uses transport equations for reactive scalars conditioned on the mixture fraction. It is argued that the fluctuations in the scalar field is linked directly to the fluctuations in the mixture-fraction field. The CMC method has been used to model several experimental burners [Roomina and Bilger, 2001; Kim and Huh, 2002]. In both these methods, an underlying density function for the mixture fraction is assumed. The PDF is described based on the local mean and variance which is then used to compute weighted source terms. The presumed-PDF approach is prescribed only for binary pure-stream mixing and cannot completely describe partially-premixed flames where the conditions of pure fuel-oxidizer mixing is not strictly met. The overall effect of the assumed sub-grid PDF shape is to decrease the reaction rate. The infinite-micromixing assumption that neglects the sub-grid fluctuation leads to short reaction zones and faster reactions [Raman et al., 2002]. Direct moment methods use first-order closures based on the Taylor expansion for the rate constant. Such schemes can become cumbersome when the number of scalars to be handled exceeds 3-4. Hence, all of the above methods perform well under restrictive conditions and are by no means a universal solution to the turbulent reactive flow problem.

This makes a case for another class of methods that solve for the underlying density function itself. The transported probability density function (PDF) method defines the transport equation of the joint PDF of the system variables (both velocity and scalar fields). The major advantage of this method is that complex reactions can be treated without approximations. However, the dimensionality of the equation depends on the number of velocity components as well as the number of scalars used [Pope, 1985]. Hence, the direct solution of the velocity-composition PDF equation is not feasible for more than 1-2 scalars. The velocity dependence can be removed resulting in the joint composition PDF. With the reduced dimensionality, simulations of several different reaction schemes have been carried out [Pope, 1976; Nguyen and Pope, 1984]. But this method requires input of the velocity and turbulence fields and also uses the turbulent-viscosity hypothesis. Pope [1981] derived a solution algorithm for the multi-dimensional velocity-composition PDF transport equation using a stochastic particle-based Monte-Carlo scheme. The linear increase in computational time with number of dimensions is an attractive feature of this algorithm. Such an algorithm also made the use of velocity-composition PDF tractable.

The velocity-PDF requires closure for the mean of the conditional pressure fluctuation and the conditional viscous term. The Generalized Langevin Model (GLM) [Haworth and Pope, 1986] uses two coefficients (a drift and a diffusion part) to jointly model the two unclosed terms. In general, the coefficients evolve with the flow. The diffusion coefficient is assumed to be a constant in the Simplified Langevin Model (SLM). Either way, the velocity model introduces turbulent dissipation in the evolution equation and hence needs an additional model for dissipation to be completely closed. Haworth and El Tahry [1991] have used an external flow solver to obtain time-scale information and solve for the velocity and fluctuations using a SLM. The joint velocity-dissipation [Pope and Chen, 1990] or velocity-frequency PDF [Van Slooten et al., 1998] will yield a closed form for the PDF transport equation. The velocity-frequency model uses the conditional mean of the turbulent frequency and is consistent in regions with no turbulence as well [Pope, 2000]. Calculations have been carried out for reacting inhomogeneous flows [Masri and Pope, 1990; Hulek and Lindstedt, 1996; Xu and Pope, 2000] based on these models. Extensions to high-speed compressible flows have also been considered [Delarue and Pope, 1997]. For simulations of practical flows, the boundary conditions for the velocity are obtained from the Eulerian equivalents like the no-slip or free-slip condition for the velocity at the wall. When the near wall region is not resolved, a wall-function approach can also be implemented with appropriate specification of the Reynolds stresses as in the Eulerian approaches. The boundary condition for the particle frequency requires that the mean increase in the log-layer and that there be a positive flux away from the wall [Dreeben and Pope, 1997a,b].

The Monte-Carlo algorithm uses stochastic particles, otherwise referred to as nodal particles that evolve according to a set of equations based on local interaction with other particles and macroscopic flow-field quantities. The stochastically equivalent system ensures that the particles evolve with the same PDF as that of the original transport equation [Pope, 1985]. The algorithm can be formulated as an Eulerian [Pope, 1985; Roekaerts, 1991; Anand et al., 1998; Grier and Fox, 1993; Jones and Khaki, 1998] or a Lagrangian [Pope, 1994a; Raju, 1996; Xu and Pope, 2000; Subramaniam and Haworth, 2000; Muradoglu et al., 1999; Tsai and Fox, 1996] scheme. In the former case, the particles reside at cell centers and a fraction of the particles are moved across each cell face based on the flux. This ensures strict particle number conservation but leads to round-off errors due to finite particle numbers. A round-off correction step [Roekaerts, 1991; Pipino and Fox, 1994] is used to keep the flux error to a minimum. In the Lagrangian scheme, particles are free to move in the domain and evolve in physical and

compositional spaces using a set of stochastic differential equations. The tracking of particles and efficient number control are key aspects to be considered in the implementation [Xu and Pope, 1999; Subramaniam and Haworth, 2000; Raman et al., 2002]. A comparison of the performance shows that the Lagrangian scheme is equivalent to a second-order finite-difference scheme under the same conditions while the Eulerian scheme is diffusive and requires a large number of particles to achieve the same level of accuracy [Mobus et al., 2001]. The Eulerian and Lagrangian methods will be used to simulate a recirculating flow and a direct comparison between the efficiency of the two methods will be made in later chapters. The low numerical diffusion of the Lagrangian scheme is the main reason for its wide usage.

The particle models are statistically stationary at steady state, meaning that the time average of the fields approach a fixed value. However, instantaneous fields contain fluctuations which can be detrimental to the numerical stability of the scheme. The velocity PDF ensures that the pressure or the gradient of the pressure appears closed through solution to a Poisson equation. However, the source term for the Poisson equation requires the gradient of the Reynolds stress and the mean velocity terms. In a statistically noisy field, a consistent stable algorithm is necessary to eliminate numerical oscillations [Fox, 2003]. This appears to be the main drawback in using *stand-alone* PDF codes. In recent years a solution methodology has emerged that seeks to use the best compromise between Eulerian flow solvers and PDF particle methods. The **hybrid** solvers use a grid-based CFD solver (finite-volume or finite-difference) to solve for the flow and turbulence fields. The PDF code is then used to handle the species transport equation with the reaction and mixing sub-processes [Correa and Pope, 1992; Wouters, 1998]. For an exothermic reaction scheme, the impact of reactions on the flow is macroscopically linked to the density change. Hence, the hybrid solvers use a density feedback [Subramaniam and Haworth, 2000] to correct the flow according to reaction. The feedback mechanism can be implemented *tightly* in the sense that the flow is corrected every PDF time step. This makes the algorithm time accurate even if the flow solver is not time accurate. On the other hand, the high computational cost of tight coupling and the ultimate goal of reaching a steady state where there are no time transients lead to a *loose* coupling scheme where the flow is corrected only after a pre-determined number of PDF time steps [Jenny et al., 2001; Muradoglu et al., 1999]. It is assumed that in the intermediate steps, the flow has not significantly changed due to reaction that it affects the overall state of the system. It is essential that *a priori* studies are done to determine the right number of time steps to be used for the feedback loop [Raman et al., 2003].

Different flow solvers have been used in the hybrid formulation. Euler equations solver [Muradoglu et al., 1999; Jenny et al., 2001], multi-block RANS scheme [Raman et al., 2002], deforming and moving mesh solvers like the KIVA code [Subramaniam and Haworth, 2000; Tolpadi et al., 1997], Large Eddy Simulation scheme [Jaberi et al., 1999; Colucci et al., 1998] are some examples. Recently Gicquel et al. [2002] have implemented a stable velocity-composition PDF scheme in LES and have compared LES and DNS simulation results, in particular, of the Reynolds stress terms. Implementation of hybrid methods to handle complex grids is computationally challenging due to the numerically intensive process of tracking of large numbers of particles [Raman et al., 2002]. Use of element-based particle tracking is found to be efficient [Subramaniam and Haworth, 2000]. The time step used depends on the mesh and local CFL criteria and could potentially make the scheme computationally intractable due to very small time steps arising from extensive grid clustering. Methods have also been devised to base the time step independent of the mesh size (in other words, not based on local CFL criteria) [Muradoglu and Pope, 2002], though their application to practical problems are yet to be seen. Independent particle iterations based on a local time step has been tested by Raman et al. [2003] and found to be very effective in dealing with heavily clustered and axisymmetric grids. The details of the implementation will be discussed in later chapters.

With the implementation of the hybrid scheme, the stiffness of the reactions is removed from the flow solver and transferred to the PDF scheme. The PDF scheme needs to tackle the problem of handling the scalar transport equations and the sub-grid mixing phenomena. Pope [1985] derived the Monte-Carlo scheme based on the particle transport equations. The method employs fractional time stepping [Oran and Boris, 2001] to separate the physical sub-processes of transport in physical space and composition space. Another split is used to separate the mixing and reaction sub-steps. The reaction sub-step then involves solving a ODE system of scalar dependent size. Structurally, the system appears simple but for a stiff kinetic scheme, numerical integration is computationally intensive. In a PDF scheme this ODE system should be integrated for each particle scalar composition every time step which could make the algorithm intractable. Without a novel scheme to solve this ODE system, the advantage of the PDF method (of using detailed chemistry) will be lost. The particle nature of the method can again be put to use. In the PDF simulation, the notional particles will evolve from nearly the same initial conditions and hence the integration process is repeated for the same trajectory in composition space. Hence, integration is reduced to the process of finding a

point in composition space given an initial composition and a time step. Pre-computed look-up tables store the results of the computation in multi-dimensional arrays [Chen et al., 1989; Maas and Pope, 1992]. Instead of integration, interpolation is used where given any initial composition, a linear interpolation on the node values in the table finds the results of the integration. The storage space required assuming a 10-point coarse grid in each dimension of the N -scalar space, is 10^N points. It is clear that beyond a small number of scalars, look-up tables will be too large to store. Maas and Pope [1992] introduced the concept of reduced-manifolds to make the storage space tractable. But in general, these methods are useful only for a particular mechanism and operating conditions. Also most of the points stored in the table may not be actually used in the simulation at all. Pope [1997] developed the In-Situ Adaptive Tabulation (ISAT) algorithm that forms the look-up table during the simulation. Simulations can now be performed on modern computers using chemistry mechanisms with 50 or more species. The ISAT algorithm uses a binary tree data storage structure for the integration results and defines a region of composition space around each stored point in which linear interpolation can be carried out. The ISAT algorithm has been used in various problems [Xu and Pope, 2000; Shah and Fox, 1999; Raman et al., 2002] and is found to yield speed-up factors of 1000 over direct integration. Several variations of the method have also been proposed [Yang and Pope, 1998b; Tang and Pope, 2003]. However, even the basic algorithm of ISAT is sufficient to ensure tractability.

The last physical sub-process to be considered is the mixing process. This involves modeling the unresolved and sub-grid length scale molecular diffusion process. The PDF description does not contain the length scale information to accurately describe the micromixing process [Fox, 1996]. The mixing is affected by a wide range of physical phenomena including turbulent stretching, molecular dissipation and changes in scalar gradients due to chemical reactions. The mixing models need to satisfy some basic criteria, two of which are that the scalar PDF should relax to a Gaussian form irrespective of the initial shape and the scalar values must remain bounded [Fox, 2003]. For multi-scalar mixing additional criteria have been proposed (Tsai and Fox [1998]). The simplest mixing model is the Curl's model (also called the Coalescence-Dispersion (CD) Model) [Curl, 1963] which describes mixing in terms of two fluid elements on initially different composition vector that mix to form two fluid elements of equal composition. However, in a stochastic simulation starting with a set of delta functions for the PDF, it can be shown that Curl's model never produces a continuous PDF [Pope, 1982]. Janicka et al. [1970] suggested a modification to the model that results in continuous PDF but still

does not converge to a Gaussian form in the asymptotic limit. Further age-based modifications [Pope, 1982] yield finite higher moments but still cannot exactly reproduce the Gaussian relaxation. The LMSE or the IEM (Interaction by Exchange with the Mean) model [Dopazo, 1979; Villermaux, 1986] is a deterministic representation of the mixing process and is the most widely used model due to the simplicity of the implementation. The IEM model is non-local in nature and does not relax the scalar PDF. In fact it retains the initial shape of the PDF [Tsai and Fox, 1998]. An improvement over this method is the Generalized IEM model proposed by Tsai and Fox [1995] which uses a reference mixture fraction PDF and a mixing time conditioned on this mixture fraction. The binomial-Langevin model [Valino and Dopazo, 1991] uses a stochastic algorithm for the particle evolution in composition space. Special boundary conditions need to be employed to maintain the boundedness of the scalars. The Euclidean Minimum Spanning Tree (EMST) [Subramaniam and Pope, 1999] satisfies the localness of the mixing process in composition space and has been used in combustion simulations [Xu and Pope, 2000]. More intensive models have been derived that take into account the length scale distributions of the scalar field [Fox, 1997] and the differential diffusion effects in multi-scalar mixing [Fox, 1999]. The development of a general-purpose mixing model is still an active area of research. Recently Fox et al. [2002] have implemented a new multi-scalar mixing model and compared with the DNS results of a three-stream mixing problem [Juneja and Pope, 1996]. Extension to inhomogeneous reacting flows will be the future direction for these models.

Much of the work in PDF modeling comes from combustion research. The need for well-tuned burners to increase performance and decrease pollutant emission has been the driving force in research in this field. Simple flames like the piloted jet diffusion flame and the bluff-body stabilized flame prove useful as test cases. The piloted Sandia D diffusion flame of Barlow and Frank [1998] and the bluff-body flame of Dally et al. [1998b] have been widely used to test models. The datasets for various scalar and velocity fields as well as the inlet conditions are well documented. Several simulations of the methane-air flame [Xu and Pope, 2000; Jones and Khaki, 1998; Masri and Pope, 1990; Chen et al., 1989; Pitsch and Steiner, 2000; Dally et al., 1998a] have shown that the turbulence-chemistry interactions are very important and need to be accounted for in the simulations. Common problems have been noted in these works. When a RANS flow model is used, the use of a $k-\epsilon$ model for the round jet leads to shorter flame lengths and needs to be corrected [Dally et al., 1998a]. The change in the source term for the dissipation equation proposed by Pope [1978] has been employed for this purpose. It has

also been noted that the use of the IEM mixing model with the standard mechanical-to-scalar time-scale ratio leads to flame extinction [Xu and Pope, 2000; Raman et al., 2002]. Lindstedt et al. [2000] have concluded that the use of a higher value than the standard is required to avoid extinction of flames. Use of a radiation model is found to improve predictions [Pitsch and Steiner, 2000; Tang et al., 2000]. It has been concluded that the Sandia D flame can be predicted quite well by existing models for the reaction and mixing processes. The bluff-body flame is more complicated due to the complex nature of the flow with a dual vortex core in the recirculation region [Dally et al., 1998b]. Velocity-composition joint PDF simulations have been successful in predicting the flow for similar flames qualitatively [Wouters et al., 1998]. Recent research [TNF, 2002; Kim and Huh, 2002] have shown that simple flow models with changes similar to that used for the piloted-flames can predict the experimental data quite accurately close to the bluff-body. However downstream predictions still need improvement and the use of unsteady methods have been discussed [TNF, 2002]. It has been argued that the use of the gradient-diffusion assumption leads to a deterioration in the predictions. Use of Reynolds stress closure is encouraged. Even with second order turbulence closures, the scalar transport uses the same gradient-diffusion assumption and its effect on predictions has not been quantified. Use of a LES scheme as a flow solver has received considerable interest [Colucci et al., 1998; Gicquel et al., 2002] particularly due to the more dependable flow-field solution. However, hybrid LES-PDF methods in 3-D flows will be computationally intensive and efficient parallel algorithms need to be implemented. In terms of modeling, it has also been recognized that the use of a more detailed model for the scalar dissipation rate improves the predictions [Pitsch, 2003].

Combustion modeling involves fast chemistry and in many cases, reactions that can be assumed to be in equilibrium. Hence the use of presumed PDF methods is considerably more accurate in predicting flames [Fox, 1996]. Unlike combustion, most chemical processes involve finite-rate chemistry. In the recent past, use of detailed simulation techniques for reactor analysis leading further to design and scale-up has become a common strategy [Tsai and Fox, 1994; Pipino and Fox, 1994; Acharya et al., 1991; Deutschmann and Schmidt, 1998; Liu and Barkelew, 1985; Tirtowidjo, 1997; Tsai and Fox, 1996; Kolhapure and Fox, 1999; Shah and Fox, 1999; Raman et al., 2001]. Simulations are complicated by the mixing effects since the nature of the chemical mechanism can lead to regimes where both micromixing and reaction have overlapping time scales. Commercial simulation codes usually neglect closure or use simple models to account for micromixing [Fluent Inc., 2001]. The effect of such assumptions need to be studied but

in general, neglecting closure leads to shorter reaction zones (with respect to the initial contact) [Peters, 2000]. Simple models that include mixing like the eddy dissipation model assume that mixing is the controlling step. This will work in combustion mechanisms where reactions are extremely fast. Finite-rate modeling with computational tractability is an essential requirement for the simulation of chemical reactors. In this context, the chlorination reactor operated by the Dow Chemical Company forms a good case study.

Chlorination involves the conversion of primary chlorinated hydrocarbons to secondary and tertiary chlorinated derivatives. Typically, a tubular reactor with feed consisting of the hydrocarbon and chlorine is used. Though the process itself is simple, the reactions span a wide range of time scales and the sub-grid mixing needs to be modeled to capture the dynamics of the reactor. Detailed kinetic schemes for chlorination involve up to 38 species and 157 reactions [Tirtowidjo, 1997; Shah and Fox, 1999], while the simplified scheme used in this work has 15 species and 21 reactions [Shah and Fox, 1999]. The flow in the coaxial reactor used for chlorination is turbulent, and direct numerical simulations of the fundamental transport equations are rendered intractable by the complex interaction between turbulence and reaction. Simulation strategies invoke one of the two simplifications - simplified flow dynamics or global chemistry schemes [West et al., 1999; Liu and Barkelew, 1985; Acharya et al., 1991]. In the past, simple flow models like the CSTR, plug-flow and combinations of such simple configurations have been used to study the reactor. These do not account for the effect of turbulence on reaction. In addition non-idealities in plant reactors, like reactants with a certain degree of non-premixedness, can only be qualitatively accounted for using non-ideal CSTR and other models with increasing complexity. Moreover, in the limit of completely non-premixed feed streams, reactor run-away is possible and can only be accurately analyzed using a detailed flow model. Gas-phase chlorination kinetics are similar to combustion chemistry and involves highly exothermic processes. Industrial reactors are usually operated at 2-3 atmospheres and in the event of a high degree of chlorination, formation of soot is possible through complete decomposition of the hydrocarbons at high temperatures. Models employing simplified flow assumptions cannot predict these conditions accurately.

In the past, use of more complete flow models has required that the numerical stiffness associated with the free-radical kinetic mechanism be reduced. Several global mechanisms involving a reduced number of species with multi-step reactions have been used in two or three dimensional geometries with limited success [Liu and Barkelew, 1985;

Acharya et al., 1991]. Most of these schemes assume that the inlet has completely pre-mixed reactants. The reaction scheme in itself exhibits non-linear behavior that can lead to instabilities. Studies using a CSTR model show the presence of a Hopf-bifurcation point near the ignition branch of the stability curve [West et al., 1999]. Such instabilities have not been corroborated by non-zero-dimensional reactor models. The presence of large numbers of radicals that are essential for accurately predicting reactor behavior also makes convergence of CFD simulations difficult due to the wide range of time scales of the reactions. This makes the chlorination reactor a very interesting test case for the hybrid simulation. The chapter on chlorination reactor will describe the different simulation techniques used. Use of the hybrid method for prediction of reaction quenching will also be discussed.

With the increasing computational expense of adding PDF modules to RANS or any other flow solver, efficient parallel implementations are necessary. For stand-alone particle codes, multiple-independent-simulations [Xu and Pope, 1999] have been used to distribute computing time across multiple processors and at the same time reduce statistical error through higher particle number density. These methods will work only for simple grids since duplication of the computation is not prohibitive and the grids are small enough to be copied across processors. For complex grids and hybrid formulation, use of domain decomposition is the only efficient paradigm to be used. Parallel particle schemes using domain decomposition have not been explored in the past. The bottle-neck is the tracking algorithm that needs to carry over particles across processors when the particles cross the domain. Parallel particle tracking techniques [Cheng and Plassman, 2002] should be adapted to the PDF framework to extend their use in the simulation of complex flows.

This review provided an overview of the current state of the art simulation techniques for turbulent reactive flows. The next section will detail the important objectives met in this thesis work.

Objectives

The following list details the important research objectives to be met in this work :

1. A hybrid computational scheme using an existing finite-volume solver cast in generalized coordinates and a grid-independent Lagrangian composition PDF method has been successfully developed

2. An object-oriented programming structure with modular operations for the physical sub-processes will be implemented to facilitate the use of different models.
3. Several non-reacting test cases will be used to demonstrate the grid-independent nature of the PDF scheme. The accuracy and efficiency of the PDF code will be established through these tests.
4. Reacting cases will be simulated with the ISAT reaction module to handle detailed kinetics. The Sandia piloted jet-diffusion flame will be simulated and the results compared with experimental data.
5. The Dow chlorination reactor will be simulated using both the uncoupled and coupled FV-PDF method and the effects of feedstream configuration on a pilot-scale reactor will be analyzed.
6. A plant-scale chlorination reactor will be simulated and exact plant quenching data will be reproduced. Simple reactor models will also be used to analyze quenching behavior and to identify stable and oscillatory reaction branches.
7. The fractional time-stepping technique will be analyzed for errors and possible pitfalls in the usage will be identified.

Overview

The remaining chapters are arranged in the following way. The next chapter discusses the mathematical models used and explains the basics of the RANS and PDF formulations. Chapter 3 discusses the numerical implementation of the hybrid algorithm and presents some simple test results. Chapter 4 details test simulations for the feedback mechanism in reactive flows. Chapter 5 is devoted to the chlorination reactor and different aspects of the reactor analysis are covered. Chapter 6 contains conclusions and directions for future work.

2 Basic Theory

Introduction

In a CFD simulation, continuity and momentum transport equations of the form

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x_i} (\rho U_i) = 0, \quad (2.1)$$

and

$$\frac{\partial \rho U_i}{\partial t} + \frac{\partial \rho U_i U_j}{\partial x_j} = \frac{\partial}{\partial x_j} \left(\mu \frac{\partial U_i}{\partial x_j} \right) - \frac{\partial P}{\partial x_i} \quad (2.2)$$

are solved to obtain the flow field, and a species transport equation of the form

$$\frac{\partial \rho \phi}{\partial t} + \frac{\partial \rho U_i \phi}{\partial x_i} = \frac{\partial}{\partial x_i} \left(\rho D \frac{\partial \phi}{\partial x_i} \right) + \rho S(\phi), \quad (2.3)$$

in a coupled or segregated manner is solved to describe the distribution of species in the domain. Here U_i is the velocity, P is the pressure and ϕ is the composition at a particular location in space and time.

The continuity (2.1) and momentum (2.2) equations characterize the flow of the system while the species conservation equation (2.3) specifies the scalar distribution field within the solution domain. The system of partial differential equations can be hyperbolic, elliptic or parabolic depending on the nature of the problem [Aris, 1989]. There are a limited number of cases for which an analytical solution can be found [Bird et al., 1960]. In most cases, this system of equations is solved by a grid-based method. The equations are discretized directly (finite-difference method) or through an integral form (finite-volume method) and then solved for the variables at each grid point. The use of a discrete space rather than the real continuous space is a source of error which can be controlled by making the grid finer. This choice of a grid with enough resolution defines the grid independence of the (CFD) solution. The development of solution algorithms and the solution of the above mentioned set of differential equations is the objective of *computational fluid dynamics*.

The flow equations for non-reacting systems have been solved for a variety of flow configurations and geometries. Such simulations have been used in a wide variety of applications and in the chemical industry, have aided in the development of more efficient reactors. As mentioned earlier, to get accurate results, the solution should be simulated on a fine grid (or mesh). By nature, turbulent flow contains a wide distribution of length and time scales of motion. To resolve all the scales, the mesh dimensions should be of the same order as the smallest of the length scale and the time step should be smaller than the smallest time scale. Direct Numerical Simulations (DNS) are aimed at resolving all these scales. However, DNS simulations are computationally expensive. For example, for homogeneous isotropic turbulence, it is known that the size of the mesh required in each direction is proportional to $Re^{3/4}$ where Re is the Reynolds number based on the velocity fluctuations and the integral length scale [Ferziger and Peric, 2002]. For high Reynolds number flow, the number of grid points required will be computationally intractable. In such cases the use of DNS fails. To solve such flows, certain approximations need to be made. This idea led to several solution techniques like the RANS (Reynolds Averaged Navier-Stokes) and LES (Large Eddy Simulation) methods. The low computational expense of the RANS method makes it the ideal choice for industrial applications. However, solution of reactive flow problems are even more challenging due to the strong coupling between the physics of flow and reaction. It will be discussed in later sections that the use of Probability Density Function (PDF) based methods offers a superior alternative to handle reactive flows. The RANS formulation will be discussed first followed by PDF formulation for reactive flows and coupling techniques between these methods.

Eulerian Approach to Solving NS Equations

There are several methods available for predicting turbulent flows. These can be divided into the following categories [Bardina et al., 1980] :

1. Use of correlations such as the friction factor as a function of the Reynolds number or the Nusselt number.
2. Use of integral equations that are derived from equations of motion by integrating over one or more coordinates. The system of equations reduces to ordinary differential equations that can be easily solved.

3. The third approach solves equations obtained by averaging the equations of motion over time (for statistically steady flows), or over a coordinate in which the mean flow does not vary. This leads to the *one point closure* and a set of differential equations called the RANS equations. This set of equations does not form a closed set and requires approximations (*turbulence models*).
4. The fourth set of methods are called the two-point closures and use equations that involve the correlation of velocity components at two distinct points or the Fourier transform of these equations. These methods are rarely used except in homogeneous turbulence.
5. The fifth method is *large eddy simulation* (LES) and solves the largest scale motions while modeling only the small scale motions. It can be regarded as a kind of compromise between RANS and DNS methods.
6. As explained earlier, in DNS the NS equations are solved for all of the motions in turbulent flow. Consequently this represents a computationally intensive and often intractable scheme.

The difference between the RANS, LES and DNS solutions lies in the *length scales* that they resolve. The length scale defines the scale at which a physical sub-process takes place in a turbulent flow. The RANS method resolves flow at the integral length scale (which defines the largest structures present in the flow) and models all processes below this length scale. LES resolves length scales up to the filter level (in most cases up to the grid level since many LES solutions use the grid as an implicit filter) and models all sub-grid scale (sub-filter scale) processes. DNS resolves all scales up to the molecular diffusion level (dissipation of energy from the smallest eddies onto the continuum). It is evident that as we move from the one point closures to the more detailed schemes, the computational expense will be higher. These methods have been used for a variety of geometries in different flow regimes. With high Reynolds number flows, the complexities and associated computational time increases making the computation impossible for reacting cases.

RANS Equations

In a fully developed flow, every variable can be written as the sum of an average value and fluctuations about that value:

$$\phi(x, t) = \langle \phi(x) \rangle + \phi'(x, t), \quad (2.4)$$

where $\langle \phi(x) \rangle$ is the ensemble average defined by

$$\langle \phi(x) \rangle = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N \phi(x), \quad (2.5)$$

where N is the number of fluid elements present in the domain. These definitions can be extended to unsteady flows as well [Wilcox, 1994]. For variable density and compressible flows, use of density weighted (or Favre-averaged) variables is convenient.

$$\langle \rho \rangle = \bar{\rho}. \quad (2.6)$$

$$\tilde{U} = \frac{\langle \rho U \rangle}{\bar{\rho}}. \quad (2.7)$$

Using these definitions, the continuity and momentum equations can be Reynolds averaged. For incompressible flows without body forces this yields the following equation in tensor form:

$$\frac{\partial(\bar{\rho}\tilde{U}_i)}{\partial x_i} = 0. \quad (2.8)$$

$$\frac{\partial(\bar{\rho}\tilde{U}_i)}{\partial t} + \frac{\partial}{\partial x_j} \left(\bar{\rho}\tilde{U}_i\tilde{U}_j + \bar{\rho}\langle \widetilde{u'_i u'_j} \rangle \right) = -\frac{\partial \langle p \rangle}{\partial x_i} + \frac{\partial \widetilde{\tau_{ij}}}{\partial x_j}, \quad (2.9)$$

where $\langle \tau_{ij} \rangle$ are the mean viscous stress tensor components:

$$\widetilde{\tau_{ij}} = \mu \left(\frac{\partial \tilde{U}_i}{\partial x_j} + \frac{\partial \tilde{U}_j}{\partial x_i} - \frac{2}{3} \delta_{ij} \frac{\partial \tilde{U}_k}{\partial x_k} \right) \quad i, j, k = 1, 2, 3. \quad (2.10)$$

On the left hand side of Eq. 2.9, terms of the form $\bar{\rho}\langle \widetilde{u'_i u'_j} \rangle$ that represent the correlation between the fluctuations, are not solved directly. These terms are instead modeled using either the gradient diffusion hypothesis [Rodi and Mansour, 1993; Wilcox, 1994] or transport equations with modeled production and dissipation terms [Launder et al., 1975].

The Reynolds-averaged scalar transport equation can be written as

$$\frac{\partial \bar{\rho}\tilde{\phi}}{\partial t} + \frac{\partial}{\partial x_j} \left(\bar{\rho}\tilde{U}_j\tilde{\phi} + \bar{\rho}\langle \widetilde{u'_j \phi'} \rangle \right) = \frac{\partial}{\partial x_j} \left(\Gamma \frac{\partial \tilde{\phi}}{\partial x_j} \right) + \widetilde{\mathbf{S}}(\phi). \quad (2.11)$$

The turbulent scalar flux, $\langle u' \phi' \rangle$, is modeled using the gradient diffusion hypothesis. This will involve the introduction of the turbulent diffusivity (which is obtained directly

from the turbulent viscosity and the turbulent Schmidt number), which will be modeled based on the local turbulent kinetic energy and dissipation. The next section will deal with models for the turbulent kinetic energy and dissipation fields.

The last term in Eq. 2.11 is the mean source term which often leads to computational intractability in turbulent flow simulations. In simpler laminar flows, the mean source term is the same as the source term computed at the mean composition since there are no fluctuations in the flow

$$\widetilde{\mathbf{S}(\phi)} = \mathbf{S}(\tilde{\phi}). \quad (2.12)$$

In turbulent flow, the presence of scalar fluctuations makes this laminar assumption erroneous. However, many RANS simulations are performed with this form for the source term and is called the no-closure approximation. It is impossible to exactly represent the sub-grid fluctuations in a Eulerian approach where the only information available is at the grid resolution level. All smaller scales have to be modeled. It is possible to extract the moments (e.g., variance of scalar) and compute second order accurate source terms. This is referred to as the moments method and requires the solution of the transport equation for the variance as well [Pope, 2000]. In binary mixing cases, a presumed PDF can be assumed for the scalar. Using the mixture fraction and its variance, the scalar PDF can be modeled. The presumed-PDF approach has been used in combustion and other fast chemistry flows where mixing determines the extent of reaction [Peters, 2000]. Only the transported PDF method solves for the complete PDF of the scalars and hence describes the exact source term without any approximation [Pope, 1976]. It has also been shown using a finite number of delta functions, the PDF of the scalar can be described [Villermaux and Falk, 1994; Fox, 1998]. In this case, the heights of the delta functions or the weights are solved directly in a Eulerian manner and the scalar PDF is approximated using the weights [Fox, 1996]. This method provides an inexpensive but less accurate alternative to full PDF schemes.

Turbulence Modeling

In laminar flows the energy dissipation and transport of mass, momentum, and energy normal to the streamlines are through viscosity (mass, momentum) or thermal conductivity (energy). Turbulence may be considered to enhance this viscosity. The eddy viscosity model can then be written as:

$$-\bar{\rho}\langle u'_i u'_j \rangle = \mu_t \left(\frac{\partial \tilde{U}_i}{\partial x_j} + \frac{\partial \tilde{U}_j}{\partial x_i} \right) - \frac{2}{3} \bar{\rho} \delta_{ij} \tilde{k}, \quad (2.13)$$

where \tilde{k} is the turbulent kinetic energy:

$$\tilde{k} = \frac{1}{2} \langle u'_i u'_i \rangle = \frac{1}{2} \left(\langle u'_x u'_x \rangle + \langle u'_y u'_y \rangle + \langle u'_z u'_z \rangle \right). \quad (2.14)$$

For the scalar transport equation the following model is used:

$$-\bar{\rho}\langle u'_j \phi' \rangle = \Gamma_t \frac{\partial \tilde{\phi}}{\partial x_j}, \quad (2.15)$$

where $\tilde{\phi}$ is the Favre-averaged mean scalar value.

Turbulence can then be described based on one or two variables. In the simplest case, a position dependent length scale L is the only description necessary. These mixing length models then can use the velocity, $q = \sqrt{2\tilde{k}}$ to obtain the turbulence energy field where q is given by:

$$q = L \frac{\partial \tilde{U}}{\partial y}. \quad (2.16)$$

Since a length and velocity scale are required to define the turbulence field, a model based on two such equations is used. In almost all these models, a transport equation for the kinetic energy of the form below is solved [Rodi and Mansour, 1993]:

$$\frac{\partial(\bar{\rho}\tilde{k})}{\partial t} + \frac{\partial(\bar{\rho}\tilde{U}_j\tilde{k})}{\partial x_j} = \frac{\partial}{\partial x_j} \left(\left[\mu_j + \frac{\mu_t}{\sigma_k} \right] \frac{\partial \tilde{k}}{\partial x_j} \right) + P_k - \bar{\rho}\tilde{\epsilon}, \quad (2.17)$$

where $\tilde{\epsilon}$ is the dissipation factor and P_k is the rate of production of turbulent kinetic energy. The production term is obtained by using the eddy viscosity model on the Reynolds stress terms.

$$P_k = -\rho\langle u'_i u'_j \rangle \frac{\partial \tilde{U}_i}{\partial x_j} \approx \mu_t \left(\frac{\partial \tilde{U}_i}{\partial x_j} + \frac{\partial \tilde{U}_j}{\partial x_i} \right) \frac{\partial \tilde{U}_i}{\partial x_j}. \quad (2.18)$$

When the rates of production and destruction of turbulence are nearly equal, the dissipation, $\tilde{\epsilon}$ and \tilde{k} and L (integral length scale) are related by

$$\tilde{\epsilon} \approx \frac{\tilde{k}^{3/2}}{L}. \quad (2.19)$$

In the $\tilde{k} - \tilde{\epsilon}$ model a transport equation for the dissipation factor is also solved [Wilcox, 1994]:

$$\frac{\partial \rho \tilde{\epsilon}}{\partial t} + \frac{\partial \bar{\rho} \tilde{U}_j \tilde{\epsilon}}{\partial x_j} = \frac{\partial}{\partial x_j} \left(\frac{\mu_t}{\sigma_\epsilon} \frac{\partial \tilde{\epsilon}}{\partial x_j} \right) + C_{\epsilon 1} P_k \frac{\tilde{\epsilon}}{\tilde{k}} - C_{\epsilon 2} \bar{\rho} \frac{\tilde{\epsilon}^2}{\tilde{k}}. \quad (2.20)$$

The eddy viscosity or the turbulent viscosity can be expressed as

$$\mu_t = \rho C_\mu \sqrt{\tilde{k}} L = \rho C_\mu \frac{\tilde{k}^2}{\tilde{\epsilon}}. \quad (2.21)$$

The problem with most of these RANS models is the determination of the constants. Experimental correlation is needed to find reasonably accurate values of the constants. In general the following values are used in most simulations [Wilcox, 1994]: $C_\mu = 0.09$, $C_{\epsilon 1} = 1.44$, $C_{\epsilon 2} = 1.92$, $\sigma_k = 1.0$, $\sigma_\epsilon = 1.3$.

The $\tilde{k} - \tilde{\epsilon}$ model is one of the most widely-used turbulence models. The model performance can be improved by tuning the constants [Udaipurwala et al., 1994]. One of the main deficiencies of the model is that it over-predicts the spreading rate of round jets. Some corrections [Pope, 1978] have been proposed. In addition to the standard $\tilde{k} - \tilde{\epsilon}$ model various modifications [Hwang et al., 1993] are available. In general, these corrections are valid only for specific flows and perform worse than the standard model for general flows.

As suggested above, the idea behind the two-equation models are to use an equation for the kinetic energy transport and use one of many other choices for the second variable (like kL , ω , ω^2 or τ). Another widely used model is the $\tilde{k} - \tilde{\omega}$ model [Wilcox, 1994]. The transport equation for ω can be written as:

$$\frac{\partial \tilde{\rho} \tilde{\omega}}{\partial t} + \frac{\partial \tilde{\rho} \tilde{U}_j \tilde{\omega}}{\partial x_j} = C_{\omega 1} \frac{P_k \tilde{\omega}}{\tilde{k}} - C_{\omega 2} \tilde{\omega}^2 + \frac{\partial}{\partial x_j} \left(\frac{\mu_t}{\sigma_\omega} \frac{\partial \tilde{\omega}}{\partial x_j} \right). \quad (2.22)$$

The $\tilde{k} - \tilde{\omega}$ model is superior in treatment of the near-wall region and in accounting for the effects of the stream-wise pressure gradients. The definition for the eddy viscosity is the same as the $\tilde{k} - \tilde{\epsilon}$ model and is rewritten in terms of $\tilde{\omega}$.

$$\mu_t = C_\mu \tilde{k} \tilde{\omega}. \quad (2.23)$$

The length scales at which different physical processes occur differ widely in fluid flow. The equations of motion for modeled turbulent flow resemble the laminar flow equations. The only difference is in the treatment of the viscosity. Hence the solution schemes are similar to the laminar case except the length and time scales for the model equations (\tilde{k} and $\tilde{\epsilon}$ or \tilde{k} and $\tilde{\omega}$) are much smaller than the flow equations themselves and hence lead to a stiff system. The solution schemes should be able to handle such stiff equations.

The reacting flow problems with detailed chemistry pose another constraint - the wide variation of the *reaction time scale*. In multi-species kinetic schemes (like the

scheme for the chlorination reactor) the time scales of the individual reactions vary by orders of magnitude. Consequently the source term in the species transport equations give rise to *stiff differential equations*. This limits the time step used for computation. DNS is computationally intractable for large kinetic mechanism based simulations. In general, the LES and RANS methods use similar modeling approaches for reactive flows. However, the additional cost of handling multiple species transport equations can make LES schemes computationally expensive. In spite of the approximations involved in the RANS based methods, they seem to offer the only tractable solution scheme. Solution of detailed kinetics is vital in studying flow based reactors especially those systems that involve comparable mixing and reaction time scales. The usual RANS approaches of presumed-PDF method or flamelet models are not general purpose solutions due to the assumptions inherent to the models. By solving for the underlying probability density function (PDF) of the joint scalar or joint velocity-scalar fields in the flow, information about the mean and higher moments of these variables can be readily obtained. For reasons to be discussed later, it may not be feasible to use a stand-alone PDF method to handle both velocity and scalar fields.

In this work, a hybrid FV - Lagrangian PDF code is used to study the flow and concentration fields in turbulent reactive flows. A brief overview of the compressible flow model used in the FV code is given below, followed by a detailed description of the composition PDF model. It should be noted that the finite-volume method used here is one of the many different schemes that can be implemented to solve the momentum and continuity equation. More details on both finite-volume and Lagrangian PDF models can be found elsewhere [Ferziger and Peric, 2002; Pope, 1981; Haworth and El Tahry, 1991; Muradoglu et al., 1999; Xu and Pope, 1999].

Compressible Flow Model

Using standard notation for mean flow quantities, $(\bar{\rho}, \bar{p}, \tilde{T}, \tilde{h}$ and \tilde{u}_β are the Reynolds-averaged mean density and pressure, and Favre-averaged temperature, enthalpy and Cartesian velocity components) and neglecting body forces and external heat sources, the conservative Reynolds-averaged Navier-Stokes (RANS) equations for the flow of a compressible fluid can be written in a generalized frame of reference $\xi_\alpha = \xi_\alpha(x_\beta)$ ($\alpha, \beta = 1, 2, 3$) as

$$J \left(\frac{\partial \mathbf{Q}}{\partial t} + \Gamma \frac{\partial \mathbf{W}}{\partial \tau} \right) + \frac{\partial \mathbf{H}_\alpha}{\partial \xi_\alpha} = 0, \quad (2.24)$$

where J is the Jacobian of the transformed coordinates with respect to the physical frame of reference, and repeated indices imply summation (e.g., $A_\alpha B_\alpha = A_1 B_1 + A_2 B_2 + A_3 B_3$). The quantity $\mathbf{Q} = [\bar{\rho}, \bar{\rho}\tilde{u}_\beta, \bar{\rho}E]^T$ represents the vector of conserved variables with E as the total specific internal energy, and $\mathbf{W} = [\bar{p}, \tilde{u}_\beta, \tilde{T}]^T$ is the vector of dependent variables which includes the pressure, velocity components and temperature. A pseudo-time term (τ) is added for low-Mach number preconditioning and Γ represents a suitable preconditioner [Shuen et al., 1993]. The vector \mathbf{H}_α is the inviscid and viscous flux difference in the ξ_α coordinate direction:

$$\mathbf{H}_\alpha = \mathbf{F}_\alpha^I - \mathbf{F}_\alpha^V = J \begin{pmatrix} \bar{\rho}\tilde{U}_\alpha \\ \bar{\rho}\tilde{u}_\beta\tilde{U}_\alpha + \bar{p}\frac{\partial\xi_\alpha}{\partial x_\beta} \\ \tilde{U}_\alpha(\bar{\rho}E + \bar{p}) \end{pmatrix} - J \begin{pmatrix} 0 \\ \sigma_{\beta\gamma}\partial\xi_\alpha/\partial x_\gamma \\ (\tilde{u}_\beta\sigma_{\beta\gamma} + q_\gamma)\frac{\partial\xi_\alpha}{\partial x_\gamma} \end{pmatrix}. \quad (2.25)$$

The stress tensor is defined as $\sigma_{\alpha\beta} = -2\mu(S_{\alpha\beta} - \frac{1}{3}S_{\gamma\gamma}\delta_{\alpha\beta})$ with

$$S_{\alpha\beta} = \frac{1}{2} \left(\frac{\partial\tilde{u}_\alpha}{\partial\xi_\gamma} \frac{\partial\xi_\gamma}{\partial x_\beta} + \frac{\partial\tilde{u}_\beta}{\partial\xi_\gamma} \frac{\partial\xi_\gamma}{\partial x_\alpha} \right). \quad (2.26)$$

The quantity \tilde{U}_α is the contravariant velocity: $\tilde{u}_\beta\partial\xi_\alpha/\partial x_\beta$. q_α is the total heat flux, written as $q_\beta = (\tilde{k} + \mu_t C_p / Pr_t)\partial\tilde{T}/\partial x_\beta$. The thermal conductivity k is calculated from the laminar viscosity and Prandtl numbers, Pr_t is the turbulent Prandtl number set equal to 0.9 for the present study, C_p is the specific heat that is calculated from polynomial fits based on temperature and μ_t is the turbulent eddy viscosity.

The inviscid flux, \mathbf{F}^I is separated into convective and pressure components using a second-order, low-diffusion flux-splitting scheme (LDFSS) [Edwards, 1997]. A second-order backward three-point physical time differencing is used for time-dependent term, an Euler differencing is used for the pseudo-time derivative and second-order central differences are used for the viscous terms. The system is closed using the ideal gas law, $\bar{p} = \bar{\rho}R\tilde{T}/W$, where W is the molecular weight of the gas mixture. It should be noted that the energy equation is not solved to determine the enthalpy or temperature, which are directly obtained from the PDF code (described below) and remain ‘‘frozen’’ during the FV iterations. The energy equation is retained as part of the original FV formulation and is not needed in the hybrid code. The temperature values are used to calculate thermo-physical properties of the fluid. Equation 2.25 is linearized and the resulting system is solved using a fully implicit multi-grid-level line relaxation scheme. In the multi-level scheme the system is solved on increasingly finer grids, interpolating the solution from coarser grids. However, the PDF code is restricted to the finest

grid. Transport equations for the Favre-averaged mean turbulent kinetic energy \tilde{k} and dissipation rate, $\tilde{\epsilon}$ are solved as a separate coupled system. The standard two-layer form of the model is used [Rodi and Mansour, 1993].

In the present work the above system is solved for a single constituent with a spatially varying molecular weight to obtain the velocity ($\tilde{\mathbf{u}}$), and turbulence quantities (\tilde{k} and $\tilde{\epsilon}$). Based on these flow quantities, the PDF code is used to obtain a new temperature (\tilde{T}), density $\bar{\rho}$ and molecular weight (W), which are fed back to the RANS code to update the velocity and turbulence fields. The PDF code models the original PDF transport equations through a stochastically equivalent system. The theory behind the PDF method is discussed next.

Probability Density Function Methods

Probability Density Function methods refer to a class of methods that model the stochastic nature of the system by computing the PDF of the underlying state variables. When more than a single variable describes the system, a joint probability density function of the variables is required. The next sub-section describes some of the basic statistical terms. It is followed by the derivation of the one-point Eulerian PDF transport equation and a description of the modeling involved. Then, stochastically equivalent systems will be introduced in terms of a Lagrangian PDF description. The physical sub-processes like mixing and reaction will be discussed with an analysis of the bottlenecks in the use of PDF methods. The numerical implementation of the PDF method and the hybrid scheme in general will be covered in the next chapter.

Fundamental Statistical Definitions

The cumulative probability Distribution Function for a random variable $\mathbf{U}(\mathbf{X}, t)$ at fixed point \mathbf{X} in space at time t is defined as

$$F_{\mathbf{U}}(\mathbf{V}; \mathbf{X}, t) = P(\mathbf{U}(\mathbf{X}, t) < \mathbf{V}), \quad (2.27)$$

where \mathbf{V} is the sample space variable of \mathbf{U} which takes all possible values of \mathbf{U} , and P is the probability that \mathbf{U} is smaller than \mathbf{V} . The *probability density function* is defined as the derivative of the distribution function with respect to the sample variable:

$$f_{\mathbf{U}}(\mathbf{V}; \mathbf{X}, t) = \frac{\partial}{\partial \mathbf{V}} F_{\mathbf{U}}(\mathbf{V}; \mathbf{X}, t). \quad (2.28)$$

This essentially means that the PDF is the probability that the random variable takes a value between \mathbf{V} and $\mathbf{V} + d\mathbf{V}$ for different realizations of the turbulent flow. The mean value or the *expectation* of a statistical quantity is defined by

$$\langle \mathbf{U}(\mathbf{X}, t) \rangle = \iiint_{-\infty}^{+\infty} \mathbf{V} f_{\mathbf{U}}(\mathbf{V}; \mathbf{X}, t) d\mathbf{V}. \quad (2.29)$$

When describing a velocity field for a turbulent flow, typically the variable U_i represents the three velocity components and is defined as a vector \mathbf{U} and the sample space by vector \mathbf{V} . The velocity fluctuations are given by

$$u_i = U_i - \langle U_i \rangle. \quad (2.30)$$

Higher-order moments like the variance can then be defined as

$$\langle u_i u_j \rangle = \iiint_{-\infty}^{+\infty} (V_i - \langle U_i \rangle)(V_j - \langle U_j \rangle) f_{\mathbf{U}}(\mathbf{V}; \mathbf{X}, t) d\mathbf{V}. \quad (2.31)$$

In turbulent flow simulations the concept of joint scalar and velocity PDF is useful. The joint cumulative distribution function is defined by Eq. 2.27 for a vector \mathbf{U} and the corresponding probability density function is given by Eq. 2.28.

In flows with multiple species a joint velocity composition PDF is defined by

$$f_{\mathbf{U}, \phi}(\mathbf{V}, \psi; \mathbf{X}, t) d\mathbf{V} d\psi = P(\mathbf{V} \leq \mathbf{U}(\mathbf{X}, t) < \mathbf{V} + d\mathbf{V}, \psi \leq \phi(\mathbf{X}, t) < \psi + d\psi). \quad (2.32)$$

In statistical terms, the conditional probability of an event $P(A|B)$ is defined as the probability that A occurs given that B occurs:

$$P(A|B) = \frac{P(A \cap B)}{P(B)}. \quad (2.33)$$

In a similar manner, a conditional PDF of a scalar conditioned on velocity can be defined as

$$f_{\phi|\mathbf{U}}(\psi|\mathbf{V}) = \frac{f_{\phi, \mathbf{U}}(\psi, \mathbf{V})}{f_{\mathbf{U}}(\mathbf{V})}. \quad (2.34)$$

Based on this equation, conditional averages can be defined for a function, say, $\langle Q(\phi, \mathbf{U}) | \mathbf{U} = \mathbf{V} \rangle$ based on the conditional PDF defined above as

$$\langle Q(\phi, \mathbf{U}) | \mathbf{U} = \mathbf{V} \rangle = \iiint_{-\infty}^{+\infty} Q(\psi, \mathbf{V}) f_{\phi|\mathbf{U}}(\psi|\mathbf{V}) d\psi. \quad (2.35)$$

All the unclosed terms in the PDF transport equation to be derived next, can be written as conditional averages of a non-linear term conditioned on one of the describing variables like \mathbf{V} .

Eulerian PDF Formulation

The Eulerian one-point one-time PDF $f_{\mathbf{U}\phi}(\mathbf{V}, \psi; \mathbf{x}, t)$ is defined as the joint probability density function of the $\mathbf{U}(\mathbf{x}, t)$ and composition $\phi(\mathbf{x}, t)$ with corresponding sample space variables \mathbf{V} and ψ . First, the transport equation for the velocity PDF will be derived. Then the scalar conservation equation will be used to arrive at the scalar PDF transport equation. More details on the derivations can be found in Pope [2000] and Fox [2003]. In order to derive the transport equation, it will be useful to define the *fine-grained* form of the Eulerian PDF. For the velocity PDF $f_{\mathbf{U}}(\mathbf{V}; \mathbf{x}, t)$ this can be written as

$$f'_{\mathbf{U}}(\mathbf{V}; \mathbf{x}, t) = \prod_{i=1}^3 \delta(U_i(\mathbf{x}, t) - V_i). \quad (2.36)$$

The δ function has certain mathematical properties that can be used in deriving the PDF transport equation. The *sifting property* is defined as

$$\int_{-\infty}^{\infty} \delta(x - a)g(x)dx = g(a). \quad (2.37)$$

This definition is used to compute the mean of the fine-grained PDF:

$$\langle f'_{\mathbf{U}}(\mathbf{V}; \mathbf{x}, t) \rangle = \int_{-\infty}^{\infty} \delta(\mathbf{V}' - \mathbf{V})f_{\mathbf{U}}(\mathbf{V}'; \mathbf{X}, t)d\mathbf{V}' = f_{\mathbf{U}}(\mathbf{V}; \mathbf{x}, t). \quad (2.38)$$

The above property can be defined for the derivatives of the δ function as well. The first derivative has the following property

$$\delta^{(1)}(v - a) = -\delta^{(1)}(a - v). \quad (2.39)$$

It can also be shown by using Eq. 2.37 that

$$f(x)\delta(x - a) = f(a)\delta(x - a). \quad (2.40)$$

Division of the δ function is not defined and hence, the above relation cannot be reduced further by cancellation. Using these properties, the time and spatial derivatives of the fine-grained PDF can be obtained:

$$\frac{\partial}{\partial t} f'_{\mathbf{U}}(\mathbf{V}; \mathbf{x}, t) = -\frac{\partial f'_{\mathbf{U}}(\mathbf{V}; \mathbf{x}, t)}{\partial V_i} \frac{\partial U_i(\mathbf{x}, t)}{\partial t}, \quad (2.41)$$

$$\frac{\partial}{\partial x_i} f'_{\mathbf{U}}(\mathbf{V}; \mathbf{x}, t) = -\frac{\partial f'_{\mathbf{U}}(\mathbf{V}; \mathbf{x}, t)}{\partial V_j} \frac{\partial U_j(\mathbf{x}, t)}{\partial x_i}. \quad (2.42)$$

Hence the transport equation of the velocity PDF can be written as

$$\frac{\partial f'}{\partial t} + V_i \frac{\partial f'}{\partial x_i} = -\frac{\partial}{\partial V_i} \left(f' \frac{DU_i}{Dt} \right), \quad (2.43)$$

where the f' refers to $f'_U(\mathbf{V}; \mathbf{x}, t)$. The right-hand side of this equation contains the substantial derivative of the velocity that can be described using the Navier-Stokes equation. The mean equation is then transformed to

$$\frac{\partial f}{\partial t} + V_i \frac{\partial f}{\partial x_i} = -\frac{\partial}{\partial V_i} \left(f \left\langle \mu \nabla^2 U_i - \frac{1}{\rho} \frac{\partial p}{\partial x_i} \middle| V \right\rangle \right). \quad (2.44)$$

The pressure term can be decomposed into the mean and fluctuating components to yield the final transport equation:

$$\frac{\partial f}{\partial t} + V_i \frac{\partial f}{\partial x_i} = \frac{1}{\rho} \frac{\partial \langle p \rangle}{\partial x_i} \frac{\partial f}{\partial V_i} - \frac{\partial}{\partial V_i} \left(f \left\langle \mu \nabla^2 U_i - \frac{1}{\rho} \frac{\partial p'}{\partial x_i} \middle| V \right\rangle \right). \quad (2.45)$$

A similar equation for the PDF of the fluctuating velocity can be derived [Pope, 2000] but is omitted here. It can be noticed from the mean transport equation that two unclosed terms arise which are conditioned on the velocity. PDF modeling aims at providing consistent closures for these two terms. The Generalized Langevin Model (GLM) [Haworth and Pope, 1986] uses two new terms G_{ij} and C_0 to jointly model the unclosed terms. The transport equation using the GLM model can be rewritten as:

$$\frac{\partial f}{\partial t} + V_i \frac{\partial f}{\partial x_i} - \frac{1}{\rho} \frac{\partial \langle p \rangle}{\partial x_i} \frac{\partial f}{\partial V_i} = -\frac{\partial}{\partial V_i} [f G_{ij} (V_j - \langle U_j \rangle)] + \frac{1}{2} C_0 \epsilon \frac{\partial^2 f}{\partial V_i \partial V_i}. \quad (2.46)$$

G_{ij} has dimensions of inverse time and C_0 is non-dimensional. They generally evolve with the flow and are a function of space and time. However, in the simplified Langevin model (SLM) C_0 is assumed to be a constant. It can be noted that the model introduces the dissipation $\epsilon(x, t)$ and hence is not completely closed. A joint velocity- frequency [Van Slooten et al., 1998] or a velocity-dissipation [Pope and Chen, 1990] formulation is used in which the transport equation for the PDF involving dissipation (or frequency) is also solved. The equations discussed above solve for the PDF of the system variables. Mean equations of the moments of the flow variables can be obtained from the PDF transport equation. For example, the PDF equation can be integrated after being multiplied by the velocity to obtain the mean momentum equation. The next section will deal with the Lagrangian PDF and will provide a practical approach to solving the PDF transport equation.

