

**Turbulent Premixed Combustion Simulation with
Conditional Source-term Estimation and Linear-Eddy
Model Formulated PDF and SDR Models**

by

Hong P. Tsui

B.Sc. (Honours), The University of British Columbia, 2008

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

in

The Faculty of Graduate and Postdoctoral Studies
(Mechanical Engineering)

THE UNIVERSITY OF BRITISH COLUMBIA
(Vancouver)

December 2016

© Hong P. Tsui, 2016

Abstract

Computational fluid dynamics (CFD) is indispensable in the development of complex engines due to its low cost and time requirement compared to experiments. Nevertheless, because of the strong coupling between turbulence and chemistry in premixed flames, the prediction of chemical reaction rates continues to be a modelling challenge.

This work focuses on the improvement of turbulent premixed combustion simulation strategies requiring the use of presumed probability density function (PDF) models. The study begins with the development of a new PDF model that includes the effect of turbulence, achieved by the implementation of the Linear-Eddy Model (LEM). Comparison with experimental burners reveals that the LEM PDF can capture the general PDF shapes for methane-air combustion under atmospheric conditions with greater accuracy than other presumed PDF models. The LEM is additionally used to formulate a new, pseudo-turbulent scalar dissipation rate (SDR) model.

Conditional Source-term Estimation (CSE) is implemented in the Large Eddy Simulation (LES) of the Gülder burner as the closure model for the chemistry-turbulence interactions. To accommodate the increasingly parallel computational environments in clusters, the CSE combustion module has been parallelised and optimised. The CSE ensembles can now dynamically adapt to the changing flame distributions by shifting their spatial boundaries and are no longer confined to pre-allocated regions in the simulation domain. Further, the inversion calculation is now computed in parallel using a modified version of an established iterative solver, the Least-Square QR-factorisation (LSQR). The revised version of CSE demonstrates a significant reduction in computational requirement — a reduction

of approximately 50% — while producing similar solutions as previous implementations.

The LEM formulated PDF and SDR models are subsequently implemented in conjunction with the optimised version of CSE for the LES of a premixed methane-air flame operating in the thin reaction zone. Comparison with experimental measurements of temperature reveals that the LES results are very comparable in terms of the flame height and distribution. This outcome is encouraging as it appears that this work represents a significant step towards the correct direction in developing a complete combustion simulation strategy that can accurately predict flame characteristics in the absence of *ad hoc* parameters.

Preface

The content of Chapter 3 is published in:

- H. P. Tsui and W. K. Bushe. Linear-eddy model formulated probability density function and scalar dissipation rate models for premixed combustion. *Flow, Turbulence and Combustion*, 93:487–503, 2014.

I was responsible for conducting all parts of the research and preparing the final manuscript in this publication. Dr. W. Kendal Bushe supervised the research, reviewed the manuscript and provided feedback as appropriate.

The content of Chapter 4 is published in:

- H. P. Tsui, M. M. Kamal, S. Hochgreb and W. K. Bushe. Direct comparison of PDF and scalar dissipation rates between LEM simulations and experiments for turbulent, premixed methane air flames. *Combustion and Flame*, 165:208–222, 2016.

I was responsible for conducting all parts of the numerical research and preparing the final manuscript in this publication. Dr. Simone Hochgreb and Mr. M. M. Kamal of University of Cambridge provided the relevant experimental data and expertise for the analysis of the measured results. Dr. W. Kendal Bushe supervised the research and provided feedback as appropriate. Dr. Hochgreb and Dr. Bushe reviewed the manuscript and provided feedback as appropriate.

The content of Chapter 5 is published in:

- H. P. Tsui and W. K. Bushe. Conditional source-term estimation using dynamic ensemble selection and parallel iterative solution. *Combustion, Theory and Modelling*, 2016. DOI: 10.1080/13647830.2016.1178811

I was responsible for conducting all parts of the research and preparing the final manuscript in this publication. Dr. W. Kendal Bushe supervised the research, reviewed the manuscript and provided feedback as appropriate.

The content of Chapter 6 has been accepted for presentation in the Combustion Symposium (2016):

- H. P. Tsui, M. M. Salehi and W. K. Bushe. LES of a turbulent premixed flame with conditional source-term estimation and linear-eddy model presumed PDF. 36th Combustion Symposium, Seoul, Korea, 2016.

I was responsible for conducting all parts of the research and preparing the final manuscript in this publication. Dr. M. M. Salehi provided feedback for the simulation methodologies as appropriate and reviewed the manuscript. Dr. W. Kendal Bushe supervised the research, reviewed the manuscript and provided feedback as appropriate.

Table of Contents

Abstract	ii
Preface	iv
Table of Contents	vi
List of Tables	x
List of Figures	xiii
List of Abbreviations	xx
Acknowledgments	xxi
1 Introduction	1
1.1 Introduction	1
1.2 Outline	3
2 Literature Survey	5
2.1 Background	5
2.2 Conservation Equations	5
2.3 Transport Properties	8
2.4 Thermodynamics	10
2.5 Chemistry	12
2.5.1 Chemical Kinetics	12
2.5.2 Reduced Chemistry	13

Table of Contents

2.6	Laminar Premixed Flame	14
2.7	Scales of Turbulence	18
2.8	Regimes of Premixed Turbulent Combustion	22
2.9	Computational Fluid Dynamics (CFD)	27