Lagrangian PDF Formulation

Pope [1985, 2000] explains the use of fluid particle terminology in the definition of the Lagrangian modeling approach. Here, instead of modeling the Eulerian PDF transport equation, the fluid is viewed in terms of particles with particular location and velocity. The forces on the particle, which in the Eulerian sense will lead to the Navier-Stokes equation, are modeled. The Lagrangian PDF is then defined in terms of a fluid particle originating from position $\mathbf{X}^p(0, \mathbf{Y}) = \mathbf{Y}$ with velocity $\mathbf{U}^p(t, \mathbf{Y})$ and is denoted by $f_L(\mathbf{V}, \mathbf{x}; t|\mathbf{Y})$. Similar to the Eulerian formulation, a fine-grained PDF $f'_L(\mathbf{V}, \mathbf{x}; t|\mathbf{Y})$ is defined in terms of delta functions:

$$f'_L(\mathbf{V}, \mathbf{x}; t|\mathbf{Y}) = \delta(\mathbf{U}^p(t, \mathbf{Y}) - \mathbf{V}) \delta(\mathbf{X}^p(t, \mathbf{Y}) - \mathbf{x}). \quad (2.47)$$

Properties of the Eulerian fine-grained PDF hold for the Lagrangian PDF as well. In addition, for incompressible flows it can be shown that [Pope, 2000]:

$$\int f'_L d\mathbf{Y} = \int \delta(\mathbf{U}^p(t, \mathbf{Y}) - \mathbf{V}) \delta(\mathbf{X}^p(t, \mathbf{Y}) - \mathbf{x}) d\mathbf{Y} \quad (2.48)$$

$$= \delta(\mathbf{U}^p(t, \mathbf{Y}) - \mathbf{V})_{\mathbf{X}^p(t, \mathbf{Y}) = \mathbf{x}} \quad (2.49)$$

$$= \delta(\mathbf{U}(\mathbf{x}, t) - \mathbf{V}). \quad (2.50)$$

The last term is the Eulerian fine-grained PDF and the mean of this term is the Eulerian PDF. By taking the mean of f'_L in the velocity space, it follows that

$$\int f_L d\mathbf{Y} = f(\mathbf{V}; x, t). \quad (2.51)$$

This relation between the Lagrangian and Eulerian PDFs ensures that the modeling of the flow based on the Lagrangian approach can be readily translated into the Eulerian form. The Eulerian formulation can be then manipulated to yield mean flow equations for the velocity and higher moments. Hence, equations for the fluid particle motion can be formulated that yield the correct Eulerian transport equation. Such models arise from the *Langevin Equation* that describes the stochastic velocity of particle in Brownian motion. It can be shown that the Langevin equation yields the correct Eulerian PDF in homogeneous turbulence (where the Lagrangian and Eulerian one-point one-time PDF are identical) [Pope, 2000]. The Langevin equation is a stochastic differential equation (SDE) which evolves a PDF that follows the *Fokker-Planck* equation [Gardiner, 1983]. The SDE is given by

$$dU^* = -U^* \frac{dt}{T_L} + \left(\frac{2\sigma^2}{T_L} \right)^{1/2} dW, \quad (2.52)$$

where T_L is a time-scale for the process and σ^2 is a positive constant. However, the original Langevin equation needs to be modified to ensure that certain fundamental turbulent concepts are obeyed. The Langevin equation evolves a statistically stationary Gaussian process for which, in isotropic turbulence the variance of the process is well-defined in terms of the local turbulent kinetic energy. The Kolmogorov hypothesis on structure function is invoked to relate the diffusion term to the dissipation of energy [Pope, 2000]. Put together, the Langevin equation that describes isotropic turbulence is given by

$$dU^* = -\frac{3}{4}C_0\frac{\epsilon}{k}dt + (C_0\epsilon)^{1/2}dW(t), \quad (2.53)$$

where C_0 is the Kolmogorov universal constant and $dW(t)$ refers to white noise with $\langle dW(t) \rangle = 0$ and $\langle dW(t)dW(t) \rangle = dt$. The GLM [Haworth and Pope, 1986] is an extension of the Langevin equation to inhomogeneous flows. As noted by Pope [2000], GLM is a class of methods that model the particle flow to obtain a closed form for the PDF transport equation. In general for a diffusion process in multi-dimensional space $\mathbf{Z} = \{\mathbf{X}, \mathbf{U}\}$, the Langevin equation can be written as:

$$d\mathbf{Z}^* = a_i(\mathbf{Z}^*, t)dt + B_{ij}(\mathbf{Z}^*, t)d\mathbf{W}(t), \quad (2.54)$$

where the drift term is specified as a vector of coefficients.

$$\begin{bmatrix} a(\mathbf{X}^*, t) \\ a(\mathbf{U}^*, t) \end{bmatrix} = \begin{bmatrix} \mathbf{U}^* \\ -\frac{1}{\rho}\frac{\partial(p)}{\partial\mathbf{x}} + G_{ij}(\langle\mathbf{U}\rangle - \mathbf{U}^*) \end{bmatrix}. \quad (2.55)$$

The diffusion term consists of the diagonal matrix,

$$B_{ij} = \begin{bmatrix} 0 & \dots & \dots & \dots \\ \dots & 0 & \dots & \dots \\ \dots & \dots & 0 & \dots \\ \dots & \dots & (C_0\epsilon)^{\frac{1}{2}} & 0 & 0 \\ \dots & \dots & 0 & (C_0\epsilon)^{\frac{1}{2}} & 0 \\ \dots & \dots & 0 & 0 & (C_0\epsilon)^{\frac{1}{2}} \end{bmatrix}. \quad (2.56)$$

Using this general form of the model, the corresponding Fokker-Planck equation for the evolution of the fluid particle PDF f^* can be written as

$$\frac{\partial f^*}{\partial t} + V_i \frac{\partial f^*}{\partial x_i} = -\frac{\partial}{\partial V_i} (a_i f^*) + B_{ij} \frac{\partial^2 f^*}{\partial V_i \partial V_j}, \quad (2.57)$$

where $B_{ij} = \mathcal{B}_{ik}\mathcal{B}_{kj}$. Direct solution to the PDF transport equation (along with a model for the dissipation term) will yield all the moments of the velocity field and will completely determine the system. However, even for simple flows, such a direct method is not feasible due to the high dimensionality of the equation. To overcome this problem, the use of the particle system based on the joint PDF of *particle* position \mathbf{X}^* and \mathbf{U}^* , $f_L^*(\mathbf{V}, \mathbf{x}; t)$ is considered [Pope, 1985]. Here the particle refers to a notional particle that evolves according to a SDE with coefficients determined by the local flow properties such that the PDF evolved corresponds to the Fokker-Planck equation for the Lagrangian one-time one-point PDF.

The marginal PDF of position of the particles is given by

$$f_{\mathbf{x}}^*(\mathbf{x}; t) = \int f_L^*(\mathbf{V}, \mathbf{x}; t) d\mathbf{V}. \quad (2.58)$$

The conditional PDF of velocity given position can be given by

$$f^*(\mathbf{V}|\mathbf{x}, t) = \frac{f_L^*(\mathbf{V}, \mathbf{x}; t)}{f_{\mathbf{x}}^*(\mathbf{x}; t)}. \quad (2.59)$$

This relation can be used to calculate the conditional mean of velocity and its moments with respect to the particle position. The particle system will correspond to the original fluid system in terms of the mean and other velocity statistics conditioned on the position. The continuity equation for the particle system can then be written as

$$\nabla \cdot \langle \mathbf{U}^*(t) | \mathbf{x} \rangle = 0 \quad (2.60)$$

The particles move with the local velocity, and the PDF f_L^* will follow the Fokker-Planck equation,

$$\frac{\partial f_L^*}{\partial t} + V_i \frac{\partial f_L^*}{\partial x_i} = -\frac{\partial}{\partial V_i} (a_i f_L^*) + \frac{\partial^2}{\partial V_i \partial V_j} \left(\frac{1}{2} b^2 f_L^* \right), \quad (2.61)$$

where $a_i(\mathbf{U}, \mathbf{x}, t)$ and $b(\mathbf{x}, t)$ are the drift and diffusion terms in the particle Langevin equations. For the particle system to evolve the PDF corresponding to the fluid system, the conditional Lagrangian velocity $f^*(\mathbf{V}|\mathbf{x}, t)$ should match $f_{\mathbf{U}}(\mathbf{V}; \mathbf{x}, t)$. Dividing the above equation by $f_{\mathbf{x}}$ we get

$$\frac{\partial f^*}{\partial t} + V_i \frac{\partial f^*}{\partial x_i} + (V_i - \langle U_i^* | \mathbf{x} \rangle) \frac{f^*}{f_{\mathbf{x}}} \frac{\partial f_{\mathbf{x}}}{\partial x_i} = -\frac{\partial}{\partial V_i} (a_i f^*) + \frac{\partial^2}{\partial V_i \partial V_j} \left(\frac{1}{2} b^2 f^* \right). \quad (2.62)$$

The third term on the left is an extra term that does not correspond to any term in the transport equation for the Eulerian PDF of velocity. To ensure that the equations

correspond, this term has to be zero which can be satisfied only if $f_{\mathbf{x}}$ is uniform. In addition,

$$\int \mathbf{V} f_L^*(\mathbf{V}, \mathbf{x}; t) d\mathbf{V} = \int \mathbf{V} f_{\mathbf{X}}^*(\mathbf{x}; t) f_L^*(\mathbf{V}|\mathbf{x}, t) d\mathbf{V} = \langle U_i^* | \mathbf{x} \rangle f_{\mathbf{X}}^*(\mathbf{x}; t). \quad (2.63)$$

Integrating Eq. 2.61 over the velocity space, it can be seen that the right hand side becomes zero and the left hand side reduces to

$$\left(\frac{\partial}{\partial t} + \langle U_i^* | \mathbf{x} \rangle \frac{\partial}{\partial x_i} \right) \ln f_{\mathbf{X}}^*(\mathbf{x}; t) = - \frac{\partial \langle U_i^* | \mathbf{x} \rangle}{\partial x_i}. \quad (2.64)$$

The continuity equation makes the right-hand-side zero which implies that the distribution $f_{\mathbf{X}}^*(\mathbf{x}; t)$ is constant. This implies that any initial uniform distribution of particles stays uniform. This conclusion is independent of the drift and diffusion terms. It will be later seen that for a composition PDF, certain additional conditions based on these coefficients are necessary to satisfy uniformity.

The particle system represents only a model of the actual system and should be compared with the fluid system based on the moments of the PDF. The Lagrangian one-time one-point PDF contains no information about the multi-point statistics and the conditional means based on particle locations of such quantities need not correspond. Comparing with the Reynolds-equation and neglecting the viscous terms, a particle pressure P can be defined such that

$$\frac{\partial \langle U_i^* | \mathbf{x} \rangle}{\partial t} + \frac{\partial \langle U_i^* U_j^* | \mathbf{x} \rangle}{\partial x_i} = - \frac{1}{\rho} \nabla P. \quad (2.65)$$

By taking the divergence of this equation, a Poisson equation for the particle pressure can be defined:

$$\nabla^2 P = - \rho \frac{\partial^2 \langle U_i^* U_j^* | \mathbf{x} \rangle}{\partial x_i \partial x_j}. \quad (2.66)$$

This equation shows one of the main reasons for the instabilities of the PDF methods. First and higher order derivatives are difficult to obtain due to the statistical fluctuations associated with the model. Use of the particle scheme hence requires a large number of particles to reduce the fluctuations and thereby leading to increased computational expense. Several flow situations have been tackled by the use of the joint velocity PDF equation along with a model for the dissipation term [Pope and Chen, 1990; Van Slooten et al., 1998]. Wall functions and near wall modeling have been studied [Dreeben and Pope, 1997b,a]. It has been shown that these models can reproduce DNS and experimental data over a wide range of operating conditions [Haworth and Pope, 1986].

Before moving onto the scalar transport equations and PDF modeling of reactive flow problems, mean estimations using particle methods is briefly considered.

It is apparent that the calculation of the mean fields and the statistics are of practical importance. Since the particles are distributed uniformly, an estimate of the mean at a particular location can be made based on a kernel estimation of all the particles centered around that point. The form of the kernel is irrelevant but should minimize the cost of mean estimation. The most common method is to use a particle-mesh method. So far, the PDF method has been considered as a grid-free technique as opposed to the Eulerian approach. When using the particle-mesh method, a grid is superimposed on the domain and estimates are made using a cloud-in-cell approach at the node points. Based on these fields, splines can be formed which can then interpolate fields onto the particle positions. Alternatively, a simple and inexpensive scheme is to use cell estimates where the particles that fall within each cell of the grid are considered to estimate the means at the cell centers. With the use of the hybrid method, to be discussed later, use of the cell estimate is found to be effective and adaptable for parallel algorithms as well. The velocity PDF formulation has been illustrated to introduce the idea of Lagrangian particle simulation and the practical implementation of numerical algorithms. However, due to the requirement of the second derivatives as the source term for the Poisson equation, the method can be unstable. The joint composition PDF is introduced with an intention of using it as a species transport solver.

Scalar Transport using PDF Method

The scalar transport equation for a single species can be written as

$$\frac{D\phi}{Dt} = \Gamma \nabla^2 \phi + S(\phi(\mathbf{x}, t)). \quad (2.67)$$

If the one time one point Eulerian PDF is denoted by $f_\phi(\psi; \mathbf{x}, t)$, the corresponding fine-grained PDF is defined by

$$f'_\phi(\psi; \mathbf{x}, t) = \delta(\phi(\mathbf{x}, t) - \psi). \quad (2.68)$$

The fine-grained PDF definition can then be manipulated to obtain the following mathematical relations:

$$\frac{\partial f'_\phi}{\partial t} = \delta^{(1)}(\phi(\mathbf{x}, t) - \psi) \frac{\partial \phi}{\partial t} \quad (2.69)$$

$$= -\delta^{(1)}(\psi - \phi(\mathbf{x}, t)) \frac{\partial \phi}{\partial t} = -\frac{\partial f'_\phi}{\partial \psi} \frac{\partial \phi}{\partial t} = -\frac{\partial}{\partial \psi} \left(f'_\phi \frac{\partial \phi}{\partial t} \right). \quad (2.70)$$

Similarly the spatial derivative can be manipulated to yield

$$\frac{\partial f'_\phi}{\partial x_i} = \delta^{(1)}(\phi(\mathbf{x}, t) - \psi) \frac{\partial \phi}{\partial x_i} \quad (2.71)$$

$$= -\delta^{(1)}(\psi - \phi(\mathbf{x}, t)) \frac{\partial \phi}{\partial x_i} = -\frac{\partial f'_\phi}{\partial \psi} \frac{\partial \phi}{\partial x_i} = -\frac{\partial}{\partial \psi} \left(f'_\phi \frac{\partial \phi}{\partial x_i} \right). \quad (2.72)$$

Differentiating both sides of the above equation for the spatial derivative, the following relation is derived

$$\nabla^2 f'_\phi = -\frac{\partial}{\partial \psi} (f'_\phi \nabla^2 \phi) + \frac{\partial^2}{\partial \psi^2} (f'_\phi \nabla \phi \cdot \nabla \phi). \quad (2.73)$$

Using the definitions of the first spatial and time derivatives, the substantial derivative of f'_ϕ can be written simply as

$$\frac{Df'_\phi}{Dt} = \frac{\partial f'_\phi}{\partial t} + U_i(\mathbf{x}, t) \frac{\partial f'_\phi}{\partial x_i} = -\frac{\partial}{\partial \psi} \left(f'_\phi \frac{D\phi}{Dt} \right). \quad (2.74)$$

In order to find the mean of this equation, the following term is modified as

$$\langle f'_\phi U \rangle = \int U_i \delta(\phi(\mathbf{x}, t) - \psi) f_\phi d\psi \quad (2.75)$$

$$= \int (\langle U_i \rangle + u'_i) \delta(\phi(\mathbf{x}, t) - \psi) f_\phi d\psi \quad (2.76)$$

$$= (\langle U_i \rangle + \langle u'_i | \phi(\mathbf{x}, t) = \psi \rangle) f_\phi. \quad (2.77)$$

The mean velocity does not depend on the scalar random field and hence appears as an unconditional mean. Using this equation, the transport equation for f_ϕ can be written as

$$\frac{\partial f_\phi}{\partial t} + \frac{\partial}{\partial x_i} [f_\phi (\langle U_i \rangle + \langle u'_i | \psi \rangle)] = -\frac{\partial}{\partial \psi} \left(f_\phi \left\langle \frac{D\phi}{Dt} \middle| \psi \right\rangle \right) \quad (2.78)$$

$$= -\frac{\partial}{\partial \psi} \{ f_\phi [\langle \Gamma \nabla^2 \phi | \psi \rangle + S(\psi)] \}, \quad (2.79)$$

which can be rewritten using Eq. 2.73 as

$$\frac{\bar{D}f_\phi}{\bar{D}t} = \Gamma \nabla^2 f_\phi - \frac{\partial}{\partial x_i} (f_\phi \langle u'_i | \psi \rangle) - \frac{\partial^2}{\partial \psi^2} (f_\phi \langle \Gamma \nabla \phi \cdot \nabla \phi | \psi \rangle) - \frac{\partial}{\partial \psi} [f_\phi S(\psi)] \quad (2.80)$$

where $\bar{D}f_\phi/\bar{D}t$ refers to the substantial derivative based on the mean velocity. The fluctuating velocity field conditioned on the scalar appears as an unclosed term in the equation and is usually modeled using a gradient-diffusion hypothesis:

$$-f_\phi \langle \mathbf{u} | \psi \rangle = \Gamma_T \nabla f_\phi \quad (2.81)$$

where Γ_T is the turbulent diffusivity.

The conditional scalar dissipation ($\langle \Gamma \nabla \phi \cdot \nabla \phi | \psi \rangle$) is the other term that needs modeling and has been the focus of research in PDF methods for reactive flows. This term is closed using a mixing model which defines the effects of molecular mixing. Though several models exist, current micromixing models are not accurate in describing the mixing process and are usually taken as a representation of the mixing process. More details about the models will be discussed in later sections. One of the other features to be noticed is that the reaction source term $S(\psi)$ appears closed in Eq. 2.80. This single feature makes the PDF method superior to existing finite-difference/finite-volume methods which assume simple closures that are not valid in the entire computational domain. To complete the derivation, the closure for the conditional turbulent scalar flux and a simple IEM (Interaction by Exchange with the Mean) mixing model [Villermaux, 1986] with constant C_ϕ are used to obtain

$$\frac{\bar{D}f_\phi}{\bar{D}t} = \nabla \Gamma_T \frac{\partial f_\phi}{\partial x_i} + \Gamma_T \nabla^2 f_\phi + \frac{\partial}{\partial \psi} \left[f_\phi \left(\frac{1}{2} C_\phi \frac{\epsilon}{k} (\psi - \langle \phi \rangle) - S(\psi) \right) \right] \quad (2.82)$$

where it is assumed that the turbulent diffusivity is much larger than the molecular diffusivity. However, the molecular diffusivity can be readily included to form the effective viscosity. The equation can be rewritten as

$$\frac{\bar{D}f_\phi}{\bar{D}t} = -\nabla \Gamma_T \frac{\partial f_\phi}{\partial x_i} + \frac{1}{2} \frac{\partial}{\partial x_i \partial x_i} (f_\phi \Gamma_T) + \frac{\partial}{\partial \psi} \left[f_\phi \left(\frac{1}{2} C_\phi \frac{\epsilon}{k} (\psi - \langle \phi \rangle) - S(\psi) \right) \right] \quad (2.83)$$

In a Lagrangian manner, if the joint PDF of position $\mathbf{X}^p(t)$ and composition $\phi^*(t)$ is denoted by $f_{x\phi}^*$, and if the distribution of f_x is considered to be uniform, the above equation can be considered as the Fokker-Planck equation for the evolution of the scalar field. The corresponding particle equations can be written as

$$d\mathbf{X}^p = [\langle \mathbf{U} \rangle + \nabla \Gamma_T] dt + (2\Gamma_T)^{1/2} d\mathbf{W}. \quad (2.84)$$

$$d\phi^* = - \left[\frac{1}{2} C_\phi \frac{\epsilon}{k} (\phi^* - \langle \phi \rangle) - S(\phi^*) \right] dt. \quad (2.85)$$

The scalar evolution equations are closed if the velocity and turbulence fields are known. Hence any Eulerian grid-based flow solver can be used to provide the coefficients of the particle equations. It is implied that the mean fields used as coefficients are evaluated at the particle positions. These two equations are part of the working set of equations. It can be shown that the drift (a_i) and diffusion (b) coefficients are implicitly related by:

$$\nabla \cdot \mathbf{a} = \frac{1}{2} \nabla^2 b^2, \quad (2.86)$$

to ensure that a uniform particle distribution will not change with time [Pope, 2000]. It can be seen that the current coefficients satisfy this equality. In variable density flows, Favre-averaged coefficients are used and the equations are altered to reflect the density changes in the domain:

$$d\mathbf{X}^p = \left[\tilde{U} + \frac{1}{\tilde{\rho}} \nabla \tilde{\rho} \tilde{\Gamma}_T \right] dt + \left(2\tilde{\Gamma}_T \right)^{\frac{1}{2}} d\mathbf{W}. \quad (2.87)$$

$$d\phi^* = - \left[\frac{1}{2} C_\phi \frac{\tilde{\epsilon}}{\tilde{k}} \left(\phi^* - \langle \tilde{\phi} \rangle \right) - S(\phi^*) \right] dt. \quad (2.88)$$

where $\langle \cdot \rangle$ indicates a Favre-averaged quantity. In RANS based models, the turbulent diffusivity is calculated from the local kinetic energy and dissipation fields. In gas-phase flows, the effect of molecular viscosity can be pronounced especially in regions of low turbulence. Hence an effective viscosity is defined by

$$\Gamma_{eff} = \Gamma + \frac{C_\mu \tilde{k}^2}{S_{C_t} \tilde{\epsilon}}. \quad (2.89)$$

The composition PDF method can be readily extended to multiple scalars. It uses an external flow solver to close the velocity and turbulence fields [Haworth and El Tahry, 1991]. The numerical implementation is based on a node-based scheme [Pope, 1985] or a Lagrangian distributed-particle scheme [Tsai and Fox, 1996]. Both these methods have been widely used for solving reacting flow problems [Jones and Khaki, 1998; Raman et al., 2002; Hulek and Lindstedt, 1996; Subramaniam and Haworth, 2000; Mobus et al., 2001]. Transport in physical space is non-trivial for complex domains but has been successfully used in several works [Subramaniam and Haworth, 2000; Raju, 1996; Raman et al., 2003]. Transport in composition space is more involved and several numerical hurdles have to be overcome in modeling mixing and treating detailed chemistry terms. In practice, a fractional time-stepping method is used [Oran and Boris, 2001] that separates the mixing and reaction processes into a set of serial steps. The mixing step is solved using a micromixing model. The reaction step requires that a highly non-linear ordinary differential equation (ODE) set be solved in every time step for every particle in the domain. The next two sections will deal with micromixing models and efficient numerical methods for implementing detailed chemistry.

Micromixing Models

The composition PDF was found to have a closed reaction source term but needed a model for the molecular mixing term. This poses a considerable modeling challenge. The main criteria for such a model are [Tsai and Fox, 1995]:

1. The scalar PDF should relax to a Gaussian form regardless of its initial shape.
2. Model bounded scalar must remain bounded.
3. Model should ensure linearity between inert scalars.
4. Model should be local in composition space. In other words, there should be no random jumps in composition due to mixing.
5. Model should ensure boundedness for joint scalar statistics.
6. Model should ensure appropriate dependence on scalar length scales.
7. Model should display appropriate dependence on the Reynolds, Schmidt and Damkohler numbers.

Below, some of the common mixing models are discussed.

Coalescence-Dispersion Model

Originally proposed by Curl [1963], this model expresses mixing in terms of pairs of fluid elements. The fluid elements mix to form two equal composition elements. The rate of mixing is determined by a mixing frequency that determines the number of pairs in an ensemble that will mix in a given time. The CD model is simple in its description and provides a qualitative description of mixing. However, it fails to predict the Gaussian relaxation of scalar PDF in homogeneous turbulence. In addition, the stochastic process does not represent diffusion but rather a jump process that can lead to mixing that is non-local in composition space. Janicka et al. [1970] and Pope [1982] have suggested improvements that provide a continuous distribution for the PDF through a random mixing frequency as opposed to a deterministic value used by Curl. These models still predict unrealistic higher-order moments.

IEM model

As a first step approximation to incorporating the micromixing process, the IEM (Interaction by Exchange with the Mean) is a suitable closure [Villermaux, 1986]. It connects the local turbulence properties, turbulent kinetic energy (k) and dissipation rate (ϵ) to the micromixing time. The IEM model proposes that the rate at which the composition of a particular particle approaches the mean is linearly proportional to

distance of the particle in composition space from the mean of the cluster - which in this case is the mean in a particular cell. Numerically this can be represented as

$$\frac{d\phi}{dt} = \frac{C_\phi}{\tau_\phi} (\langle\phi\rangle - \phi), \quad (2.90)$$

where ϕ is the composition vector associated with a stochastic particle. The IEM model can be seen as an improvement over the coalescence-dispersion model where two particles at random are paired and the particles in the composition of the particles are set to the mean of the pair. This clearly would lead to random jumps in the composition space and consequently violates the model criteria. However the IEM model can be seen as representative of the micromixing process which traditional FV codes neglect. It will be shown in Chapter 4 that even this simple mixing model provides excellent agreement with experimental reacting flow data.

GIEM model

The Generalized IEM model introduced by Tsai and Fox [1995] tends to alleviate some of the problems of the IEM model by tracking the evolution of a shadow scalar. The individual scalars mix at different rates and the shape of the PDF relaxes to a Gaussian in the absence of chemical reactions (as it should for a homogeneous non-reacting passive scalar). The rate of mixing in each direction of the composition space is linearly related to the rate of mixing in the shadow scalar space. The GIEM model has been extended to multiple scalars [Tsai and Fox, 1998] as well.

Binomial Langevin Model

This model proposed by Valino and Dopazo [1991] uses a stochastic model for mixing in scalar space. The random walk has a drift term similar to the IEM model and a diffusion term that is supposed to model turbulent dispersion. The scalar increments can then be written as:

$$d\phi^* = A_\phi (\phi^* - \langle\phi\rangle) dt + B_\phi^{1/2} dW^b \quad (2.91)$$

where A_ϕ and B_ϕ are the coefficients for the drift and diffusion terms. These coefficients are specified such that the mean term reduces to the IEM model and the diffusion term approaches zero quadratically at scalar boundaries. dW^b indicates binomial Wiener increment instead of the usual Gaussian increment. The definition of the model requires

that a parameter that determines the amount of turbulent dispersion be specified. This is usually tuned to match homogeneous turbulence data but is by no means a universal value. Other modifications for non-equal weight based particle simulations are required to enforce the correct decay of the variance [Wouters et al., 1998; Wouters, 1998].

EMST model

The Euclidean Minimum Spanning Tree (EMST) model proposed by Subramaniam and Pope [1999] tries to address the local mixing criteria for mixing models. It is argued that IEM model predicts extinction events incorrectly in the limit of fast chemistry due to the lack of the localized mixing behavior. In the EMST model, the evolution of the scalars is governed by compositions in the scalar neighborhood. The EMST model uses a tree structure to form a composition sub-space in which each particle is linked to at least one neighbor particle. The mixing effect of each particle on the other is determined by a set of coefficients which are fixed in such a way as to obtain the desired decay rate for the variance. The EMST model has been tested for reacting flows involving flame chemistry and has been used to predict flames with high extinction probability [Xu and Pope, 2000; Masri et al., 1996]. The EMST model is an extension of the mapping closure [Pope, 1991] to multiple scalars [Xu and Pope, 2000].

Other Mixing Models

Several other models like the mapping closure [Valino, 1995], the Multi-scale IEM [Heinz and Roekaerts, 2001] have been proposed that in general are valid for certain test flows. Their wide applicability has not been tested so far. The Lagrangian Spectral Relaxation model [Fox, 1995, 1999] tends to address most of the issues associated with the mixing model criteria but its use in practical CFD simulations is yet to be demonstrated.

For this study, the IEM model was found to be satisfactory as a representation of the unmixedness. Since no current model can predict mixing accurately, the IEM model was used due to the simplicity in implementation. One of the contentious issues in using the IEM model is the mechanical-to-scalar time-scale ratio that appears as the coefficient in the model. Usually, this value is set at 2.0 indicating that the velocity and scalar length scales are identical. It has been experimentally verified that the ratio is close to 1.5 for free jets [Panchapakesan and Lumley, 1993]. Theoretical derivations [Peters, 2000] show that the ratio is not a constant in reacting flows and the dependence of reaction for arbitrary mechanisms is still unanswered. It should be noted that the choice of the

micromixing time scale is central to most mixing models. This ends the discussion on mixing models.

Reaction Computation

The reaction process is reduced to a set of ODE after the time-splitting is applied to the transport in composition space. This can simply be written as:

$$\frac{d\phi}{dt} = \mathbf{S}(\phi). \quad (2.92)$$

In spite of the apparent simplicity of the equation, numerical integration of a stiff set of differential equations is computationally challenging [Oran and Boris, 2001]. For a PDF simulation involving particle numbers of the order of a million that need to be integrated over several thousand time steps, such integrations can make the simulation intractable. In finite-volume or finite-difference grid based methods, several simplifications can be employed. In many cases, the reaction set is reduced and the number of differential equations to handled are consequently lower. In fast chemistry limits, use of equilibrium chemistry is common [Bilger, 1980]. In this method, a conserved scalar, namely mixture fraction is tracked in the computational domain. It is assumed that the reactions have reached equilibrium and by knowing the fraction of component from the two streams, the mean composition of all the scalars can be obtained. Since RANS and LES methods do not resolve all the length scales associated with the conserved scalar, the mean and variance of the mixture fraction are determined at each location in the domain. By assuming an underlying distribution function for the mixture fraction using its mean and variance, the composition can be calculated. Since binary mixing can be closely approximated by a beta-PDF, almost all known applications use this function form for the PDF of the mixture fraction. The equilibrium compositions are calculated *a priori* and stored in a *look-up* table.

Finite-rate chemistry on the other hand is difficult to handle because it cannot be described by a conserved scalar alone. The reaction path lies on a multi-dimensional space and needs to be tracked in all directions as opposed to the equilibrium assumption where the pathway lies on a two dimensional surface. Hence finite-rate chemistry calculations require that the transport equations are solved for all the species and temperature - together termed as scalars. Also, the reaction itself will not proceed to equilibrium in each time step. The look-up table then can be extended to include all the species which can be represented by a multi-dimensional array. Similar to the equilibrium chemistry method, calculations can be done to store the results of integration for a given initial

composition, time step, and reaction conditions. However for a set of initial conditions, depending on the time step, a finite extent of reaction is possible. Then the integration itself is simplified to linear interpolation where, given the look-up table and the initial condition, the results of integration can be found. It can be seen that for a N -species chemistry, there should be N dimensions in the table. In addition each dimension needs to be discretized with M points. In general, M can vary for each dimension. The total storage memory required is of the order of N^M which can very easily become intractable for mechanisms with 8-10 species. Attempts to simplify the table, like the Intrinsic Low-Dimensional Manifold (ILDm) methods [Maas and Pope, 1992] are useful but are chemistry dependent.

However, the look-up table approach is an attractive feature, especially in PDF methods, where nearly identical integrations have to be repeated for several thousand particles. Pope [1997] developed the *in-situ* adaptive tabulation (ISAT) method to overcome the computational issues with the look-up table approach. One of the main defects is that the look-up table stores points in the entire composition space while a particular computation may use only a portion of the table. Hence a distinction was made between the *realizable* and *accessed* space. Pope [1997] concluded that the accessed space in a particular simulation was much smaller compared to the realizable space and hence a table that is built during the simulation will be more efficient in conserving storage space. The ISAT algorithm uses a binary tree data structure which store nodes correspond to the integration points. An ellipsoid of accuracy is defined around the node within which linear interpolation is carried out. A mapping matrix that contains the Jacobian evaluated at the node point is also stored. The algorithm uses an efficient search routine in the binary tree data structure (Fig. 2.1). Given any initial composition as a query point, a search is carried out over the entire set of stored points until the region in composition space corresponding to the query point is reached. If the query point lies within the ellipsoid corresponding to a previously stored node, linear interpolation is carried out based on the mapping matrix stored at that node :

$$\phi_{in} \xrightarrow{\mathbf{R}} \phi_{out}. \quad (2.93)$$

The interpolation is first order accurate and the final composition is determined by:

$$\phi_{out} = \phi_{node} + \mathbf{J}(\phi_{node} - \phi_{in}). \quad (2.94)$$

If the query point is not within the ellipsoid of any other node, then two possibilities exist. It is possible that the new point can be included within the domain of another node

by “growing” the ellipsoid corresponding to that node. This is the growth algorithm that potentially reduces the number of storage nodes in the tree by orders of magnitude [Fox, 2003]. If a growth is not possible, then a new node is created with the query point. Integration is carried out to find the composition after reaction and the mapping matrix is computed from the Jacobian. The interpolation is first-order accurate and is able to reproduce direct integration results accurately [Yang and Pope, 1998a]. Use of multiple binary trees reduce the storage requirement and increase the speed of retrieval. The ISAT algorithm has been used in several PDF calculations and has been found to yield speed-ups of over 500 compared to direct integration [Shah and Fox, 1999; Kolhapure and Fox, 1999; Xu and Pope, 2000; Raman et al., 2001]. The ISAT-CK version used in this work also has a optically thin radiation model incorporated into the algorithm to handle flame calculations [Tang et al., 2000]. Later chapters contain numerical results of the speed-up for particular chemical mechanisms.

Hybrid Method

The velocity-composition-frequency method can be applied to a reacting flow system to solve for all the variables in the multi-dimensional solution space. As noted earlier, the PDF scheme requires derivatives of the velocity field for the solution of the Poisson equation for pressure. This can be computationally unstable and hence an alternative approach stemming from the external flow closure applied by Haworth and El Tahry [1991] is considered. Figure 2.2 shows the flowchart of the method [Jenny et al., 2001]. Typically an Eulerian flow solver like a finite-difference or finite-volume scheme is implemented to handle the momentum, continuity, and in the case of a RANS method, the turbulence equations. The PDF method is then formulated as a joint composition scheme and the closure coefficients in the particle transport equations are obtained from the mean flow fields. For a variable density reacting flow case, the PDF method feeds back the temperature and mean molecular weight to the flow solver. Using the pressure field, the new density corresponding to reaction is calculated and the mean flow adjusted accordingly. This iterative process is repeated until convergence is reached. The algorithm can be tightly coupled in which case the feedback is initiated after every PDF time step or it can be loosely coupled where the PDF scheme iterates over a fixed number of time steps before feeding back the mean profiles. The latter approach is computationally less expensive for steady-state problems. If time-accurate solutions are needed, the first approach should be used.

Summary

It can be seen that the hybrid scheme provides a viable and stable alternative to the joint velocity-composition-frequency PDF. It is a marked improvement over traditional grid-based techniques that use *ad hoc* closures. However, the computational expense of the hybrid method is much higher than the traditional Eulerian methods and efficient implementation of the particle algorithms are key to achieve a tractable scheme. The next chapter will discuss the numerical implementation of the methods discussed with the central idea of achieving grid independent efficiency.

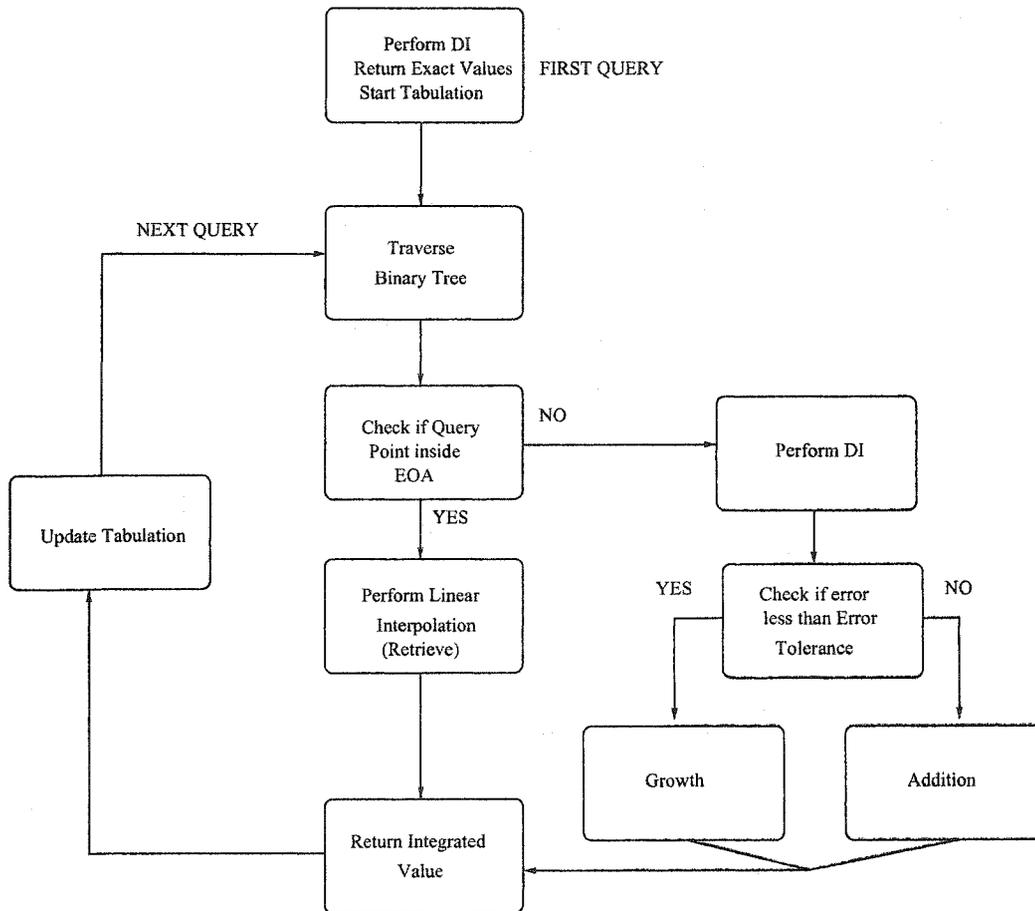


Figure 2.1 ISAT algorithm.

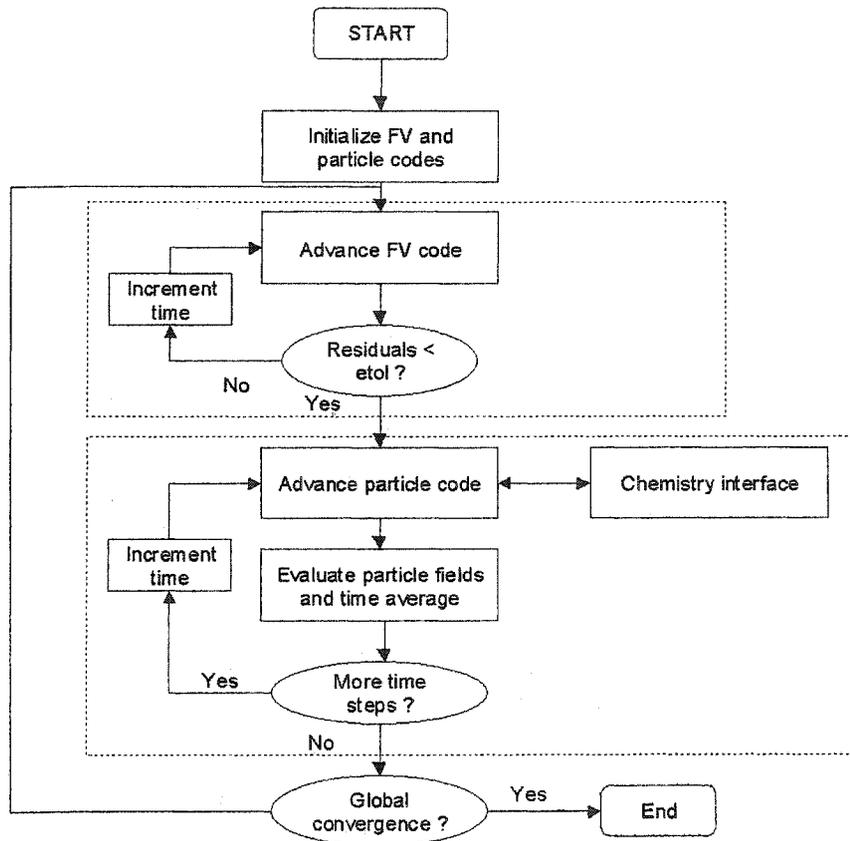


Figure 2.2 Flowchart of hybrid Eulerian flow-solver/ Transported PDF method. [Jenny et al., 2001].

3 Numerical Methods

Introduction

Numerical algorithms and their efficient implementation are important in making a model practically useful. Reacting flows are often modeled using equilibrium assumptions or simplified reaction mechanisms. In the recent past, flamelet models and conditional moment closure (CMC) methods with detailed chemistry have been used extensively in combustion simulations [Kim and Huh, 2002; Roomina and Bilger, 2001; Pitsch and Steiner, 2000]. In spite of their popularity, the underlying assumptions restrict their applications. For example, use of flamelet models for recirculating flow that involves partially premixed reaction zones violates the basic assumption of reaction in thin layers. In this regard, the PDF transport equation contains closed reaction source term and can be handled directly. Use of novel algorithms for handling large kinetic mechanisms (like ISAT [Pope, 1997]) have made full-PDF simulations tractable. The need for a theoretically sound micromixing model has often been cited as the main drawback of the PDF method [Roomina and Bilger, 2001; Peters, 2000]. However, use of existing micromixing models have yielded good agreement with experimental data [Xu and Pope, 2000; Muradoglu et al., 1999]. Even in the limit of premixed reactions, use of currently available models can provide a qualitative insight into the dynamics of the system. The main drawback of the method lies elsewhere.

To be useful as a computational tool in the industry, PDF methods should be able to operate on arbitrary geometries and irregular grids. Traditional PDF methods have been formulated on Cartesian orthogonal grids and extensions to complex grids has not been pursued actively. Use of a particle-based scheme requires an efficient management of particles [Raman et al., 2002]. Non-linear sorting algorithms for example, can quickly make the scheme intractable. *An ideal PDF implementation should be able to operate on any domain of arbitrary grid structure.* In other words, none of the features of the grid itself can be used for optimization since this will limit the generality of the code.

This chapter outlines the details of such a scheme and explains the individual algorithms required to achieve this generality.

Eulerian NS Solver

For the hybrid scheme using a joint-composition PDF, a flow solver is required to supply the velocity and turbulence fields. Since reactions can induce density gradients, the only requirement is that a variable-density formulation of the flow solver be used. In this work, a compressible formulation cast in generalized coordinates is utilized for the Lagrangian particle solver. The CHEM3D code [Harvey, 1998] uses a multi-block format to handle structured non-orthogonal grids. The finite-volume code can also handle multi-species transport equation with reaction source terms but contain no closure for these terms [Harvey and Edwards, 1999]. It is designed to use flamelet models [Williams, 1975] which have a feedback structure similar to the PDF code [Harvey and Pitsch, 2000]. Particle schemes that can work on such a generalized flow structure can be assumed to be easily portable to any other flow solver with minimal changes.

In the feedback mechanism employed here, the Eulerian solver is iterated first to obtain a steady-state flow field with cold-flow conditions. This flow field is used to start the PDF iterations, initially without feedback. Once a steady state is obtained, the feedback loop is initiated. This involves transferring the mean temperature, molecular weight and density to the flow solver. Here, the static pressure changes are assumed to be negligible and hence the feedback of any two of the above three fields will close the ideal gas equation. The instantaneous mean fields contain statistical fluctuations and hence are time-averaged over a small number of PDF iterations. The number of PDF time steps and the time-averaging steps are problems dependent and needs to be tested and to a large extent depends on the reaction associated density changes in a given time step.

PDF Scheme

As explained in Chapter 2, the high dimensionality of the PDF transport equation requires that an alternative procedure with reduced computational time be used. Particle-based Monte-Carlo algorithms are found to scale linearly with number of dimensions and provide a tractable scheme for solving the PDF transport equation [Pope, 1985]. The particle-based simulation can be described as a Lagrangian or Eulerian

system. The Lagrangian scheme needs efficient management of the particles while the Eulerian system contains a simplified system with fixed particle position. Use of Eulerian PDF scheme where particles are located at grid-nodes with a deterministic position have been considered [Roekaerts, 1991; Pipino and Fox, 1994]. The simplified particle tracking aids in using complex geometry. In this algorithm, only the face fluxes of a cell are required and are directly computed by any finite-volume code. However it has been shown that the use of the Eulerian scheme leads to large numerical diffusion and can seriously degrade the simulation results [Mobus et al., 2001]. For this reason, the Eulerian PDF scheme will not be considered here. Mobus et al. [2001] have also shown that the use of the Lagrangian PDF scheme can be implemented in such a way that the numerical diffusion can be limited. Similar techniques have been analyzed elsewhere [Jenny et al., 2001; Muradoglu et al., 1999].

The Lagrangian particle scheme consists of several sub-steps to handle transport in either physical space or composition space. Figure 3.1 details the sub-steps for each time iteration in a PDF algorithm. In the context of hybrid schemes (Fig. 2.2), this flowchart refers to the particle side of the method. Data interpolation and particle tracking are key elements of the Lagrangian PDF scheme. The only consistency condition imposed here is the equivalence of the particle mean density and the Eulerian mean density fields. Since the Eulerian flow solver obtains the density fields from the particle field, consistency is naturally imposed. The velocity fields are then adjusted to satisfy the continuity equation.

At the start of the iteration, the notional particles are distributed uniformly in the computational domain. If continuity is satisfied, the particles will stay uniformly distributed [Subramaniam and Haworth, 2000]. The particles are assigned a weight such that the sum of the weights of the particles in a computational cell is proportional to the cell mass. In flows involving axisymmetric domain or high velocity gradients, cell masses can vary several orders of magnitude. In the Lagrangian method, due to the random movement of particles, particle-based cell weights can be disrupted instantaneously, though the average cell mass should equal the cell volume times density of that cell. Clustering routines that reduce the fluctuations in the particle weights are implemented [Haworth and El Tahry, 1991]. Once the mean fields are obtained from the flow solver, the data is interpolated onto the particles using a bilinear interpolation scheme. This marks the beginning of the time iteration for the PDF code. It is followed by a half-step physical transport. For computational expense and stability considerations, reaction calculations are applied for the full time step [Oran and Boris, 2001].

The physical transport is carried out for another half time step to complete the iteration. Data output is carried out at a pre-determined frequency. At the end of the time step, depending on the frequency of feedback, particle mean fields of temperature and mixture molecular weight are fed back to the flow solver. The feedback loop itself is iterated many times until convergence is attained. These are essential features of the algorithm and depending on the nature of the PDF scheme, the computational expense of each sub-step will vary. However, in the presence of a complex chemistry mechanism, the ISAT subroutine takes-up nearly 95% of the computational time [Raman et al., 2003]. The next section will explain in detail the sub-steps of the algorithm for the Lagrangian scheme. The algorithms are explained for a 2-dimensional grid but can be readily extended to 3-dimensional geometries.

Lagrangian Particle Scheme

In this version of the PDF method, the notional particles are free to move in the domain and are governed by the particle transport equation (Eq. 2.84). The algorithm is based on object-oriented methodology with a hierarchical data structure. The grid is defined as the super-structure that contains all the details of the simulation. The grid is composed of blocks which correspond to the multi-block structure of the flow solver. Each block contains cells, faces and nodes which are at the same level of the hierarchy. All faces contain information about the nodes they are composed of and all cells contain information about the faces and the nodes that are part of the cell. In addition each face stores the cell numbers that it links. Particles are owned by cells and ownership is transferred when particles cross a face. Due to the finite-volume nature of the code, ghost cells are defined on the outer boundaries of the original grid. In case the ghost cell defines a block boundary, faces of the cell are duplicated and owned by the adjacent block. If the block boundary defines a physical boundary like an inlet or a wall, the faces of the ghost cell are left undefined. Particles contain all the information about the scalar fields through the individual composition vectors. When particles exit the physical domain they are discarded and the particle number is left vacant. When new particles come into the domain through an inlet or due to the clustering operations, the vacant array positions are utilized. No sorting of the particles arrays are done and all operations are carried out based on integer tags given to the individual topological groups (like faces or cells). The following sub-sections will describe the steps involved in the PDF scheme.

Particle Initialization

The hybrid simulation is started from a cold-flow solution where the flow solver is iterated to obtain a divergence-free (relative to a error-tolerance) flow field. The particles are uniformly distributed in the domain based on a nominal particle number density defined as the number of particles per cell. Each particle is assigned a number or a tag which indicates their position in the particle array. Also each particle carries the following information:

1. Location coordinates x_1, x_2, x_3 . The location information is updated each time step by the particle tracking algorithm. Particle positions are initialized using a mapping routine similar to the particle interpolation algorithm discussed next. Each cell is transformed into a unit square using a bilinear mapping function. Then a random position inside the unit square is chosen and the coordinates in the computational domain are obtained using a inverse mapping function.
2. Weight w_p . The sum of the particle weights in a given cell corresponds to the cell mass. The stochastic nature of the algorithm will cause fluctuations in the number density and control algorithms are employed to decrease the fluctuations and thereby minimize errors in mean field computations due to low sample numbers.
3. Velocity U_p, V_p, W_p . The velocity fields are interpolated from the data provided by the flow solver using the bilinear mapping technique discussed next.
4. Composition array ϕ_1, \dots, ϕ_N . Each particle carries a vector of composition and temperature. For non-premixed flows, the vector is extended to include mixture fraction as well. In general, no distinction is made between temperature and scalars, and they are transported using the turbulent Schmidt number. This is equivalent to assuming identical turbulent Prandtl and Schmidt numbers in a Eulerian grid-based method.
5. Cell number C_p . The cell number is required to associate a particle to the cell. The object-oriented structure makes this information redundant but periodic checking is carried out to ensure that there is a one-to-one mapping between particles and owner cells.

For reacting flows, it may be essential to artificially ignite the flow. In such cases, either the entire domain or an appropriate block is set to a pre-determined ignition

temperature by forcing the particle temperature. It was also found that for the methane-air flame discussed in Chapter 5, artificial ignition was not necessary due to the presence of a pilot flame. However, the chlorination scheme requires a ignition temperature of 800 K to initiate sustainable reactions.

Mean Field Interpolation

Particle transport equations in both physical and compositional spaces require the mean fields from the flow solver interpolated onto the particle position. Based on the transport equation in physical space (Eq. 2.84), a random velocity can be defined as

$$\mathbf{U}_p = \left[\tilde{\mathbf{U}} + \frac{1}{\bar{\rho}} \nabla \bar{\rho} \Gamma_{eff} \right] + \sqrt{2\Gamma_{eff}} \frac{d\mathbf{W}}{dt}. \quad (3.1)$$

It should be noted that this is not a mathematical identity that can be derived from the equation for the differential motion of the the particle. This definition simplifies the particle tracking procedure explained next. Strictly speaking, this definition of \mathbf{U}_p is invalid for infinitesimal dt since the derivative of the Weiner diffusion process does not exist. For the case of finite time step, the Wiener diffusion term is replaced by a finite-difference approximation of the derivative :

$$\frac{d\mathbf{W}}{dt} = \frac{d\mathbf{W}}{\Delta t} \quad (3.2)$$

where the $d\mathbf{W}$ term is sampled from a normal distribution with variance equal to the time step Δt .

Interpolation on regular grids can be carried out using “global” techniques like the spline method. But on irregular grids with arbitrarily oriented blocks, these methods are either not defined or cumbersome to use. Here a cell-by-cell interpolation is carried out. As indicated in Fig. 3.2 each cell contains information about the neighborhood defined as the 9 cells (in 2-dimension) that surround the current cell (including the cell itself). For the sake of simplicity, the cells are denoted by rectangles but they can be of any shape. The inner boxes joining the cell centers encompasses the entire query cell (i, j). Hence any particle in this cell can be mapped onto one of inner boxes. Since the finite-volume scheme solves for the cell-centered variables, the mean fields can be directly obtained at the cell centers of all the cells in the domain. To avoid redundant searching and interpolation, an integrated algorithm is used (Fig. 3.3).

Given the vertices of a cell, it is possible to map the cell onto a unit square (in 2-d) or a unit cube (in 3-d). A bilinear function is defined based on the vertices as:

$$l_1 = x_2 - x_1; l_2 = x_4 - x_1; l_3 = x_3 + x_1 - x_2 - x_4, \quad (3.3)$$

and

$$m_1 = y_2 - y_1; m_2 = y_4 - y_1; m_3 = y_3 + y_1 - y_2 - y_4, \quad (3.4)$$

where (x_i, y_i) are the vertices in cyclic order. The bilinear function used here defines only the inverse mapping explicitly on the coordinates. The transformed coordinates (a, b) can then be related to the actual coordinates (x, y) as

$$\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} x_1 + al_1 + bl_2 + abl_3 \\ y_1 + am_1 + bm_2 + abm_3 \end{bmatrix}. \quad (3.5)$$

Given any (x, y) it is then possible to map the point to the unit square using an inverse relation. Since the inverse relation contains implicit variables, an iterative procedure is used to obtain the mapped coordinates. Similar definitions of mean field variables can be derived to define the interpolation based on the properties at the nodes. For any scalar field γ , with values γ_i at vertex i , the following definition holds:

$$\gamma_{(x,y)} = \gamma_1 + a(\gamma_2 - \gamma_1) + b(\gamma_4 - \gamma_1) + ab(\gamma_3 + \gamma_1 - \gamma_2 - \gamma_4). \quad (3.6)$$

Since the coefficients in this method (namely, the differences of the scalar fields at the vertices) do not change due to the PDF step, re-computation of the coefficients are carried out only at the beginning of every PDF iteration.

Interpolation is carried out multiple times for the same particle in each time step. Every time step iteration starts with interpolation over all the particles. During particle tracking, mean-field coefficients are adjusted when the particle crosses a face. In axisymmetric flows, a correction algorithm is applied that is similar to tracking and hence may involve additional interpolation to correct the particle positions. Particles entering the domain through an inlet face also need mean-field interpolation. Hence an efficient implementation is necessary to reduce computational expense. In simulations with nearly orthogonal grids, cell-centered values were first interpolated using a first-order scheme onto the cell vertices. The particle interpolation algorithm was used in a similar manner using the vertices of the cell itself as the corner points. This reduced the computational load in retrieving from memory cell-centered values of the entire neighborhood. Comparisons with the original algorithm showed no significant difference in the results.

The mapping routine is also used to randomly distribute particles. If the cells are composed of squares, it can be seen that a random choice of (a, b) will uniformly distribute the particles in the square. However, for skewed cells, a uniform distribution in unit square will not correspond to a uniform distribution in the original coordinates.

Hence a weighting function based on the Jacobian of the transformation is used [Papoulis, 1991]

$$J_p = \begin{vmatrix} \frac{\partial x}{\partial a} & \frac{\partial x}{\partial b} \\ \frac{\partial y}{\partial a} & \frac{\partial y}{\partial b} \end{vmatrix} \quad (3.7)$$

Particle Tracking

Each time step consists of two half-steps in physical space where particles are moved according to Eq. 2.84. As explained earlier, each cell should contain information about the particles inside its domain in order to evaluate the particle mean fields. During the transport in physical space, particles can reach boundaries, and special conditions are applied to ensure that the Eulerian boundary conditions are satisfied. Hence particle tracking involves the combined task of moving particles in physical space, applying boundary conditions and updating the particle pointers of a cell.

For non-orthogonal irregular grids, use of simple mapping routines to determine particle location will fail. Hence the element-to-element strategy proposed by Subramaniam and Haworth [2000] is implemented here. In this algorithm, the particle location at the start of the time step is identified by its coordinates (x_p, y_p) . For a given cell, the face normals, n_i corresponding to each of the cell faces is obtained (See Fig. 3.4). Using simple geometrical technique it can be shown that for every normal, n_i with directional coefficients (l, m) , the equation of a perpendicular line passing through vertices (x_1, y_1) and (x_2, y_2) is given by

$$lx + my + c = 0, \quad (3.8)$$

where the coefficient c can be determined from the coordinates of the vertices. This definition will provide a consistent way of defining the equation of the line and normal. It should be noted that the face normals may themselves be oriented in different directions and is not required to be outward facing. The coefficients are normalized in such a way that the equation represents a unit normal. Given the equation of the line, the distance from a point to the line, s_i , can be calculated by substituting the coordinates of the point in the equation for the line.

$$s_i = lx_p + my_p + c. \quad (3.9)$$

If the velocity of the particle is known, the component in the direction normal to the

face is given by

$$v_i = IU_p + mV_p, \quad (3.10)$$

where (U_p, V_p) is the particle velocity. If the particle is moving towards the face, v_i and s_i will be of opposite signs, otherwise it can be concluded that the particle is moving away from the face. If the particle is indeed moving towards the face, then the time it would take to reach the face is given by:

$$t_i = -\frac{s_i}{v_i}. \quad (3.11)$$

This process is repeated for the all the faces and minimal time step is found based on the global time-step, the local time-step and the face-minimal time-step:

$$\Delta t_p = \min_{i=1}^4 (\Delta t, t_i). \quad (3.12)$$

The global time step, Δt , is fixed at the start of the simulation based on the smallest mixing time, $\tilde{\epsilon}/\tilde{k}$ in the entire domain.

If the global is the smallest time-step determined by Eq. 3.12, the particle remains in the cell. The final position of the particle is given directly by

$$\mathbf{x}_f = \mathbf{x}_p + \mathbf{U}_p \Delta t_p. \quad (3.13)$$

The owner cell retains the particle and the control moves to the next particle in the array. On the other hand, if the face-minimal time-step is the smallest time scale, then the particle will cross a face and enter a new cell. In this case, the particle is moved to the face corresponding to the minimal-time and the global time step reduced by the time-step used.

$$\Delta t \Rightarrow \Delta t - \Delta t_p. \quad (3.14)$$

Taking into account the fact that the particle now lies on the face of the adjoining cell, control is transferred to this new cell. The particle velocity \mathbf{U}_p is recomputed using the current coordinates. A Gaussian random number is synthesized to obtain the Wiener diffusion term. The above procedure is repeated to find the minimal time based on all the other faces of this cell. Since the global time step reduces every time the particle moves in the physical domain, a final state corresponding to zero global time step is reached at which point the particle controls are transferred to the final cell and the iteration for the next particle begins.

It is found that the CFL criteria normally used with Eulerian flow solvers is restrictive for clustered grids where the local time step can be very small. Since the Eulerian solvers

use this local step for all the cells, the number of iterations required to reach steady state can be very high. In the present scheme, this problem is averted by arbitrarily choosing a global time step and modifying the time step of the particle based on the local conditions. Sharp velocity gradients or extensive grid-clustering are local properties and may not extend over the entire domain. Such a cell-based scheme ensures that a truly grid-independent time step can be chosen for the simulation. It should be noted that though steady-state adaptive time-stepping methods [Muradoglu and Pope, 2002] have been formulated, a time-accurate scheme has not been formulated for realistic flows.

Boundary Conditions

Particle equivalent Eulerian boundary conditions are simple to apply with this tracking scheme. In general, the boundary conditions can be classified into three categories - inlet, outlet and reflecting boundary conditions. These conditions are handled as follows:

1. Inlet boundary conditions directly correspond to the Eulerian inlet conditions, and are treated as a separate fractional step in the algorithm. However, the scheme is only a slight modification of the particle tracking scheme. For every cell having an inlet face, the mass flux is computed based on the cell centered velocity, density, and the area of the face. Then the number of particles to be injected is determined based on the mean particle weight in the cell. The particles are evenly distributed within the inlet cell. It can be seen that the maximum weight that can be input is the weight of the cell itself. Hence the time step will be limited by the cell volume:

$$\Delta t_{inp} = \frac{Vol_c}{U.A}, \quad (3.15)$$

where Δt_{inp} is the inlet time step, Vol_c is the cell volume and $U.A$ is the flux normal to the inlet face. In effect this is a CFL type of criteria that restricts the time step. If diffusion flame configurations are being simulated, heavy grid clustering is employed near the inlet to capture the large gradients in the velocity. Hence the global time step will be restricted by the inlet even though an adaptive tracking algorithm has been implemented elsewhere. To accommodate arbitrary time steps, a modification to the particle injection is necessary. This is accomplished by randomizing the inlet time step. The modified algorithm first places the particles at arbitrary positions on the inlet face. The number of particles is chosen based on a lower threshold, the inlet cell weight and an upper limit. The time step of

injection is randomized based on the global time step to yield:

$$\Delta t_{inp} = 2r\Delta t \quad (3.16)$$

where r is a uniform random number between 0 and 1. Once the time step is determined, the particle is moved using the tracking algorithm using the inlet time step as the global time step. This will ensure that the additional weight that may be input is distributed beyond the first cell and transported into the domain.

2. Outlet boundary condition corresponds to the outlet condition in the Eulerian sense. Particles that leave the domain are removed from the tag list of the owner cell and the corresponding particle array is marked vacant. When new particles are added in the inlet or created during clustering operations (discussed next), these vacant arrays are used.
3. Reflecting boundary conditions are used for treating wall boundaries and axes. Reflection of the particle ensures that the flux across the boundary is zero. Particle velocities are reflected with respect to the boundary. If the normal to the boundary is defined as n , then the reflected velocity \mathbf{U}_r of a particle starting with velocity \mathbf{U}_p is given by

$$\mathbf{U}_r = \mathbf{U}_p - 2(\mathbf{U}_p \cdot n)n. \quad (3.17)$$

4. All equivalent Eulerian boundary conditions can be treated using these three particle boundary conditions. For example, a freestream boundary can be treated as an inlet with particle injection at the freestream velocity. Wall boundaries are treated as adiabatic and no attempt has been made to include isothermal or heat-loss criteria.

Axisymmetric Correction

The particle transport equations in physical space are cast in Cartesian coordinate frame corresponding to (z, y, x) . However, for an axisymmetric domain, the NS equations are solved using a polar coordinate frame corresponding to (z, r, θ) . Hence the random walk needs to be corrected. Figure 3.5 represents the coordinate frames. To apply the correction, the w component of the random walk corresponding is modeled as

$$w = \sqrt{2\Gamma_{eff}dt}. \quad (3.18)$$

Since the Cartesian y component is the radial projection of the random walk in polar coordinates, the following correction in the y coordinate of the particles is applied:

$$y^* = \sqrt{y^2 + w^2}. \quad (3.19)$$

where y^* is the corrected y position. It can be noted that the effect of this correction will decrease with distance from the origin. Particle tracking can be applied by using an artificial velocity to move the particle to correct the radial position. In all simulations, an arbitrary velocity of $U = 1 \text{ m/s}$ normal to the axis is used to move particles for the time duration given by:

$$\Delta t_c = \frac{y^* - y}{U}. \quad (3.20)$$

The corrected position is always further away from the axis than the original particle position. The particles are thus kept away from reaching the axis. This algorithm was found to retain the uniform particle distribution. The correction algorithm is applied after the normal particle tracking has been completed.

Particle Clustering

Practical flows often involve steep velocity and scalar gradients that need to be captured using a clustered grid structure. In some cases, axisymmetric domains can be used to reduce computational cost. In either of these cases, the cell volumes can vary over a large range in the domain. For example, in axisymmetric grids, cell volumes are dependent on the distance of the cell from the axis. Hence cells near the axis will have very small volumes and those farther away will be heavier. Reacting flows like flame combustion can cause steep density gradients that can amplify further the differences in the cell masses. Since the Lagrangian particles represent mass density weighted PDF, in the stochastic simulation, the sum of the particle weights in any cell is proportional to the cell mass. Hence particles in different sections of the domain can have widely varying weights. When these particles move in the domain following prescribed local velocities, cell-particle weight mismatch can occur. This refers to the intrusion of a “heavy” particle in a cell where the mean particle weight is much smaller than the incoming particle. Similarly, a host of “light” particles can enter a cell with a much higher mean particle weight. These two situations will lead to unfavorable fluctuations in particle number density that can result in errors in the estimation of the mean fields. Hence a particle clustering-breakup procedure is implemented to better control the particle number density [Haworth and El Tahry, 1991].

The particle-clustering algorithm has two parts - cloning and break-up. After the physical transport including inlet conditions have been handled, the mean particle weight \bar{w}_i in a cell i with M particles is found by:

$$\bar{w}_i = \frac{\sum_{j=1}^M w_j}{N}, \quad (3.21)$$

where w_j is the weight of j^{th} particle in the cell. N is the nominal particle number density that needs to be preserved within statistical accuracy. For every particle in the cell with properties $(\mathbf{x}_j, \phi_j, w_j)$, the following rules are applied [Fox, 2003]:

1. If $w_j > 2\bar{w}_i$ then the particle is split into two identical particles or the single particle can be “cloned”.

$$\begin{pmatrix} \mathbf{x}_j \\ \phi_j \\ w_j \end{pmatrix} \longrightarrow \begin{pmatrix} \mathbf{x}_j \\ \phi_j \\ w_j/2 \end{pmatrix} \text{ and } \begin{pmatrix} \mathbf{x}_j \\ \phi_j \\ w_j/2 \end{pmatrix} \quad (3.22)$$

2. If $w_j < \bar{w}_i/2$, then the particles are lighter than the average particle in the cell. There are several ways to handle this scenario. Usually, the particle is added to another randomly selected particle keeping the mean composition unaltered.

$$\begin{pmatrix} \mathbf{x}_j \\ \phi_j \\ w_j \end{pmatrix} \text{ and } \begin{pmatrix} \mathbf{x}_k^* \\ \phi_k^* \\ w_k^* \end{pmatrix} \longrightarrow \begin{pmatrix} \mathbf{x}_k^* \\ \phi_k^+ \\ (w_j + w_k^*) \end{pmatrix} \quad (3.23)$$

where $\phi_k^+ = \frac{w_j \phi_j + w_k^* \phi_k^*}{w_j + w_k^*}$. In grids with heavy clustering such a “passive” algorithm will not be sufficient to maintain the particle number density. The alternative “aggressive” clustering procedure is used where heavy particles are cloned into a number of particles such that the weight of each of the clone is same as the nominal particle weight. In a similar way, all light particles are identified and sorted. Then the particles are combined with other light particles in the list until the nominal particle weight is reached. This can lead to spatial errors if the final position of the clustered particle is chosen based on the weighted average of the individual particles. This is similar to a mixing in physical space that will pull particles towards the cell center and can distort the spatial uniformity of the particle distribution. In general, the clustered particle position can also be set to

the position of a random particle in the sorted list to minimize the spatial error. The effect of clustering on the correlation between particle weight and position was not tested here but has been considered elsewhere [Wouters, 1998]. In the simulations carried out here, no major difference in mean fields was found when the clustering schemes were changed.

As mentioned earlier, in axisymmetric cases, particle weights are dependent on their distance from the axis. Hence particles near the axis will invariably have a lower weight compared to particles situated further away. Hence use of the weights directly to determine the nominal particle weight can lead to a bias towards the heavier particles. To rectify this bias, the weights are first divided by their radial position (i.e. their distance from the axis) to create importance weights \hat{w}_j [Pope, 1994b]:

$$\hat{w}_j = \frac{w_j}{y_j}, \quad (3.24)$$

where y_j is the particle position relative to the axis. The clustering algorithm is now applied on the importance weights using the procedure describe above. As the cell location moves farther from the centerline the difference between using particle weights and importance weights will be minimal. However, near the axis, even minor changes in the radial coordinate can alter the weights drastically. Particle clustering ends the process in physical space for the half-step.

Mean Field Estimation

The conditional mean of composition conditioned on position can be obtained from the particle composition vectors. To compute the mean at any point \mathbf{x} , all the particles in the domain need to be used along with a weighting function or a kernel that increases the weight in the vicinity of \mathbf{x} [Pope, 2000]. Usually a bandwidth is defined such that the kernel value is zero beyond this distance. In the hybrid simulations, a mesh is defined for the Eulerian flow solver and the feedback variables need to computed on the same mesh. Several techniques including spline estimates [Jenny et al., 2001] and a three-stage algorithm for structured grids [Subramaniam and Haworth, 2000] have been developed. Fox [2003] details many of these schemes. Here a simple cloud-in-cell estimate [Hockney and Eastwood, 1981] is used. The means can then be defined in a cell as:

$$\langle \phi | x \rangle = \frac{\sum_{i=1}^{N_p} w_i \phi_i^*}{\sum_{i=1}^{N_p} w_i}, \quad (3.25)$$

where $i = N_p$ is the number of particles in the cell each with weight w_i . The mean is a piecewise constant in a given cell. The mean field is also required by the micromixing model and the present definition of the mean ensures that it is conserved by mixing in a cell.

For the hybrid simulation, the cell-estimate at a single time step can be statistically noisy and can lead to numerical instability. In general, the mean is averaged over multiple time steps before being fed to the flow solver. This statistical variance reduction technique has been tested elsewhere [Xu and Pope, 1999; Muradoglu et al., 1999].

Mixing and Reaction

Transport in composition space is through the mixing and reaction sub-steps. Fractional time-stepping is used to form a series of sequential steps that simplifies the handling of these two processes [Pope, 1985; Yang and Pope, 1998a]. The mixing process can be simulated using any of the several known micromixing models [Curl, 1963; Villermaux, 1986; Valino and Dopazo, 1991; Subramaniam and Pope, 1999; Fox, 1997]. Here the Curl's model and IEM model are implemented. For the models except the last one listed above, the scalar dissipation time scale is directly linked to the turbulence time scale using the mechanical-to-scalar time-scale ratio.

$$\frac{\tau}{\tau_\phi} = C_\phi = 1 \quad (3.26)$$

It is noted that in later chapters, this definition will be altered to match the usual definition of the ratio in as twice the current definition. Hence C_ϕ will be set to 2 normally. Irrespective of the definition, the assumption is that the scalar and velocity dissipation fields have the same time scale.

1. In the Curl's model, the N particles in a cell are first paired at random. Then the number of pairs that will mix is decided based on the following relation:

$$N_{pair} = \beta \Delta t \omega N, \quad (3.27)$$

where Δt is the time step and ω is the turbulence frequency. β is a mixing parameter that is set to 2.0 for the Curls' model. Mixing involves modifying the composition in each of these N_{pair} pairs to the mean composition of the pair. So for two particles i and j , mixing will alter the composition as:

$$\phi_i^{t+\Delta t^*} = \phi_j^{t+\Delta t^*} = \frac{\phi_i^t + \phi_j^t}{2}. \quad (3.28)$$

where $\phi^{t+\Delta t^*}$ is the composition vector of the particle after the mixing sub-step. It should be noted that a reaction sub-step follows mixing that will update the composition of the particle to the next time-step. Janicka et al. [1970] and Pope [1982] suggested modifications to the Curls' model in order to produce the right higher order moments of a conserved scalar in homogeneous turbulence. In particular, the modified Curl's model of Janicka *et al.*, can be implemented for particle pair (i, j) as

$$\phi_i^{t+\Delta t^*} = \frac{\alpha\phi_j^t + (1-\alpha)\phi_i^t}{2}, \quad (3.29)$$

$$\phi_j^{t+\Delta t^*} = \frac{\alpha\phi_i^t + (1-\alpha)\phi_j^t}{2}, \quad (3.30)$$

where α is a random number chosen from a uniform distribution. Also the value β is set to 3.0 for this model. Curls' model in its original and modified form has been used in turbulent diffusion flame simulations Lindstedt et al. [2000] and is found to yield reasonable agreement with experimental data.

2. The Interaction by Exchange with the Mean (IEM) model is the simplest and most common model used in PDF simulations. It is a deterministic model that relaxes the particle composition towards the mean at the same linear rate in the directions in composition space. The particle implementation is as follows:

$$\phi^{t+\Delta t^*} = \langle \phi^t \rangle + [\phi^t - \langle \phi^t \rangle] e^{-C_\phi \Delta t \omega}. \quad (3.31)$$

Here $\langle \phi^t \rangle$ refers to the cell-mean at time t . In all the simulations to be discussed in later chapters, only the IEM model has been used. The assumption of constant C_ϕ in reacting flows can be source of error [Peters, 2000] and will be further studied in the next chapter.

The reaction sub-step will then be reduced to a set of ODE. For every particle,

$$\phi^{t+\Delta t} = \int_0^{\Delta t} \mathbf{S}(\phi) dt \quad \text{with} \quad \phi(0) = \phi^{t+\Delta t^*}. \quad (3.32)$$

The reaction source term requires no closure as it is computed directly. In addition, wide separation of reaction time scales can make the above equation very stiff numerically and will require special integrators that can handle such equations. Even with efficient integrators, numerical integration for the large ensemble of particles can be intractable for large chemistry mechanisms. Reduced chemistry techniques [Petzold and Zhu, 1999] or pre-computed look-up tables [Maas and Pope, 1992] can be used. However the generality

of these methods is questionable. The ISAT scheme developed by Pope [1997] provides a more tractable alternative. The only input required is the chemistry mechanism in a prescribed CHEMKIN format [Kee et al., 1989].

The ISAT algorithm was explained in Chapter 2. For a particular mechanism, *a priori* studies can be used to determine the optimum size of the ellipsoid of accuracy. This is in effect to specify the maximum radius of the ellipsoid in any given direction and is parameterized by the error tolerance ε . For the thermo-chlorination reaction to be discussed in Chapter 5, such preliminary tests were carried out using a zero-dimensional reactor. The pair-wise mixing stirred reactor (PMSR) Pope [1997] is a non-ideal CSTR that has the additional parameter of mixing time τ that represents a microscopically non-mixed reactor. A Lagrangian micro-mixing model can then be used to simulate mixing. The pair-wise reactor uses a mixing model similar to the Curl's model that leads to a large *accessed* region in composition space (refer to Chapter 2). Hence this reactor is suitable for tests on ISAT speed-up and effect of error tolerance.

The results of the tests are shown in Tables 3.1, 3.2, and 3.3. In the first test, the error tolerance parameter was varied from 10^{-4} to 10^{-5} and the final steady-state temperature was compared with the output from direct-integration. The number of particles in the reactor was set at 100. It is seen (Table 3.1) that as the error tolerance parameter is increased, the final temperature deviates more from the actual result. However, with higher error tolerance, the computational time is drastically reduced with the highest error tolerance providing a speed-up ratio of over 17. The second column in the table refers to the number of binary trees used to store the look-up table and is set a constant value of 4. The number of records indicate the number of nodes in the look-up table which directly translate into the storage memory used by the table. Again, higher error tolerance value leads to lesser storage requirement. The number of direct-integrations are considerably reduced even with the lowest error tolerance. However, an order of magnitude increase in error tolerance reduces the number of direct integrations only by a factor of 4. It is seen that around an error tolerance of 7×10^{-5} the error in the result is only 0.3 K but at $\varepsilon = 1 \times 10^{-4}$ the error shoots by more than 10 times this value. It is concluded that the use of error tolerance value between 7×10^{-5} and 1×10^{-4} will yield reasonably accurate results.

The second table (Table 3.2) shows the scale-up with particle numbers. The PMSR reactor can be considered as a single cell in a full-PDF simulation. It is interesting to note that when particle numbers are increased, the computational time does not increase linearly. This is an important property of ISAT that arises from repeated interpolations

Table 3.1 ISAT performance with variation in error tolerance parameter.

Error tolerance	Ntrees	Records	DI	Time	Outlet temperature
1×10^{-8} (DI)	4	-	1.7×10^7	35 hrs	705.2 K
1×10^{-5}	4	11994	48512	6 hr 44 min	705.1 K
5×10^{-5}	4	3896	20664	4 hr 12 min	705 K
7×10^{-5}	4	2978	15968	2 hr 14 min	704.9 K
1×10^{-4}	4	2635	12561	1 hr 56 min	701.1 K

Table 3.2 ISAT performance with change in the number of particles in the PMSR. Ntree = 8 and Error tolerance = 7×10^{-5} .

No. of particles	Records	DI	Time
500	3468	15320	3 hr 42 min
1000	4737	22424	8 hr 17 min
2000	7931	35411	9 hr 10 min

that replace integrations. The ISAT method can be thought of as a neural network that trains itself based on the particles that it “sees”. As more and more particles are passed through the algorithm, a larger portion of the composition space is “mapped”, thereby reducing the need for direct integration. Theoretically, once the entire region of the accessed space is mapped, an asymptotic retrieval rate should be observed. However, in all simulations carried out, a finite number of records were being added until the steady state is observed [Pope, 1997]. In a full-pdf simulation that uses particle numbers of the order of 10^6 , a sub-linear scale-up time will lead to very high computational speed-up. In fact, without such a look-up table, full-pdf calculations with detailed chemistry are intractable even with most modern computers available today.

The third set of simulations are aimed at testing the multiple-binary tree structure used by ISAT. The binary tree consists of leaves which contain the mapping matrix and nodes that contain the cutting plane vector defined as the bisector the perpendicular line between two node point values. A decision criteria is formed to determine the closest node to the query point at that level of the tree. The query point will go through several such decisions before reaching the leaf. This can lead to large retrieval times if the tree has too many nodes. Additionally the problem of an “unbalanced” tree need to be considered. This is a common problem with binary tree data structures where one side

Table 3.3 ISAT performance with change in the number of binary trees.

Ntrees	Error tolerance	Records	DI	Time
2	1×10^{-5}	19188	81388	10 hr 14 min
4	1×10^{-5}	11994	48512	6 hr 44 min
6	1×10^{-5}	9957	40171	5 hr 43 min
8	1×10^{-5}	8034	29885	4 hr 36 min
10	1×10^{-5}	7396	27589	4 hr 23 min

of the tree has more nodes than the other side. Ideally, the height of a tree is defined as the farthest leaf from the root or the largest number of nodes to be traversed from the root to the leaf. For a completely balanced tree with n nodes, the height will be $\log_2(n+1)$. In the worst case scenario the height can be n itself. Certain initial conditions and reaction parameters can cause nodes to be added preferentially on one side of the tree. However, it can be seen that if particles from the same region in composition space are mapped to the tree, there is a higher probability of the tree being balanced (though fully balanced trees are not always possible). The multiple-binary tree based ISAT uses this concept by dividing different sections of the composition space among the different trees, thereby maintaining a more balanced and smaller structure with shorter retrieval times. From the tests (Table 3.3 it can be found that as the number of trees are increased there is a significant drop in the number of records stored and the number of direct integrations as well. This arises from the fact that ISAT allows overlapping ellipsoids of accuracy and with well maintained trees, such overlaps are reduced. The number of direct integrations also decreases but the effects are less pronounced as the trees are increased. At around $Ntrees = 8$, the computational time reaches a minimum and further increase in number of trees does not yield appreciable speed-up or reduction in storage space.

After the reaction step, the second half-step of the transport in physical space is carried out to complete the PDF time step. The algorithms implemented here has been individually tested for simple flow configurations and found to yield consistent results. The PDF scheme has been implemented using a modular concept which ensures that individual steps and algorithms can be changed without affecting the entire scheme. The next section uses this algorithm to simulate simple shear flow cases will demonstrate the generality of the method.

Test Case

Simple test cases are performed to verify the accuracy of the PDF scheme. This configuration reduces the errors associated with the flow-field computation. To aid comparison with the FV code, only non-reacting cases with a single mixture fraction are treated here. In addition, use of mixing model is avoided to ensure direct computation of variance profiles using the mean mixture fraction values. It should be noted however, that the use of a mixing model will not affect the mean profile and will only change the variance field. The next chapter will detail all aspects of the mixing model and will be tested against experimental data.

The shear layer geometry is shown in Fig. 3.6. The velocity of the two streams are specified by U_a and U_b . The mixture fraction mf is set to 0 for the lower stream and 1 for the upper stream. The dotted lines indicate a rough plot of the shear layer development. Three different test cases are performed. For the first two cases, a simple grid with two blocks (as defined by the solid rectangles in the Fig. 3.6) each with 53×41 grid points is used. The third case uses a complex grid structure. No feedback loop is necessary due to the absence of reactions. Initially, the finite-volume code is iterated to generate the flow field as well as the finite-volume predictions of the mixture fraction and variance profiles. Then the same flow fields are used with the PDF scheme to obtain the conserved scalar mean and variance profiles from the particle values. The following cases were studied.

Uniform Velocity

Both U_a and U_b were set to a constant velocity of 2 m/s. The upper and lower boundaries were set to free-stream conditions with identical velocities as the inlet. Physical properties of air was used to simulate the scalar shear layer. Figure 3.7 shows the comparison for the mean and variance obtained from the FV and PDF schemes at different downstream locations. Since no mixing model was used, the variance field can be obtained analytically as:

$$\langle \xi'^2 \rangle = \langle \xi \rangle (1 - \langle \xi \rangle). \quad (3.33)$$

It should be noted that the variance profiles presented for the PDF scheme has been time-averaged. This implies that the right hand side function in the above equation needs to be time-averaged so that an exact comparison can be made. In the case of the FV code, a steady-state solution is obtained and the analytical expression can be used for with the steady-state mixture fraction value. A variance transport equation with no

dissipation term is solved using the FV code and compared with the PDF results. The profiles of the mean indicate that the PDF code and the FV code predictions are in excellent agreement thereby indicating that the particle tracking scheme is functioning properly for convection dominated flows. The turbulent diffusivity is of the same order as molecular diffusivity in this case and the random velocity component of the particles are at least two orders of the magnitude smaller than the mean convection component. The variance profiles show that the PDF scheme is able to predict the analytical result with excellent accuracy while the FV method seems to contain excessive numerical diffusion in the variance transport equation. This feature has been noted by Subramaniam and Haworth [2000] where it is attributed again to the numerical dissipation associated with the implementation of the variance equation. However, it is seen that the FV predictions improve in accuracy farther downstream.

Non-Uniform Velocity

In the second test case, the grid used before is retained. Only the velocities of the streams are changed to create a velocity-shear layer with high production of the turbulent kinetic energy and consequently, turbulent diffusivity. The inlet velocities are set at 20 m/s and 10 m/s for the two streams. The mixture fraction inlet conditions are identical to the previous case. The mean and variance of the mixture fraction from the PDF and FV schemes are shown in Fig. 3.8. It can be seen that the presence of the velocity gradient increases the width of the shear layer. The mean profiles from both the schemes are still in good agreement indicating that the random velocity component of the particle model is functioning properly. The turbulent diffusivity based velocity is of the same order as the mean convective velocity in this case. It can be seen that the variance profiles from the PDF code are in very good agreement with the analytical values and that the FV results still show substantial numerical dissipation of the variance.

From these two tests, it is concluded that the particle tracking scheme works efficiently on orthogonal grids in the absence of a velocity gradient and performs equally well in velocity-shear layers. The variance profiles are in good agreement with analytical values. It is also noted that the FV scheme underestimates the variance in both the cases. No attempt has been made to analyze the accuracy of the FV code with respect to scalar transport.

Arbitrary Grid Structure

The main objective of this work is to extend the PDF schemes to arbitrary grids with irregular structure and non-orthogonal faces. In this context, the shear flow geometry was modified to accommodate a more rigorous test case with separated inlets. The geometry used is part of a grid used in full-scale RANS simulations of industrial plants [Harvey, 2002]. Figure 3.9 shows the block structure of the domain. The coordinate axes of each block has been marked to indicate the principal indexing order. It can be seen that the blocks are arranged almost at random and no particular feature of the grid can be used to simplify particle tracking. In addition, the face-to-face strategy in the particle tracking should now be able to handle faces that may be differently oriented in each block. There are four inlet faces (with respect to the grid) that represent two inlets. The centerline is specified as an axis in order to test the reflecting boundary conditions. The inlets near the axis have a fluid inlet velocity of 0.5 m/s with mixture fraction of 0. The outer inlets have a fluid inlet velocity of 1.0 m/s with mixture fraction of 1. These values are only representative and several tests with higher velocities have been performed and no change in the results were observed.

Figure 3.10 shows the grid structure within the blocks close to the initial contact of the two streams. The rounded wall zone has been enlarged to indicate the extensive clustering near the walls required to capture the gradients without the use of a wall function. The mean and variance of the mixture fraction at different downstream positions are shown in Figs. 3.11 and 3.12. No mixing model was used and the analytical expression used in the above cases is plotted with the variance profiles. It can be seen that the results from the FV and PDF codes agree very closely for all axial positions. The deviations in the mean profile indicate that the scalar shear-layer spread is smaller than the FV code. The variance profiles follow the trends in the previous cases. The FV code seems to underpredict the variance at all the positions shown. In addition, the FV code does not correctly predict the location of the peak and is considerably more symmetric than the PDF predictions given that the velocity ratio is 2.0 based on the inlet velocity. However, similar to past trends, the FV predictions seem to improve in quality downstream.

Summary

This conclusively proves that the PDF scheme is able to operate on arbitrary irregular grids and provide accurate predictions in terms of the mean and the variance of the mixture fraction. This concludes the test for numerical accuracy of the PDF schemes. Several other tests have also been conducted to analyze the effect of particle number, time step and feedback iteration steps on the results. These tests are particular to the flow problem considered and will be discussed in later sections in the context of the application. The next section will validate the feedback scheme and the detailed chemistry implementation using experimental data.

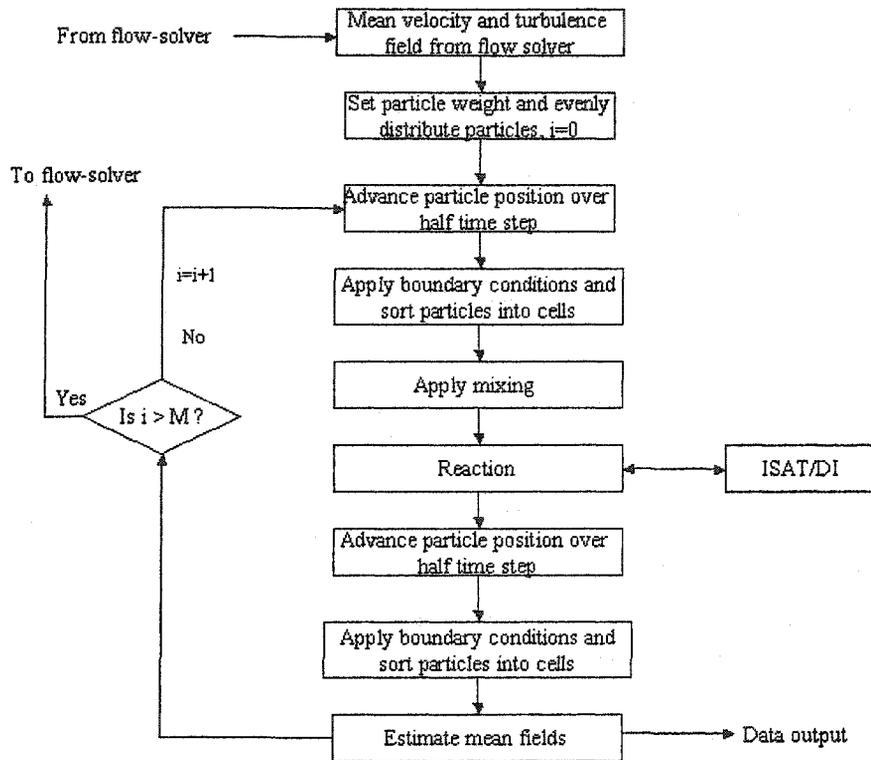


Figure 3.1 Flowchart of transported PDF method.

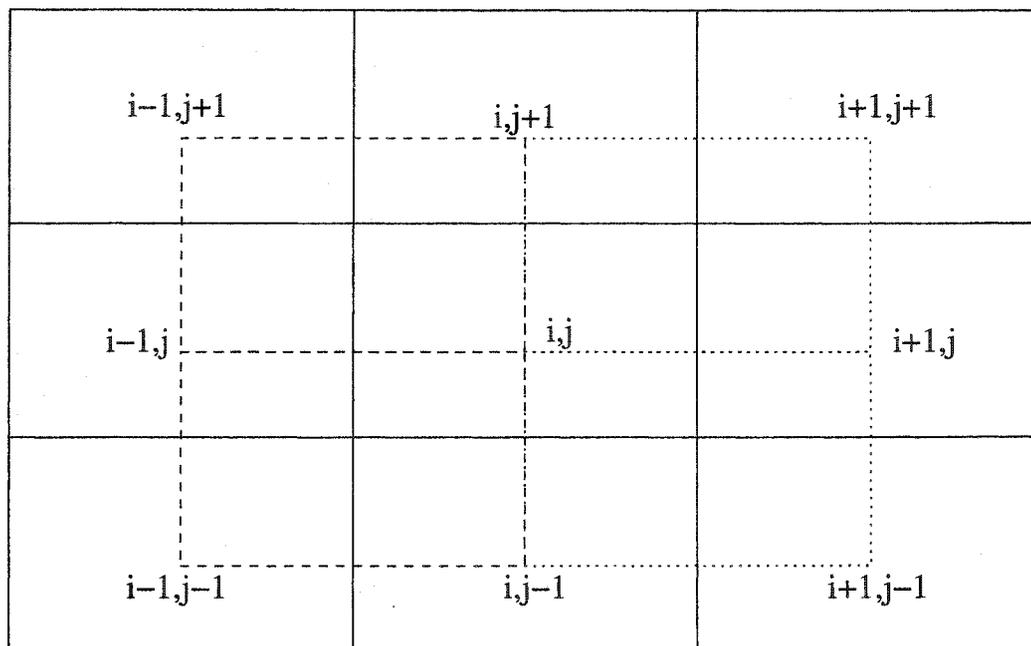


Figure 3.2 Cell neighborhood for interpolation.

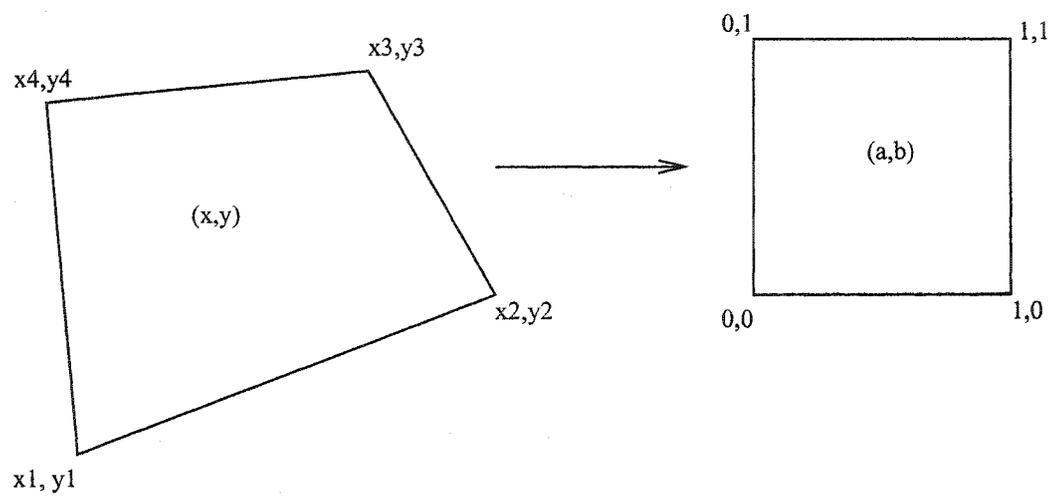


Figure 3.3 Bilinear mapping used to transform to unit square.

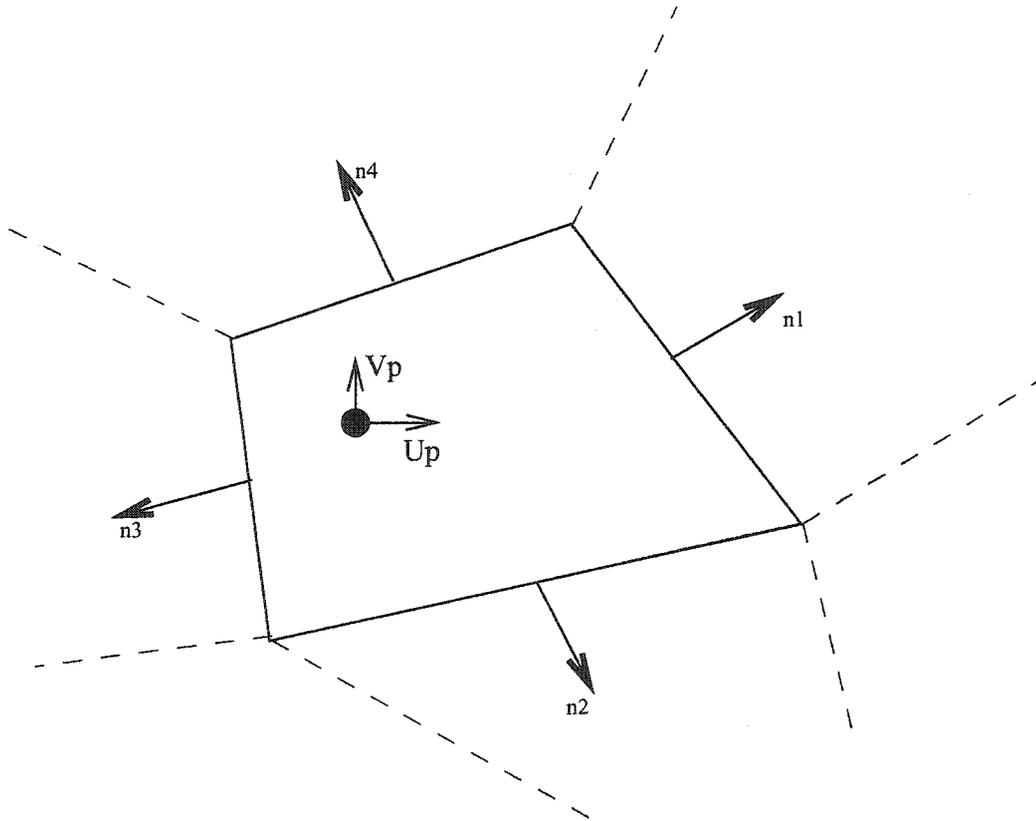


Figure 3.4 Single cell tracking procedure. Cell normals are indicated by n_i .

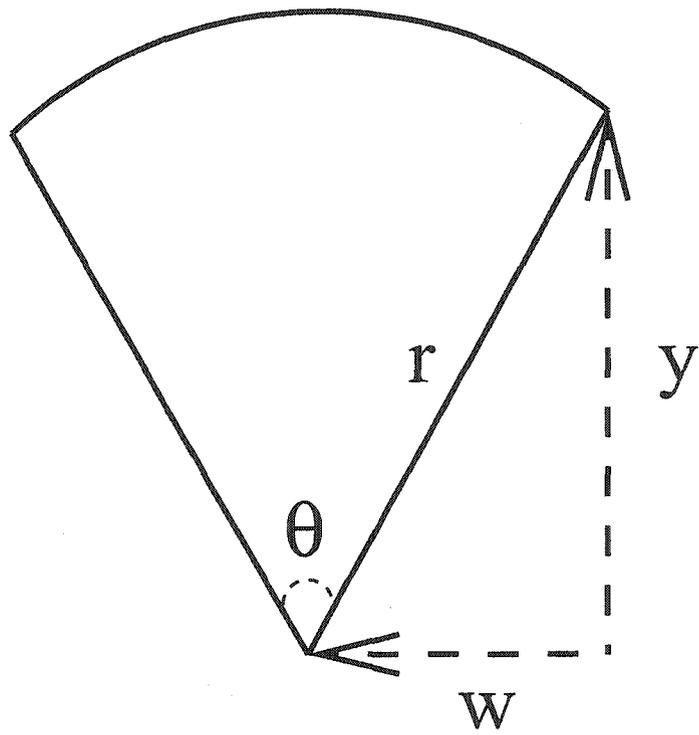


Figure 3.5 Correction applied to axisymmetric domains.

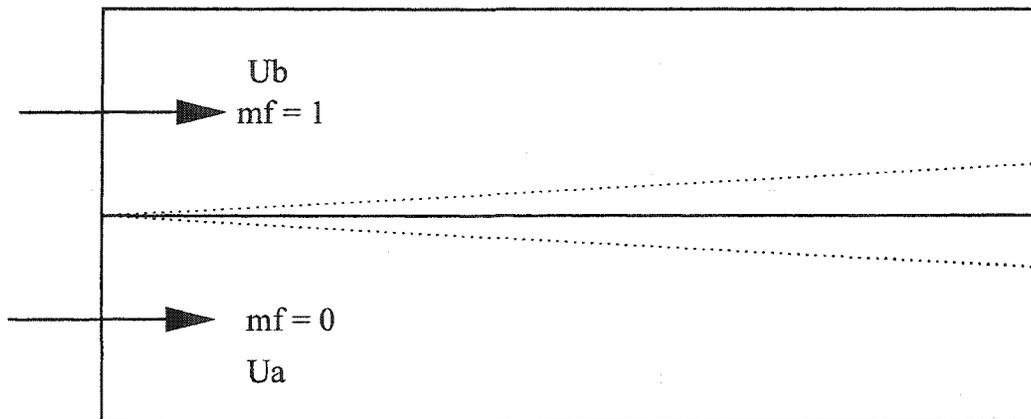


Figure 3.6 Shear layer domain used for the test cases. The dotted lines show a rough outline of the shear layer.

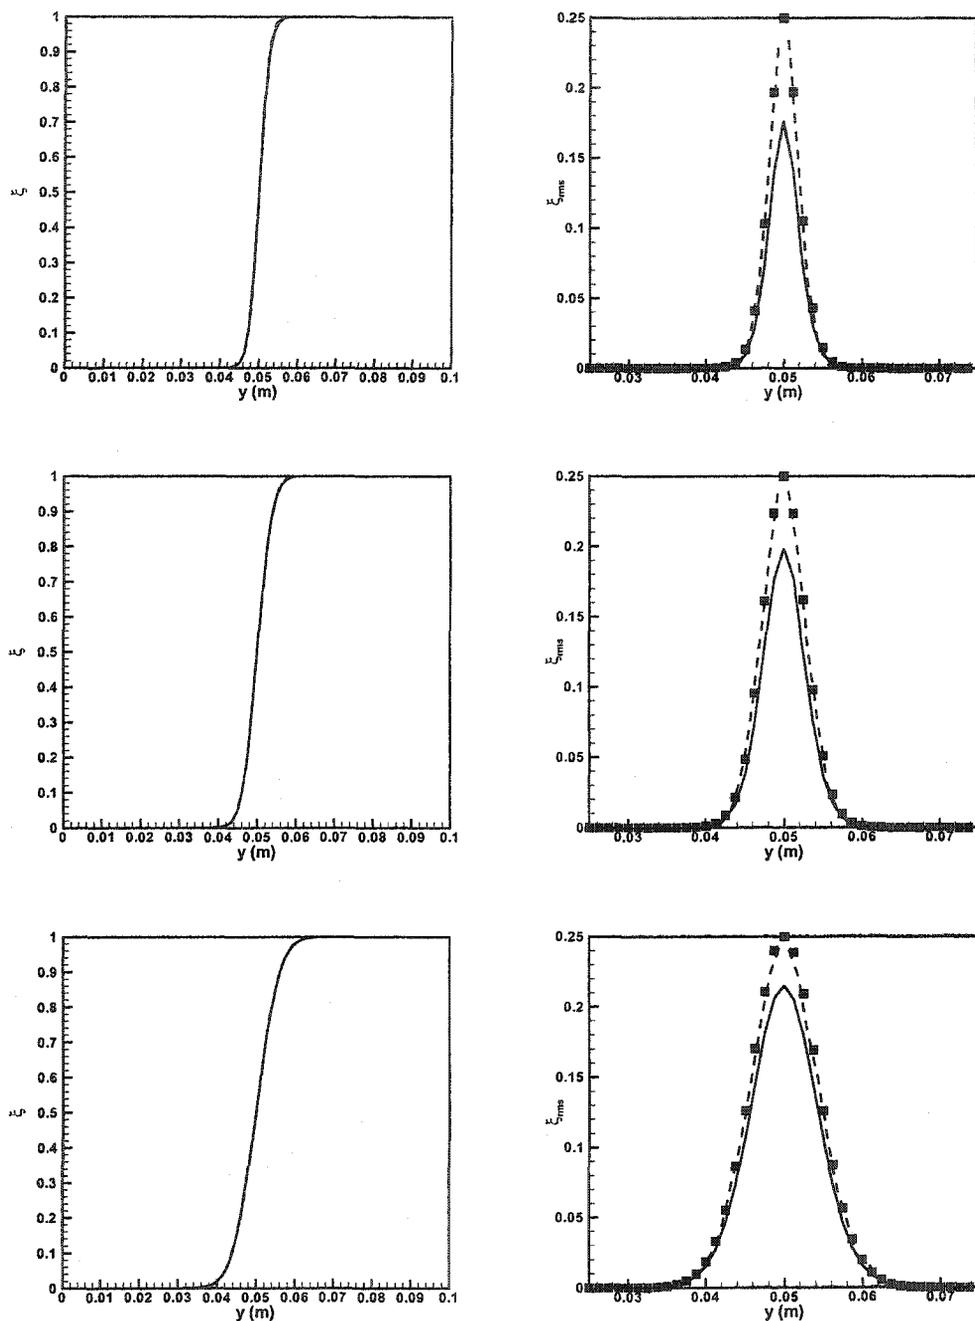


Figure 3.7 Comparison of Finite-Volume/ PDF results of mixture fraction and variance profiles for test case with constant velocity. Dashed line - PDF. Solid line - FV. Symbols - Theoretical value.

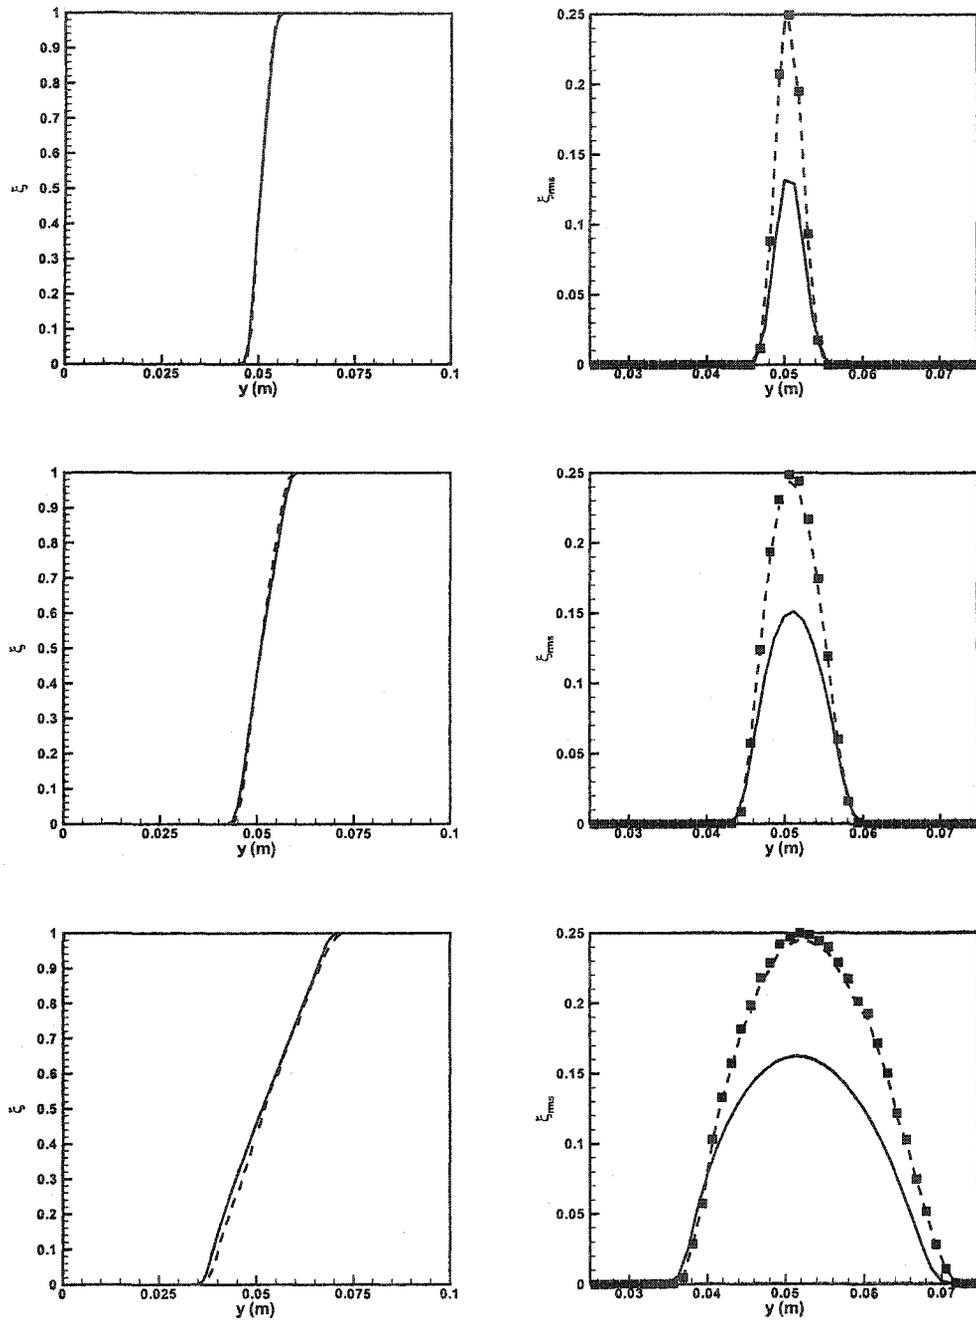


Figure 3.8 Comparison of Finite-Volume/ PDF results of mixture fraction and variance profiles for test case with velocity-shear. Dashed line - PDF. Solid line - FV. Symbols - Theoretical value.

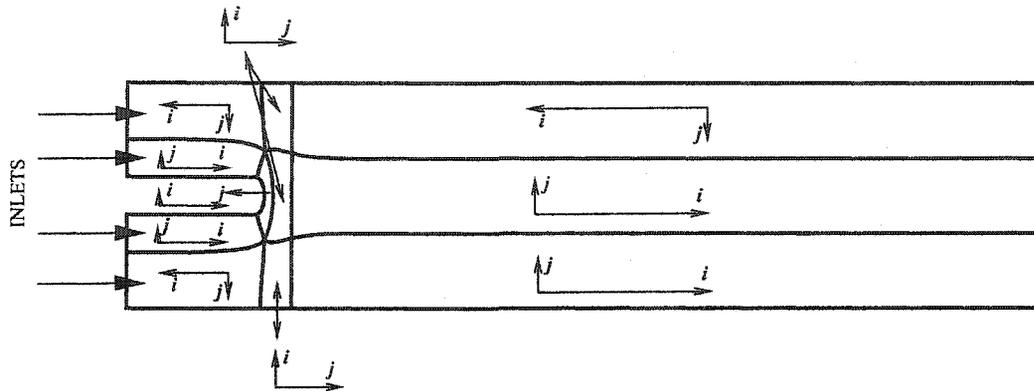


Figure 3.9 Grid structure for the test case. The individual block orientations have been marked with respective axes.

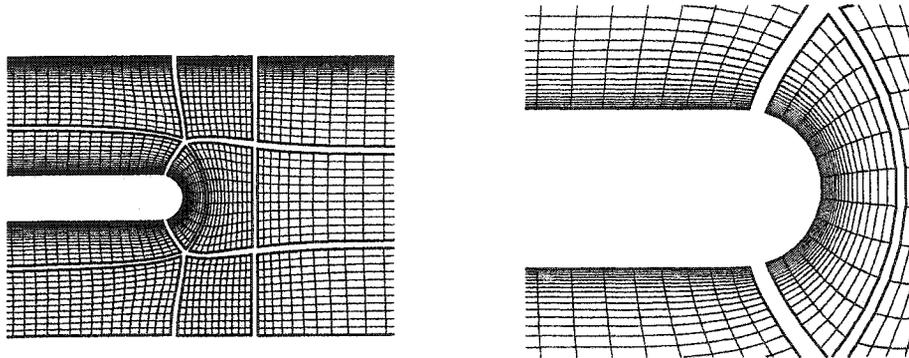


Figure 3.10 Enlarged view of the grid. The blocks have been slightly displaced to distinguish block boundaries.

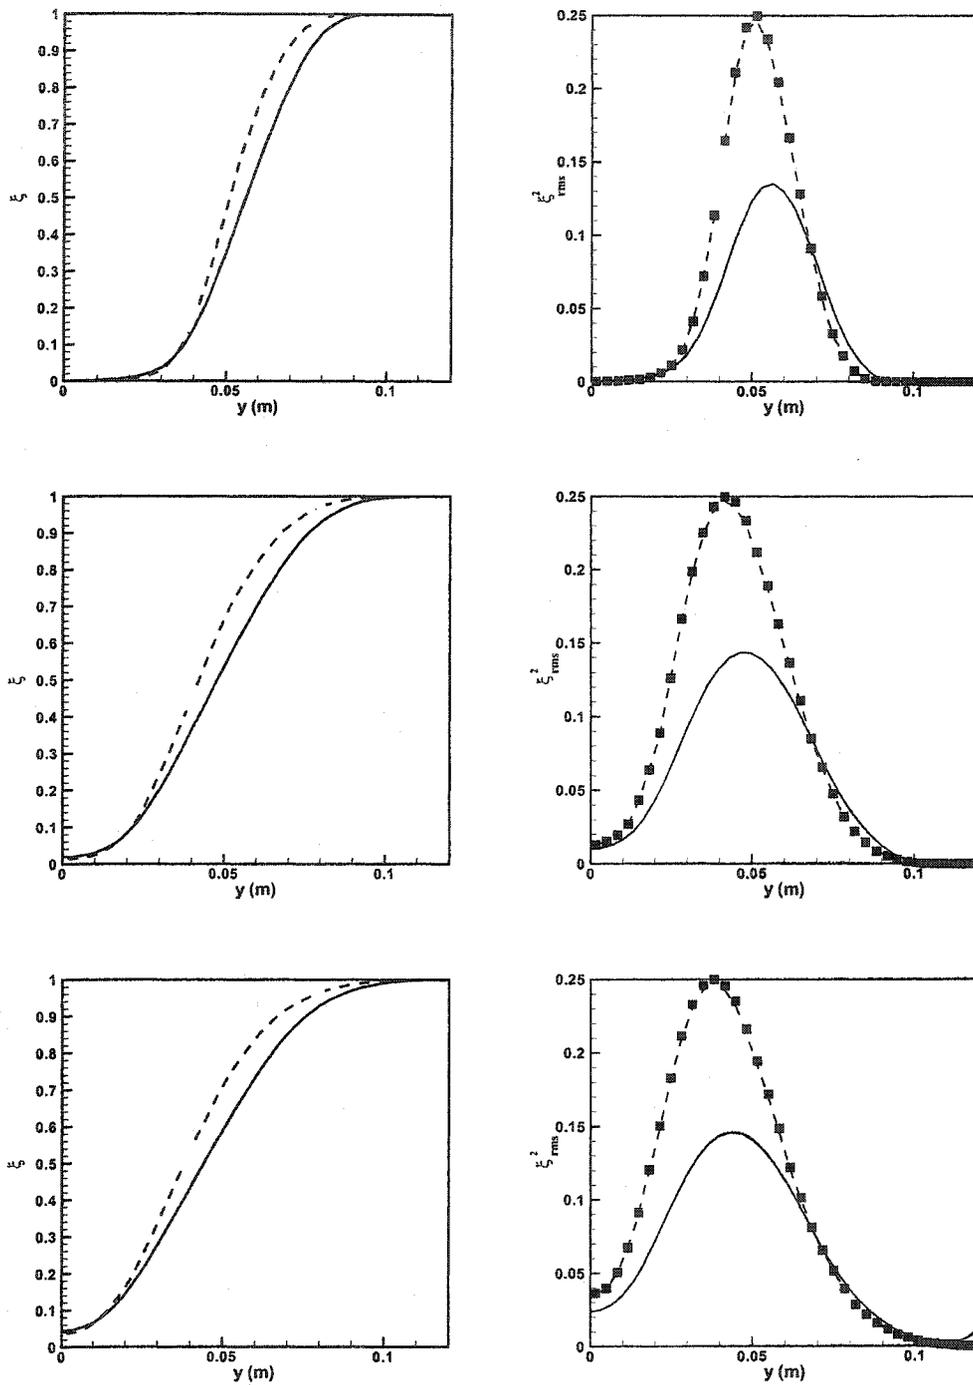


Figure 3.11 Comparison of Finite-Volume/ PDF results of mixture fraction and variance profiles for the arbitrary test grid. Dashed line - PDF. Solid line - FV. Symbols - Theoretical value.

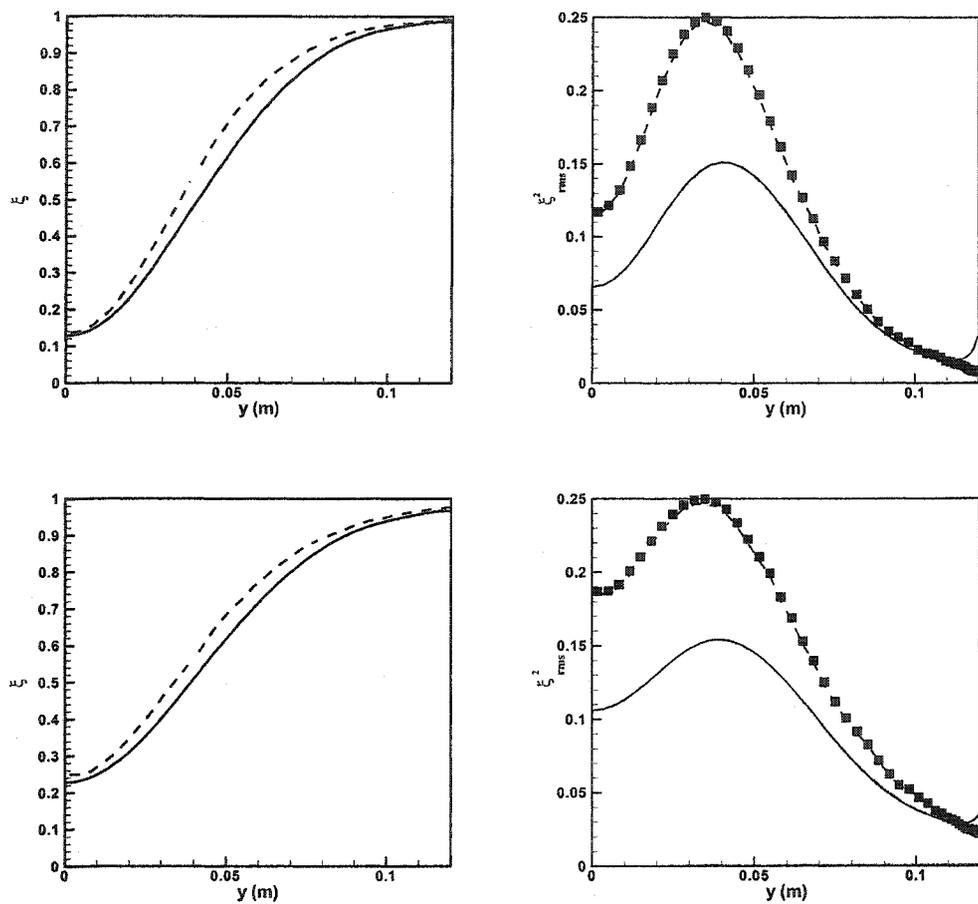


Figure 3.12 Comparison of Finite-Volume/ PDF results of mixture fraction and variance profiles for the arbitrary test grid. Dashed line - PDF. Solid line - FV. Symbols - Theoretical value.

4 Methane-Air Diffusion Flame Simulation

Introduction

The methane-air jet diffusion flame configuration of Barlow and Frank [1998], termed Sandia D flame, is an ideal candidate for testing the numerical stability and accuracy of the hybrid algorithm. The experimental set-up is simple and the data set is quite detailed to allow modeling with few assumptions. The inlet conditions are measured and readily available and the boundary conditions can be easily incorporated into FV codes. LES calculations [Pitsch and Steiner, 2000] and other PDF calculations [Xu and Pope, 2000] have yielded good predictions. Pitsch *et al.* [Pitsch and Steiner, 2000] have used the flamelet model with a reduced form of the GRI-2.11 mechanism to obtain good agreement with experimental data. Most simulations have been carried out with such skeletal mechanisms involving 4-16 species. It has been concluded that the modified $\tilde{k} - \tilde{\epsilon}$ or Reynolds-Stress based turbulence models are quite good at predicting the mean flow quantities. Eulerian as well as Lagrangian composition PDF codes coupled with a finite-volume code are found to yield reasonably accurate predictions over most of the simulation domain. Some discrepancies have been noted especially in the radial profiles of the mixture fraction and temperature. The different reaction mechanisms were found to yield qualitatively similar results [Pitsch and Steiner, 2000].

The main advantage of using PDF methods is that detailed chemistry can be implemented along with fast source term integrators to minimize the effect of the reduced chemistry mechanism on the solution. The GRI reaction schemes provide the most comprehensive and standardized set of mechanisms for natural gas combustion. The GRI-2.11 mechanism has a list of 49 species and 277 elementary reversible reactions [Bowman et al., 1995]. Several of these reactions may be unimportant to the present case but the mechanism itself has been optimized including these reactions. They are hence included in the calculations as well. The rate constants and the sensitivity of the parameters have been extensively studied. The GRI-3.0 mechanism has 53 species and

325 reactions [Smith et al., 2000]. A reduced scheme, the 16-species Augmented Reduced Mechanism (ARM) is also used to compare the effect of chemistry on flame prediction [Sung et al., 1998]. The ARM mechanism does not contain NO species. The GRI-3.0 release contains modifications that produce higher values for the prompt NO reaction and leads to a higher NO mass fraction. To demonstrate the effect of the chemistry, we have used the GRI-2.11 mechanism in its entirety along with the GRI-3.0 scheme and the ARM. The choice of mixing model and its parameters were found to affect the stability of the flame [Xu and Pope, 2000]. Here the mechanical-to-scalar time-scale ratio, which appears as a parameter in the IEM model, is varied to understand its effect on the mean and variance of the scalar fields.

The FV-PDF hybrid algorithm is used to simulate this diffusion flame using detailed chemistry. In the next section details of the numerical procedure relevant to this case is detailed. A FV simulation of the reacting flow with no sub-grid closure will be presented which will illustrate the need for micromixing models. Several cases using different chemistry mechanisms and model constants will be discussed in the later part of the chapter.

Numerical Procedure

The FV simulation of a flame involves the solution to the entire set of flow and species transport equations. The detailed chemistry source term and the Jacobian of the reaction matrix are directly incorporated into the implicit solver. The FV code does not contain a closure for the source term and assumes complete scalar mixing at sub-grid scales. It was found that the turbulence fields solved using the \tilde{k} - $\tilde{\epsilon}$ equations were sensitive to inlet boundary conditions. For the flame simulations in this work, the axial velocity and the turbulent kinetic energy at the inlet were specified from the available experimental data. It is also known that the \tilde{k} - $\tilde{\epsilon}$ turbulence model overpredicts the axial velocity decay of the round jet [Pope, 1978]. To correct for this defect, the dissipation equation model constants have to be tuned. Usually, the value of $C_{\epsilon 1}$ is increased or the value of $C_{\epsilon 2}$ is decreased corresponding to an increase in the source term for the dissipation equation [Dally et al., 1998a]. Preliminary tests with non-reacting cases showed that for the present implementation of the turbulence model, $C_{\epsilon 1}$ needs to be set to 1.53 or $C_{\epsilon 2}$ has to be changed to 1.85. For the non-reacting case, both these changes yielded the same result for the axial and radial spread of the jet.

The FV code is formulated for complex multi-block grids. Though the code can

handle multi-species transport equations with stiff source terms, the need for sub-grid closure has been shown elsewhere [Raman et al., 2003]. Here, the transported PDF scheme is used as an alternate species handler that provides sub-grid closure through a mixing model. The detailed chemistry can be effectively handled through the ISAT algorithm. However, the generality of the PDF code implementation is limited due to the inherent non-linear sub-models. Traditionally, PDF methods have been formulated for Cartesian grids where searching, tracking, and interpolation of particles are trivial. For complex grids, such grid-dependent methods cannot be scaled linearly with domain size. Tracking particles across a complex domain will need specialized algorithms that are independent of the grid structure and particle number density. The algorithmic modifications needed to attain this flexibility are explained below.

The PDF part of the solution uses the finest grid in the multi-block FV simulation. The grid as well as the cell-centered velocity and turbulence fields are transferred to the PDF code. The initialization involves uniform distribution of the notional particles and assignment of composition vectors to particles. Usually the particles are set to a high temperature to initiate combustion. In the case of piloted flames, such artificial ignition is not required. The PDF particles represent the mass of the cell. This is marked using a weight such that the sum of the weights of all the particles in a cell correspond to the mass of the cell. But in axisymmetric simulations or cases with non-Cartesian and clustered grids, the volume range and consequently the cell-mass range may span 2-3 orders of magnitude. Consequently, bigger cells end up having large numbers of particles and the smaller cells have very few particles. This introduces statistical error and bias in the calculation of the means. To correct this defect, variable particle weights are used. Each particle is initialized to a certain weight as follows:

$$wt_{p,i} = \frac{vol_i \rho_i}{N_p \sum_{j=1}^{N_c} vol_j \rho_j}, \quad p = 1, \dots, N_p \quad (4.1)$$

where $wt_{p,i}$, vol_i and ρ_i are the particle weight, cell volume, and density. N_c is the total number of cells and N_p is the nominal number of particles in that cell. As the particles evolve in physical space they will move into bigger or smaller cells. Accordingly the particles are combined or split to maintain the particle number density [Haworth and El Tahry, 1991]. The sum of the particle weights is always equal to the cell mass (within statistical errors) which is ensured by the continuity mass balance. Use of variable weights is a common practice [Xu and Pope, 2000] and by combining this with variable particle number density, heavier grid clustering can be used. In the grid used here, the ratio of the smallest to the largest cell mass is of the order of 10^{-6} , but the error in

continuity has been reduced to less than 5% of the cell mass.

For particle transport, bilinear interpolation is used to compute the velocity and turbulent diffusivity at the particle position. The Wiener diffusion term in the transport equation is represented by a set of standard normal distributed random numbers with variance Δt . Due to the large number of particles involved, tracking through non-Cartesian cells is non-trivial. The PDF code is based on an object-oriented methodology that links the particles to the cells. When particles move out of cells, the owner cell of the particle transfers control to the new cell. This eliminates the need for sorting of particle arrays through non-linear algorithms. Also we use a cell-face strategy similar to the algorithm in [Subramaniam and Haworth, 2000] to move particles. By this method, the face normals are used to compute the shortest time required for a particle to reach any of the faces in the cell. The particle is moved to this face and the procedure repeated until the total motion time is equal to the time step. This method is grid independent and eliminates the need for particle searching that is computationally expensive on non-Cartesian grids.

For a Cartesian grid, the time step is chosen based on a local CFL type condition:

$$\Delta t = C \min_{i,j} \left(\frac{\Delta x}{u}, \frac{\Delta y}{v}, \frac{\tilde{k}}{\tilde{\epsilon}} \right), \quad (4.2)$$

where $C < 1$ for a grid of size of $i = 1, imax$ and $j = 1, jmax$. Clustering leads to reduction in grid spacing thereby reducing the time step. For the type of grids used in this work, this minimum time step is of the order of 10^{-6} s and makes the computation intractable. Variable time-stepping techniques have been formulated for “stretched” grids [Muradoglu and Pope, 2002], but the time accuracy of the PDF method is lost. Though we use a steady-state approach in this work, coupling for an unsteady FV-PDF code is numerically similar and an attractive method that we are developing. To maintain the time accuracy of the PDF method and to effectively use grid clustering we have incorporated multiple-stepping for particles. The particles are moved in a cell according to the local time step determined by the minimum of the flow time scale and the mixing time scale. The global time step is achieved by multi-step movement of the particle. The cell values determine the particle velocities at each fractional step. This makes the transport numerically accurate. Theoretically, the global time step can then be set to any value. But the serial nature of the fractional time-stepping method leads to errors in composition space as particle mixing is done only after complete particle stepping. Since the time step is now independent of the grid, numerical tests need to be performed to determine an optimum global time step.

The performance of the PDF algorithm was analyzed for the flow configuration (Fig. 2.1) used for the flame calculations. The results are compared in Fig. 4.2. Here, different tests are performed to estimate the increase in computational load with the number of particles per cell. The times shown indicate the time taken for a tracking step for the entire ensemble averaged over 10 time steps. Curve (a) is the base case that represents the time taken using a time step $\Delta t = 2.5 \times 10^{-4}$ s for a grid size $N = 2000$ (number of cells). Curve (b) presents results for the same grid but tracking with double the time step. This increases the number of particle interpolations and face normal calculations. However, it can be seen that curve (a) and (b) differ by approximately a factor of two for all particle numbers considered. Curve (c) represents the tracking time for a grid of size $4N$ but with the original time step Δt . It can be seen that the times recorded are almost five times the tracking time for case (a). In all the cases studied, the tracking time increases linearly with number of particles per cell and consequently, the number of particles in the domain. From the results in Fig. 4.2 it can be concluded that the tracking time scales linearly with particles, grid size and global time step. It should be noted that when chemical reactions are added, the ISAT subroutines use 93% of the computational time. However, it is expected that for very complex grids, the increase in the computational efficiency of the tracking algorithm will have a significant impact on the total simulation time.

Simulation Conditions

The Sandia D flame is a piloted methane/air jet diffusion flame [Barlow and Frank, 1998]. The fuel consisting of 25 % methane and 75 % air forms the inner fuel jet with a diameter of $D = 7.2$ mm. The flame is stabilized using a pilot with a diameter of $D_p = 18.2$ mm. The pilot has a composition corresponding to that of a burnt fuel mixture. Barlow and Frank [Barlow and Frank, 1998] have experimentally studied this flame. The jet configuration and inlet conditions for velocity, turbulent kinetic energy, temperature and species concentrations are identical to the experimental setup. The central jet has a bulk velocity of 49.6 m/s . The pilot issues at 11.4 m/s and the coflow velocity is set to 0.9 m/s . The inlet velocity profile was obtained from experiments.

The inlet conditions for turbulence are more difficult to prescribe since for the $\tilde{k} - \tilde{\epsilon}$ model, the turbulence intensity and the length scale are required. The turbulence intensity can be obtained from the velocity fluctuation data provided by Sandia. Most RANS simulations use a constant length scale across the inlet but a linear variation of

the form

$$L_m = C y_w, \quad (4.3)$$

where y_w is the distance to the wall, is used here [TNF, 2000]. The value of C can be estimated by solving a simple pipe flow with the appropriate jet diameter and flow conditions and using the exit value of kinetic energy and dissipation factor:

$$C = \frac{C_\mu^{\frac{3}{4}} k^{\frac{3}{2}}}{\epsilon y}. \quad (4.4)$$

Over the jet cross-section, the value of C is found to be almost a constant value of 0.12. This inlet profile is computed for the main fuel jet and the pilot. The coflow has very low radial variation in velocity and the length scale is assumed to be a constant.

Non-reacting flow cases with inlet conditions identical to the experimental setup for the reacting case [Barlow and Frank, 1998] were used to study the grid dependency of the solution. Freestream boundary conditions were applied at the outer boundary of the coflow and an axisymmetric centerline was used. The freestream was set to the coflow conditions. The odd-shaped domain (Fig. 2.1) was used so that at the inlet and the exit, the cells are rectangular. The other axial locations were determined by solving for the roots of a cubic polynomial. This minimizes the flux error in the FV code. The size of the domain was chosen such that the freestream boundary is not affected by the jet reactions. It was concluded that a 53×63 3-block grid yields grid-independent solutions and is able to resolve the velocity gradients. The grid points are clustered at the entrance and near the shear layer between the jets to resolve the gradients. For the PDF code, 50-70 particles per cell were found to yield particle-number independent results. The time step was set to 5×10^{-5} s based on the lowest mixing time scale.

The following simulation strategy is adopted to minimize computational time:

1. A FV flow solution is obtained with the prescribed boundary conditions. The pilot temperature is set to 1880 K.
2. With constant flow profile, a PDF solution is obtained using either the ARM or the GRI mechanism, and with infinite micromixing rate. This is similar to obtaining a finite-volume solution.
3. Starting from the solution in the previous step, the mixing model is turned on with the desired C_ϕ value.
4. Once a steady-state solution is reached, the feedback is initiated. The PDF code then runs for M steps and feeds temperature back to the finite-volume code. The

FV code runs for N steps and transfers back converged flow and turbulence fields. The value of M is arbitrarily fixed to be 30. N is found using trial and error to ensure that the finite-volume code converges to preset residual levels before the flow field is transferred.

To ensure that the above strategy does not lead to systematic errors, several independent simulations with finite-rate mixing were carried out starting from non-reacting initial conditions. These were found to converge to the same solution but were much slower. The computational expenses in some of these simulations compared to the ones done with the above methodology were more than a factor 10 larger.

One of the advantages of the present formulation of the PDF method is that variable time steps can be used without dependence on grid. However, this affects the mixing model. Even though mixing is linear in time, due to particle movement through multiple cells in the same time step, finite error is introduced in the sub-grid mixing process. In particular, particles can now pass through cells without actually affecting the means of the scalar composition in those cells. Several independent simulations were carried out with both reacting and non-reacting cases to study this error. In the first set of the simulations, the mixture fraction and its variance profile along the axis and in the radial direction at selected axial positions were compared for different time steps. The simulations were performed under non-reacting conditions. The mixing constant was set to $C_\phi = 2.5$ after it was determined that this parameter has no effect on the time step chosen. It can be seen (Fig. 4.3) that the time step has no effect on the axial mean profile of the mixture fraction. The radial profiles (Fig. 4.4) show some minor differences but in general no extraneous numerical diffusion is visible. However, the RMS mixture fraction shows a linear trend of increase in the peak value with decrease in time step. This can be directly attributed to the mixing model. The decay of the fluctuations is exponential which means that an increased time step causes a larger decay. The change in the peak RMS value is negligible for time steps smaller than $\Delta t = 5 \times 10^{-5}$ s. This is chosen as the time step based on the non-reacting flow conditions.

The second set of test cases uses the GRI-2.11 chemistry under reacting conditions. The C_ϕ value is set to 2.5. The time steps were set at 5×10^{-5} s and 1×10^{-4} s. The axial temperature profiles (Fig. 4.5) and select radial profiles (Fig. 4.6) along with the RMS value of the temperature have been plotted. Similar profiles for O_2 and CO_2 have been plotted to verify scalar independence. It can be seen that the time step has no effect and the solutions agree closely. The minor differences in the RMS values are due to the statistical nature of the solution scheme. Reactions tend to decrease scalar fluctuations

thereby reducing the effects of the mixing model. All the other simulations in this work with the ARM chemistry are computed with a time step of 5×10^{-5} s and the GRI mechanism calculations were done with a time step of 1×10^{-4} s.

Details of Test Cases

In the following sections we will discuss the nature of the solutions and their dependence on kinetic schemes and model parameters. The interaction of turbulence and chemistry is an important aspect of this study. The tests here were designed to test the sensitivity of the various parameters in the turbulence, reaction and mixing processes. The model constants in the turbulence dissipation equation, $C_{\epsilon 1}$ and $C_{\epsilon 2}$, are used to study the dependence of reaction on the dissipation rate. The chemical mechanisms are themselves subject to different approximations and hence need to be compared against each other to quantify their performance. The mixing model constant $C_{\phi} = 2.0$ is again based on non-reacting flows and needs to be tested for reacting flows. All the data presented have been time averaged after reaching steady state to reduce statistical fluctuations. All the calculations were averaged for 100 time steps corresponding to $1/3$ of the residence time of a particle moving along the centerline. It was found that the resulting profiles did not change with longer time-averaging.

The simulations are organized in the following manner. The finite-volume simulation of the flame will be discussed first and a case made for the importance of the closure of chemical source terms. Case A is the base calculation that agrees the best with experimental data. It will be discussed first to show the full PDF closure effect. Case B is used to test the sensitivity to the turbulence model constant. Cases C and D illustrate the effect of the mixing model constant. Cases E and F compare the effect of chemistry on the flame prediction, using the same model constants as Case A. Case G shows the effect of neglecting the radiation model and is simulated with constants identical to Case A. The parameters used in each case are given in Table 1 for future reference.

The experimental data of Barlow and Frank [1998] contain several species radial profiles at select axial locations. In addition, mixture fraction profiles calculated from species data by setting the fuel mixture fraction to 1.0 and the coflow mixture fraction to 0.0, have also be tabulated. Here, only select profiles are used for the comparison with simulation.

Finite Volume Simulation with Detailed Chemistry

The finite-volume code is used first without the PDF code to understand the need for closure of the source term. The complete GRI-2.11 mechanism with 49 species is used to simulate the methane-air flame. Standard model constants are used for the turbulence equations. A constant turbulent Schmidt number of 0.7 is used for computing the diffusivity of the scalars. The axial velocity and RMS velocity profile along the centerline are shown in Fig. 4.7. It can be seen that there is close agreement for the mean profile but the RMS fluctuations are underpredicted along the entire length of the domain. It can also be seen that there is a rapid decay of kinetic energy appearing at the inlet indicated by the dip in the RMS velocity. The mean profile shows that the core of the profile extends up to $x/d = 10$ and decays rapidly afterwards. Figure 4.8 shows the centerline mean temperature, mixture fraction and scalar profiles. The peak in temperature is around 250 K higher than the experimental value. Also the location of the peak is farther upstream indicating faster reaction rates. However, the quenching after the reaction zone is predicted quite well and the outlet temperature is close to experimental values. The mixture-fraction profile indicates that the flame length is considerably underpredicted and the stoichiometric mixture fraction ($\xi_{st} = 0.351$) is observed at around $x/d = 30$. The decay of mixture fraction itself is underpredicted at upstream positions and overpredicted in the downstream positions especially near the location of the peak flame temperature. Again, the outlet mixture fraction is predicted very well by the simulation. The species profiles reflect these trends with peak values shifted upstream. It can be noted that the peaks are also higher than experimental values, which is directly linked to higher temperatures predicted by this scheme. Discussion of the radial profiles is postponed at this point as a similar case (Case D) will be analyzed later.

The finite-volume simulations point out the inadequacies in the present formulation. The use of a completely mixed model for the sub-grid scalar field leads to a short flame with high peak temperatures. The scalar profiles indicate higher production due to this excess heat. The mean profile of the mixture fraction has large errors basically due to the failure of the turbulence model to predict the correct spreading rate. It is not known whether the correct spreading rate prediction in itself will improve species fields and the temperature field. However, it can be concluded that a sub-grid closure is required and the effect of the turbulence model needs to be studied to better characterize the flame. The following sections will discuss the FV-PDF predictions and further analyze the effects of the other sub-models.

FV-PDF Predictions of the Sandia D Flame

As discussed before, simulation of the methane-air flame is quite widespread as a test case for the verification of the accuracy of a reactive flow code. The composition PDF code used here solves the species transport equation and provides the density correction to the FV code. It is known that the simulation of a round jet in an axisymmetric domain will overpredict the decay rate of the jet [Pope, 1978]. Various corrections have been proposed but the simple change of the dissipation model constant is found to be effective [Dally et al., 1998a]. In this case, the value of $C_{\epsilon 1}$ is set to 1.53 instead of the normal value of 1.44. It should be noted that a value of 1.6 has been suggested for predicting the correct spread. However, this value can be used only for a free jet and the present value of 1.53 was found sufficient to predict the correct spreading rate as compared to the experiment. The value of the mixing model constant was fixed at 2.4 to ensure that the flame stays lit. The complete GRI-2.11 mechanism with 49 species and 277 reactions was used. The optically thin radiation model implemented in the ISAT-CK algorithm was used [Tang et al., 2000].

The axial profiles of temperature and selected chemical species are shown in Fig. 4.9. The comparison with experimental data is excellent and the quenching of the flame downstream is predicted quite accurately. The air entrainment in the latter sections of the flame is predicted quite accurately as shown by the oxygen profile. CO_2 profiles show overprediction near the end of the flame but it is not reflected in the CO profiles. The H_2O profile follows a similar trend to the CO_2 profile while the prediction of H_2 is almost exact. It can also be noted that the error in oxygen prediction is roughly translated into the CO_2 and H_2O overprediction. The axial NO prediction is very accurate with the GRI-2.11 mechanism and the hydroxyl profiles are also reproduced with minimal errors. The effect of the chemistry mechanism on the NO prediction will be discussed later.

The radial profiles (Fig. 4.10) of temperature show good agreement with experimental data in predicting the spread of the flame. It can be seen that the RMS profiles of the temperature fluctuations are also in close agreement. The mixing model does not seem to adversely affect the prediction of the subgrid fluctuations. This is again verified by the radial profiles of the mixture fraction and its variance. However, the variances are underpredicted significantly after $x/d = 45$. The peak in the variance is predicted significantly upstream. However, experimental uncertainties also need to be taken into account as the inlet variance itself seems to be non-zero from the given data. The validity of the turbulence models in regions with sharp density gradients needs to be further examined. It can also be seen that the H_2O profiles are slightly overpredicted

in the radial direction as well, although minor species like the hydroxyl radical are accurately predicted. The CH_4 radial profile shows good agreement with experimental data at all locations considered. However, the radial profile of O_2 show lower values than experiment and is consistent with the overprediction observed in the H_2O profiles.

To illustrate the interaction of turbulence and chemistry, conditional means of scalars conditioned on mixture fraction are plotted (Figs. 4.11 and 4.12). The conditional mean of O_2 , CH_4 , H_2O , H_2 , CO_2 , CO , NO , OH are shown at different axial positions. It can be observed that the conditional means are quite accurate for all the major species. There are some minor errors in the CO and H_2 profiles at $x/d = 30$ though the mean profiles at these locations are in close agreement with the experimental data. The NO profiles show good agreement but are consistently underpredicted which could be due to the inclusion of the radiation model that quenches the flame by nearly 100-200 K. In general, it is seen that the mixing model combined with the detailed chemistry is able to predict the flame behavior in extended mixture fraction-composition space. The mixing branches close to pure fuel and pure oxidizer conditions are predicted very well. The observed difference could also be numerical, since the number of particles at the peak are comparatively fewer, and could lead to slight errors in the calculation of the means. No effort has been made to test this artifact as the comparisons seem to be in good agreement in the rest of the mixture-fraction space.

Figure 4.13 shows the axial velocity profile and the RMS of velocity. The decay rate of the velocity is predicted quite accurately. The RMS values normalized by the centerline velocity are higher than experimental values indicating that the decay of the kinetic energy is underpredicted. Though no direct conclusion can be drawn about the decay of kinetic energy from this plot, the correct spreading rate prediction indicates that the turbulent kinetic energy along the centerline is overpredicted. The effect of turbulent model parameters will be looked at in detail next.

Effect of Model Constants in Turbulence Energy Dissipation Equation

As noted above, the standard $\tilde{k} - \tilde{\epsilon}$ equation overpredicts the decay rate along the centerline. Consequently flames are shorter and broader. In this case, it was found that the fast decay rate causes flame extinction due to rapid dissipation of the pilot flame enthalpy. To compensate for this dissipation, the mixing model constant had to be set to very high values which led to flame lengths of the order of 25-30 jet diameters. It is known [Pope, 1978] that the simulation of jets with 2-d axisymmetric conditions overpredicts the decay rate. The so-called round-jet/plane-jet anomaly is common to almost

all the turbulence models except the $k - \omega$ model and the enstrophy-equation model. Corrections to the model equations have been proposed [Pope, 1978] that decrease the spreading rate. Here the model constants are modified to produce the same effect. There are two different ways - either by increasing $C_{\epsilon 1}$ [Dally et al., 1998a] or by decreasing $C_{\epsilon 2}$ [Jones and Khaki, 1998]. It has been argued that both these changes will produce similar effects. Here we compare these changes. Case B is simulated with $C_{\epsilon 2} = 1.85$ as opposed to a standard value for round jets (1.80). The present implementation was found to produce the correct axial decay (Fig. 4.13) and further decreases in the constant led to longer flame lengths and large errors in the prediction of the radial spread of the flame. Further increase in the parameter value led to shorter flames and eventual extinction.

The axial temperature profile (Fig. 4.14) shows close agreement with experimental data. It can be seen that the quenching near the end of the domain is less than Case A, suggesting reduced entrainment from the outer coflow. Also the peak in the temperature is attained at a shorter length than the experimental data. The radial profiles of temperature (Fig. 4.15) show a wider spread of the flame. Even at $x/d = 15$, the flame is broader than the experimental data. The radial temperature profile at $x/d = 30$ is overpredicted and the RMS value of temperature is underpredicted at the peak. The predictions are much better downstream of the stoichiometric flame length. The predictions in the mixture fraction profile follow a similar trend with a larger radial spread than that observed in Case A. The discussion of other species profiles have been omitted due to the similarity in the predictions with the temperature profile.

However, the conditional means (Fig. 4.16) indicate very close agreement with the experimental data, principally because the spatial dependence has been eliminated. Since the variance predictions of this simulation are similar to Case A, very small variation is noted in the conditional means. It can be noted that the NO profile indicates slightly higher levels, especially for axial locations further downstream. This can be directly attributed to the higher temperatures predicted at these locations. However, for the hydroxyl radical, the conditional means have been quite accurately predicted as before. The axial velocity and its RMS value compared against the experimental data shows very little difference from Case A. There is a slight overprediction of the kinetic energy values towards the end of the domain. It might be possible to control the turbulence intensity as well as the velocity decay rate by a combined variation of the two model constants. This test clearly shows that the $\tilde{k} - \tilde{\epsilon}$ model is flawed in the prediction of round jets but still can be used with minor corrections even for reacting cases. The effect

of the model parameters on simulation results are quantitative and the jet behavior in terms of the axial and radial profiles are not affected qualitatively.

Effect of Mixing Model

Sub-grid scale mixing is one of the key physical processes that is difficult to model due to the lack of clear experimental techniques to study scalar fluctuations at the smallest scales. Several mixing models have been developed mostly based on non-reacting scalars, for which DNS simulations of homogenous turbulence are available [Juneja and Pope, 1996]. Models like the EMST [Subramaniam and Pope, 1999], BMC/GIEM [Tsai and Fox, 1998] that provide a more detailed mixing model are found to be better approximations for reacting flows but still depend on the mixing parameter C_ϕ for the description. This mechanical-to-scalar time-scale ratio has been found to be around 1.5-2.0 for non-reacting passive scalars in constant density flows [Panchapakesan and Lumley, 1993] and the assumption of equal scales for momentum and mass transfer. In premixed combustion, Corrsin [Peters, 2000] has deduced that for linearly reacting scalars, this ratio decreases with reaction. On the contrary, experimental data of Gagnepain *et al.* [Gagnepain et al., 1998], show that this ratio actually increases with reaction rate. More study is required before a conclusive model can be formulated. In spite of their limitations, mixing models quantify non-mixedness and the reduced reaction rates, to a large extent, mimic the flame behavior. Hence, model parameters can be adjusted to predict the reactions more accurately.

It has been noted elsewhere that the IEM model leads to extinction of the flame [Xu and Pope, 2000]. It was found that the standard value of 2.0 for the mixing constant led to detachment of the flame near $x/d = 15$ and consequent blow-out. The parameter was varied from a value of 2.4 (where the flame was found to be stable) to a value of 3.0. In addition a case corresponding to infinite micromixing was also considered. The results (Fig. 4.17) show that for the case with $C_\phi = 3.0$, there is very little variation in the axial and radial mean profiles of the temperature and different species as compared to Case A. It can also be seen that the axial variance profiles are not significantly lower for the $C_\phi = 3.0$ case. This suggests that the value of C_ϕ is of local importance in regions of the flame close to extinction. Numerically, at around 15 jet diameters, the pilot is not strong enough to initiate a sustainable ignition rate. It was observed that extinction leading to flame blow-off initiated around this point for values of $C_\phi < 2.4$. For the infinite-rate micromixing case, the variances are near zero and the peak axial temperature is significantly higher than Cases A and C. The artificially increased mixing

rate due to the higher C_ϕ values induces ignition and keeps the flame lit. This observation is not universal, in that increased mixing rate can also lead to extinction. It should also be noted that as the velocity of the inlet is increased, extinction probability increases around $x/d = 15$ [Barlow and Frank, 1998].

The centerline profiles (Figs. 4.17 and 4.18) of temperature as well as species for the infinite micromixing case (Case D) show good agreement with the finite-volume simulations without sub-grid closure. Both these cases neglect sub-grid fluctuations and have similar impact on the shortened flame length. The higher temperatures lead to higher mass fractions of CO and NO with peaks shifted upstream. The oxygen profile shows overprediction at the end of the flame while complete methane consumption occurs well before the experimental data. The radial profiles (Fig. 4.19) show that changing C_ϕ by small amounts (from the nominal value of 2.0) has very less impact on the mean but significantly changes the variance profile. It can be seen that the error in prediction of the RMS temperature is close to 50% at $x/d = 45$. The infinite rate mixing based simulation shows the extreme case where the variances are suppressed to nearly zero. The finite values arise from the nature of the algorithm. The data output is carried out after a convection step which marks the end of the time step and the variance is forced to be zero only after mixing. The mixture fraction profile (Fig. 4.19) shows that the assumption of complete micromixing changes the flow field only slightly from the finite-mixing rate cases. The mean profiles are predicted very well up to the largest downstream positions considered. However, the radial spread is found to be higher than the other cases. The conditional means of temperature and chemical species (Fig. 4.20) for Case C show very little impact of the micromixing constant. The plots are almost identical to Case A indicating that in this parameter range, the C_ϕ constant does not yield significant quantitative differences. This study clearly shows that the use of the sub-grid scale mixing model is critical in accurately predicting the means. In addition, the nature of the mixing model itself will affect the prediction of the second-order moments. However, the variation of the model constant in the range of 2.4-3.0 does not affect the conditional or unconditional means significantly.

Effect of Chemistry Mechanism and Radiation

The axial comparison with experiments (Fig. 4.21) for Cases E and F show that the chemistry has very little effect on the axial temperature and mixture fraction profiles. The results also compare well with Case A using the GRI-2.11 mechanism. The variance predictions are unaltered from Case A. It can be noted that the mass fractions of CO_2 and

H₂O are overpredicted and the profile of O₂ underpredicted. The difference in the axial temperature profile after the peak is higher than Case A and the quenching due to radiation seems to be less severe unlike Case A. However, the minor species are still predicted quite accurately. It can be seen that the NO concentration is extremely high compared to experiments and the calculations of Case A. This can be a direct result of the mechanism change in GRI-3.0 chemistry that produces higher prompt NO. The higher temperatures near the end of the reactor will also add to the production rate. It is known that even an increase of 50-100 K in temperature can increase NO production rates dramatically resulting in such overpredictions. Both mechanisms predict the OH radical quite accurately. The radial predictions (Fig. 4.22) of mixture fraction and temperature confirm that the jet spread and flame width are in close agreement with the experiment and do not differ significantly from Case A predictions. The conditional mean of temperature indicates a good prediction by the GRI-3.0 mechanism at all the axial locations. The ARM mechanism seems to attain lower means at $x/d = 30$ but both mechanisms are able to correctly predict the mixing branches.

The conditional mean (Fig. 4.23) of NO shows a higher peak reflecting the higher mass fraction in the means. Even in the presence of radiation quenching, it is consistently overpredicted. This is in agreement with the changes in the prompt NO chemistry. The hydroxyl radical is predicted accurately indicating that the effect of the hydroxyl mechanism is unchanged from the GRI-2.11 mechanism. The GRI-2.11 mechanism was shown to underpredict the conditional mean of NO at all the axial positions considered. The possible effect of radiation is considered next in Case G. This calculation was started from the Case A steady-state solution. After the radiation model was turned off, the simulation was continued for 2.5 residence times and the solution was found to reach a steady state. The axial temperature profiles (Fig. 4.24) show an increase in the peak value that quenches very slowly. Consequently the NO mass fractions are quite high after the temperature peak. The predictions before $x/d = 45$ are quite good but deteriorate rapidly in a overpredicted temperature field. The conditional mean (Fig. 4.25) of the temperature shows a slight increase in the temperature mean from Case A. Interestingly, this effect seems to be prominent at $x/d = 15$ where the axial temperature is quite accurately predicted. The NO profile also shows a pronounced peak at this location. In general, the conditional mean values are higher than Case A for the three scalars shown, but the hydroxyl radical profile seems to be less affected by the radiation model.

Conclusions

A stable hybrid FV-PDF algorithm has been implemented and tested for accuracy. The modified tracking algorithm is found to scale linearly with grid size and particle number density. The time step size can be chosen independent of the grid and was found to be accurate for higher time steps than that stipulated by the local CFL criteria. The time step however affects the mixing process and step sizes higher than the local mixing time can lead to systematic errors. Very large time steps increase the computing time since a larger fraction of the time is spent on particle tracking and interpolation. Optimal time steps are problem dependent and *a priori* tests need to be done. For such tests, a conservative starting point can be fixed based on the smallest mixing time in the entire domain. Since the mixing time might itself change during the simulation due to changes in the turbulence fields caused by density changes, reacting cases also need to be tested. For the methane-air flame studied here, a time step of 5×10^{-5} s is found to yield consistent results for both reacting and non-reacting cases.

The feedback algorithm for the FV-PDF scheme has been tested with both density and molecular weight changes supplied to the FV code. The feedback frequency is another tunable parameter that is dependent on the problem. Smaller numbers of iterations will also require less time-averaging steps and will increase the noise in the feedback fields. The FV code might not be able to handle such statistically noisy fields and consequently diverge. On the other hand, very large iteration numbers will lead to a larger lag between the flow and scalar fields. As a result, the convergence-rate will have large fluctuations in the local flow and scalar fields. In flame simulations, such fluctuations might lead to extinction. For the system studied here, a feedback frequency of 30 is used and is found to yield stable and accurate results.

The methane-air piloted flame (Sandia D) was then simulated based on the above results. Excellent agreement was found with the experimental data using the model parameters in Case A. The model constant for the turbulence dissipation equations had to be changed to account for the correct decay rate of the jet. Several tests were performed for non-reacting jets to correct the spreading rate. Finally a choice of $C_{\epsilon 1} = 1.53$ was made. The mechanical-to-scalar time-scale ratio was set at 2.4 to keep the flame lit. It was found that the extinction probability of the flame due to flow fluctuations was especially high at around 15 jet diameters from the inlet. The higher mixing rate was found to result in a stable flame. The centerline profiles of the velocity, mixture fraction and species means were found to agree well with the experimental values. The RMS axial velocity was over-predicted especially near the end of the jet. The change to the

other model constant, C_{ϵ_2} , did not improve the centerline RMS predictions. However, the radial profiles were more spread out indicating a wider flame zone and consequent differences in the species profiles.

The choice of the chemistry mechanism used was also found to affect the predictions. Though major species were predicted consistently for both the reduced 17-species mechanism and the extended 53-species GRI-3.0 scheme, the prediction of minor species showed differences. Especially, the GRI-3.0 mechanism shows higher NO production and a significantly higher peak than the 49-species GRI-2.11 mechanism as noted elsewhere [Pitsch and Steiner, 2000]. The conditional means reflect this trend. The temperature profiles show higher levels especially along the centerline. The radiation mechanism was also tested and is found to decrease temperature profiles by 200-250 K near the end of the flame. The higher temperature in simulations with no radiation model leads to significantly higher NO prediction even for the GRI-2.11 mechanism and it can be assumed that the GRI-3.0 simulations will result in even higher NO levels. It is concluded that the optically radiation mechanism employed in ISAT-CK is effective in predicting the correct temperature profile but its assumptions need to be tested based on experiments for different combustion mixtures. It was found that the use of ISAT-CK led to speed-up of up to 500 compared to direct integration in simulations using the ARM mechanism. Direct integration based simulations were intractable for the GRI mechanisms.

The direct comparison with experiments have conclusively shown that the FV-PDF hybrid algorithm is a viable scheme for simulating turbulent reactive flows with detailed chemistry. The next chapter will deal with the simulation of chlorination reactors using pilot and plant-scale configurations.

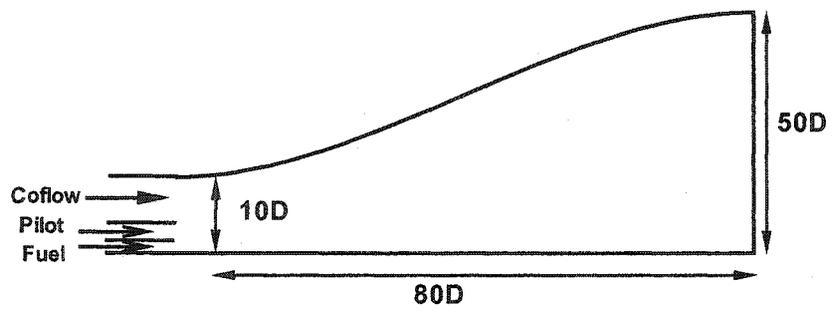


Figure 4.1 Simulation domain used for Sandia D flame.

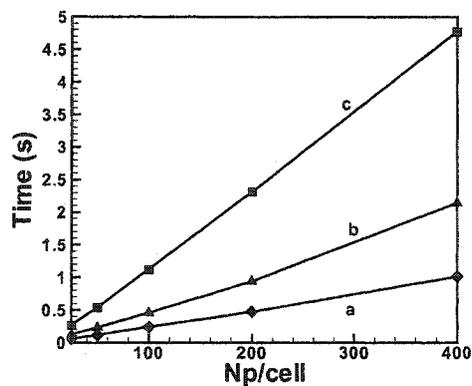


Figure 4.2 Comparison of ensemble tracking time per iteration with different particle numbers per cell for varying grid size. (a) 2000 cells and time step of 2.5×10^{-4} s. (b) 2000 cells and time step of 5×10^{-4} s. (c) 8000 cells with time step of 2.5×10^{-4} s.

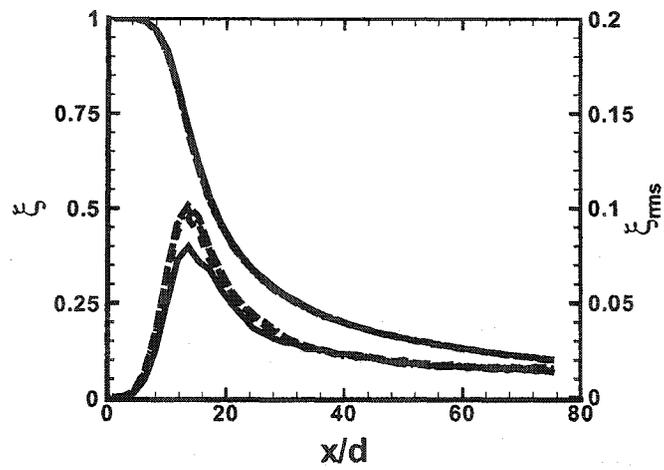


Figure 4.3 Axial mean mixture fraction and variance for test case used for optimal time step determination. Solid lines use a time step of 10^{-4} s. Dashed lines: 5×10^{-5} s. Dashed-dotted lines: 2.5×10^{-5} s.

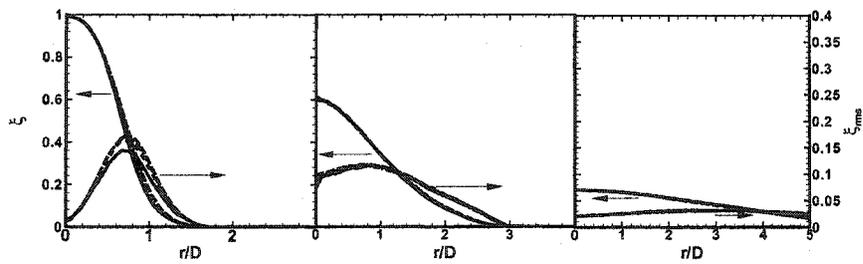


Figure 4.4 Radial mean mixture fraction and variance at (left) $x/d = 15$ (middle) $x/d = 30$ and (right) $x/d = 45$ for test case used in time step determination. Solid lines: 1×10^{-4} s. Dashed lines 5×10^{-5} s. Dashed-dotted lines: 2.5×10^{-5} s.

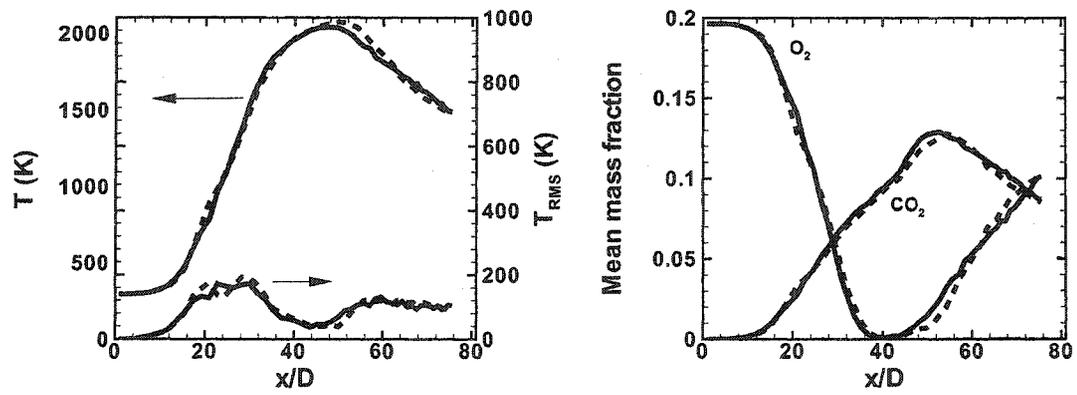


Figure 4.5 (Left) Axial mean temperature and RMS of temperature for two different time steps. (Right) Axial mean species profiles for the two cases. Solid lines: 1×10^{-4} s. Dashed lines 5×10^{-5} s.

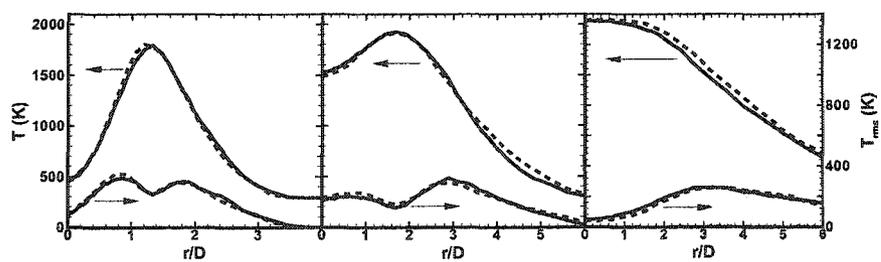


Figure 4.6 Radial mean temperature and RMS of temperature for reacting flow case at axial positions of (left) $x/d = 15$ (middle) $x/d = 30$ and (right) $x/d = 45$ for two different time steps. Solid lines: 1×10^{-4} s. Dashed lines: 5×10^{-5} s.

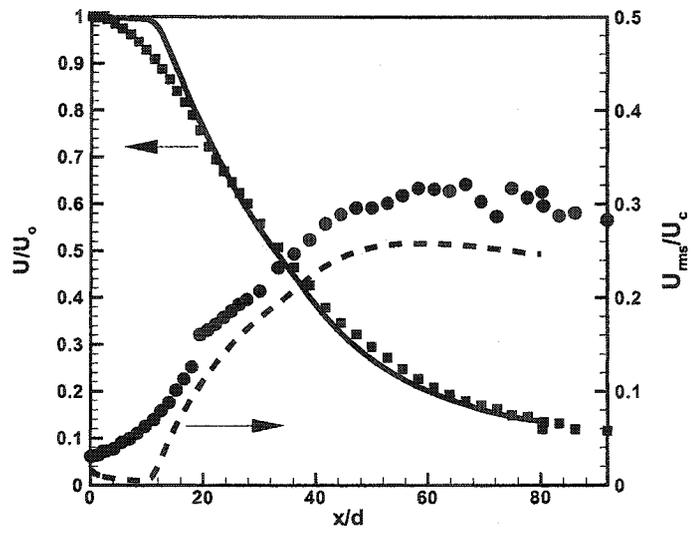


Figure 4.7 Centerline profile of axial velocity and RMS axial velocity obtained without a micromixing model in the FV code compared with experimental data. Lines show simulation results and the symbols are experimental data.

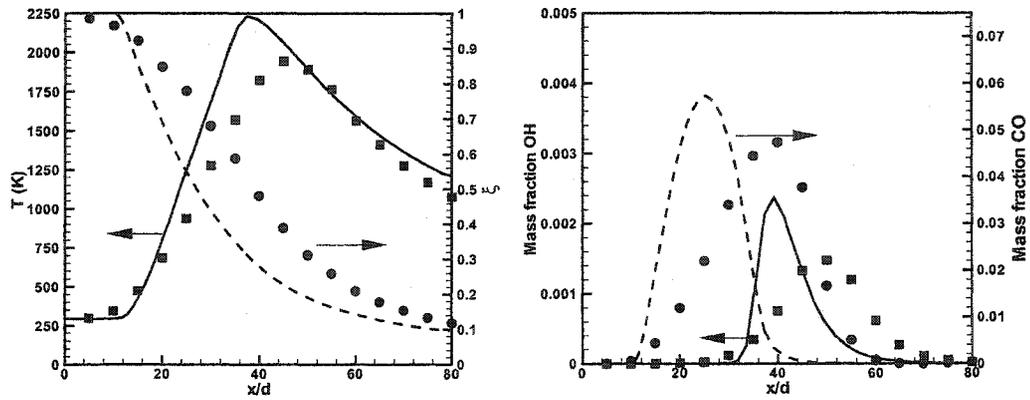


Figure 4.8 Centerline profile of (left) temperature and (right) chemical species obtained without a micromixing model in the FV code compared with experimental data. Lines show simulation results and the symbols are experimental data.

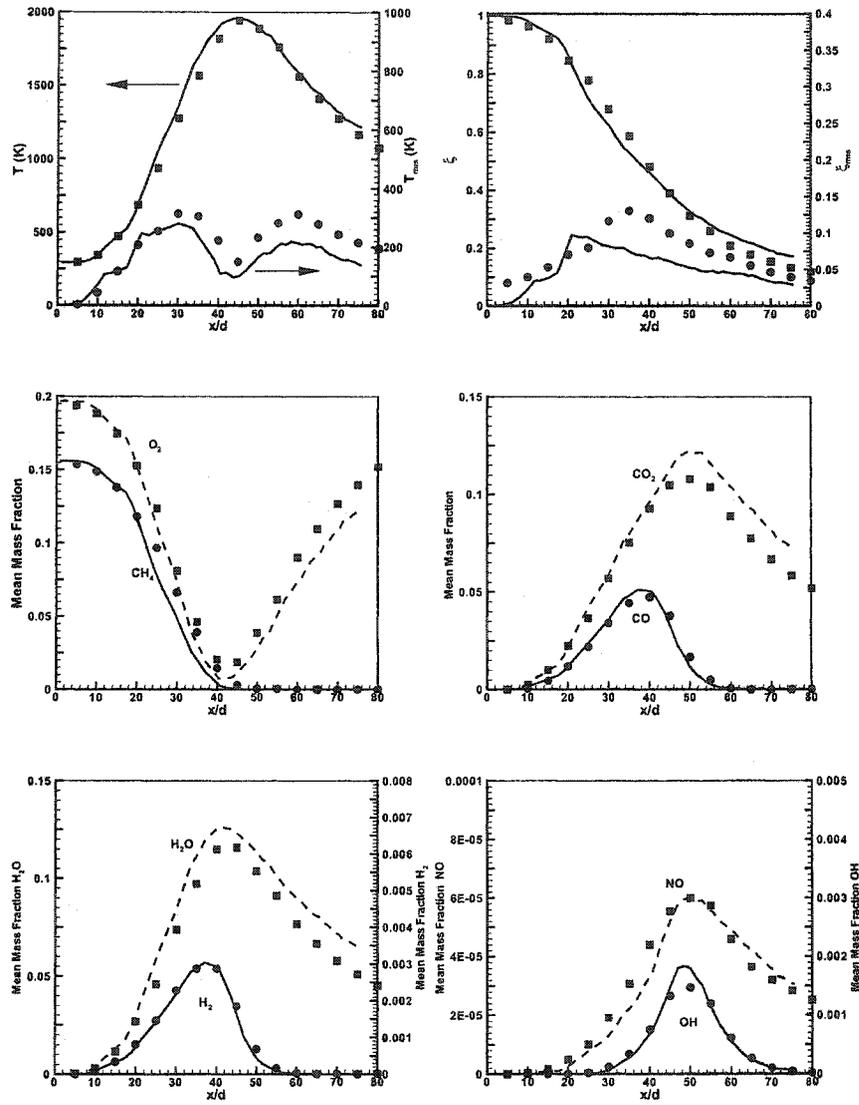


Figure 4.9 Axial profile of chemical species (lines) and experimental data (symbols) for Case A.

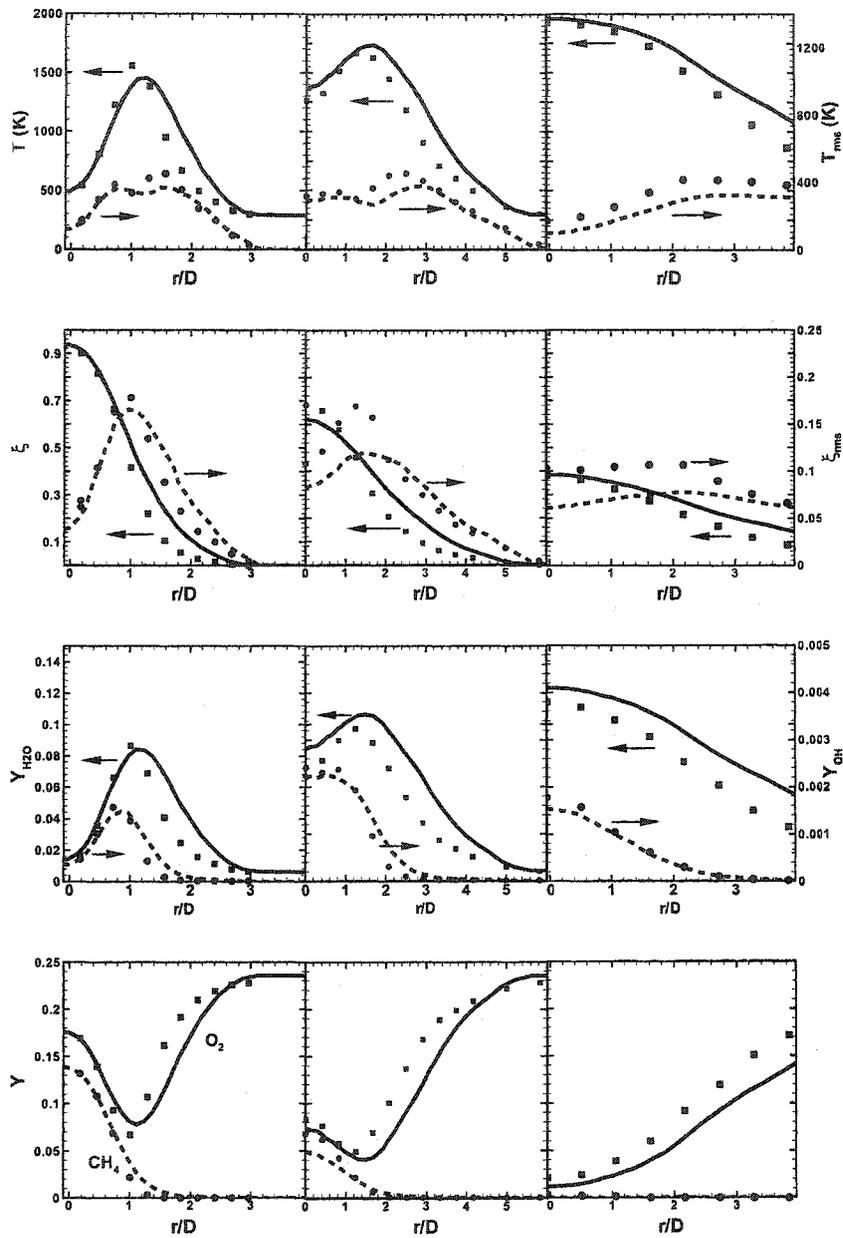


Figure 4.10 Radial profiles at axial positions of (left) $x/d = 15$, (middle) $x/d = 30$ and (right) $x/d = 45$ for temperature, mixture fraction and chemical species for Case A. Symbols are experimental data and lines are simulation results.

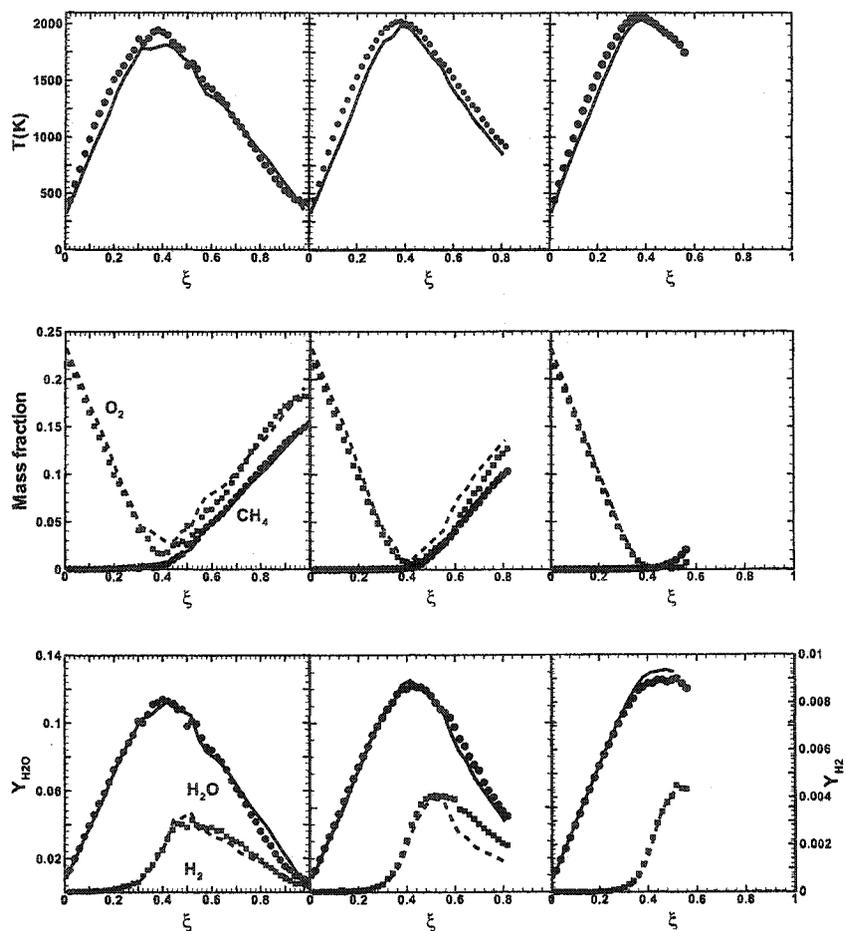


Figure 4.11 Conditional means at axial locations of (left) $x/d = 15$, (middle) $x/d = 30$ and (right) $x/d = 45$ for selected chemical species for Case A. Symbols are experimental data and lines are simulation results.

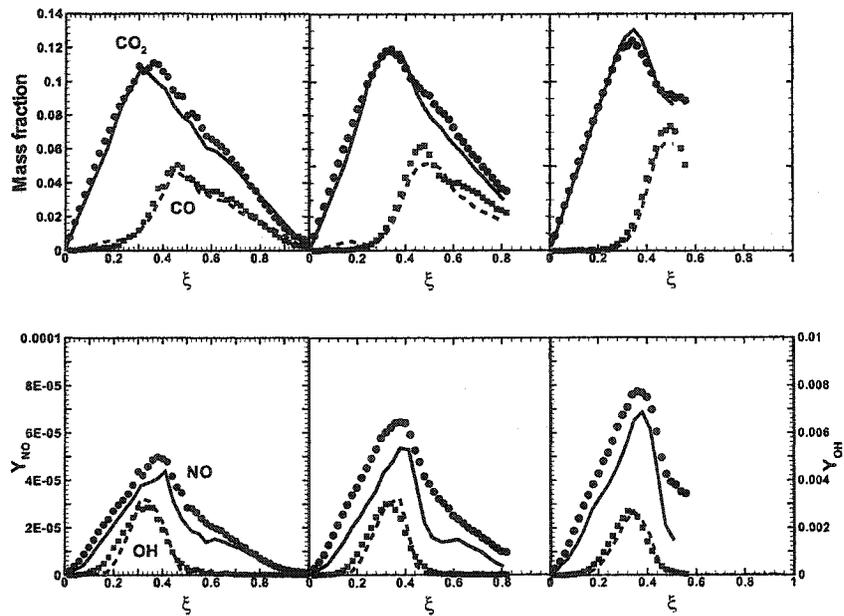


Figure 4.12 Conditional means at axial locations of (left) $x/d = 15$, (middle) $x/d = 30$ and (right) $x/d = 45$ for selected chemical species for Case A. Symbols are experimental data and lines are simulation results.

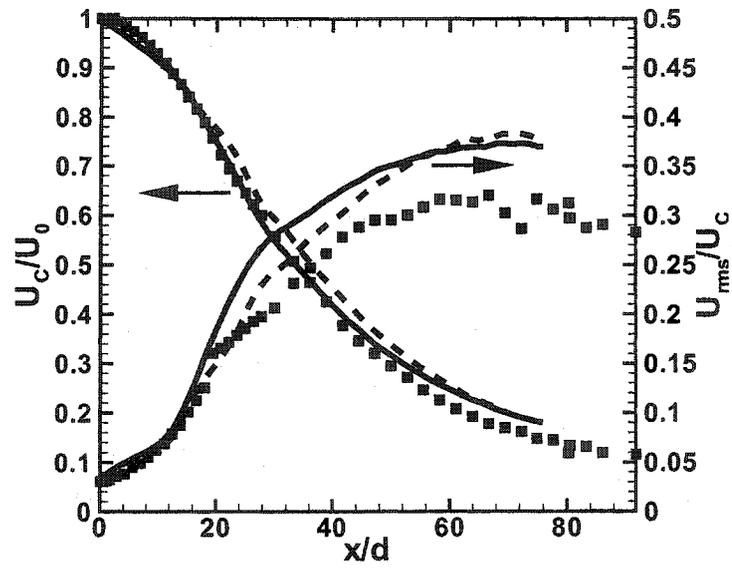


Figure 4.13 Mean and RMS axial velocity for Case A (solid line) and Case B (dashed line) compared with experimental data (symbols).

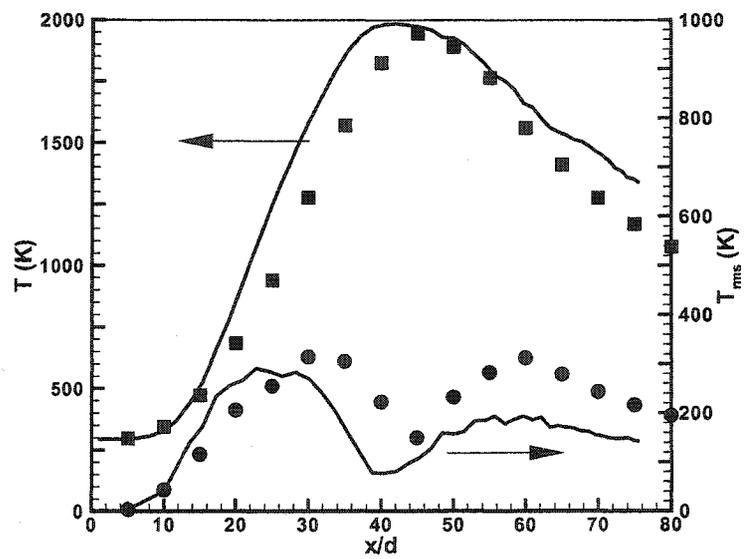


Figure 4.14 Centerline mean and RMS profile for Case B (solid line) compared with experimental data (symbols).

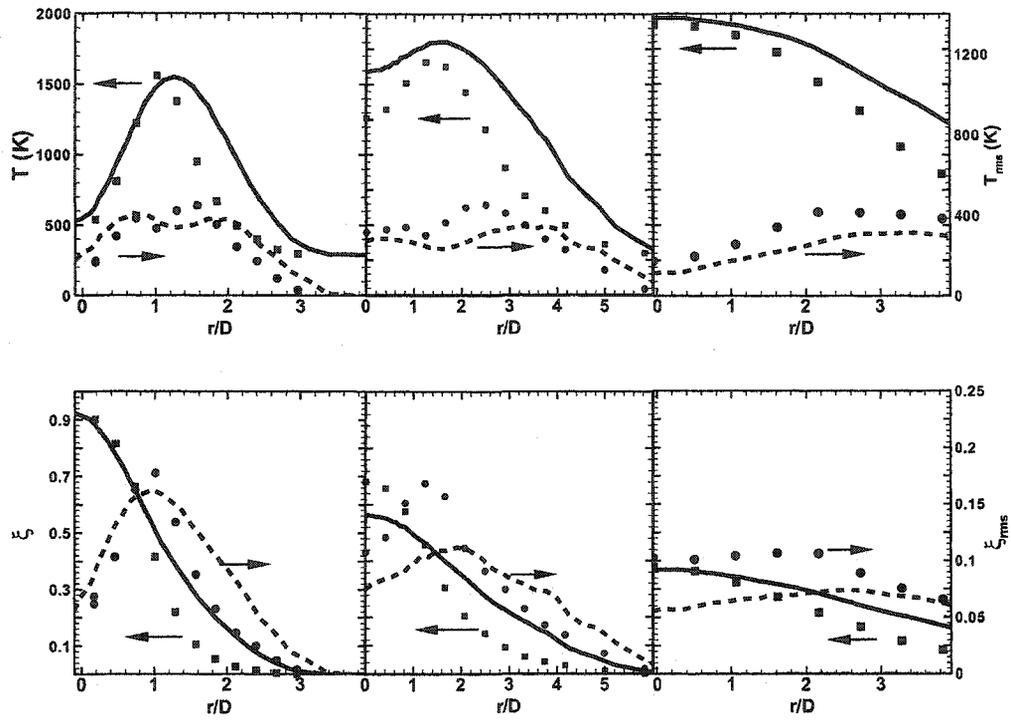


Figure 4.15 Radial mean temperature and mean mixture fraction profiles for Case B at axial positions of (left) $x/d = 15$, (middle) $x/d = 30$, and (right) $x/d = 45$. Lines are simulation results and symbols are experimental data.

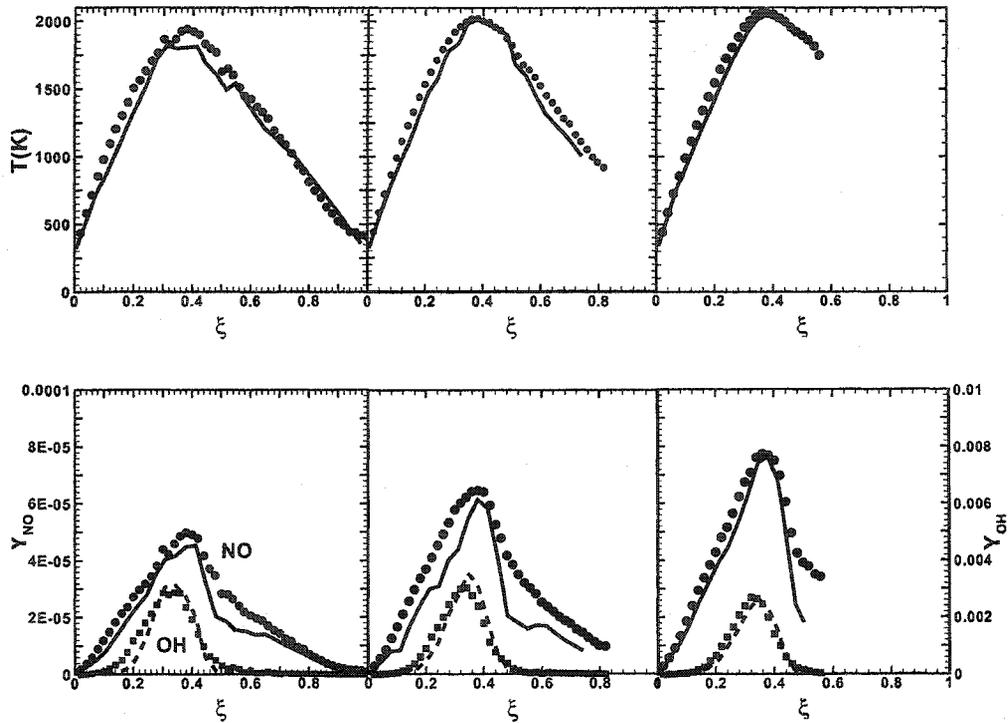


Figure 4.16 Conditional means of temperature and selected chemical species for Case B at axial positions of (left) $x/d = 15$, (middle) $x/d = 30$, and (right) $x/d = 45$. Lines are simulation and symbols are experimental data.

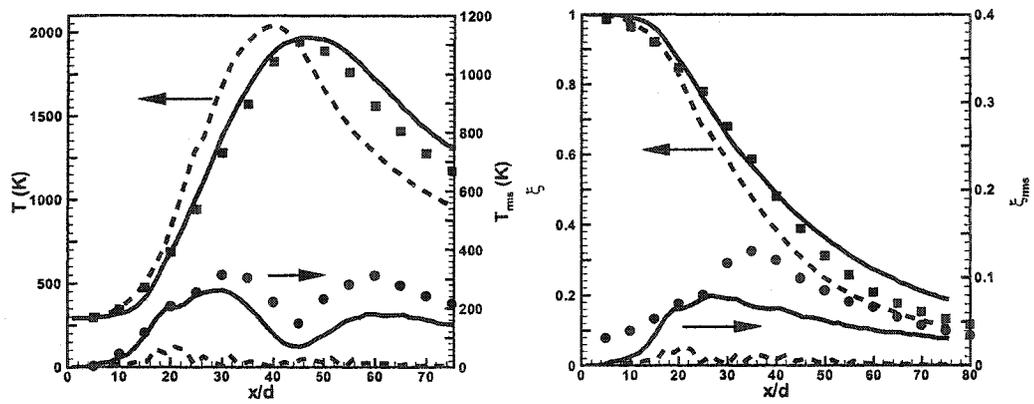


Figure 4.17 (Left) Mean and RMS centerline temperature profile for Case C (solid line) and Case D (dashed line) compared with experimental data (symbols). (Right) Centerline mean and RMS mixture fraction profile for the same cases.

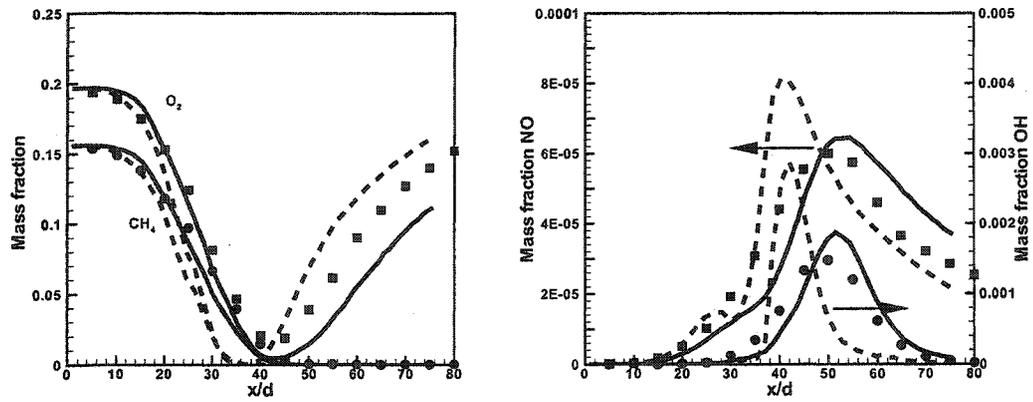


Figure 4.18 (Left) Mean methane and oxygen profile for case C (solid line) and Case D (dashed line) compared with experimental data (symbols). (Right) Centerline NO and OH profile for the same cases.

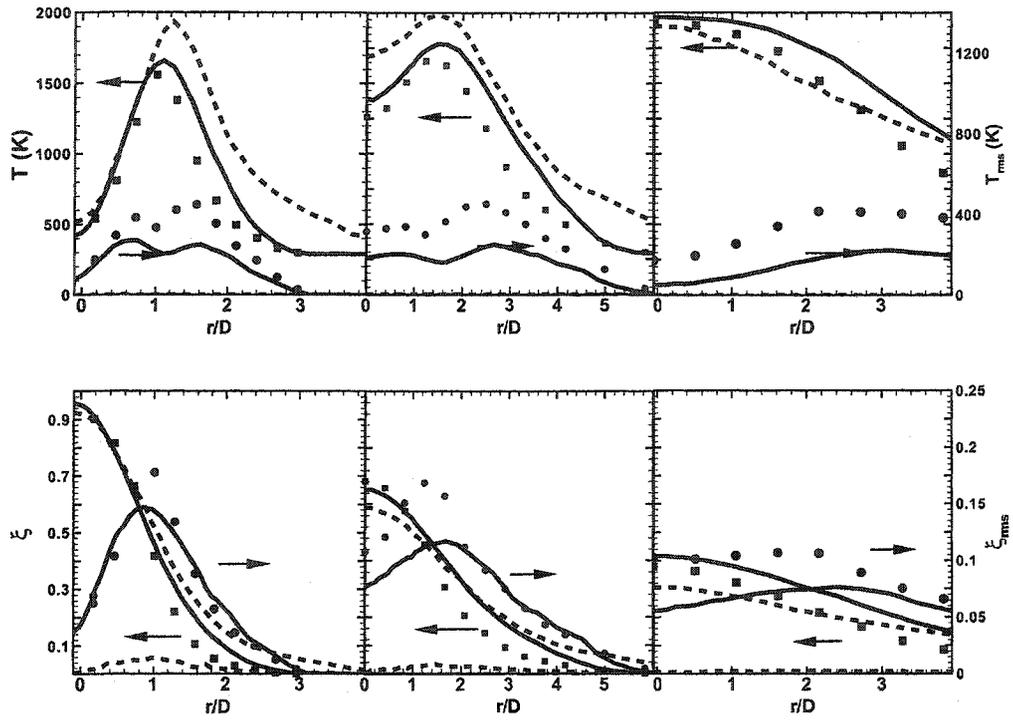


Figure 4.19 Radial profiles of means of temperature and mixture fraction for Case C (solid line) and Case D (dashed line) at axial positions of (left) $x/d = 15$, (middle) $x/d = 30$, and (right) $x/d = 45$. Lines are simulation results and symbols are experimental data.

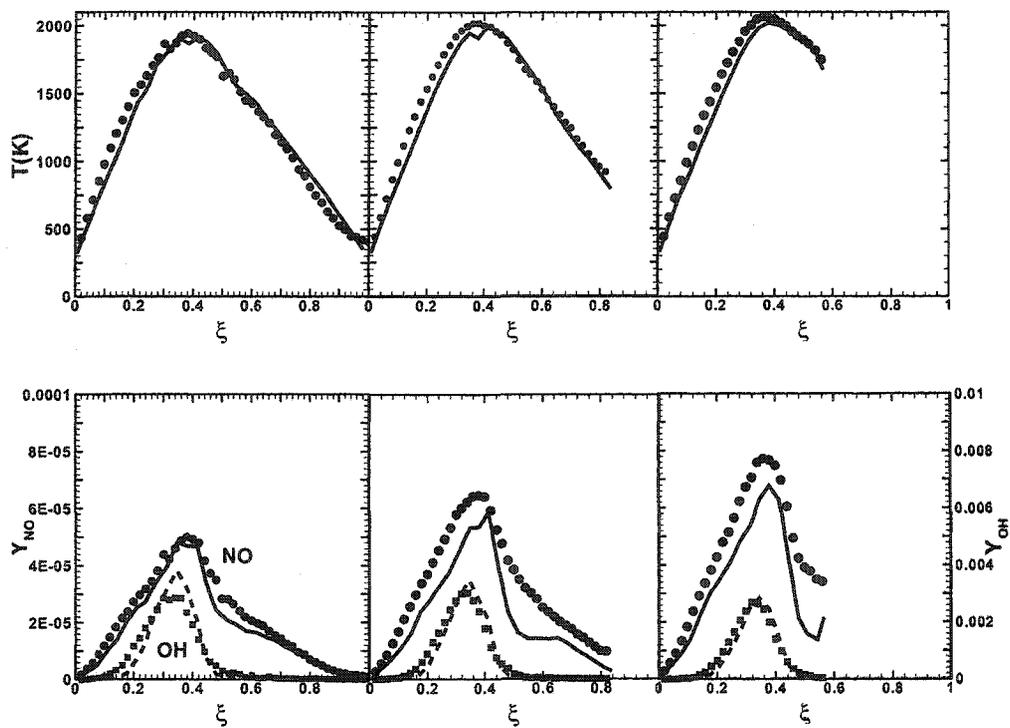


Figure 4.20 Conditional means of temperature and selected chemical species for Case C at axial positions of (left) $x/d = 15$, (middle) $x/d = 30$, and (right) $x/d = 45$. Lines are simulation and symbols are experimental data.

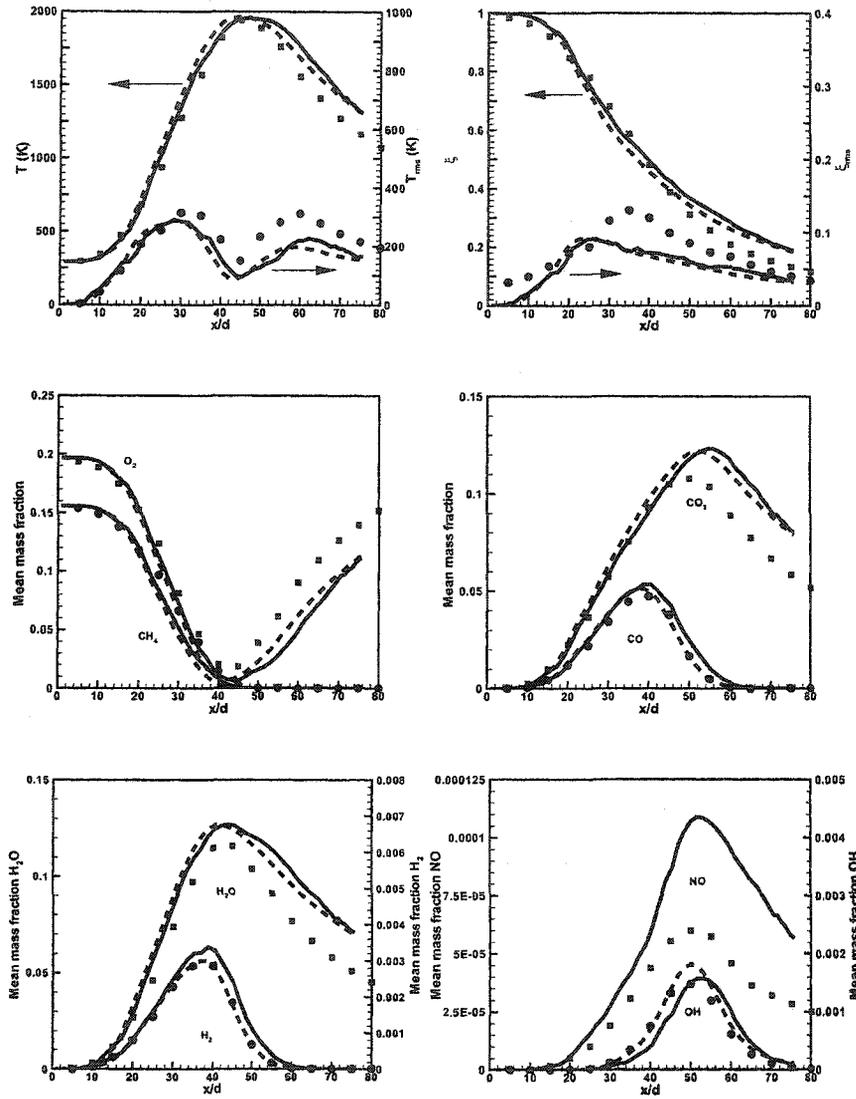


Figure 4.21 Axial profiles of selected chemical species for different mechanisms compared with experimental data. Solid lines are GRI-3.0 mechanism results and dashed lines are ARM results. Symbols are experimental data.

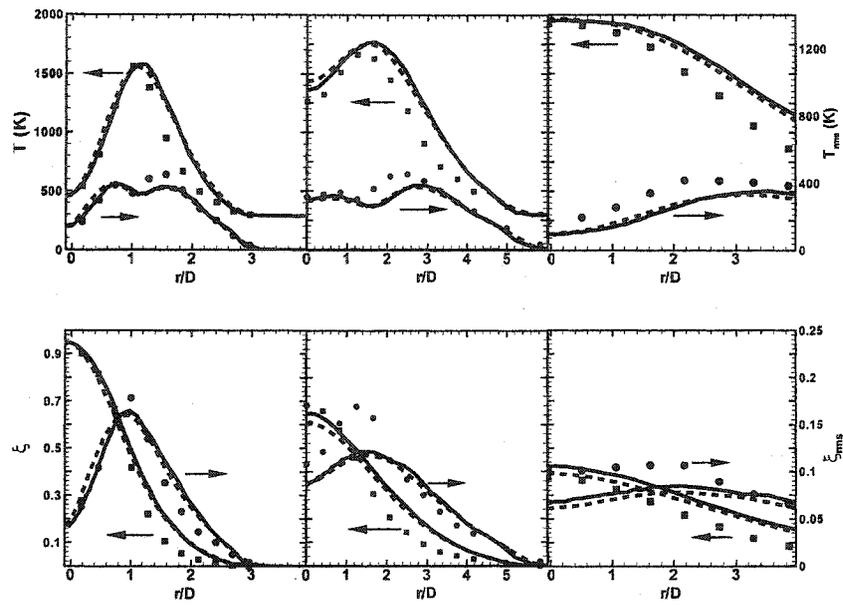


Figure 4.22 Radial profile of mean temperature and mixture fraction for different mechanisms at axial positions of (left) $x/d = 15$, (middle) $x/d = 30$, and (right) $x/d = 45$. Solid lines are GRI-3.0 mechanism (Case E) results and dashed lines are ARM (Case F) results. Symbols are experimental data.

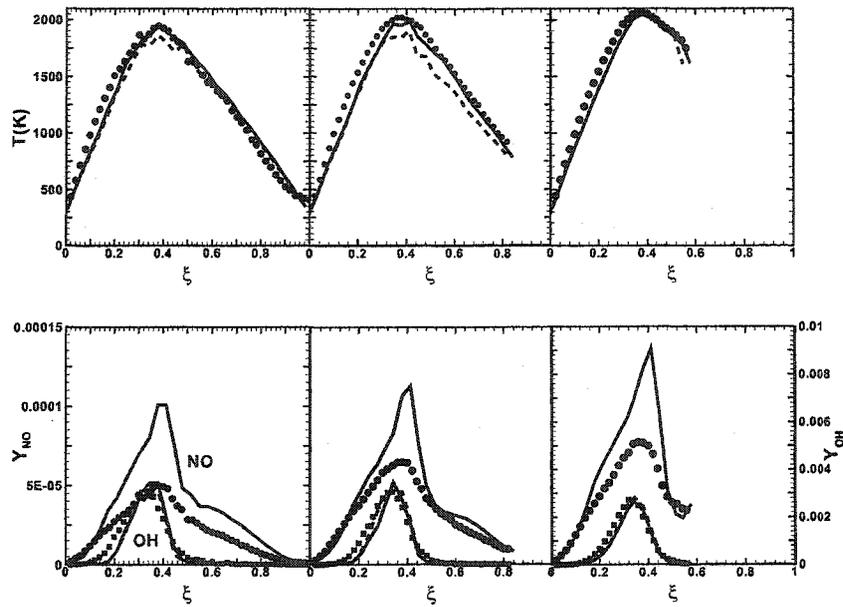


Figure 4.23 Conditional means of temperature and selected chemical species for different mechanisms at axial positions of (left) $x/d = 15$, (middle) $x/d = 30$, and (right) $x/d = 45$. Solid lines are GRI-3.0 mechanism (Case E) results and dashed lines are ARM (Case F) results. Symbols are experimental data.

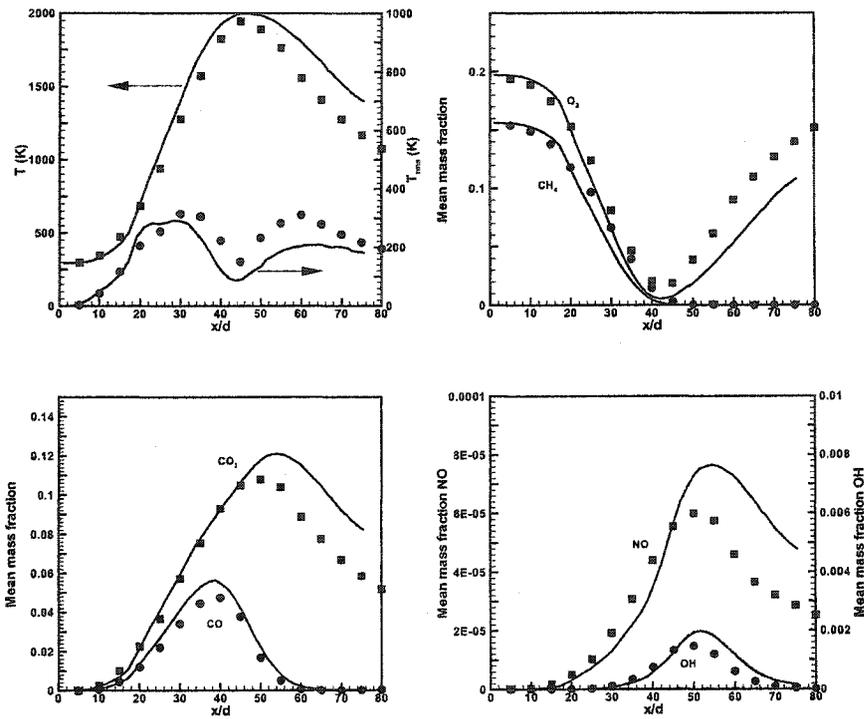


Figure 4.24 Axial profiles of scalars for Case G that neglects radiation. Solid lines are simulation result and symbols are experimental data.

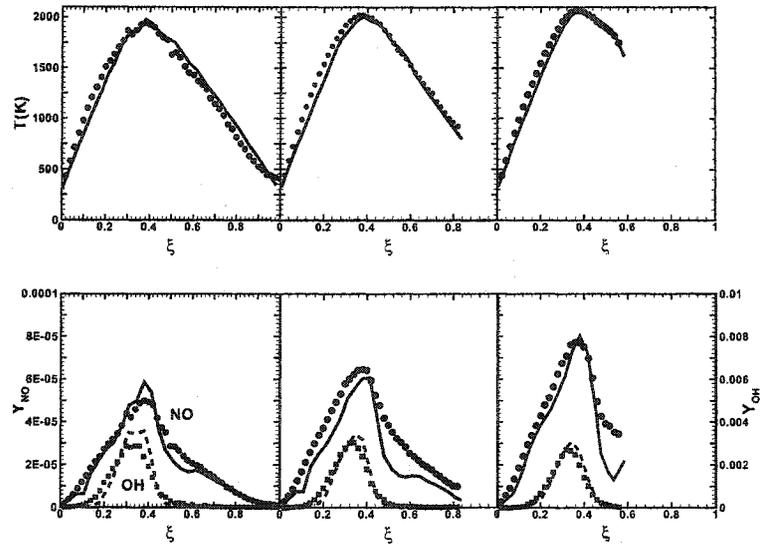


Figure 4.25 Conditional means of temperature and selected chemical species for Case G at axial positions of (left) $x/d = 15$, (middle) $x/d = 30$, and (right) $x/d = 45$. Solid lines are simulation results and symbols are experimental data.

5 Simulation of Gas-phase Chlorination Reactor

Introduction

Gas-phase chlorination of hydrocarbons is a commercially important process used for the production of chlorinated derivatives (e.g., methyl chloride, methyl dichloride, and chloroform) [Wiseman, 1972]. These chlorinated products have a variety of applications as industrial solvents, intermediates and environmentally-friendly refrigerants [Reed, 1991]. In the present study, methyl chloride is used as the primary feed and chloroform is the desired product. Plant-scale reactors typically operate in either premixed or partially premixed, adiabatic mode with the inlet feed jets used to enhance turbulent mixing. The free-radical reactions, however, give rise to a host of secondary products which are difficult to separate. Using methyl chloride as a feed, the maximum yield of desirable products occurs around 700-740 K. At higher temperatures, the primary products rapidly pyrolyze to a wide range of chlorinated alkenes and eventually reactor fouling due to carbon deposition can become a problem. The detailed kinetic scheme [Tirtowidjo, 1997; Shah and Fox, 1999] used in this work involves 38 species and 152 reactions [Tirtowidjo, 1997]. A 21-step reduced scheme with 15 species is also used in the CFD simulations (see Tables. 5.1 and 5.2). The modeling of thermochlorination reactors is difficult because of the strong coupling between the turbulent flow and reaction processes. However, faster computers and more efficient algorithms have made such simulations considerably easier [Deutschmann and Schmidt, 1998; Tsai and Fox, 1996]. While successful CFD simulations of methane chlorination reactors have been reported [Acharya et al., 1991], most studies have been limited to smaller reaction schemes or one-dimensional models [West et al., 1999]. In general, previous studies have concluded that unless a detailed kinetic scheme is employed, CFD does a poor job of predicting finite-rate chemistry effects, minor species formation and reactor extinction. In addition, most CFD simulations have not accounted for micromixing effects which are known to affect reaction rates and hence the reactor performance. Thus, for the accurate predic-

tion of minor species and reactor stability studies, full PDF methods should be employed due to their ability to account for turbulence-chemistry interactions by treating complex reactions without modeling assumptions [Tsai and Fox, 1996; Pope, 1985].

The chapter is organized in such a way that the simulations performed go from a simple zero-dimensional reactor to progressively complex configuration and solution methodology. First, a description of the reactor configuration is given. Then CSTR and PMSR results are discussed. This is followed by a discussion of the constant density simulation of a 2-dimensional chlorination reactor using FV and PDF methods and the effect of the constant-density assumptions. These simulations will prove the need for the Lagrangian PDF scheme to get an accurate description of the reactor dynamics. Then the hybrid-scheme discussed in previous chapters will be used to study different reactor configurations and operating conditions. Finally, as a note of caution on the implicit assumptions of the hybrid scheme, a study on fractional time stepping will be analyzed.

Reactor Configuration

Industrial gas-phase chlorination processes typically use coaxial, jet-stirred reactors similar in geometry to the one represented in Fig. 5.1. This reactor has two inlet streams which feed a fixed ratio of organics to chlorine. Although complete premixing is desired, in industrial reactors the two streams are not completely premixed. The degree of mixing is thus an important aspect of the reactor performance. Here, however, complete premixing is assumed as an ideal case and reactor performance analyzed based on this assumption. (Studies with non-premixed feed streams will be discussed later in this chapter) In this study the premixed inlets have identical inlet velocities of 3.11 m/s which corresponds to a reactor residence time of 7.5 s. The feed streams have 25% chlorine, 36% methyl chloride and 39% methyl dichloride by mass. The inlet temperature is set at 323 K. The feed jet diameters are adjusted to get equal velocities. The flow fields exhibit special characteristics that are critical to sustaining the reaction. The incoming coaxial jets create a flame-like structure that transports the reactants inside the reacting zone. Mixing occurs primarily in the high-shear region surrounding the interface of the primary jets. The penetration of the reactants deep inside the reactor facilitates secondary and tertiary chlorination. The chlorination reaction is highly exothermic and like combustion flames, stays lit only above a critical temperature. Hence the feedback of enthalpy to the incoming cold feed streams is an essential feature of the flow. The rate of heat transfer as well as the location of the reattachment point of the recirculation

Table 5.1 Chemical species.

Primary species	HCl	Cl ₂	CH ₃ Cl
	CH ₂ Cl ₂	CHCl ₃	CCl ₄
Secondary species	CHCl ₂ – CH ₂ Cl	CCl ₃ – CCl ₃	CCl ₃ – CH ₂ Cl
	CCl ₃ – CHCl ₂	CHCl ₂ – CHCl ₂	Cl [•]
	CHCl ₂ [•]	CCl ₃ [•]	CH ₂ Cl [•]

zone are important parameters in determining the stability of the reactor. The accurate determination of the low Reynolds number recirculating flow is critical for obtaining accurate predictions for chemical species distributions inside the reactor. Indeed, it has been found that relatively small variations in the degree of flow recirculation will result in reactor extinction. In order to better understand the role of micromixing, results for simplified reactor models which ignore spatial gradients at large scales are discussed first.

Simplified Reactor Models

Traditionally, a simple continuous stirred tank reactor (CSTR) model and/or a plug-flow reactor (PFR) model with recycle have been used to understand chlorination reactor performance. In fact, under certain operating conditions, these models are quite accurate in providing outlet species concentrations. In these studies, highly-simplified reaction mechanisms (typically involving a chlorine disassociation step) have been used to determine the extent of reaction [West et al., 1999]. In this study detailed chemistry is implemented in two such simple models (namely the CSTR and the pairwise mixing stirred reactor (PMSR) models) and density variations neglected. When comparing with the CFD results presented later, the reader may note that the CSTR model is analogous to a single-particle PMSR model and both the CSTR and PMSR models neglect spatial variations in the mean species concentrations.

The CSTR model for constant density flow is given by

$$\frac{d\phi}{dt} = \frac{1}{\tau}(\phi_{in} - \phi) + \mathbf{S}(\phi) \quad (5.1)$$

where ϕ_{in} is the inlet mass fraction/enthalpy vector and τ is the reactor residence time. In order to be consistent with the CFD simulations, operator splitting is used to separate

Table 5.2 21 step reaction scheme.

No.	Reaction	A_n	E_n
1	$\text{Cl}_2 \rightleftharpoons 2\text{Cl}^\bullet$	0.185E+13	0.580E+5
2	$\text{CH}_2\text{Cl}_2 \rightleftharpoons \text{CH}_2\text{Cl}^\bullet + \text{Cl}^\bullet$	0.1E+17	0.811E+5
3	$\text{CHCl}_2 - \text{CH}_2\text{Cl} \rightleftharpoons \text{CHCl}_2^\bullet + \text{CH}_2\text{Cl}^\bullet$	0.1E+16	0.884E+5
4	$\text{CHCl}_3 \rightleftharpoons \text{CHCl}_2^\bullet + \text{Cl}^\bullet$	0.1E+17	0.776E+5
5	$\text{CHCl}_2 - \text{CHCl}_2 \rightleftharpoons 2\text{CHCl}_2^\bullet$	0.1E+16	0.824E+5
6	$\text{CCl}_3 - \text{CCl}_3 \rightleftharpoons 2\text{CCl}_3^\bullet$	0.1E+16	0.714E+5
7	$\text{CCl}_3 - \text{CH}_2\text{Cl} \rightleftharpoons \text{CCl}_3^\bullet + \text{CH}_2\text{Cl}^\bullet$	0.1E+16	0.827E+5
8	$\text{CCl}_3 - \text{CHCl}_2 \rightleftharpoons \text{CCl}_3^\bullet + \text{CHCl}_2^\bullet$	0.1E+16	0.759E+5
9	$\text{CCl}_4 \rightleftharpoons \text{CCl}_3^\bullet + \text{Cl}^\bullet$	0.1E+17	0.709E+5
10	$\text{CH}_3\text{Cl} + \text{Cl}^\bullet \rightleftharpoons \text{CH}_2\text{Cl}^\bullet + \text{HCl}$	0.1E+14	0.460E+3
11	$\text{CH}_2\text{Cl}_2 + \text{Cl}^\bullet \rightleftharpoons \text{CHCl}_2^\bullet + \text{HCl}$	0.1E+14	0.460E+3
12	$\text{CH}_2\text{Cl}_2 + \text{CH}_2\text{Cl}^\bullet \rightleftharpoons \text{CHCl}_2^\bullet + \text{CH}_3\text{Cl}$	0.3E+12	0.110E+5
13	$\text{CH}_2\text{Cl}_2 + \text{Cl}^\bullet \rightleftharpoons \text{CH}_2\text{Cl}^\bullet + \text{Cl}_2$	0.1E+15	0.224E+5
14	$\text{CHCl}_3 + \text{Cl}^\bullet \rightleftharpoons \text{CCl}_3^\bullet + \text{HCl}$	0.1E+14	0.460E+3
15	$\text{CHCl}_3 + \text{CH}_2\text{Cl}^\bullet \rightleftharpoons \text{CCl}_3^\bullet + \text{CH}_3\text{Cl}$	0.3E+12	0.110E+5
16	$\text{CHCl}_3 + \text{CHCl}_2^\bullet \rightleftharpoons \text{CCl}_3^\bullet + \text{CH}_2\text{Cl}_2$	0.3E+12	0.110E+5
17	$\text{CHCl}_3 + \text{Cl}^\bullet \rightleftharpoons \text{CHCl}_2^\bullet + \text{Cl}_2$	0.1E+15	0.189E+5
18	$\text{CHCl}_3 + \text{CH}_2\text{Cl}^\bullet \rightleftharpoons \text{CHCl}_2^\bullet + \text{CH}_2\text{Cl}_2$	0.3E+12	0.110E+5
19	$\text{CHCl}_2 - \text{CHCl}_2 + \text{Cl}^\bullet \rightleftharpoons \text{CHCl}_2 - \text{CCl}_2^\bullet + \text{HCl}$	0.1E+14	0.460E+3
18	$\text{CCl}_4 + \text{Cl}^\bullet \rightleftharpoons \text{CCl}_3^\bullet + \text{Cl}_2$	0.1E+15	0.122E+5
20	$\text{CCl}_4 + \text{CH}_2\text{Cl}^\bullet \rightleftharpoons \text{CCl}_3^\bullet + \text{CH}_2\text{Cl}_2$	0.3E+12	0.110E+5
21	$\text{CCl}_4 + \text{CHCl}_2^\bullet \rightleftharpoons \text{CCl}_3^\bullet + \text{CHCl}_3$	0.3E+12	0.110E+5

the reaction term from the inflow-outflow term in Eq. 5.1. The residence time and inlet conditions of the CSTR were first set to be the same as for the coaxial jet-stirred reactor ($\tau = 7.5$ s, $T_{in} = 323$ K). However, for this inlet temperature, increasing the residence time had no effect on reactor performance. (In fact, the inlet temperature $T_{in} = 323$ K is too low to sustain reactions in a CSTR for this residence time.) The inlet temperature was thus increased, and it was found that at $T_{in} = 365$ K (Fig. 5.2) the reactor ignites and remains lit. Using this inlet temperature, increasing the residence time yielded similar results with an ignition at $\tau = 3.85$ s (Fig. 5.3). Although the assumption of complete backmixing is far from reality, these results illustrate the effect of the residence time on reactor stability. In the full-scale reactor, fluid particles have long sojourn times in the recirculation region, and short sojourn times in the entrance region. Whether these regions will remain lit will depend on the extent of backmixing and the local micromixing times. However, pilot-scale reactors with $\tau = 7.5$ s have been found to remain stable at an inlet temperature of $T_{in} = 323$ K. Thus, it can be concluded that local variations in the extent of mixing will affect reactor dynamics, and these effects must be studied using more detailed flow models.

The PMSR model [Pope, 1997] is an extended version of the CSTR model with a finite number of fluid particles that evolve in composition space by reaction and micromixing. The particles are paired and the micromixing model is applied between the partners in every pair. The particles exchange partners based on a fixed pairing time and relax to the mean composition (of the partners) based on a pre-determined micromixing time scale τ_m . The PMSR model reduces to a CSTR as the micromixing and pairing time approach zero. (Further details on the PMSR model can be found elsewhere [Pope, 1997; Shah and Fox, 1999].) For the jet-stirred reactor, the “average” micromixing time (volumetric average of k/ϵ) was found from the turbulence model to be around $\tau_m = 0.558$ s. The pairing time was (arbitrarily) set equal to the micromixing time. With the residence time set to $\tau = 7.5$ s and $T_{in} = 323$ K, it was found that the PMSR model yields a stable reacting steady state. The assumption of finite-rate micromixing (as opposed to infinite-rate micromixing) thus suffices to change the behavior of the reactor. The micromixing time was then varied and it was found that the reacting steady-state temperature decreased with decreasing micromixing time (Fig. 5.4) until the reactor eventually extinguishes. Although the PMSR model does predict a reacting steady state at the desired inlet temperature and residence time, it cannot predict changes in reactor stability due to variations in the flow field (e.g., changes in local shear rate due to variations of the feed jet radius). In order to address such questions, it is necessary

Table 5.3 ISAT performance statistics.

Method	Error tol.	CPU (min)	No. of records	No. of DI	Outlet Temp (K)
DI	10^{-8}	78	-	240,000	718
ISAT	10^{-3}	25	8056	23,453	719
ISAT	10^{-2}	10	2002	11,244	719.5
ISAT	10^{-1}	3.5	336	3,179	716

to predict both the detailed flow field and species concentration fields. This can be done using CFD simulations of the jet-stirred reactor with detailed chemistry.

Constant-density PDF Simulation of Premixed-feed Reactor

In this section, CFD simulations of the jet stirred chlorination reactor are done with the Lagrangian PDF method. The composition PDF code requires a stationary flow field obtained from a CFD turbulence model. Here, the flow is assumed to be at steady state and incompressible and hence the feedback portion of the hybrid code is not used. (Effects of density variations on reactor performance will be discussed in the next section.) The velocity, pressure, and turbulent kinetic energy fields are pre-computed using the CHEM3D code [Harvey and Edwards, 1999]. The flow is solved for an axisymmetric case using a $k-\omega$ turbulence model. The cell-centered velocity, turbulent kinetic energy and energy dissipation fields are input to the composition PDF code. From these fields the turbulent diffusivity and the local mixing times are computed (Fig. 5.5). In the composition PDF codes, the flow domain is assumed to contain a finite number of stochastic particles, which move in physical space using the flow profiles from the turbulence models and evolve in composition space by reaction and micromixing. Usually, the particles evolve over 3–4 residence times to obtain steady-state profiles. The micromixing in each cell is accounted for by a simple interaction by exchange with the mean (IEM) model [Villermaux, 1986]. This model assumes a linear rate for micromixing that is proportional to the distance between the mean composition vector in a cell and the particle composition vector. The micromixing rate involves a time constant which is taken to be proportional to the turbulence time scale τ_u .

Table 5.4 PDF simulation statistics. Wall time refers to real time (as if measured by a wall clock) required for the simulation.

Method	No. of CPUs	Source calls per Δt	Wall time	C_ϕ	No. of records	n_s
DI	1	560	23.5 hrs	∞	-	15
ISAT	1	560	7.4 hrs	∞	707	15
ISAT	1	25,000	18.1 hrs	1	930	15
ISAT	1	25,000	4 days	1	2,106	38

Simulation Conditions

A fixed time step is used with fractional time stepping to facilitate the use of ISAT to calculate the reaction source term. On average 80–120 particles per grid cell are employed on a 40×28 axisymmetric, Cartesian grid. The grid has non-uniform cells with denser grid points near the entrance region to capture the interface of the two feed jets. The simulations were performed on a 266 MHz Alpha workstation. Simulation of the initial transient regime requires a large number of records in the ISAT binary tree and consequently the storage memory requirement is relatively large. Individual studies with a PMSR model were carried out to test the validity of the parameters used in ISAT (see Table 5.3). It was found that even at a error tolerance value of $\varepsilon_{tol} = 10^{-1}$, the results from interpolation are quite accurate. The use of ISAT was found to yield significant speedups (see Table 5.4). It was also found that the speedup is non-linear with respect to the number of chemical source term calls. The use of the detailed 38-species mechanism further slowed down the simulation as compared to the reduced 15-species mechanism.

Simulation Results

Three different simulation conditions were tested. Two of these used the detailed 38-species mechanism, while the third used the reduced 15-species mechanism. The mixing time was found to be low near the entrance region due to higher velocities (Fig. 5.5). In the detailed chemistry simulations, the micromixing parameter (C_ϕ) was varied to analyze the effect of micromixing on reaction. Large C_ϕ forces the composition of the particles inside a grid cell to the mean value, thereby approximating a finite-volume code which neglects micromixing. (Note, however, that the spatial transport algorithm used in the Lagrangian PDF code is more accurate for the same grid resolution than the

method used in the CFD code [Mobus et al., 2001].) It was found that (Fig. 5.6) the use of a micromixing model results in a higher outlet temperature (714 K) as compared to a completely micromixed reactor (703 K). This showed that the implementation of even a simple micromixing model makes a definite difference in the reactor performance predictions. It can also be seen that although the two different mechanisms (detailed and reduced chemistry) yield similar temperature and species profiles, reduced chemistry produces a higher temperature in the diffuse reaction zone. Nevertheless, the predicted reactor outlet temperature and species concentrations agree closely, although simulation with detailed chemistry resulted in a slightly more diffuse reaction zone. The computational time for detailed chemistry was approximately five times higher than that for reduced chemistry. In both cases the outlet temperature is around 715 K, and the concentration of chloroform in the product is between 12–14% by mass. The temperature and species profiles (Fig. 5.7) show the presence of a characteristic reaction zone that surrounds the inlet jet. The chlorine profile and its penetration into the reactor determines the nature and extent of the reaction. Rapid depletion of chlorine can lead to regions of high temperature that promote disassociation of chloroform and methyl dichloride. However, a higher chlorine content near the outlet of the reactor can result in further chlorination of the desired products, resulting in increased production of carbon tetrachloride and other secondary products. Maximum chloroform yield is known to occur around 700–740 K, and the premixed jet-stirred reactor configuration is quite suitable for obtaining this temperature range.

These simulations clearly show that Lagrangian composition PDF method is able to predict outlet temperatures and species distributions with reasonable accuracy when compared to data from industrial chlorination reactors. The reduced mechanism captures the essential dynamics of the reactor and agrees closely with results obtained using detailed chemistry. Hence the reduced chemistry is useful in estimating reactor performance. The reactor flow configuration is critical in sustaining the reactor. The proper choice of inlet velocities is necessary to keep the recirculation large enough to increase enthalpy back-flow, but at the same time avoid forming a large dead zone that might result in either a reaction runaway or extinction due to the presence of stagnant reacting fluid that could be considered more like a batch reactor. Although not studied in detail in this work, the interface between the jets improves mixing and hence the effective ratio of the jet velocities is an important factor in determining the mixing zone. It is evident that the reaction zone itself is situated in this high-shear layer. The jet velocity ratio of 1.0 used in this work results in a restricted mixing zone as well as a limited reaction zone

that is characterized by a sharp temperature gradient. The jet velocities also depend on the inlet pipe diameters which indirectly affect the positioning of the mixing zone since the interface of the jets themselves is moved by changing the pipe diameter. The concentration of methyl dichloride was found to increase locally near the mixing zone indicating that most of the primary and secondary chlorination occurs in the mixing zone. Again these reactions are highly exothermic and the presence of the backmixing zone helps carry back the enthalpy to the inlet jets. The flow profiles compared with the final species profiles show the distinct tendency of the reaction to follow the fluid flow. This hints at the use of flow to control the reaction.

Industrial chlorination reactors have been found to exhibit non-linear behavior like extinction, re-ignition and sustained oscillations which are not predicted by this study. This is probably due to the assumption of constant-density flow in an overall exothermic system. To understand the effect of density, the results from the finite-volume code simulations with the same chemistry mechanism but with no sub-grid closure was used. These calculations indicate an interesting phenomena. If flow and reaction are decoupled for a theoretical explanation, it has been found that the highly exothermic reaction causes a large density gradient which causes the flow to accelerate. However, this acceleration causes the reactants inside to be pushed out thereby decreasing reaction and consequently the temperature and density gradients. This slows down the flow, thereby forming a closed negative feedback loop that leads to a steady state. Two different cases were studied with the FV code. In the first case the flow was non-reacting and only the mean flow quantities were evolved. In the second case, the flow was reacting with a 15-species kinetic mechanism, but neglected the concentration fluctuations. The center-line axial velocity profile (Fig. 5.8) for the non-reacting case closely follows the reacting case profile. The density profiles have been added to emphasize the large gradients in the reactor. The difference in profiles can be seen from the contour plots of velocity (Figs. 5.9 and 5.10). It can be seen that the length of the flame or the penetration distance is higher for the reacting flow. However, the structure of the flow remains the same with a large recirculation zone near the entrance and nearly plug flow zone at the end. The constant density case thus provides the starting point for simulating turbulent flow reactors with complex chemistry.

These results clearly show the applicability of the Lagrangian PDF method to the chlorination reactors. However, the use of a hybrid-method is essential to be assured that the simulations are accurate. The following sections will detail the hybrid-simulation of the chlorination reactor in the context of feed-stream configuration effects on reactor

behavior.

Effect of Feed-stream Configuration

Simulation Conditions

In this section, the hybrid-scheme is employed with density correction. The coaxial jet-stirred configuration (Fig. 5.1) was found to provide good mixing of the reactants with a high shear layer at the interface of the jets. This reactor is used to study the effect of feed-stream mixing on the exothermic chlorination reactor. Methyl chloride and methyl dichloride are grouped as organics and are considered separated from the chlorine stream. As in the previous section, the reactor has two coaxial inlets. The two streams can be premixed or passed separately in each of the inlets. This leads to three different configurations that are analyzed. For all the simulations, the inlet organic stream has a composition of 48% methyl chloride and 52% methyl dichloride with a mass flow rate of 0.02844 kg/s. The inlet chlorine stream has a flow rate of 0.00936 kg/s. In the premixed case, the reactants are assumed to be completely mixed at the above mass flow ratios before entering the reactor. In this work, the jet diameters are adjusted such that the inlet velocity of the streams are equal to 3.01 m/s for both premixed and non-premixed cases. The inlet temperature of both streams is set to 323 K. Although pilot-scale experimental studies show reactor stability, non-ideal CSTR studies with the same residence time (≈ 7.5 s) and inlet temperature show a high probability for reactor extinction. The coaxial reactor has a characteristic recirculation zone between the inlet and the wall that provides near complete mixing. Also the recirculation zone helps to transport back enthalpy to the incoming cold fluid to sustain the reaction. (In the limit of plug flow, it can be easily shown that the reactor will not stay lit.)

Orthogonal uniform grids were used for the premixed case. A grid-independence study was carried out using a 53×33 mesh and a 53×79 mesh. Radial temperature profiles at different axial positions for the two cases using infinite rate mixing condition are shown in Fig. 5.11. The time-averaged fields for different grid resolutions show very little difference. The smaller grid size was chosen for all cases reported in this paper. The non-premixed cases are found to have steep temperature gradients and hence grid clustering was necessary near the inlets. However, the overall mesh size was retained for the sake of comparison.

The reduced mechanism with 15 species and 21 reactions (Table 5.2) was used in these

simulations. It has been noted before that the use of a larger 38-species mechanism does not fundamentally alter the behavior of the reactor, but offers some quantitative differences. Since the thrust of these simulations are to study the effect of inlet configuration, the reduced mechanism was found to be sufficient. It should be noted that the simulations could be repeated for the detailed chemistry mechanism without modifying the hybrid FV, PDF code. The only appreciable difference would be the proportion of CPU time needed by ISAT to treat the chemistry.

For the non-premixed case with chlorine in the inner stream, the jet diameters are the same as that for the premixed case. In the other non-premixed case where the organics with a higher mass flow rate are in the inner stream, the inner-jet diameter was increased to 0.034 m and the outer-jet diameter was kept the same (0.04 m). For each of these cases two conditions were studied. In the first (finite-rate micromixing), the IEM model with the standard value of the mixing constant ($C_\phi = 1.0$) was used based on the local turbulence field. In the other (infinite-rate micromixing), complete micromixing was assumed ($C_\phi = \infty$). The latter is equivalent to a FV solution with no sub-grid closure for micromixing. The overall time step was set based on the minimum micromixing time for a non-reacting flow field. Chemical reactions usually lead to an increased rate of micromixing (due to changes in the flow field) and a consequent decrease in the variance of the scalars. However, it was noticed that the change in the minimum micromixing time was quite small, thereby validating the choice of the time step. In the finite-rate micromixing cases, the time step was set to 2×10^{-3} s. For studies with infinite-rate micromixing, the time step was increased to 1×10^{-2} s. Due to the absence of micromixing, the time step for this case is controlled by the reaction and the error caused by fractional time stepping. It was found that smaller time steps yielded no differences in the solution. It will be noted in the last section that the simulation of the infinite-rate mixing case exhibits multiple-steady states in part due to the fractional time-stepping algorithm.

Simulation Results

The jet-stirred reactor simulations were carried out on a single processor on a 900 MHz SUNFIRE machine. It was found that reaction source term calculations took nearly 93% for the computational time. To minimize CPU time, the simulations were carried out in three stages. Initially, a non-reacting flow field was generated and the PDF code was run with complete micromixing ($C_\phi \gg 1$) with no temperature feedback. Using the converged scalar fields, feedback was initiated and the flow field converged.

In the final stage, the micromixing model was turned on ($C_\phi = 2$) and the simulation continued for 1.5-2.0 residence times. It was found that this strategy provides a stable solution with an approximately 45% reduction in computational time as compared to starting the simulation with all physical processes activated. The streamlines for the three configurations are shown in Fig. 5.12. The effect of recirculation on flow will be discussed for each inlet configuration below.

For the non-premixed case, the scatter plots of scalars value conditioned on the mixture fraction are used for analysis. The mixture-fraction definition used is

$$\xi = \frac{Y_C - Y_{C_1}}{Y_{C_2} - Y_{C_1}} \quad (5.2)$$

where Y_C refers to the mass fraction of elemental carbon, and the subscripts refer to the mass fraction of carbon in the two inlet streams. In the case of a non-premixed inlet, a value of 0.1753 is used for the organic stream carbon content (Y_{C_2}) and the pure chlorine stream (Y_{C_1}) has a carbon mass fraction of 0.0. This normalizes the mixture-fraction definition to vary from 1.0 for the organic stream to 0.0 for the pure chlorine stream.

Premixed Case

For this case, both inlet streams contain 25% chlorine, 36% methyl chloride and 39% methyl dichloride (% based on total mass). The profile of the temperature (Fig. 5.13a) shows a flame-like structure with a maximum temperature of 752 K. Since chlorination involves formation of free chlorine atoms, the penetration of chlorine (Fig. 5.13b) inside the reactor to a large degree determines the efficiency of the reactor. The residual chlorine free radical shows the region of reaction to be situated in the thin diffusion layer surrounding the jet. The chlorine molecules disassociate due to the higher temperature in this region. There is a competitive reaction between carbon tetrachloride and chloroform with the concentration profiles following the temperature profile (Fig. 5.13c). The formation of organic free radicals leads to quenching in the outer layers of the reaction zone. This contains product contamination by tertiary chlorination. The highly exothermic nature of the reactions leads to a large temperature rise that gets carried back into the recirculation zone (Fig. 5.12a). The increase in temperature also tends to push the gases out of the reactor. This reduces the effective residence time of the reactants in the main stream thereby reducing reaction and bringing down the temperature. Eventually an equilibrium is reached where chlorine conversion is complete and the enthalpy increase is kept in check by expanding gases that increase the mean flow velocity.

It has been observed in the previous section that even when the temperature feedback is absent, the structure of the flame can be predicted quite accurately. The constant-density simulation done previously produces very similar temperature and product profiles. However, the feedback to the FV solver changes the re-attachment point and thereby affects the volume of fluid trapped in the recirculation region. Longer re-attachment lengths may lead to dilution of enthalpy and result in global extinction. In essence, the stability of the reactor is controlled by the residence time in the recirculation region rather than the global flow rate. The premixed case is stable to the point that changes in re-attachment length are compensated by reaction dilution or increase to maintain the ignited state of the reactor.

The case with infinite-rate micromixing does not significantly change the product distribution (Fig. 5.13d) and shows a similar flame structure. It should be noted that the chlorine is not completely consumed and residual fraction of the order of 10^{-4} is present at the outlet. This explains the decrease in outlet temperature to about 720 K (Fig. 5.14). The chloroform yield is only slightly different mainly arising from decreased chlorination in the reaction layer (Figs. 5.13d and 5.14). The excess chlorine free radical is now limited to the region of interaction with the recirculation zone. The end of the flame has no free radicals showing complete consumption of disassociated chlorine. Also the reaction layer is very diffuse and extends deeper into the recirculation region. Chloroform yield is uniform across the outlet at 14% by mass. This confirms the results observed in the previous section about the effect of density on reaction. For the premixed-case, the feedback loop has not changed the behavior of the reactor.

Chlorine in the Inner Jet

In this case, the reactants are separated and the chlorine stream is fed through the inner jet. The temperature profile (Fig. 5.15a) shows that the reaction zone has now been shifted to the end of the jet core. This is quite evident considering the fact that a high degree of mixing is required to get sustainable reaction. The chlorine profile (Fig. 5.15b) follows that of a simple jet but is shortened due to the high temperature gradients arising from reaction. The end of the jet has almost completely mixed particles that yield temperatures as high as 1500 K. The reaction zone, which is detached from the jet core, has a "two-pronged" structure due to the entrainment of the outer cold streams in the recirculation zone. It can be noted that the peak in the chloroform mass fraction is found downstream of the peak in the temperature. This is mainly due to the higher decomposition rates that increase with temperature. One of the distinct

differences from the premixed case is that the recirculation zone does not carry all the enthalpy generated from the reaction. The reaction layer is located at the end of the entrainment zone leading to limited heat recirculation (Fig. 5.12b). In fact this acts as the stabilizing factor preventing excessive reaction in the jet core that might lead to reactor run-away. Movement of the reaction further downstream will cut off enthalpy feedback leading to reactor extinction. Hence this configuration is inherently unstable and minor disturbances in inlet stream properties may lead to extinction. The reaction layer itself is drawn out axially due to the expansion caused by the temperature rise. The higher temperatures and consequent product decomposition decreases chloroform yield to 10% (Fig. 5.15c).

The carbon tetrachloride yield is also restricted to around 1.2%. In the reaction zone, the concentration of chlorine free radical is very high (more than six orders of magnitude greater than the premixed case) and leads to the formation of carbon-carbon bonds with tertiary chlorinated derivatives. It should be noted that this free radical is the excess fraction after consumption in primary chlorination reactions. It indicates the scarcity of organic free radicals that can consume chlorine atoms. The higher temperature leads to further decomposition of chlorine molecules and a relatively high mass fraction (17%) of hydrogen chloride is formed. The outer zone of the reaction layer also carries a relatively high fraction of organic species which causes the carbon-carbon bond formation to be favored.

The conditional scalar scatter plots (Fig. 5.16) based on the mixture fraction (Eq. 5.2) indicate an interesting feature of the chemistry. The reactions are essentially separated into two branches. The fast reaction rates ensure that the scalars are near equilibrium in the reaction zone. In the zone near the inlet (at $x/r = 1$), mixing is limited and the reaction is limited to particles with high chlorine content. As the jets mix more downstream, higher temperatures (partially due to the feedback enthalpy) lead to local ignition and almost all the particles reach the equilibrium curve. It can also be seen that certain particles in the carbon rich zone fall below this curve suggesting that the chlorine levels in these particles are not high enough to cause ignition. They are still controlled by the mixing process and will ignite as they move downstream. The chloroform production is at a peak and the maximum yield is around 20% by mass. It can be seen that the maximum temperatures are found around a mixture fraction of 0.25 which is close to the stoichiometric mixture fraction for maximum chloroform conversion. The chloroform peak lags slightly behind temperature with the maximum observed in two separate branches in a chlorine rich and carbon rich zone, respectively. At the end of the reaction

zone, the chlorine fraction is completely consumed. It should be noted that due to the ignited state of the reactor, the particle temperatures collapse onto the reacting branch. The chloroform fraction is reduced through secondary reactions to its outlet value.

The infinite-rate micromixing (Fig. 5.17) shows similar behavior to that of the premixed case. The diffusive nature of the infinite-rate micromixing case combined with the shear layer at the interface of the jets lead to a nearly premixed condition and the rest of the reactor behaves in the same way as the premixed reactor. However, the recirculation zone is cooler than the premixed case due to less chlorine content in the feedback stream. The organic layer shields the high chlorine jet from entering the dead zone and keeps the chloroform yield at almost the same level as the premixed case (Figs. 5.15d and 5.17). Even though faster mixing leads to higher reaction rates, the rate of dilution of enthalpy prevails balancing any exothermic release of heat with a decrease in temperature through mixing with colder inlet streams. Comparing Figs. 5.15c and 5.15d, it is evident that this configuration provides an excellent example of the need for micromixing closures in reacting flow simulations.

Chlorine in the Outer Jet

In this case, the organics stream is in the inner jet and chlorine in the outer jet. The jet diameters have been changed to ensure that inlet velocities are the same as the previous cases. The presence of a pure chlorine stream near the recirculation region changes the zone of reaction. The reactions are now confined to the recirculation region and the jet interface that interacts with this zone (Fig. 5.18a). The chlorine stream (Fig. 5.18b) is partially entrained into the recirculation flow and hence creates a "bifurcated" front. Also complete chlorine consumption is observed within the first 10% of the reactor length. Such short penetration lengths have an adverse effect of tertiary chlorination that produces carbon tetrachloride. The depletion of chlorine leads to reaction quenching through carbon-carbon bond formation. The position of the reaction zone causes much of the organic stream to pass through the reactor without reaction (Fig. 5.12c). Even though high temperatures are observed in the reactor, conversion of primary chlorinated derivatives is very low. The chloroform yield (Fig. 5.18c) is around 4%, which is the lowest among the three configurations. The strong mixing conditions present near the chlorine-organics interface keep the reactor lit. The recirculation zone will cause a positive feedback loop with increasing temperatures. However, higher temperatures will lead to faster depletion of chlorine and a reduction in the reaction layer thickness. This will reduce heat release into the recirculation zone and will reduce re-

action rates. Such a mechanism is inherently unstable with slight changes in mixing condition between streams leading to reactor run-away. It should also be noted that the mechanism used here does not account for soot formation. In pilot scale reactors, at such high temperatures deposition of carbonic soot on the reactor surface has been observed.

There is a high level of residual chlorine free radical that is produced at the interface of the chlorine stream and the recirculating flow. However, the maximum organic free radical is produced slightly upstream leading to high reaction rates as these free radicals move into the chlorine stream. The resulting temperatures propel the reactions towards the high temperature branch locally. However, such strong reactions are confined to the recirculation zone and the rest of the reactor is extinct as there is very little transfer of enthalpy and chlorine molecules to sustain reaction. The reaction zone is highly localized and the reactor state cannot be observed from the centerline species or temperature profiles, the latter being the only data available from a plant scale reactor. Hence this reactor constitutes a risk in an industrial scenario.

The scatter plots (Fig. 5.19) show the existence of a distinct reaction branch and a mixing branch. The reactions occur at a higher temperature and are restricted to particles with high chlorine fraction. The chloroform profiles show that a maximum of 28% mass fraction is attained for particles in zones where mixing with cooler non-reacting particles is more probable. This keeps the reacting particles lit and feeds reactants as well. Further downstream, there is a complete consumption of chlorine and the reacting branch collapses to the mixing branch with global extinction around $x/r = 6$. The chloroform plots show that the preferred mass fraction is the lower peak observed in the reacting zone. Higher mass fractions are decomposed to other carbon-carbon bonded chlorine derivatives. Near the end of the reactor, the reaction branches have collapsed completely onto the mixing branch. The complete depletion of chlorine upstream leaves no reactants, but a high temperature that leads to further decomposition of product.

The infinite-rate micromixing case (Figs. 5.18d and 5.20) shows results different from the previous non-premixed case. The faster mixing provides for better mixing conditions thereby spreading the core of the inner jet. This leads to higher temperatures with shorter reaction zones. The peak in temperature occurs near the inlet and away from the centerline with global extinction observed farther downstream. The final outlet temperatures remain almost the same as the finite-rate micromixing case. As before, product decomposition is very high and a surge of reactant mass fraction is seen at the end of the reactor. It can be noticed that the centerline temperature profile looks very

similar to the premixed reactor. CFD simulations with no sub-grid scale closure for the source term will thus yield similar centerline profiles for both the premixed and non-premixed cases. Evidently, closure in the form of a mixing model is required, though the form of the model that will be the most suitable should be decided based on experimental data.

The centerline temperature for the finite-rate micromixing case shows a sudden increase towards the end of the reactor. This is due to the inherent instabilities in the reactor that causes the re-attachment point to vary with time. Such changes lead to “peel-off” of high temperature layers that are then pushed into the plug-flow zone and carried out of the reactor. However, they are found to have very little effect on the product composition. The frequency of “peel-off” varied with time and may depend on the frequency of feedback (i.e. M). These results were found with $M \leq 10$ while for $M > 100$, the reactor became extinct.

These results indicate that the feed-stream configuration is very important in determining product yield. In addition, the use of a micromixing model to account for the sub-grid fluctuations is indispensable in simulating the chlorination reactor. It can be concluded that the use of the premixed-case leads to the maximum product yield and that the other configurations can be inherently unstable. Non-premixed inlets can also pose operational hazards unless monitoring systems can keep track of all zones of the reactor (as opposed to measuring centerline temperatures or outlet chloroform yield). Since the premixed case offers the most stable configuration, the rest of the chapter will deal only with this inlet configuration. The following sections will analyze plant-scale reactors that use premixed feed-streams. The last section will revisit the non-premixed configuration in the context of numerical-accuracy of the time-splitting techniques.

Simulation of Plant-scale Chlorination Reactor

The FV-PDF simulations of the pilot-scale reactors have led to a better understanding of the chlorination reactors. Here a pilot-scale reactor geometry is used to verify if the dynamics of the pilot-scale reactor remain unchanged with scale-up. The reactor geometry details were provided by The Dow Chemical Company and are confidential. Here scaled variables will be used to denote the dimensions of the reactor. The geometry considered is shown in Fig. 5.21. This reactor is a slightly modified version of the coaxial reactors considered before. A sparger of length L_s penetrates the tubular domain. The reactor of length L and radius R is assumed to have adiabatic walls. The radius of the

sparger is set at R_s . The sparger location creates a large dead zone where the fluid is almost at rest.

The inlet conditions are similar to the premixed conditions considered before. A mixture of organics and chlorine is fed at 333 K into the reactor. The feedstream mixture consists of CH_3Cl , CH_2Cl_2 , and Cl_2 with a molar ratio of 2:1:1. The operating pressure is considerably lower at 2.36 atm. Hybrid simulation procedure similar to the previous cases is adopted. The domain consists of 4 blocks with a total of 9700 control volumes. 50 particles per cell is used with half a million particles throughout the domain. The time step is initially set at 1×10^{-2} s and after 2 residence times reduced to 5×10^{-3} s. The simulations are run for a total of 3 residence times (calculated based on the inlet velocity) after which steady state was observed. The reduced chemistry scheme used before is utilized here. Acharya et al. [1991] have simulated the same configuration using a moment-closure method with a much coarser grid. The CFD results show reasonable agreement with data collected from the plant, though the effect of the grid resolution on simulation results has been left unanswered. In this work, the simulation of this plant-scale reactor is the first step towards simulating unstable configurations.

The axial velocity profile with streamlines are shown in Fig. 5.22. It can be seen that a large recirculation zone is present near the exit of the sparger. The recirculation zone is extended inward into the dead zone and traps some of the incoming cold fluid. The turbulence time scale (Fig. 5.22) indicates very low mixing rate in the recirculation zone and the mixing time is of the same order as the residence time. It was found that average velocity in this zone is at least 2 orders of magnitude smaller than the inlet velocity. Time dependent simulations of such configurations with widely varying flow-times can be expensive. Here a larger time step is used initially to speed-up the convergence and after the initial conditions have been convected out, the time-step is decreased and simulation continued until steady-state is attained. The scalar profiles are shown in Fig. 5.23. It can be seen that most of the reactions takes place at the exit of the sparger. There is near-uniform temperature in the rest of the reactor. The chlorine consumption is complete and the decomposition of chloroform is negligible. This reactor behaves in the same way as the premixed pilot-scale coaxial reactor considered before. The highly exothermic chlorination reactions ignite around 600 K and hence the inlet stream needs to be heated. The recirculation region provides this extra enthalpy to increase the fluid temperature. Reactions become progressively faster as fluid travels across the reaction zone defined as the thin region at the end of the sparger with large temperature gradients. Fully-reacted mixture attains a final state of around 780 K

with a chloroform yield of 14% (Fig. 5.24). The composition is uniform in the region downstream of the reaction zone. This is reflected in the axial composition plot which shows that all scalar gradients are associated with the temperature gradients. The plug-flow region following the reaction zone contains a homogeneous product and very low decomposition is observed. It was also found that a very low number of ISAT nodes are stored in the tree supporting the idea that the reactions are very fast and reach a single steady state. Hence the region of composition space to be mapped is very small and can be described by a small number of nodes.

This simulation shows that the plant-scale and pilot-scale reactors have similar flow patterns and dynamics. It is seen that the product yield depends only on the inlet composition indicating very fast chemistry that consumes the entire chlorine in the inlet stream. The reactor stays lit as long the recirculation zone is strong enough to feed back enthalpy to the cold inlet stream. The dead-zone is of no particular importance in this reactor. However, the length of the sparger can play an important role in determining the recirculation rate. The next set of simulations are aimed at understanding the effect of inlet configuration and sparger length on reactor performance.

Model Performance under Plant-operating Conditions

Plant operations are subject to a certain degree of uncertainty in the control of operating parameters. For example, inlet feed composition can contain a varying percentage of chlorine in the chlorination reactor. In general, reactor designs take into account these uncertainties. In continuous flow reactors, one of the optimization parameters is the flow rate - to achieve the maximum throughput without reactor quenching. From the CSTR studies in the first section of this chapter, it can be seen that for the current mechanism and operating conditions, reactor extinction occurs at a residence time close to 6 s. The flow-rates used in the coaxial reactors and plant-scale reactor considered before, correspond to a residence time of around 8 s. Theoretically this implies that an increase of flow-rate to 1.3 times the current flow-rate should be possible without causing reactor quenching.

The four-different simulations performed here try to highlight the fact that such direct scale-up from a zero-dimensional reactor study is not accurate. It will also show that even in the limit of premixed feed, the reduced heat transfer induced by the micromixing model will alter the behavior of the reactor. The reactor geometry is similar to the configuration used for the simulation of plant-scale reactor in the previous sec-

Table 5.5 Feed stream composition and operating conditions for the plant-scale operation

Case	Y_{Cl_2}	$Y_{\text{CH}_3\text{Cl}}$	$Y_{\text{CH}_2\text{Cl}_2}$	Flow rate (kg/s)	Op. pressure (atm.)
A	0.2724	0.3819	0.3457	7.38	2.36
B	0.2184	0.4025	0.3791	8.84	2.02

tion. However, the sparger length is one-third of the length used in that case. The two cases considered correspond to one stable and one unstable configuration [Tsai, 2002]. The details of the feedstream configurations and other parameters are given in Table 5. Case A represents a stable configuration with maximum product yield and is the desired set-point for the control system. Case B is an unstable operation in which reactor quenching has been observed. For each of the cases two simulations were performed - one using the IEM micromixing model and the other assuming complete mixing as is usually done in FV simulations without sub-grid closure.

The results of Case A simulation with and without micromixing model are shown in Fig. 5.25. It can be seen that the micromixing model has very little effect on the temperature field. The addition of micromixing lengthens the reaction zone as seen in the simulations with the pilot-scale reactors. However, quantitatively there is little difference in the final product yield with both simulations predicting 14.3% chloroform mass fraction at the outlet. The streamline plot (Fig. 5.25) is a superimposed field with the top half indicating the infinite-rate mixing case and the bottom half representing the finite-rate mixing case. It can be seen that qualitatively, the recirculation zone appears identical and the re-attachment point is located at almost the same axial location. It can be concluded that in this case, the micromixing model has no effect on reaction.

Simulations of Case B corresponding to an unstable reactor configuration yielded very interesting results. Similar to Case A, two different simulations with and without a micromixing model were carried out. The temperature plots in Fig. 5.26 show the results of the simulations. The infinite-rate mixing case shows reduced temperatures in the reaction zone but the enhanced mixing rate tends to keep the reactor lit. On the other hand, when finite-rate mixing is used, the reactor quenches with a blob of unreacted inlet stream spreading into the reaction zone. The quenching process is initiated by the lowering of the temperature in the dead-zone above the sparger. This leads to progressive cooling as observed in the infinite-rate mixing case. However, slower heat transfer rates fail to sustain the reaction. The ignition temperature for this mixture is around 650 K.

It can be seen that the reaction zone near the sparger exit is significantly cooler than the ignition temperature. In addition, the lower temperature gradients lead to weaker recirculation zones that do not penetrate the dead-zone. This leads to a significantly lesser amount of enthalpy being fed back to the cold inlet stream. These effects together are sufficient to cause global extinction.

The axial plot of chlorine composition (Fig. 5.27) shows increasing mass fraction in the outlet stream. This indicates that the chlorination reactions are not complete as opposed to previous cases where the chlorine consumption was close to 100%. The plot at the longest time corresponds to the time at which extinction was observed. The reduction in reaction rate has led to a residual chlorine mass fraction in the outlet indicating that extinction phenomena is not local and the process was initiated long before the unreacted fluid entered the reaction zone. It rules out extinction due to numerical instabilities which would have caused extinction to occur locally and would have occurred simultaneously with the cold fluid entrainment. The progressive cooling can be attributed only to the weak recirculation and slower heat-transfer rates.

This concludes the discussion on the plant-scale reactors. It has been observed that the use of the micromixing model is essential in predicting extinction. The mixing model is accurate enough to predict reactor behavior accurately. It should be noted that the mixing model is not strictly valid for the premixed case and describes the mixing of pure binary streams only.

Numerical Accuracy of the Time-splitting Scheme

In all the simulations discussed in the chapter, the numerical accuracy of the fractional time-stepping method or the time-splitting technique has been implicitly assumed to be first-order accurate. By chance, an interesting phenomena was observed in the simulation of the non-premixed configuration explained in previous sections with chlorine in the inner stream. This leads to a lifted flame profile for the temperature distribution. The normal process of simulation proceeds as follows:

1. Simulate cold flow with no reaction.
2. Use the cold flow profiles as starting conditions for reacting flow with infinite-rate mixing.
3. Use steady state infinite-rate mixing solution to start finite-rate mixing simulation.

The above steps were carried out and the final steady state for the finite-rate mixing was obtained for the non-premixed inlet configuration. However, to ensure that the solution is accurate, infinite-rate mixing was used and simulated until steady state was reached. Surprisingly, the final state of the system was not similar to that obtained in the forward loop of the cycle (Fig. 5.28). The temperature profiles in the three simulations are shown in Fig. 5.29. It can be seen that the high temperature steady-state corresponding to infinite-rate mixing (from here-on referred to as state A) exhibits a higher peak temperature than the corresponding finite-rate mixing (state B). However, the lower steady state exhibits no lifted flame profile and is similar to the non-premixed configuration. The chlorination scheme is known to exhibit oscillatory behavior [West et al., 1999], but no multiple-steady state has been observed. Hence further analysis is required before the reason for this behavior is determined.

It was noted that for the finite-rate mixing condition, the final steady state was independent of the initial conditions. When the mixing model was turned on at state A, the final solution was identical to the state B. To verify that the state B was grid independent an independent simulation with twice the number of grid points in each direction was carried out. The axial profile of temperature (Fig. 5.30) shows very little difference between the two grids indicating that the solution with the coarser grid is well-resolved. Since infinite-rate mixing is independent of the mixing model used, the other main source of error could be the chemistry scheme. It is possible that model reduction from the full-chemistry with 38 species to 17 species had resulted in a scheme that leads to multiple-steady states. Simulations with the full-chemistry model yielded identical results for all the cases considered and the figures are hence omitted. The only other variable used in the simulation is the time step. Since the fractional time-stepping scheme is an explicit first-order scheme [Pope, 2000], with larger time-steps, the associated error in the solution will be higher. Hence the time step was varied from the original value of 1×10^{-3} s by factors of 1/2 and the axial profiles of temperature compared. The results (Fig. 5.31) show that with decreasing time step, the location of the peak temperature moves towards the inlet. It can be seen that no asymptotic-limit is observed and with smaller steps, the peak temperature will move closer to the inlet.

It is known that for the fractional-time stepping scheme to be first-order accurate, the time step chosen should be of the same order as the time scale of the fastest process in the sequence of time-steps. Since the mixing-rate is infinite, the time-step is limited by this physical process and should be set to a near-zero value to increase the accuracy of the simulation. Since such a limitation will make the scheme intractable, the use of the

any finite step in this configuration leads to a systematic error. In the finite-rate mixing case, the time step then will be decided by the fastest mixing time and the time-steps used in the studies have been chosen based on this condition. The reaction sub-step does not play a role since the ISAT integration ensures that through a stiff-integration process, the fastest steps in the chemistry scheme are accurately tracked. Hence the finite-rate mixing case yields the same steady-state irrespective of the starting condition, time step or chemistry scheme used.

Summary

This section has provided some key insights into the design of the chlorination reactor. It has been showed that the premixed configuration in an industrially dependable design for maximizing product yield. Through studies of different non-premixed configurations, it has also been established that the fast chemistry of the chlorination reactor combined with strong interaction between turbulence and reaction necessitates the use of a micromixing model and detailed chemistry. The Monte-Carlo based PDF code along with the FV code to handle flow field solutions is found to be the best solution algorithm for such reactors. The plant-scale reactor simulations showed that the PDF code is able to predict accurately the outlet temperature and product yields. It was also shown the use of the micromixing model is essential in predicting quenching of reactors. Use of infinite-rate micromixing led to a stable configuration while it is known (through real-life plant operations) that under the given operating conditions, the reactor quenches. This capability can be used to test the stability of the reactor design. It should be noted though that the micromixing model is not accurate for premixed flows. However, the reduced reaction rates resulting from slower heat transfer simulates the effect of micromixing even though the accuracy of this assumption needs to be tested using more extensive experimental data. Several simulation parameters including grid size, time-step size and feedback iteration frequency were studied and optimum values were determined for the chlorination scheme. The final study showed that the time-splitting technique should be handled with care especially in the limit of infinite-rate micromixing where any finite-time step induces error.

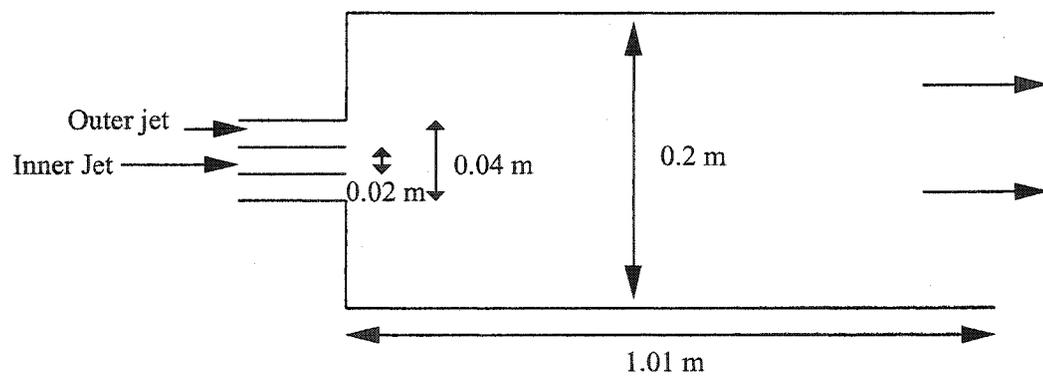


Figure 5.1 Schematic diagram of jet-stirred reactor.

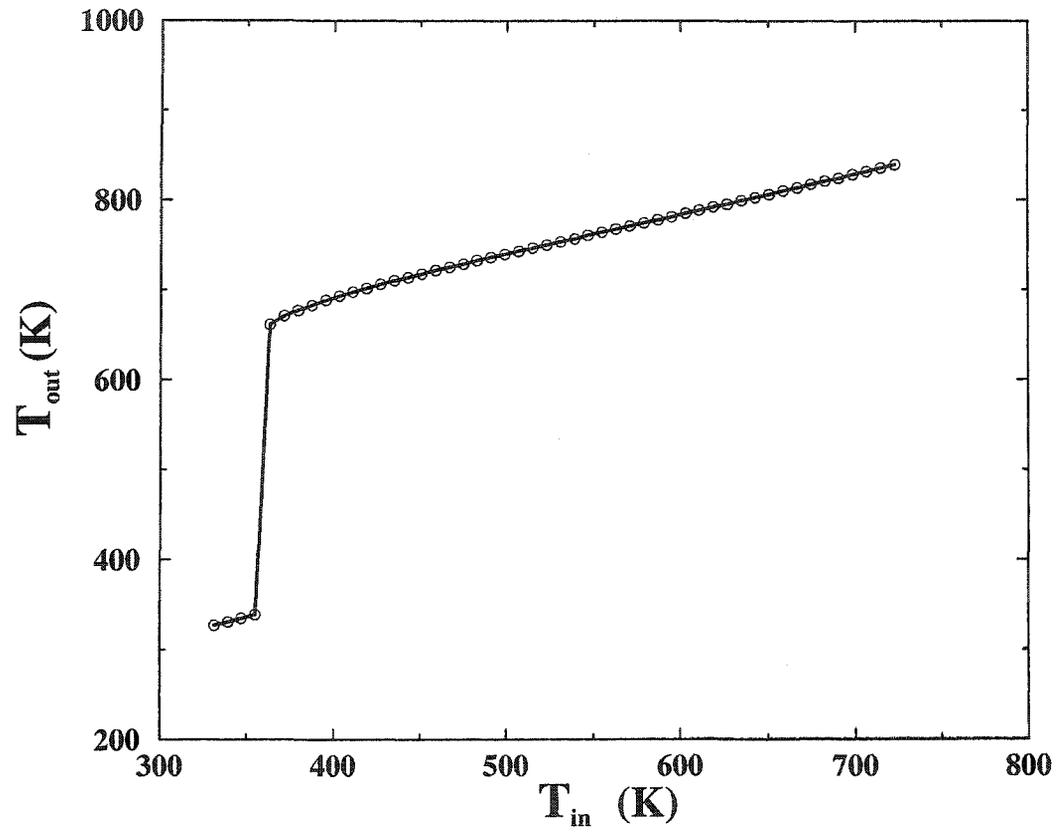


Figure 5.2 CSTR model - Variation of outlet temperature with inlet temperature at a constant residence time of 7.5 s.

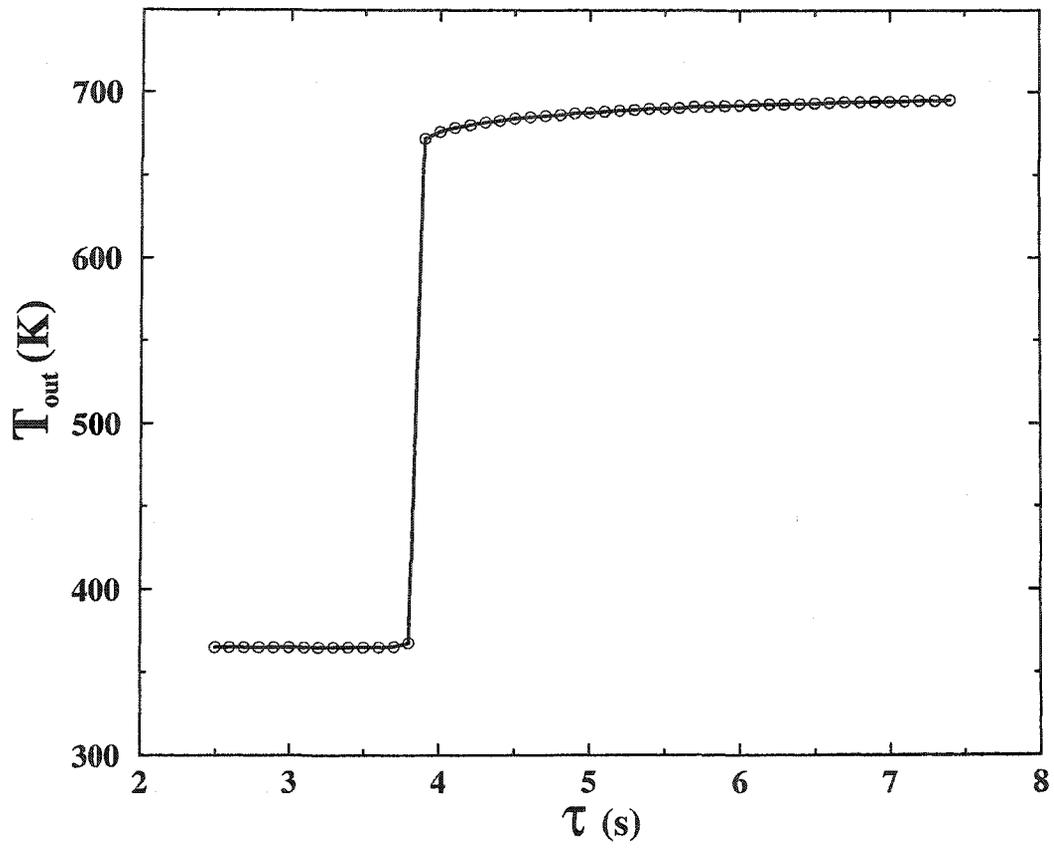


Figure 5.3 CSTR model - Variation of outlet temperature with residence time at inlet temperature of 365 K.

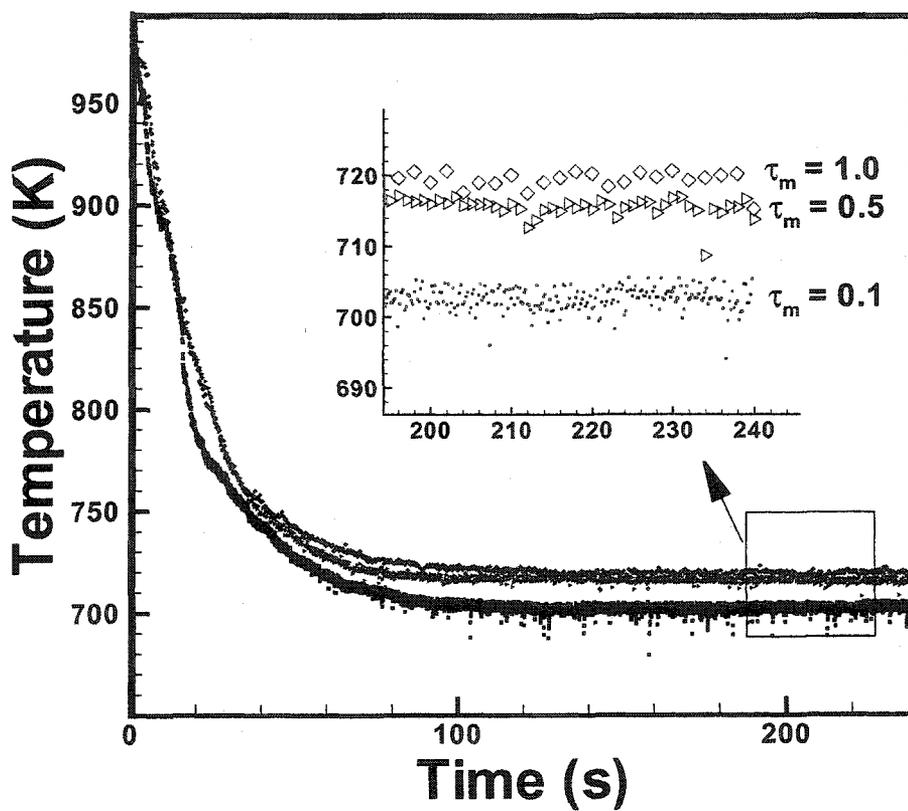


Figure 5.4 PMSR model - Effect of micromixing time on final temperature. Residence time is 7.5 s with inlet temperature of 323 K.

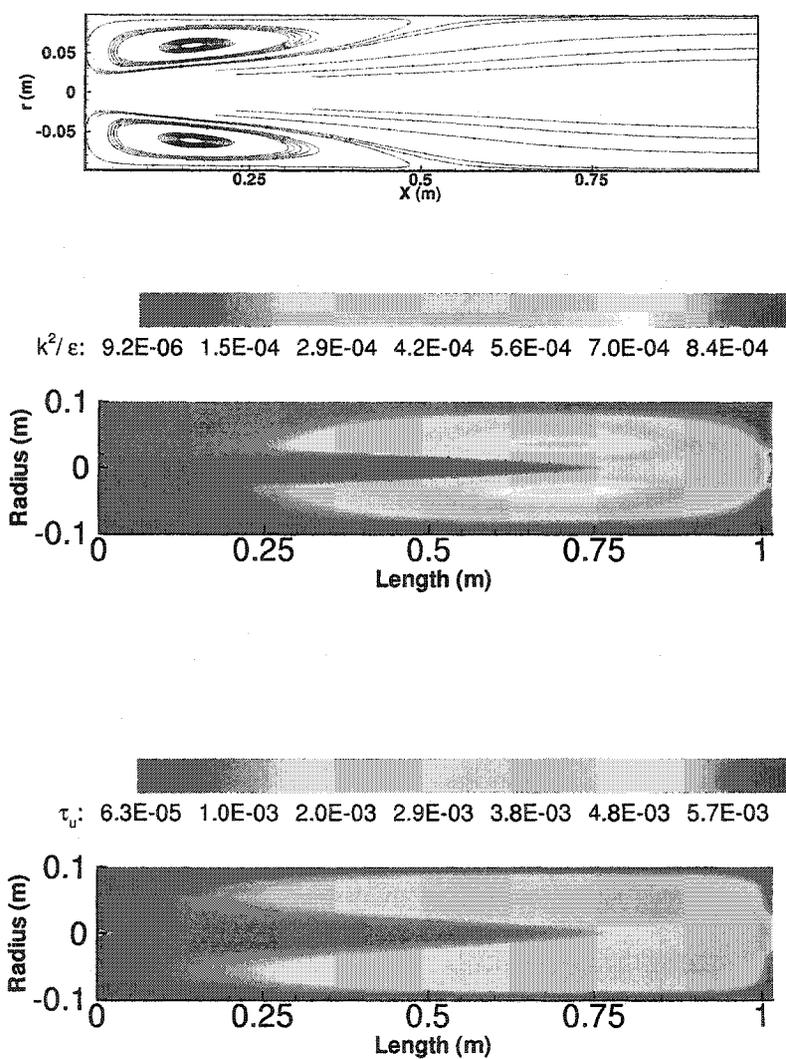


Figure 5.5 Mean velocity (top), turbulent diffusivity (middle), and local mixing time (bottom) profiles obtained from a $k-\omega$ turbulence model with a 15-species reaction mechanism.

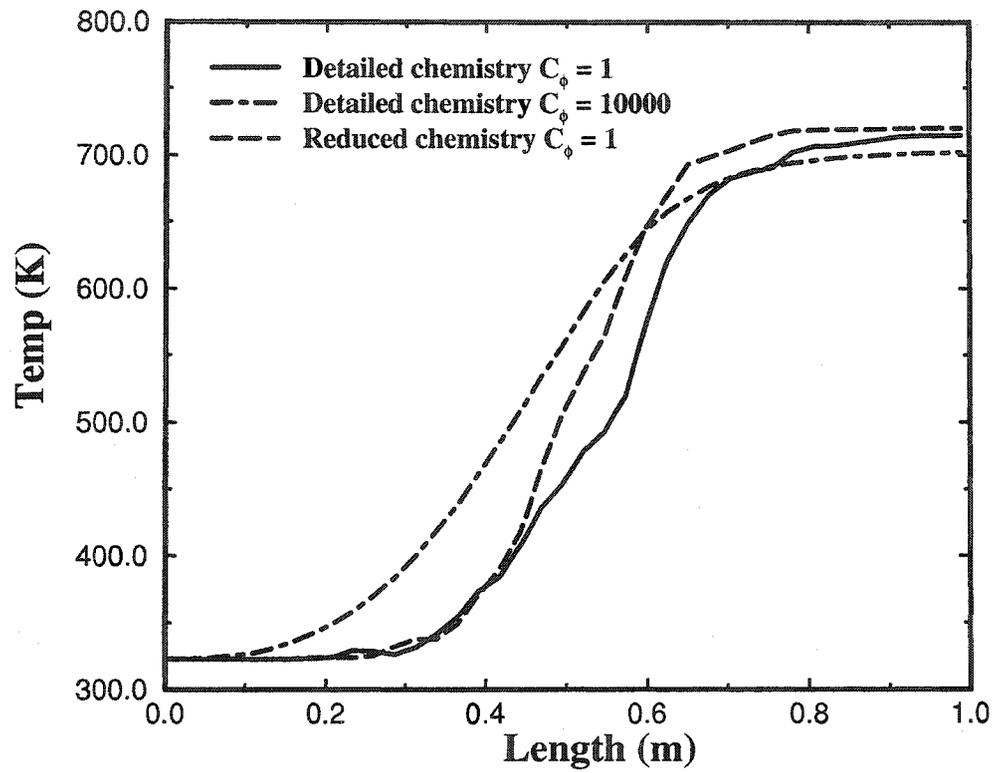


Figure 5.6 Comparison of time-averaged centerline temperature profiles for detailed and reduced chemistry with different micromixing parameters after three residence times.

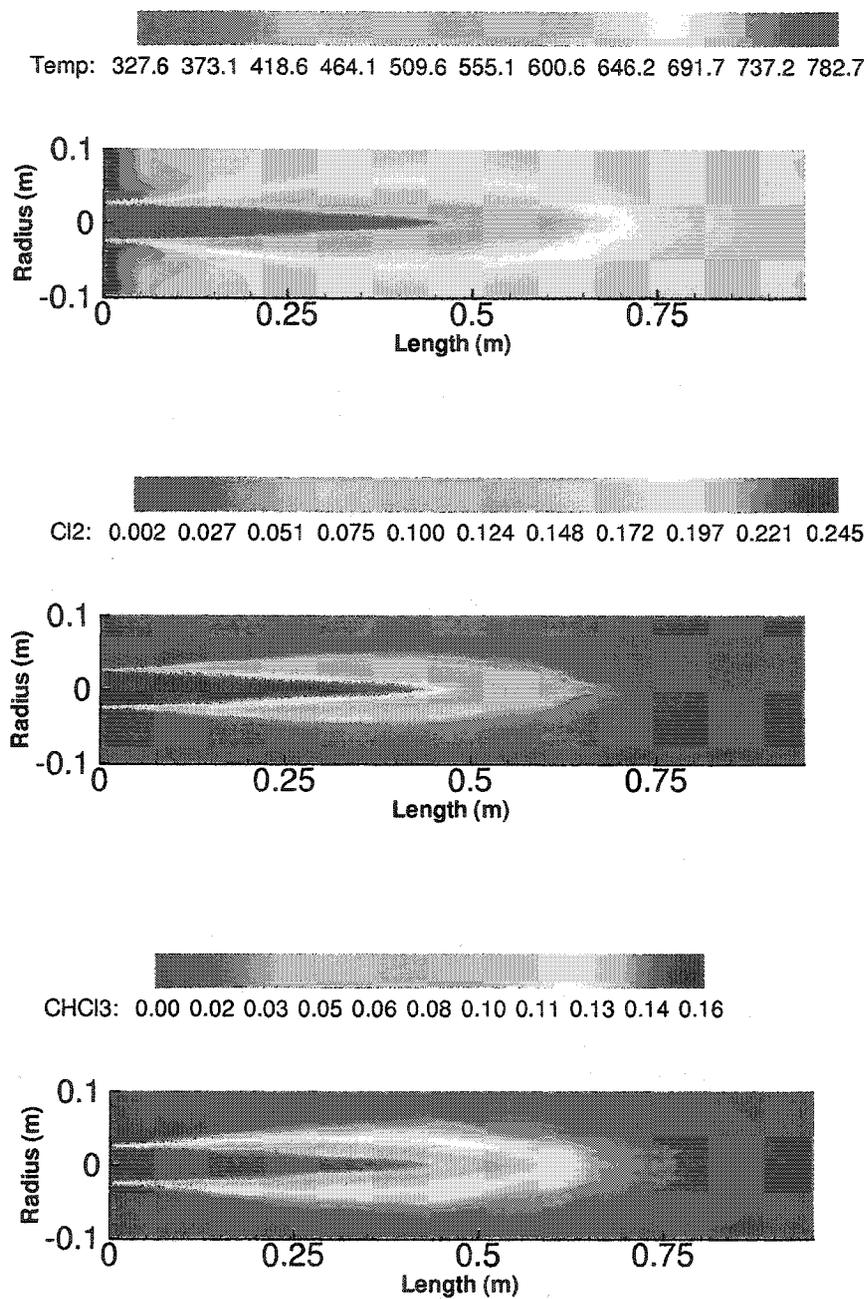


Figure 5.7 Temperature (top), chlorine (middle), and chloroform (bottom) profiles with detailed chemistry simulated using a composition PDF method.

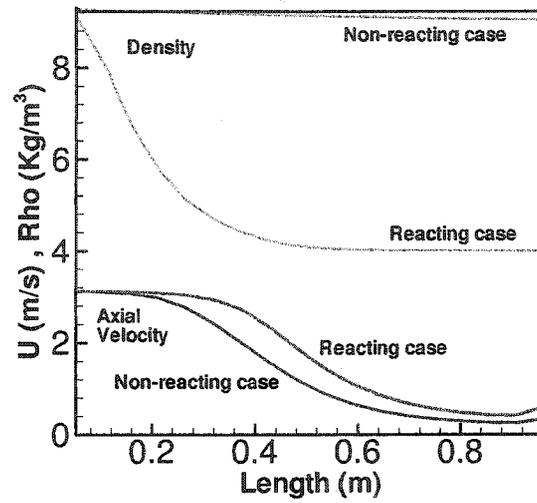


Figure 5.8 Effect of density on flow conditions.

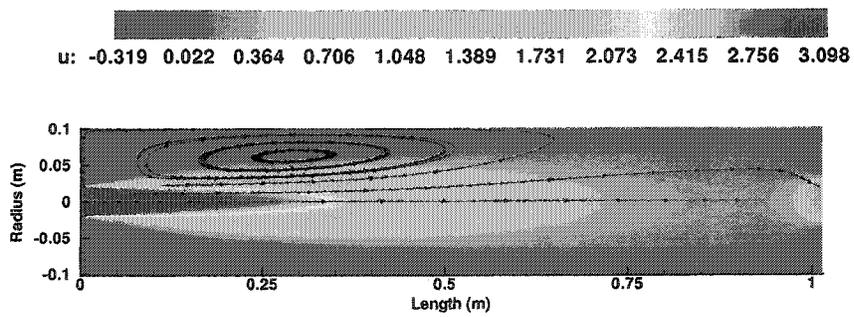


Figure 5.9 Velocity profile for cold flow condition.

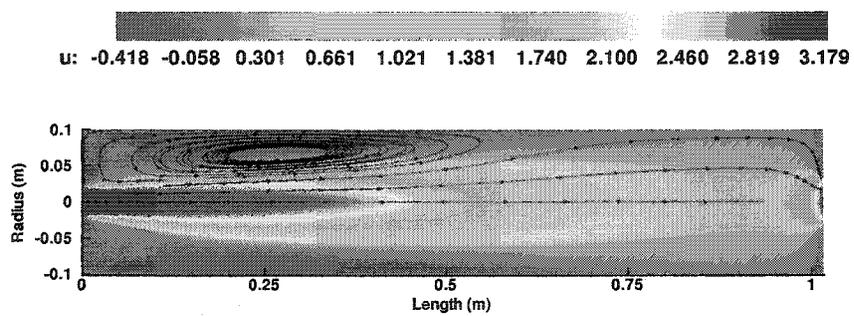


Figure 5.10 Velocity profile for reacting condition.

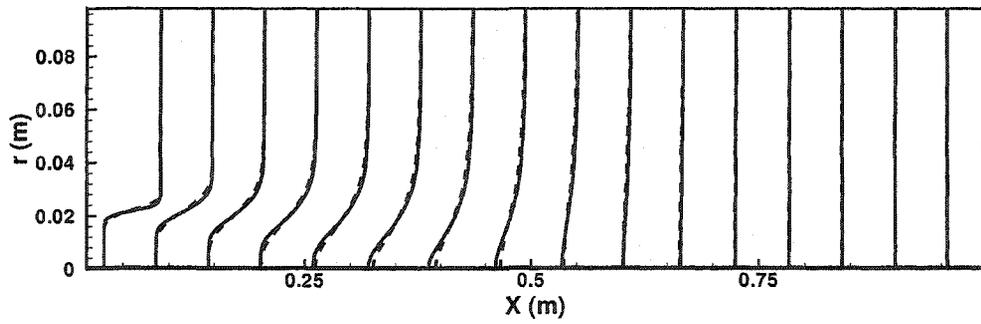


Figure 5.11 Radial temperature profiles for different grid sizes at various axial positions. The solid lines represent a 53×33 grid and the dashed lines represent 53×79 grid. Temperature has been normalized with respect to an arbitrary value of 750 and the maximum axial shift is 0.05 m for this temperature.

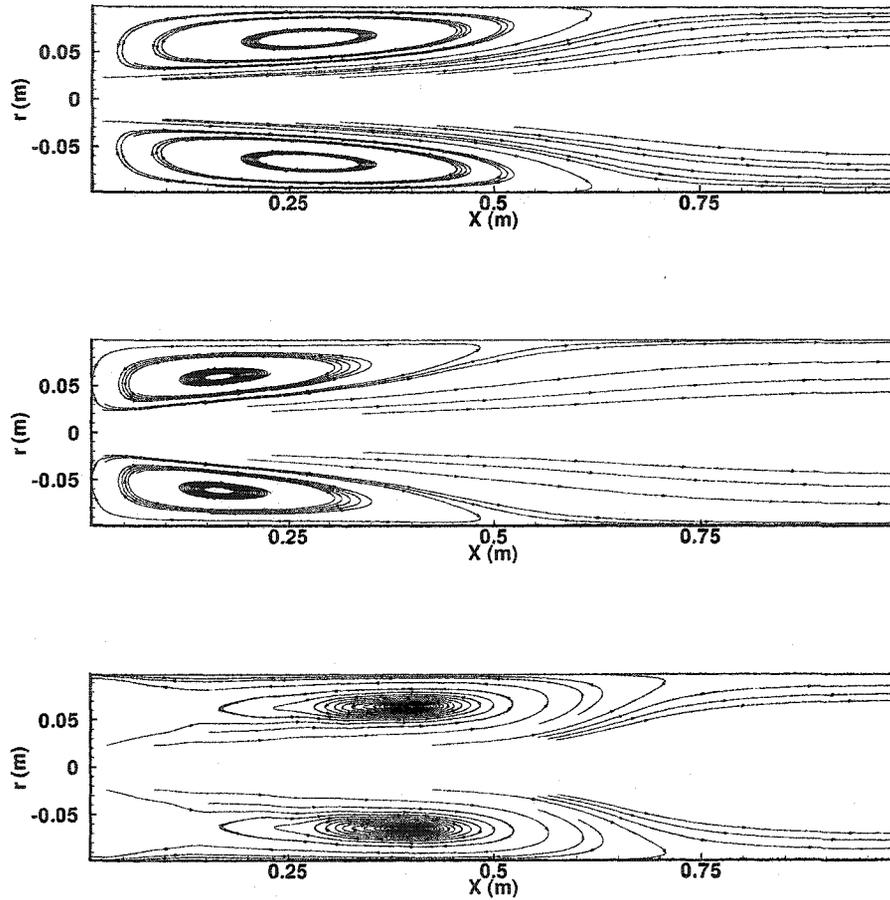


Figure 5.12 Streamlines illustrating recirculation zones for (top) Premixed (middle) non-premixed with chlorine inner jet and (bottom) non-premixed with chlorine outer jet configurations.

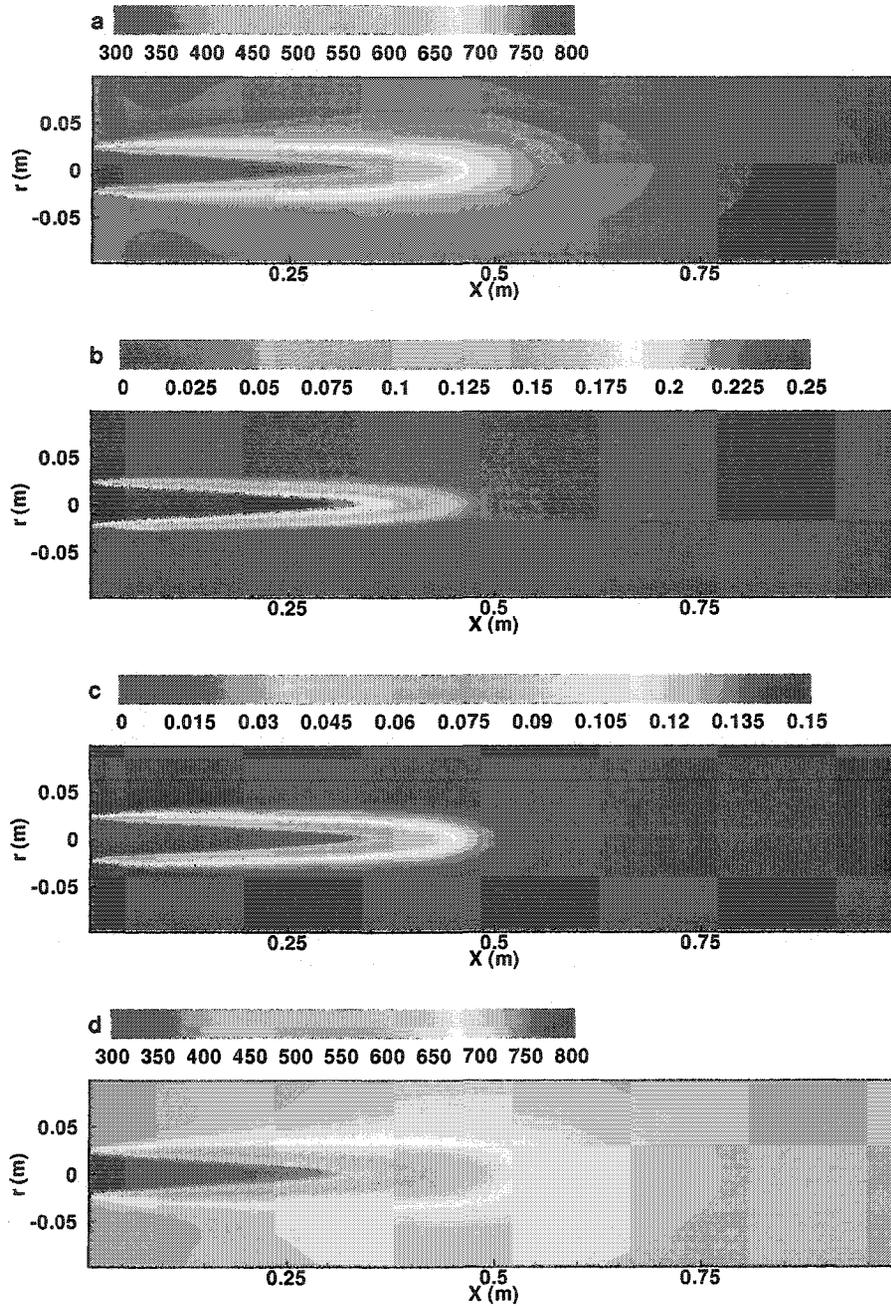


Figure 5.13 Profiles of (a) temperature, (b) chlorine and (c) chloroform for the premixed case. Plot (d) is the Temperature profile with infinitely fast micromixing.

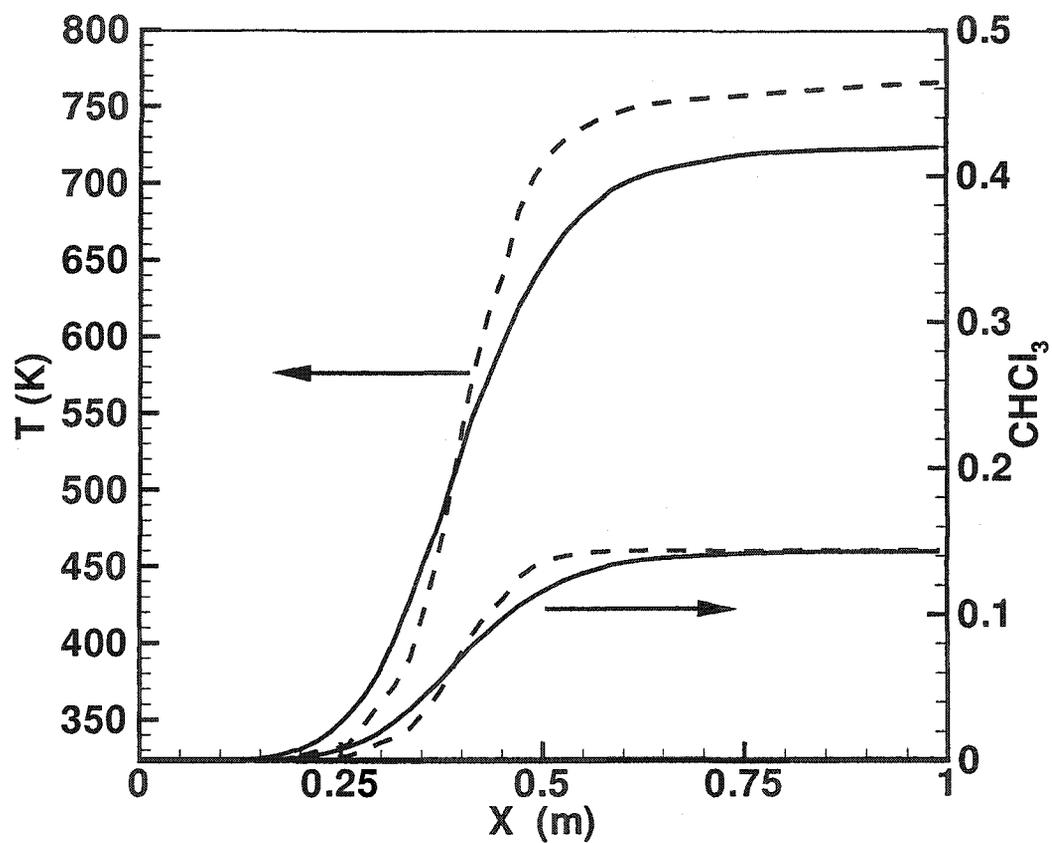


Figure 5.14 Centerline temperature and product profile for (solid) infinitely fast micromixing case and (dashed) IEM mixing model with premixed inlets.

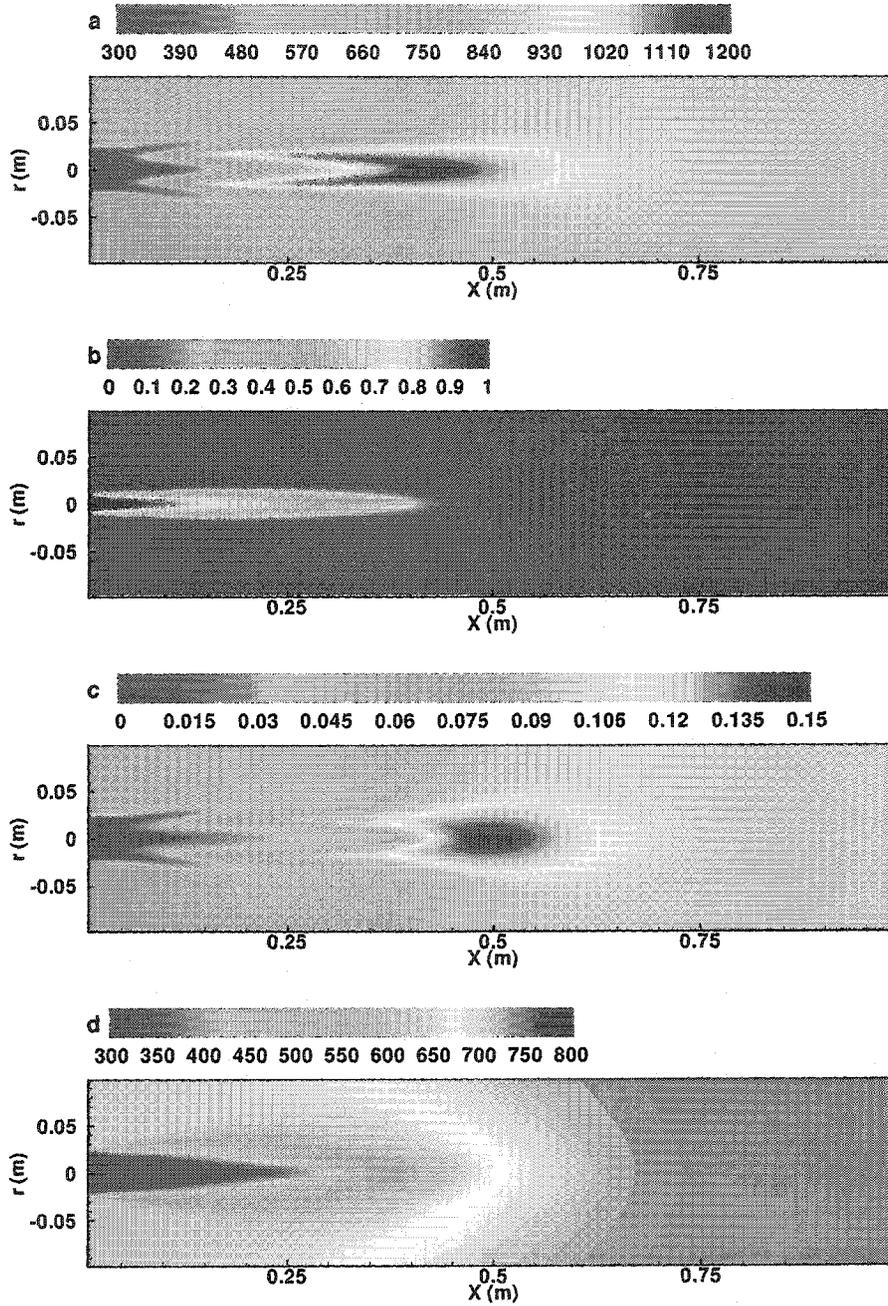


Figure 5.15 Profiles of (a) temperature, (b) chlorine and (c) chloroform for the non-premixed case with chlorine as the inner jet. Plot (d) is the Temperature profile with infinitely fast micromixing.

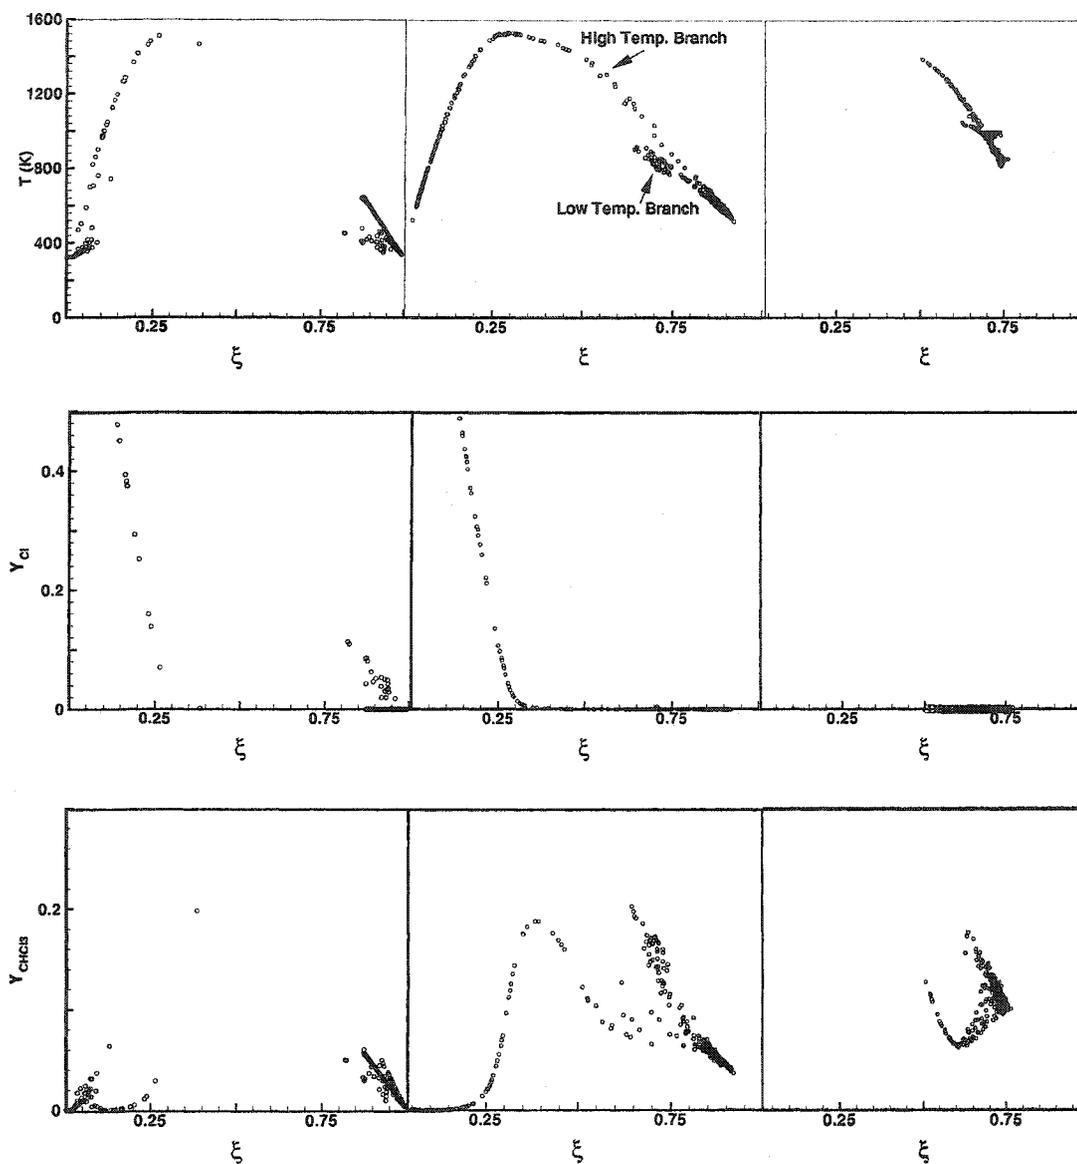


Figure 5.16 Radial scatter plots at different axial distances of (a) temperature, (b) chlorine and (c) chloroform for the non-premixed case with chlorine as the inner jet. Starting from left, the figures in each column correspond to axial positions of $x/r = 1$, $x/r = 3$, and $x/r = 6.7$, respectively.

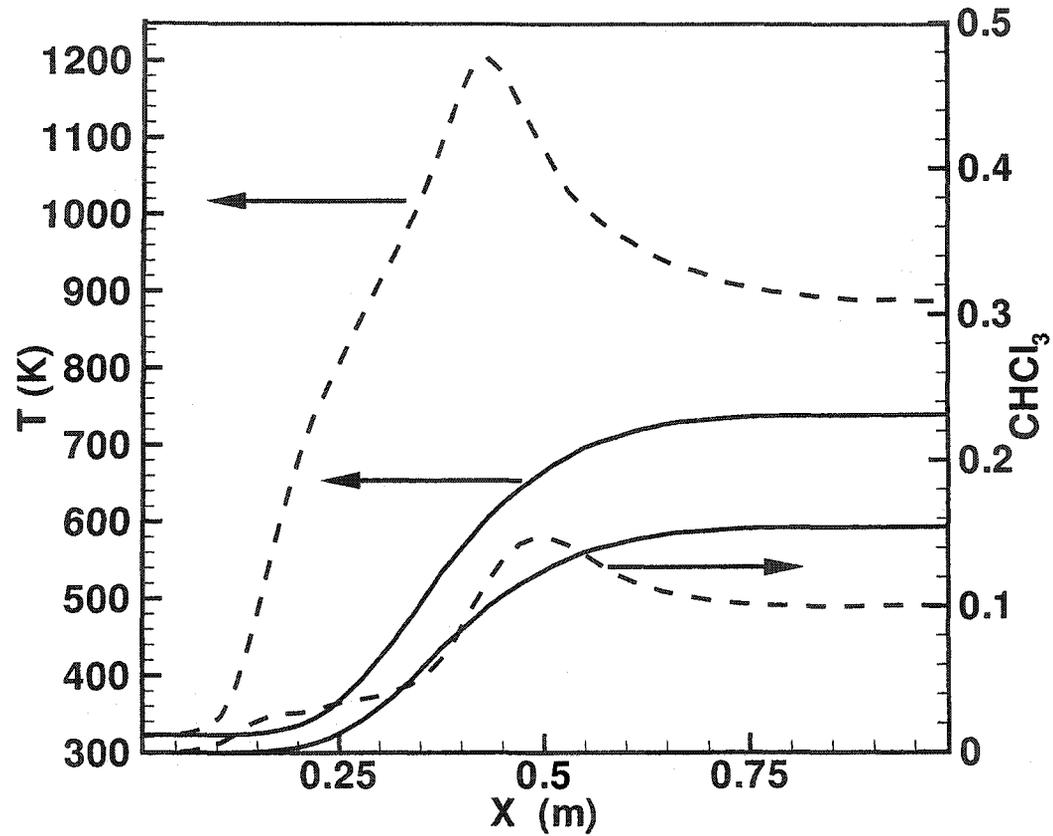


Figure 5.17 Centerline temperature and product profile using (solid) infinitely fast micromixing condition and (dashed) IEM mixing model for non-premixed inlets with chlorine in the inner stream.

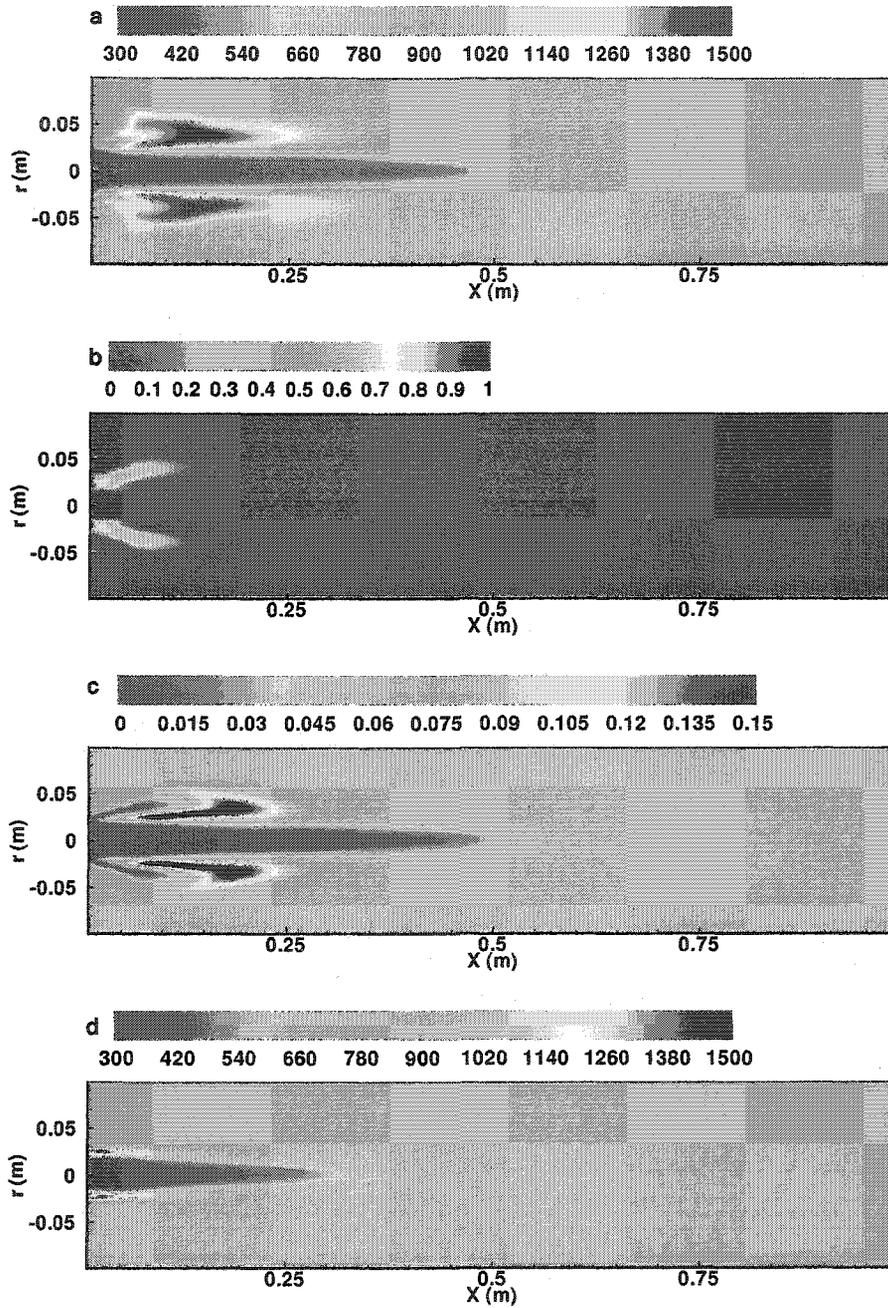


Figure 5.18 Profiles of (a) temperature, (b) chlorine and (c) chloroform for the non-premixed case with chlorine as the outer jet. Plot (d) is the Temperature profile with infinitely fast micromixing.

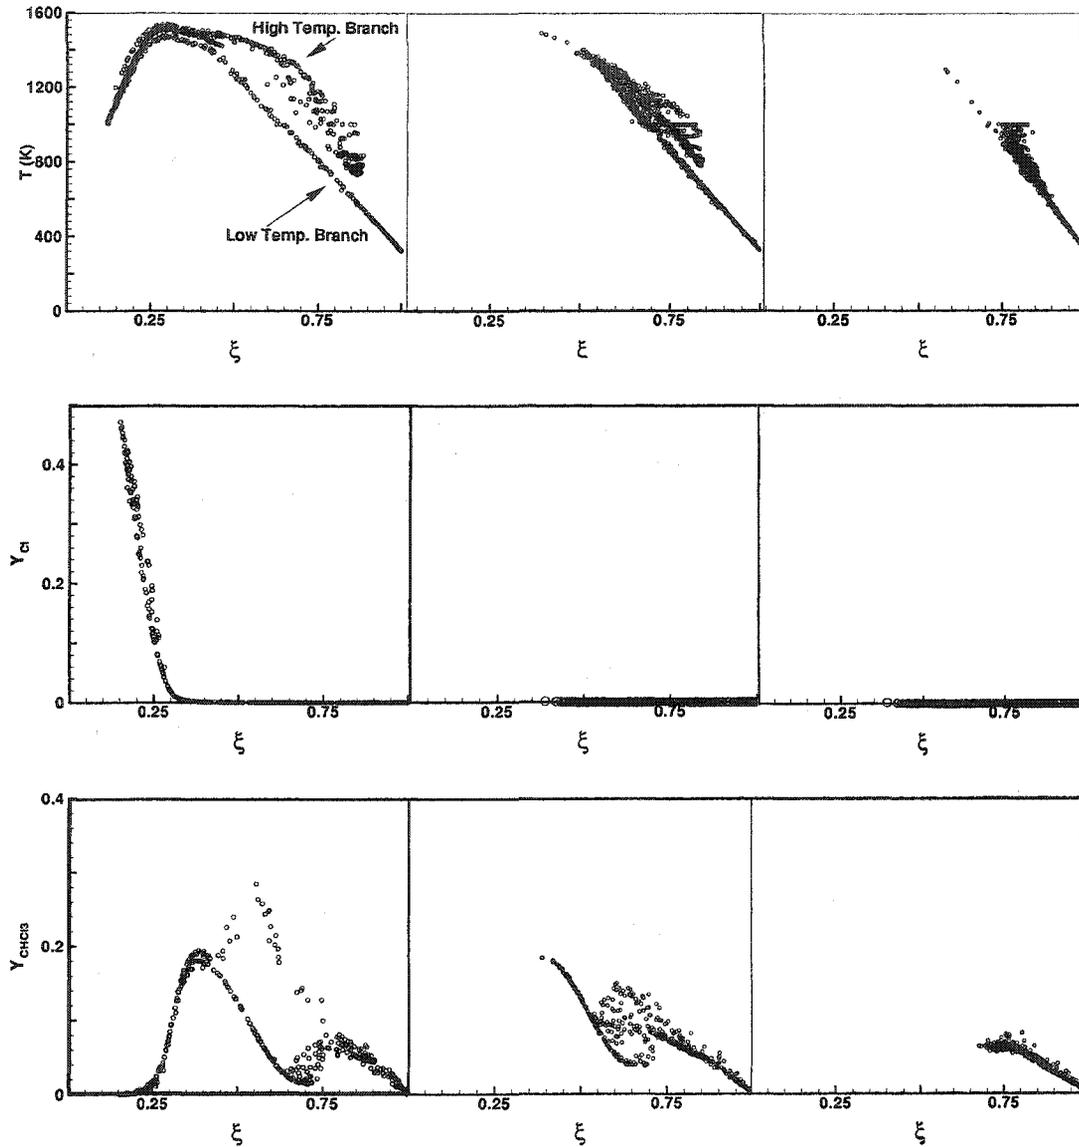


Figure 5.19 Radial scatter plots at different axial distances of (a) temperature, (b) chlorine and (c) chloroform for the non-premixed case with chlorine as the outer jet. Starting from left, the figures in each column correspond to axial positions of $x/r = 1$, $x/r = 3$, and $x/r = 5$, respectively.

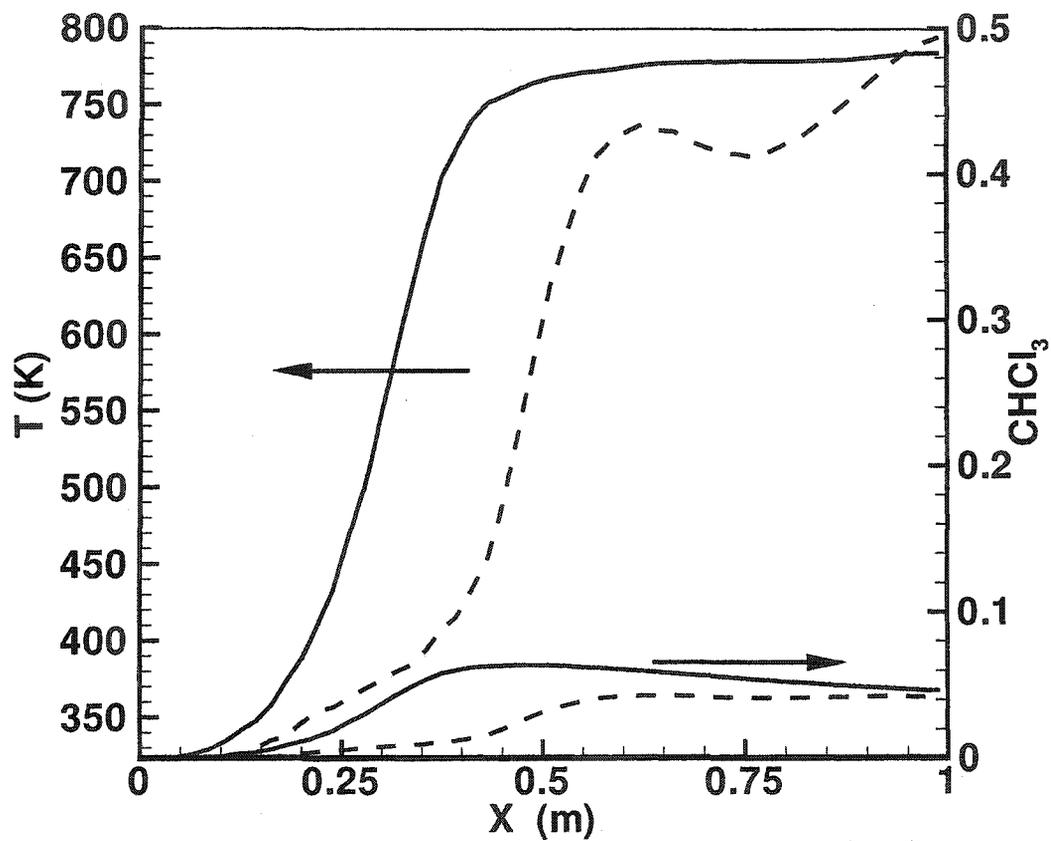


Figure 5.20 Centerline temperature and product profile using (solid) infinitely fast micromixing condition and (dashed) IEM mixing model for non-premixed inlets with chlorine in the outer stream.

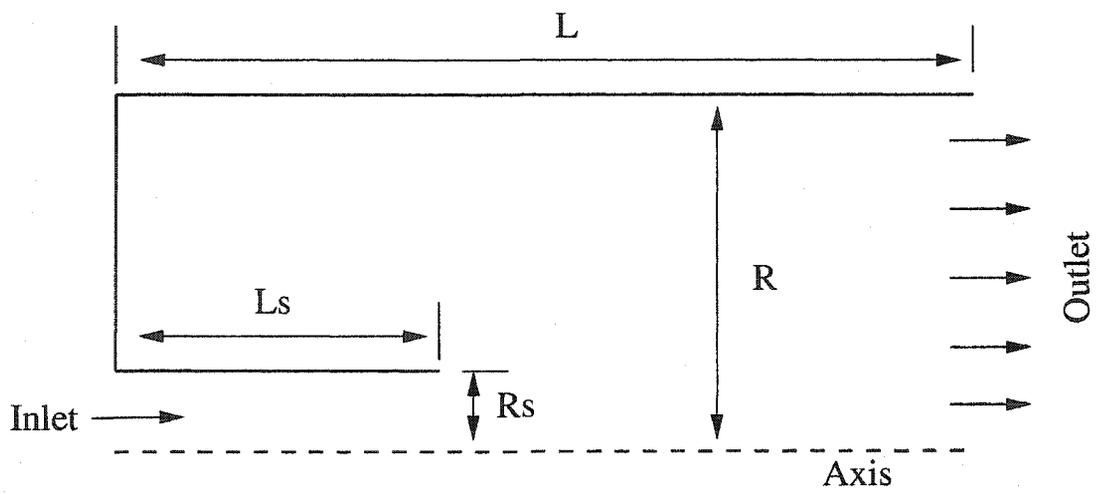


Figure 5.21 Configuration for the plant-scale chlorination reactor with sparger. L_s is the sparger length. R_s is the sparger radius. L and R are the external dimensions of the reactor.

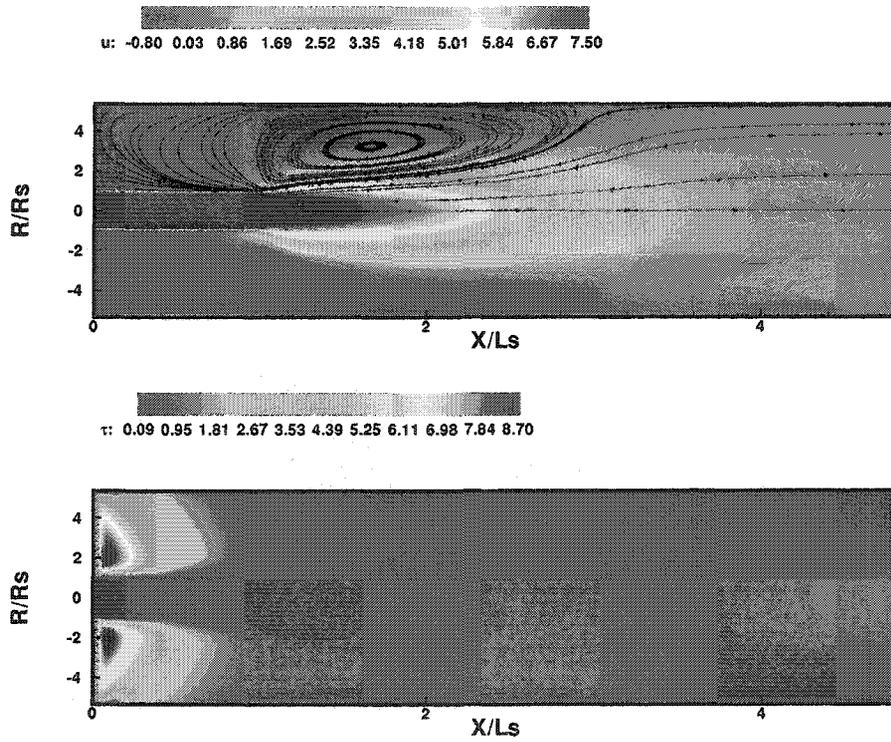


Figure 5.22 (Top) Velocity-streamline profile and (bottom) mixing time profile for the plant-scale chlorination reactor.

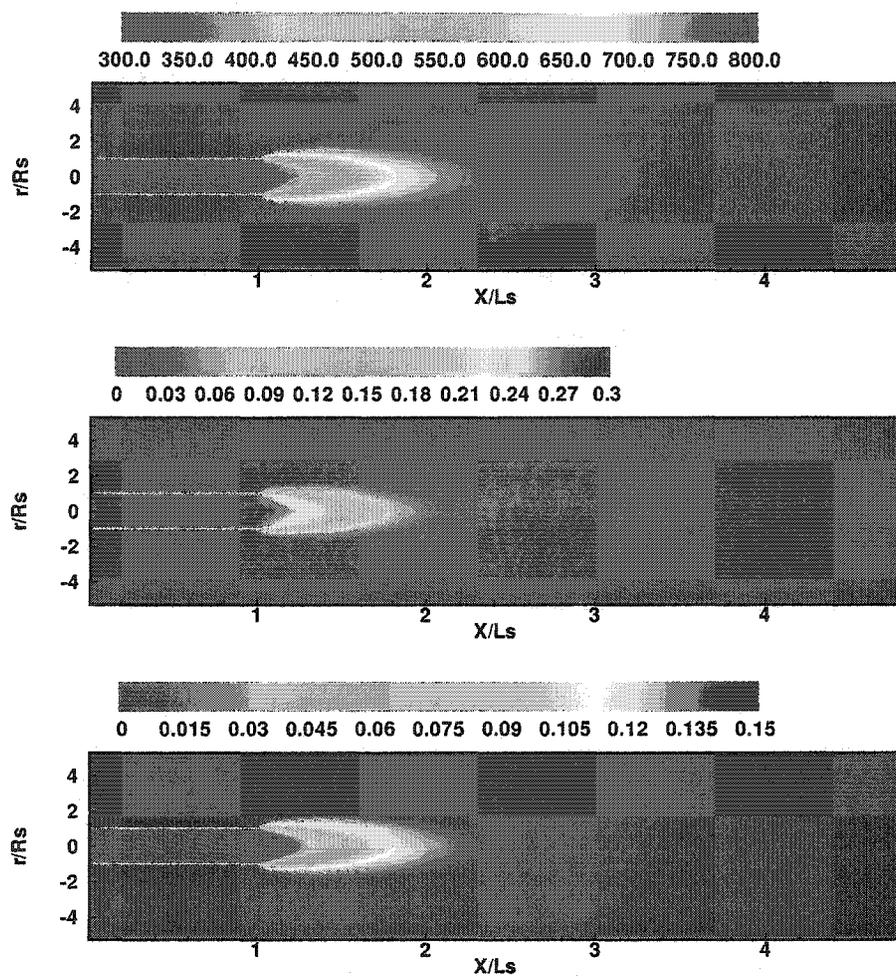


Figure 5.23 (Top) Temperature, (middle) chlorine, and (bottom) chloroform profiles for the plant-scale reactor.

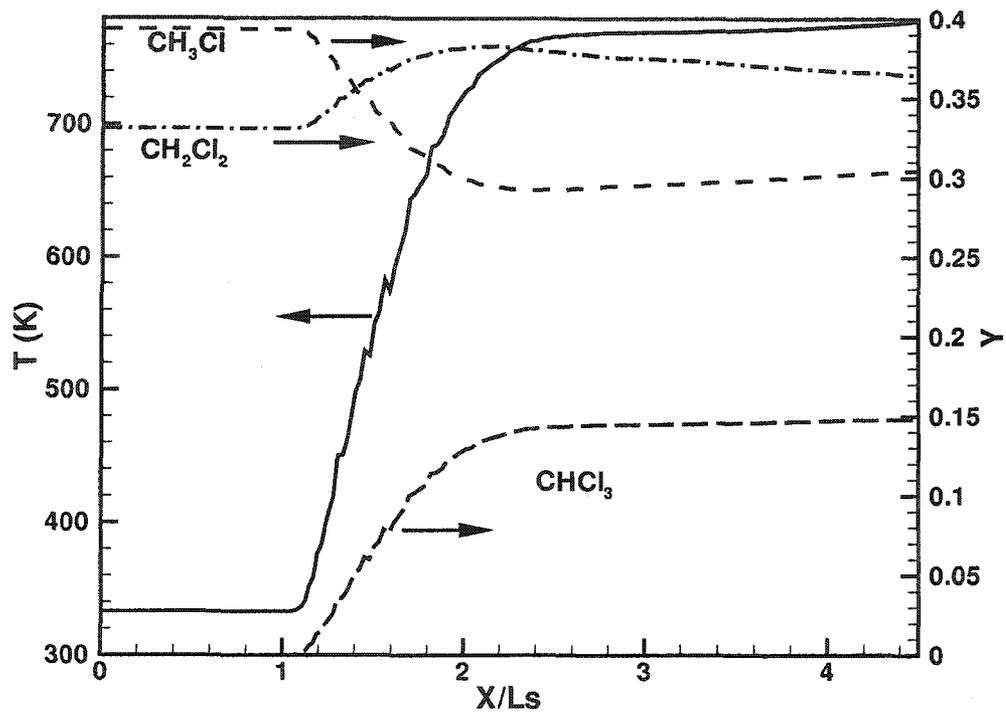


Figure 5.24 Axial profile of temperature and selected scalars for the plant-scale chlorination reactor.

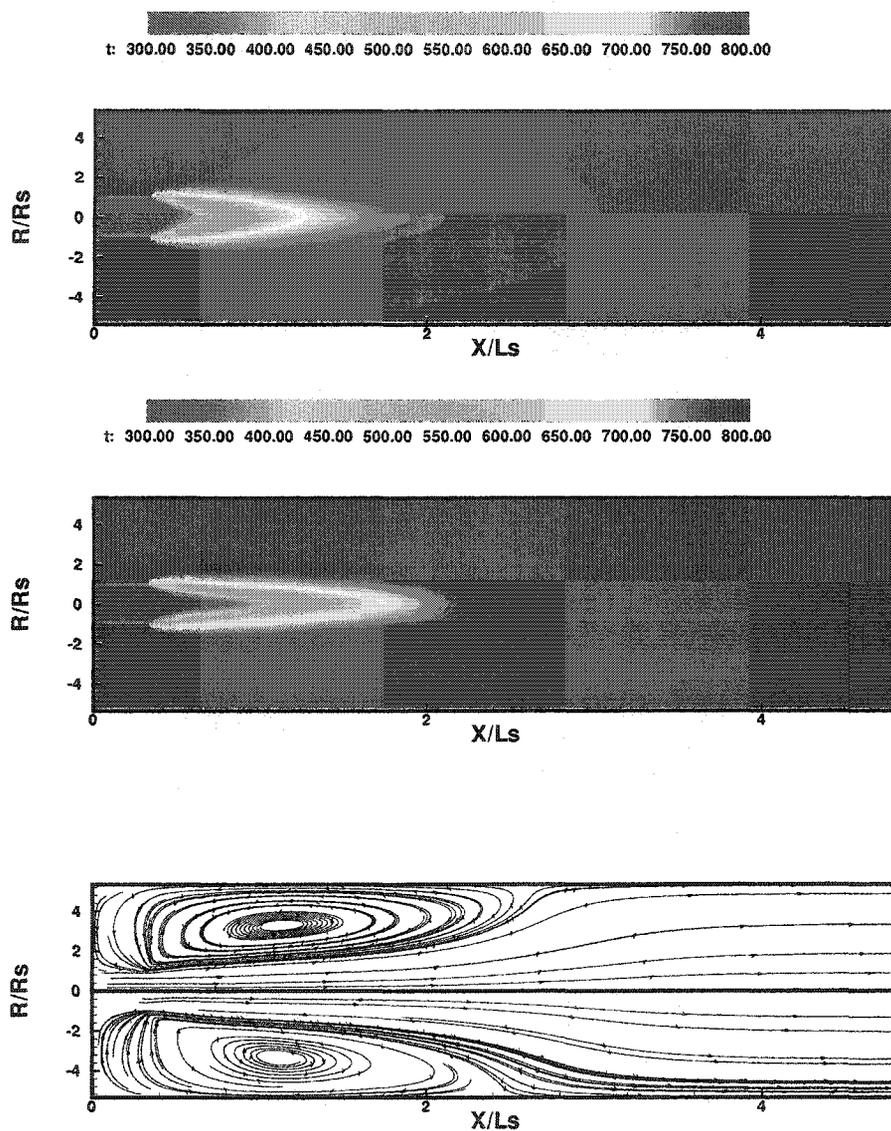


Figure 5.25 Temperature profile and streamline velocity plot of plant scale reactor simulation under stable operating conditions. (Top) Neglecting micromixing, (middle) using micromixing. The top half of the streamline plot corresponds to the infinite-rate mixing case and the bottom half represents the finite-rate mixing case.

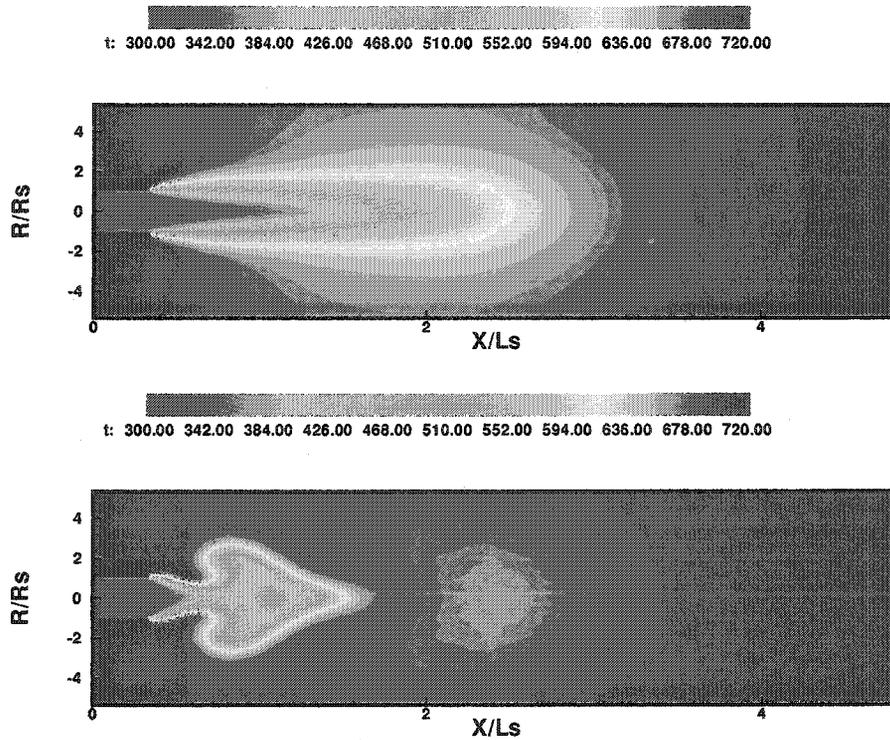


Figure 5.26 Temperature profiles of plant scale reactor simulation under unstable operating conditions. (Top) Neglecting micromixing and (bottom) considering micromixing.

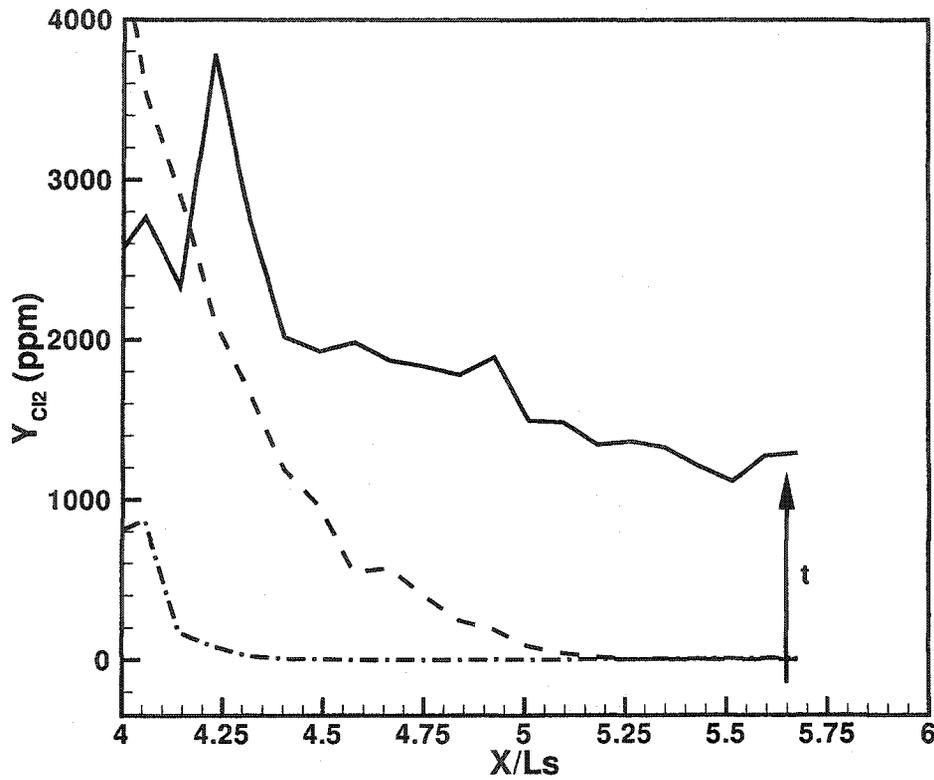


Figure 5.27 Instantaneous centerline profiles of chlorine mass fraction in ppm near the exit of the reactor. The arrow indicates the direction of increasing simulation time relative the starting solution from a infinite-rate mixing case. Dashed-dotted line - $t = 0.01s$. Dashed line - $t = 3s$. Solid line - $t = 9s$

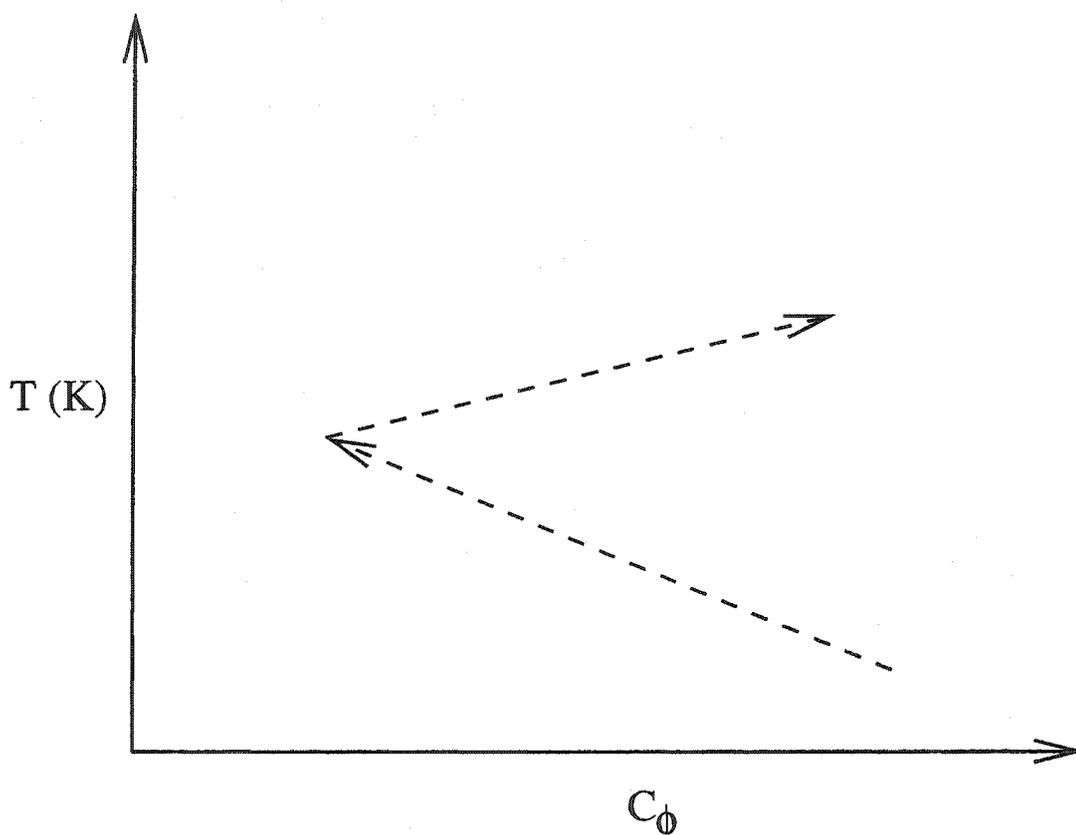


Figure 5.28 Simulation conditions for the test of numerical-accuracy. The higher C_ϕ corresponds to infinite-rate mixing and the lower value corresponds to finite-rate mixing. The temperature axis represents the peak temperature found along the centerline in each of the simulations.

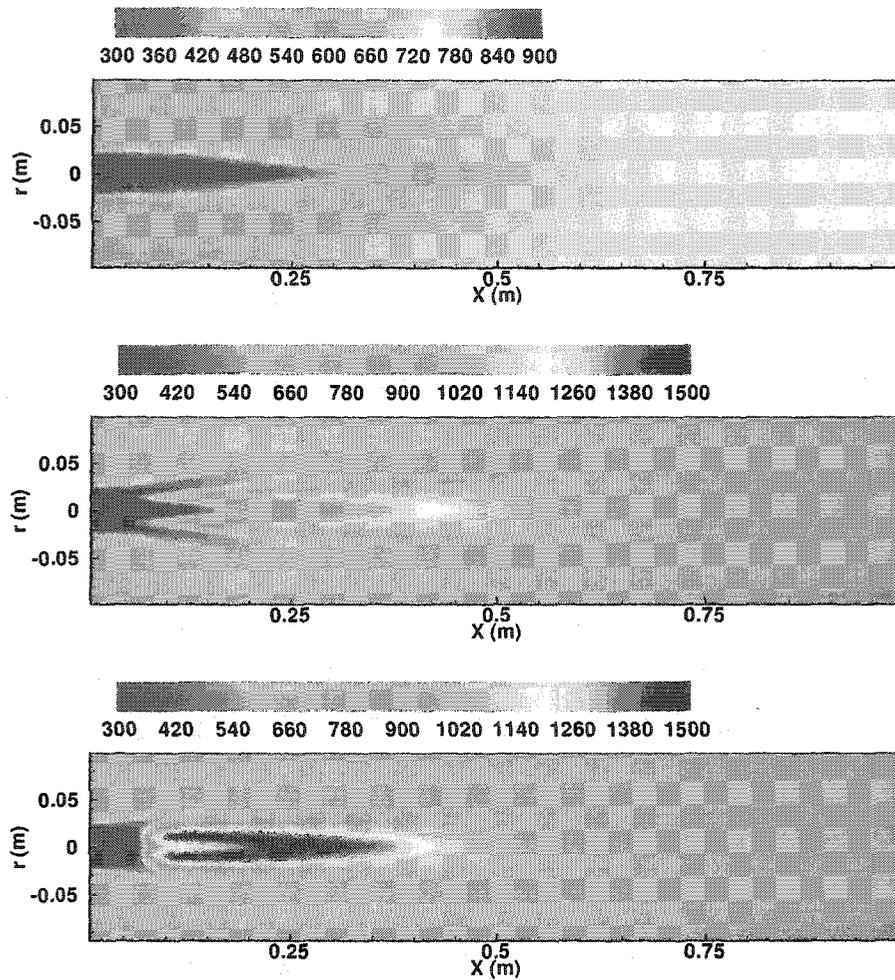


Figure 5.29 Final temperature profiles for the three cases corresponding to the reference points in Fig. 5.28. (top) Infinite-rate mixing lower steady-state, (middle) finite-rate mixing steady-state, and (bottom) Infinite-rate mixing higher steady-state.

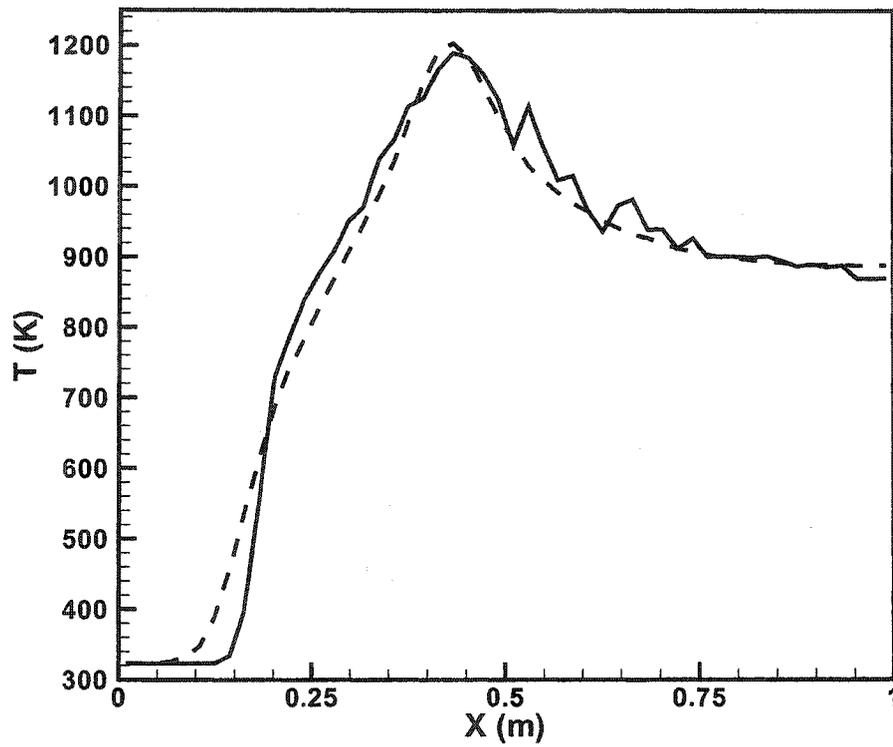


Figure 5.30 Test for grid-independence of finite-rate mixing case. Dashed lines represent the results from the original grid while solid lines show the results from the grid with twice the number of grid points in each direction.

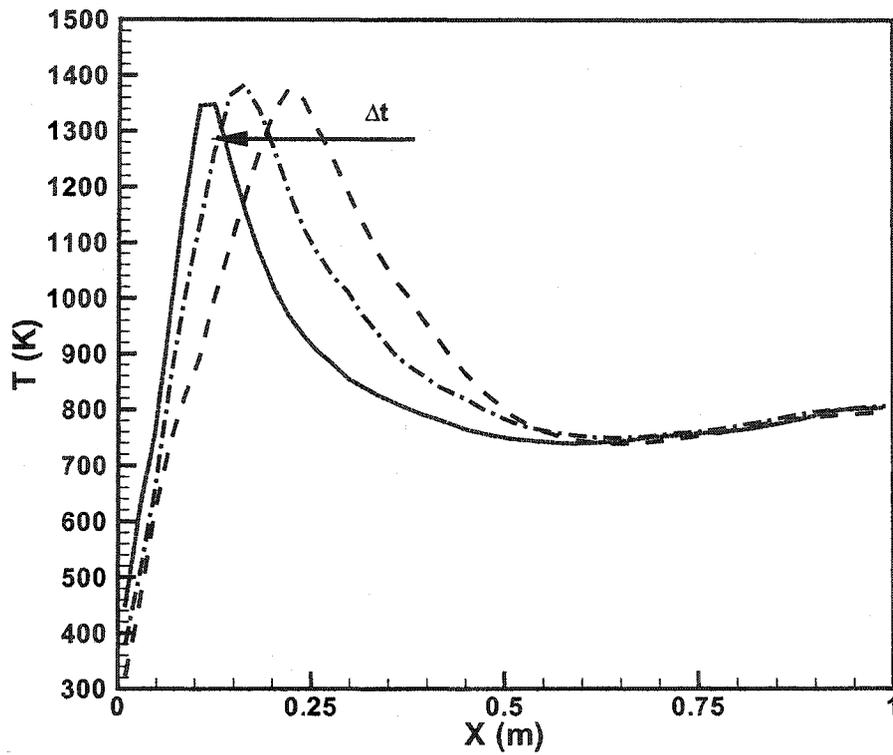


Figure 5.31 Effect of time step on axial temperature profile. The arrow points to the direction of decreasing time steps. Dashed line - $\Delta t = 10^{-3}$. Dotted line - $\Delta t = 5 \times 10^{-4}$. Solid line - $\Delta t = 2.5 \times 10^{-4}$.

6 Conclusions and Future Work

In this chapter, the major objectives met in this thesis work are discussed. Further, the aspects of the chlorination reactor studied in the preceding chapters are summarized. Future directions with regard to the numerical methods and reactor design are identified.

Hybrid Scheme

Use of PDF methods aid in providing a more detailed description of reaction without simplifying assumptions about reaction mechanisms or the turbulent flow. However, the deficiencies of the method identified in Chapter 2 necessitate the use of a hybrid scheme. Several case studies discussed in the previous chapters have conclusively proved that the hybrid scheme can be used as a general tool for all turbulent flows. The state-of-the-art algorithms used for the Monte-Carlo scheme has expanded its capability to handle complex geometries and large chemistry mechanisms. The limitation of the hybrid scheme is based on the limitation of the flow solver only. Use of such a hybrid scheme ensures that any advances made in the solution methods for the flow-field can be readily used for reacting flows. In a similar way, the micromixing model, which is an active area of research and development, can be replaced with better models as and when they are available. The ISAT scheme for handling chemistry mechanism is found sufficient to handle some of the largest mechanisms available.

The non-reacting flow test cases showed that the PDF results agree well with FV profiles for the mean flow. However the FV code underpredicts the variance while the PDF code agrees well with the analytical profiles for the variance. The test case with the arbitrary grid showed that the PDF code is able to handle complex grid structures without using any of the properties of the grid itself. The algorithms employed were found to scale linearly with number of particles and number of cells. This indicated that the non-linear sorting algorithms used in traditional PDF codes were one of the major pitfalls towards adopting such codes for general use. It should however be noted that

the PDF code based profiles are only as good as the underlying flow-field and hence the flow-solver should be able to handle complex geometries with limited numerical diffusion arising from algorithms for such grid structures.

The hybrid scheme was used to simulate a turbulent diffusion flame stabilized by a pilot. The Sandia D flame has been studied using several computational techniques with regard to the flow field. However, the reactions are invariably handled using reduced mechanisms with 4-10 species. The only other study using a detailed chemistry has been reported by Xu and Pope [2000]. Flamelet model has been used to produce reasonable agreement with experimental data [Pitsch and Steiner, 2000]. However none of the simulations are able to predict the major as well as minor species and the flow field accurately. The hybrid scheme used here was shown to predict the flame accurately in all aspects. It was also observed that the prediction of minor species is still dependent on the mechanism used. Many important features of the hybrid algorithm like grid-independent time-stepping and scale-up of tracking algorithms were studied. It was found that while choosing the time step for the simulation, the minimum local mixing time should be taken into consideration. The feedback mechanism involving the transfer of the temperature and mean mixture molecular weight were found to accurately evolve the flow field. The use of the temperature field as the only feedback variable led to mass imbalance in the flow-solver resulting in non-uniform particle distribution in the domain. These results validate the FV-PDF scheme and the comparison with experimental data show that even with simple mixing models, the hybrid scheme is more accurate than FV schemes with no closure.

Thermochlorination Reactor

Chlorination schemes are industrially important reactions and are similar to several other processes used in the chemical process industry. The simple reactor models used traditionally fail to predict the correct trends observed from plant-scale reactor. Use of a comprehensive scheme that takes into account the turbulence-chemistry interaction using a turbulence model (as in a RANS solver) and detailed chemistry (as in a PDF scheme) is essential in understanding the dynamics of the reactor. The chlorination chemistry is similar to many combustion processes. Similar to combustion phenomena, prediction of ignition is a key aspect of the simulation. Since micromixing plays an important role during ignition, use of simple CSTR models will lead to large errors. Even non-ideal zero-dimensional reactors are not sufficient since these models neglect

the spatial dependence of flow.

Use of a constant-density PDF method with frozen flow fields proved that the Monte-Carlo scheme can be used to simulate the chlorination reactor using detailed chemistry. The simple IEM mixing model had been used in the test cases and was found to be sufficient representation of the micromixing phenomena. The premixed reactor is a very stable configuration. Subsequent simulations with non-constant density FV scheme showed that the density changes did not affect the nature of the reaction zone though minor quantitative changes were observed in terms of product yield. It was also observed that while the CSTR model predicts extinction under the same circumstances, the PDF scheme that takes into account the inhomogeneous nature of both turbulence and chemistry was found to yield a stable reactor. This clearly showed that zero-dimensional models overpredict reactor extinction.

The premixed case offers the best performance with high chloroform yield and maximum reactor stability. In the case of segregated inlets, it is safer to use chlorine as the inner stream. This prevents a high mass fraction of chlorine or chlorine free radicals in the recirculation zone. Though reactor performance in terms of product yield is only slightly lower than the premixed case, the reactor operation itself is unstable and is prone to extinction. Accurate inlet control needs to be exercised to keep the reactor lit. Also with such high temperatures, the formation of carbon soot is more probable and may lead to deposition of particles on the reactor surface. This will limit the industrial throughput of the coaxial reactor. The last case that was studied used chlorine in the outer stream. This conceivably led to high temperatures with much of the enthalpy trapped in the recirculation zone. The dynamics of the reactor are inherently unstable with a high probability of reactor run-away. The reactor stability is determined only by the fraction of organic stream that can get inside the recirculation zone. Higher mass fractions will eventually lead to complete hydrocarbon decomposition and formation of soot particles. Hence, the interaction of jet mixing and reaction is important in predicting the reactor performance.

The scatter plots showed that the reaction equilibrium is attained at very high temperatures for high chlorine fractions. It also showed that the chloroform yield can be as high as 28% by mass while even the premixed case can produce only 14%. The reaction branch also showed that it is detrimental to have high chlorine or organic mass fraction as this might lead to product decomposition at high temperatures. The conditioning on the mixture fraction reduces the dimensionality of the problem. However, simple hydrodynamic models can be used to study the equilibrium surface for different inlet

feed compositions. It was also seen that in the non-premixed case with chlorine as the inner stream, the reactor stays lit and the particles near the outlet lie on the reaction branch. However, in the case with chlorine as the outer jet, the particles near the outer jet collapse onto the mixing branch showing global extinction. However, particles near the reacting zone were found to produce higher mass fraction of chloroform. Longer residence time in the high temperature recirculating flow eventually led to product decomposition.

It was found that closure at the sub-grid level for the scalars changes the structure and location of the reaction zone. For non-premixed inlets, use of completely micromixed conditions leads to higher temperatures and lesser product yield than those cases with a sub-grid level micromixing model. However, for the premixed case, faster mixing widens the jet core and dilutes the recirculation region. This leads to lower temperatures and a decrease in the chlorine consumption. Further study is required that may involve other micromixing closures to determine the effect of the choice of the model on reaction. The present study provides an understanding of the fundamental dynamics of the reactor. It is seen the recirculation zone helps to sustain the reactor by enthalpy feedback. The different feed-stream configurations change only the location of the reaction zone. The stability is still provided by limited heat recirculation back to the cold inlet stream to pre-heat the reactants until ignition. It is also seen that when the enthalpy is trapped in a limited reaction zone (as in the case with chlorine in the outer jet), global extinction might occur even though hot spots are present near the inlet zone.

Further simulations were carried out on plant-scale reactors to demonstrate the practical use of the simulation scheme. The plant reactor with an inlet sparger was simulated with premixed inlet conditions and was found to yield results very similar to the premixed pilot-scale reactor. The hybrid scheme can then effectively be used to study scale-up of industrial reactors. Plant-operating conditions were then directly used to simulate stable and unstable operations. It was observed that the hybrid scheme using a micromixing model was able to predict extinction while the simulations using infinite-rate mixing predicted a stable reactor. This demonstrated that the use of the micromixing model, though inaccurate in the premixed limit, is important to represent the finite-rate mixing process. It should also be noted that the region where mixing is important is the reaction zone which is partially-premixed due to recirculating "burnt" feed. Hence the use of the mixing model is not completely inaccurate. These simulations have shown the hybrid scheme is an efficient and accurate way to capture the dynamics of the reactor.

Future Directions

As mentioned earlier, the hybrid scheme is limited only by the capability of the flow solver. All RANS schemes are limited by their turbulence models which are not universal and have been shown in previous chapters to involve tuning based on problem conditions. This suggests that a more reliable flow-solver is required. In this context, use of LES schemes as flow solvers can increase the accuracy and reliability of the hybrid scheme. Recent work by Colucci et al. [1998] have demonstrated such techniques for 2-dimensional configurations. Another alternative will be to use velocity-composition PDF scheme where both the velocity and scalar fields are solved by the PDF scheme. The hybrid scheme then uses Reynolds-stress terms calculated from the particle code. This will provide a better alternative to using the turbulent-viscosity hypothesis but problems arising from statistically noisy mean fields need to be handled consistently. Time averaging techniques similar to the methods used in this work can be employed. Simulation of 3-d flows with the hybrid scheme will be computationally expensive if not intractable. Use of parallel algorithms are a requirement to handle such large calculations. Parallel tracking schemes have been explored elsewhere [Cheng and Plassman, 2002] but the speed-up on large-processor machines have not been tested.

The micromixing model used in the PDF scheme is the subject of intense research and progress towards a model-constant free scheme is yet to be formulated. However, promising results [Fox et al., 2002] in this direction are being pursued currently. Similarly, detailed kinetic schemes have been formulated for several reactions but the accuracy and validity of these mechanisms over a wide range of operating conditions need to be tested. It was seen that optimized chemistry schemes like the GRI mechanisms are not necessarily accurate in predicting experimental results. The PDF schemes can also be extended to include particle precipitation and soot formation [Wang and Fox, 2003]. Fine particle formation in flames and the prediction of soot formation in fuel-rich flames are of industrial importance and reliable techniques using PDF schemes have been developed similar to the scheme used here. The application of such extensions to practical flows aimed towards design of the industrial reactors is an active area of research [Fox, 1996]. These developments are geared towards industrial reactors signaling a change from combustion driven research to industrial process design.

In general, the composition PDF method used in this work has widespread applications in the field of engineering. It can be used in conjunction with existing numerical methods to improve their reliability and accuracy. This particular aspect of the PDF method will be a key feature that will be utilized in future applications.

Bibliography

- Acharya, S., Jang, D. S., West, D. H., and Hebert, L. A. (1991). A moment closure method for modeling a multistep chlorination reaction. In *AIChE Spring National Meeting*, Houston. Paper 33a.
- Akselvoll, K. and Moin, P. (1996). Large eddy simulation of turbulent confined coannular jets. *Journal of Fluid Mechanics*, 315:387–411.
- Anand, M. S., James, S., and Razdan, M. K. (1998). A scalar PDF combustion model for the national combustion code. In *34th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, Cleveland, Ohio.
- Aris, R. (1989). *Vectors, Tensors, and the Basic Equations of Fluid Mechanics*. Dover Publications Inc.
- Baldwin, B. S. and Lomax, H. (1978). Thin-layer approximation and algebraic model for separated turbulent flows. Paper 78-257 AIAA.
- Bardina, J., Ferziger, J. H., and Reynolds, W. C. (1980). Improved subgrid models for large eddy simulation. AIAA paper. 80-1357.
- Barlow, R. S. and Frank, J. H. (1998). Effects of turbulence on specific mass fractions in methane/air jet flames. *Proceedings of the Combustion Institute*, 27:1087–1095.
- Bilger, R. W. (1980). *Turbulent Reacting Flows*, chapter Turbulent flows with Non-premixed reactants, pages 65–113. Springer, Berlin.
- Bilger, R. W. (1993). Conditional moment closure for turbulent reacting flow. *Physics of Fluids*, 5(2):436–444.
- Bird, R. B., Stewart, W. E., and Lightfoot, E. N. (1960). *Transport Phenomena*. John Wiley and Sons.

- Boersma, B. J. (1999). Direct numerical simulation of a turbulent reacting jet. In *Center for Turbulence Research Annual Research Briefs*, pages 59–72.
- Boussinesq, J. (1877). Essai sur la theorie des eaux courantes. *Mem. Presentes Acad. Sci.*, 23:46.
- Bowman, C. T., Hanson, R. K., Davidson, D. F., and *et al.* (1995). Gri-mech 2.11. http://www.me.berkeley.edu/gri_mech/.
- Chen, J. Y., Kollmann, W., and Dibble, R. W. (1989). Pdf modeling of turbulent non-premixed methane jet flames. *Combustion Science and Technology*, 64:315–346.
- Cheng, J.-R. C. and Plassman, P. E. (2002). A parallel particle tracking framework for applications in scientific computing. Submitted to *Journal of Supercomputing*.
- Colucci, P. J., Jaber, F. A., and Givi, P. (1998). Filtered density function for large eddy simulation of turbulent reacting flows. *Physics of Fluids*, 10(2):499–515.
- Correa, S. M. and 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*, pages 279–285, Pittsburgh. Combustion Institute.
- Curl, R. L. (1963). Dispersed phase mixing I. theory and effects in simple reactors. *AIChE Journal*, 9:175–181.
- Dally, B. B., Fletcher, D. F., and Masri, A. R. (1998a). Flow and mixing fields of turbulent bluff-body jets and flames. *Combustion Theory and Modelling*, 2:193–219.
- Dally, B. B., Masri, A. R., Barlow, R. S., and Fietchner, G. J. (1998b). Instantaneous and mean compositional structure of bluff-body stabilized nonpremixed flames. *Combustion and Flame*, 114:119–148.
- Delarue, B. J. and Pope, S. B. (1997). Application of PDF methods to compressible turbulent flows. *Physics of Fluids*, 9:2704–2715.
- Deutschmann, O. and Schmidt, L. D. (1998). Modeling the partial oxidation of methane in a short-contact-time reactor. *AIChE Journal*, 44:2465.
- Dopazo, C. (1979). Relaxation of initial probability density function in the turbulent connection of scalar fields. *Physics of Fluids*, 22:20.

- Dreeben, T. D. and Pope, S. B. (1997a). Probability density function and Reynolds-stress modeling of near-wall turbulent flows. *Physics of Fluids*, 9(1):154.
- Dreeben, T. D. and Pope, S. B. (1997b). Wall-function treatment in PDF methods for turbulent flows. *Physics of Fluids*, 9(9):2692–2703.
- Edwards, J. R. (1997). A low-diffusion flux-splitting scheme for Navier-Stokes calculations. *Computers and Fluids*, 26:645–659.
- Eswaran, V. and Pope, S. B. (1988). Direct numerical simulations of the turbulent mixing of a passive scalar. *Physics of Fluids*, 31:506.
- Ferziger, J. H. and Peric, M. (2002). *Computational Methods for Fluid Dynamics*. Springer, 3rd edition.
- Fluent Inc. (2001). *FLUENT Manual Version 6.1*. Fluent Inc., Lebanon, New Hampshire, USA.
- Fox, R. O. (1995). The spectral relaxation model of the scalar dissipation rate in homogeneous turbulence. *Physics of Fluids*, 7(5).
- Fox, R. O. (1996). Computational methods for turbulent reacting flows in the chemical process industry. *Revue de l'Institut Francais du Petrole*, 51:215.
- Fox, R. O. (1997). The Lagrangian spectral model of the scalar dissipation rate in homogenous turbulence. *Physics of Fluids*, 9:2364.
- Fox, R. O. (1998). On the relationship between Lagrangian micromixing models and computational fluid dynamics. *Chemical Engineering Processing*, 37:521–535.
- Fox, R. O. (1999). The Lagrangian spectral relaxation model for differential diffusion in homogeneous turbulence. *Physics of Fluids*, 11(6).
- Fox, R. O. (2003). *Computational Models for Turbulent Reactive Flows*. Cambridge University Press, To be published.
- Fox, R. O., Cha, C. M., and Trouillet, P. (2002). Lagrangian mixing model. In *CTR Summer Program*. 369-380.

- Gagnepain, L., Chauveau, C., and Gokalp, I. (1998). A comparison between dynamic and scalar timescales in lean premixed turbulent flames. In *Twenty-Seventh Symposium (International) on Combustion*, pages 775–787, Pittsburgh, PA. The Combustion Institute.
- Gardiner, C. W. (1983). *Handbook of Stochastic Methods*. Springer-Verlag.
- Germano, M. (1991). Turbulence : The filtering approach. *Journal of Fluid Mechanics*, 286(229-255).
- Gicquel, L. Y. M., Givi, P., Jaber, F. A., and Pope, S. B. (2002). Velocity filtered density function for large eddy simulation. *Physics of Fluids*, 14(3):1196–1213.
- Grier, M. R. and Fox, R. O. (1993). A particle-conserving composition joint probability density function code. Technical Report 260, Kansas State University.
- Harvey, A. D. (1998). Chem3d: A compressible formulation Navier-Stokes equation solver. Unpublished FORTRAN code.
- Harvey, A. D. (2002). Private communication.
- Harvey, A. D. and Edwards, J. R. (1999). Preconditioned Navier-Stokes solution technique for laminar methane-air non-premixed combustion. Unpublished.
- Harvey, A. D. and Pitsch, H. (2000). Modeling turbulent reacting methane thermochlorination flows. In *Proceedings of the Summer Program 2000*, pages 181–191. Center for Turbulence Research.
- Haworth, D. C. and El Tahry, S. H. (1991). Probability density function approach for multidimensional turbulent flow calculations with application to in-cylinder flows in reciprocating engines. *AIAA Journal*, 29(2):208–218.
- Haworth, D. C. and Jansen, K. (2000). Large-eddy simulation on unstructured deforming meshes: Toward reciprocating IC engines. *Computers and Fluids*, 29:493–524.
- Haworth, D. C. and Pope, S. B. (1986). A generalized Langevin model for turbulent flows. *Physics of Fluids*, 29:387–405.
- Heinz, S. and Roekaerts, D. (2001). Reynolds number effects on mixing and reaction in a turbulent pipe flow. *Chemical Engineering Science*, 56:3197–3210.

- Hockney, R. W. and Eastwood, J. W. (1981). *Computer Simulation Using Particles*. McGraw-Hill.
- Hulek, T. and Lindstedt, R. P. (1996). Computations of steady-state and transient premixed turbulent flames using PDF methods. *Combustion and Flame*, 104:481–504.
- Hwang, C. C., Zhu, G., Massoudi, M., and Ekmann, J. M. (1993). A comparison of the linear and nonlinear k - ϵ turbulence models in combustors. *Journal of Fluids Engineering*, 115:93.
- Jaberi, F. A., Colucci, P. J., James, S., Givi, P., and Pope, S. B. (1999). Filtered mass density function for large-eddy simulation of turbulent reacting flows. *Journal of Fluid Mechanics*, 401:85–121.
- Janicka, J., Kolbe, W., and Kollmann, W. (1970). Closure of the transport equation for the PDF of turbulent scalar fields. *Journal of Non-Equilibrium Thermodynamics*, 4:47–66.
- Jenny, P., Pope, S. B., Muradoglu, M., and Caughey, D. A. (2001). A hybrid algorithm for the joint PDF equation of turbulent reactive flows. *Journal of Computational Physics*, 166:218–252.
- Jones, W. P. and Khaki, M. (1998). Pdf modeling of finite-rate chemistry effects in turbulent nonpremixed jet flames. *Combustion and Flame*, 115(210-229).
- Juneja, A. and Pope, S. B. (1996). A DNS study of turbulent mixing of two passive scalars. *Physics of Fluids*, 8(8):2161.
- Kee, R. J., Rupley, F. M., and Miller, J. A. (1989). *Chemkin-II: A Fortran Chemical Kinetics Package for the Analysis of Gas Phase Chemical Kinetics*. Sandia Report.
- Kim, J., Moin, P., and Moser, R. D. (1987). Turbulence statistics in fully developed channel flow at low Reynolds number. *Journal of Fluid Mechanics*, 177:133–166.
- Kim, S. H. and Huh, K. Y. (2002). Use of conditional moment closure model to predict NO formation in a turbulent CH_4/H_2 flame over a bluff-body. *Combustion and Flame*, 130(1-2):94–111.
- Klimenko, A. Y. (1995). Note on the conditional moment closure in turbulent shear flows. *Physics of Fluids*, 7(2):446–448.

- Kolhapure, N. and Fox, R. O. (1999). CFD analysis of micromixing effects on polymerization in tubular low-density polyethylene reactors. *Chemical Engineering Science*, 54:3233.
- Launder, B. E., Reece, G. J., and Rodi, W. (1975). Progress in the development of a Reynolds-stress turbulence closure. *Journal of Fluid Mechanics*, 68:537–566.
- Lilly, D. K. (1967). The representation of small-scale turbulence in numerical simulation experiments. In Goldstine, H. H., editor, *Proceedings of the IBM Scientific Computing Symposium on Environmental Sciences*, pages 195–210, Yorktown Heights, NY. IBM.
- Lilly, D. K. (1992). A proposed modification of the Germano subgrid-scale closure method. *Physics of Fluids A*, 4:633–635.
- Lindstedt, R. P., Louloudi, S. A., and Vaos, E. M. (2000). Joint scalar probability density function modeling of pollutant formation in piloted jet diffusion flames with comprehensive chemistry. *Proceedings of the Combustion Institute*, 28:149–156.
- Liu, C. and Barkelew, C. H. (1985). Numerical analysis of jet-stirred reactors with turbulent flows and homogeneous reactions. *AIChE Journal*, 32:1813.
- Maas, U. A. and Pope, S. B. (1992). Simplifying chemical kinetics: Intrinsic low dimensional manifolds in composition space. *Combustion and Flame*, 88:239.
- Masri, A. R. and Pope, S. B. (1990). Pdf calculations of piloted turbulent nonpremixed flames of methane. *Combustion and Flame*, 81(1):13–29.
- Masri, A. R., Subramaniam, S., and Pope, S. B. (1996). A mixing model to improve pdf simulation of turbulent diffusion flames. *Twenty-Sixth Symposium (International) on Combustion*, 26:49.
- Mobus, H., Gerlinger, P., and Bruggemann, D. (2001). Comparison of Eulerian and Lagrangian Monte Carlo PDF method for turbulent diffusion flames. *Combustion and Flame*, 124:519–534.
- Moin, P., Squires, K., Cabot, W., and Lee, S. (1991). A dynamic subgrid-scale model for compressible turbulence and scalar transport. *Physics of Fluids A*, 3:2746–2757.
- Muradoglu, M., Jenny, P., Pope, S. B., and Caughey, D. A. (1999). A consistent hybrid finite-volume/particle method for the PDF equations of turbulent reactive flows. *Journal of Computational Physics*, 154:342–371.

- Muradoglu, M. and Pope, S. B. (2002). A local time stepping for solving the PDF equations of turbulent reacting flow. *AIAA Journal*, 40:1755–1763.
- Nguyen, T. V. and Pope, S. B. (1984). Monte Carlo calculations of turbulent diffusion flames. *Combustion Science and Technology*, 42:13–45.
- Oran, E. S. and Boris, J. P. (2001). *Numerical Simulation of Reactive Flow*. Cambridge University Press, 2nd edition.
- Panchapakesan, N. R. and Lumley, J. L. (1993). Turbulence measurements in axisymmetric jets of air and helium. part 2. helium jet. *Journal of Fluid Mechanics*, 246:225.
- Papoulis, A. (1991). *Probability, Random Variables, and Stochastic Processes*. McGraw-Hill, Inc.
- Peters, N. (2000). *Turbulent Combustion*. Cambridge University Press.
- Petzold, L. and Zhu, W. (1999). Model reduction for chemical kinetics: An optimization approach. *AIChE Journal*, 45(4):869.
- Pipino, M. and Fox, R. O. (1994). Reactive mixing in a tubular jet reactor: A comparison of PDF simulations with experimental data. *Chemical Engineering Science*, 49(24B):5229–5241.
- Pitsch, H. (2003). Improved pollutant predictions in large-eddy simulations of turbulent non-premixed combustion by considering scalar dissipation rate fluctuations. *Proceedings of the Combustion Institute*, 29:1971–1978.
- Pitsch, H. and Steiner, H. (2000). Large-eddy simulation of a turbulent piloted methane/air diffusion flame (sandia flame D). *Physics of Fluids*, 12(10):2541–2554.
- Pope, S. B. (1976). The probability approach to the modelling of turbulent reacting flows. *Combustion and Flame*, 27:299–312.
- Pope, S. B. (1978). An explanation of the turbulent round-jet/plane-jet anomaly. *AIAA Journal*, 16:279–281.
- Pope, S. B. (1981). A Monte Carlo method for the PDF equations of turbulent reactive flow. *Combustion Science and Technology*, 25:159–174.
- Pope, S. B. (1982). An improved turbulent mixing model. *Combustion Science and Technology*, 28:131–135.

- Pope, S. B. (1985). PDF methods for turbulent reactive flows. *Progress in Energy and Combustion Science*, 11:119.
- Pope, S. B. (1991). Mapping closures for turbulent mixing and reaction. *Theoretical and Computational Fluid Dynamics*, 2:255–270.
- Pope, S. B. (1994a). Lagrangian PDF methods for turbulent flows. *Annual Review of Fluid Mechanics*, 26:23–63.
- Pope, S. B. (1994b). Pdf2dv: A code to solve the velocity-turbulence frequency-composition joint pdf equation. Unpublished FORTRAN code.
- Pope, S. B. (1997). Computationally efficient implementation of combustion chemistry using in-situ adaptive tabulation. *Combustion Theory Modelling*, 1:41.
- Pope, S. B. (2000). *Turbulent Flows*. Cambridge University Press.
- Pope, S. B. and Chen, Y. L. (1990). The velocity-dissipation probability function model for turbulent flows. *Physics of Fluids A*, 2:1437–1449.
- Raju, M. S. (1996). Application of scalar Monte Carlo probability density function method for turbulent spray flames. *Numerical Heat Transfer, Part A*, 30:753–777.
- Raman, V., Fox, R. O., and Harvey, A. D. (2002). Hybrid finite-volume / transported PDF simulations of a partially premixed methane-air flame. Submitted to *Combustion and Flame*.
- Raman, V., Fox, R. O., Harvey, A. D., and West, D. H. (2001). Cfd analysis of premixed methane chlorination reactors with detailed chemistry. *Industrial and Engineering Chemistry Research*, 40:5170–5176.
- Raman, V., Fox, R. O., Harvey, A. D., and West, D. H. (2003). Effect of feed-stream configuration on gas-phase chlorination reactor performance. *Industrial and Engineering Chemistry Research*, 42:2544–2557.
- Reed, D. J. (1991). *Encyclopedia of Chemical Technology*. Wiley, New York, 4 edition.
- Reveillon, J. and Vervisch, L. (1998). Subgrid-scale turbulent micromixing: Dynamic approach. *AIAA Journal*, 36(3):336–341.
- Rodi, W. and Mansour, N. N. (1993). Low Reynolds number $k - \epsilon$ modeling with the aid of direct simulation data. *Journal of Fluid Mechanics*, 250:509–529.

- Roekaerts, D. (1991). Use of a Monte Carlo PDF method in a study of the influence of turbulent fluctuations on selectivity in a jet-stirred reactor. *Applied Scientific Research*, 48:271.
- Rogers, M. M. and Moser, R. D. (1994). Direct numerical simulation of a self-similar turbulent mixing layer. *Physics of Fluids*, 6:903–923.
- Roomina, M. A. and Bilger, R. W. (2001). Conditional moment closure (CMC) predictions of a turbulent methane-air flame. *Combustion and Flame*, 125(3):1176–1195.
- Shah, J. J. and Fox, R. O. (1999). CFD simulation of chemical reactors: Application of in-situ adaptive tabulation to methane thermochlorination chemistry. *Industrial and Engineering Chemistry Research*, 38(11):4200–4212.
- Shuen, J. S., Chen, K. H., and Choi, Y. (1993). A coupled implicit method for chemical non-equilibrium flows at all speeds. *Journal of Computational Physics*, 106:306–318.
- Smagorinsky, J. (1963). General circulation experiments with the primitive equations: I. the basic equations. *Mon. Weather Review*, 91:99–164.
- Smith, G. P., Golden, D. M., Frenklach, M., and *et al.* (2000). Gri-mech 3.0. http://www.me.berkeley.edu/gri_mech/.
- Subramaniam, S. and Haworth, D. (2000). A PDF method for turbulent mixing and combustion on three-dimensional unstructured deforming meshes. *International Journal of Engine Research*, 1:171–190.
- Subramaniam, S. and Pope, S. B. (1999). Comparison of mixing model performance for non-premixed turbulent reactive flow. *Combustion and Flame*, 117:732.
- Sung, C. J., Law, C. K., and Chen, J. Y. (1998). An augmented reduced mechanism for methane oxidation with comprehensive global parametric validation. *Twenty-Seventh Symposium (International) on Combustion*, 27:295–304.
- Tang, Q. and Pope, S. B. (2003). Implementation of combustion chemistry by in situ adaptive tabulation of rate-controlled constrained equilibrium manifolds. *Proceedings of Combustion Institute*, 29:1411–1417.
- Tang, Q., Xu, J., and Pope, S. B. (2000). Probability density function calculations of local extinction and no production in pilot-jet turbulent methane/air flames. *Proceedings of Combustion Institute*, 28:133–139.

- Tirtowidjo, M. (1997). Fundamental kinetic modeling of industrial reactors. In *AICHE Annual Meeting*, Los Angeles. paper 78b.
- TNF (2000). *Fifth International Workshop on Measurement and Computation of Turbulent Non-premixed Flames*, Delft, The Netherlands. Page 226-227.
- TNF (2002). *Sixth International Workshop on Measurement and Computation of Turbulent Non-premixed Flames*, Sapporo, Japan.
- Tolpadi, A. K., Correa, S. M., Burrus, D. L., and Mongia, H. C. (1997). Monte Carlo probability density function method for gas turbine combustor flowfield predictions. *AIAA*, 13(2):218-225.
- Tsai, K. (2002). Private communication.
- Tsai, K. and Fox, R. O. (1994). Pdf simulation of a turbulent series-parallel reaction in an axisymmetric reactor. *Chemical Engineering Science*, 49(24B):5141-5148.
- Tsai, K. and Fox, R. O. (1995). Modeling multiple reactive scalar mixing with the generalized IEM model. *Physics of Fluids*, 7(11):2820-2830.
- Tsai, K. and Fox, R. O. (1996). PDF modeling of turbulent-mixing effects on initiator efficiency in a tubular LDPE reactor. *AICHE Journal*, 42:2926.
- Tsai, K. and Fox, R. O. (1998). The BMC/GIEM model for micromixing in non-premixed turbulent reacting flows. *Industrial Engineering Chemistry Research*, 37(6):2131-2141.
- Udaipurwala, A., Fox, R. O., and Chapman, K. S. (1994). Comparative study of turbulence models for predicting turbulent flow in a confined coaxial jet reactor. Technical Report 261, Kansas State University.
- Valino, L. (1995). Multiscalar mapping closure for mixing in homogeneous turbulence. *Physics of Fluids*, 7:144.
- Valino, L. and Dopazo, C. (1991). A binomial Langevin model for turbulent mixing. *Physics of fluids A*, 3:3034.
- Van Slooten, P. R., Jayesh, and Pope, S. B. (1998). Advances in PDF modeling for inhomogeneous turbulent flows. *Physics of Fluids*, 10:246-265.

- Villermaux, J. (1986). *Micromixing phenomena in stirred reactors in Encyclopedia of Fluid Mechanics*, chapter 27. Gulf, Houston, TX.
- Villermaux, J. and Falk, L. (1994). A generalized mixing model for initial contacting of reactive fluids. *Chemical Engineering Science*, 49(24B):5127–5140.
- Wang, L. and Fox, R. O. (2003). Application of in-situ adaptive tabulation to CFD simulation of nano-particle formation by reactive precipitation. Accepted for publication in *Chemical Engineering Science*.
- Warnatz, J., Maas, U., and Dibble, R. W. (1996). *Combustion*. Springer.
- West, D. H., Hebert, L. A., and Pividal, K. A. (1999). Detection of quenching and instability in industrial chlorination reactors. In *Fall AIChE Annual Meeting*, Dallas.
- Wilcox, D. C. (1994). *Turbulence Modeling for CFD*. DCW Industries, Inc.
- Williams, F. A. (1975). Recent advances in theoretical descriptions of turbulent diffusion flames. In Murthy, S. N. B., editor, *Turbulent Mixing in Nonreactive and Reactive Flows*, pages 189–208. Plenum Press, New York.
- Wiseman, P. (1972). *An Introduction to Industrial Organic Chemistry*. Wiley, 1 edition.
- Wouters, H. A. (1998). *Lagrangian models for turbulent reacting flows*. PhD dissertation, Delft University of Technology.
- Wouters, H. A., Nooren, P. A., Peeters, T. W. . J., and Roekaerts, D. (1998). Effects of micro-mixing in gas-phase turbulent jets. *International Journal of Heat and Fluid Flow*, 19:201–207.
- Xu, J. and Pope, S. B. (1999). Assessment of numerical accuracy of PDF/Monte Carlo methods for turbulent reacting flows. *Journal of Computational Physics*, 152:192–230.
- Xu, J. and Pope, S. B. (2000). Pdf calculations of turbulent nonpremixed flames with local extinction. *Combustion and Flame*, 123:281–307.
- Yang, B. and Pope, S. B. (1998a). An investigation of the accuracy of manifold methods and splitting schemes in the computational implementation combustion chemistry. *Combustion Flame*, 112:16.

Yang, B. and Pope, S. B. (1998b). Treating chemistry in combustion with detailed mechanisms - in situ adaptive tabulation in principal directions - premixed combustion. *Combustion and Flame*, 112:85–112.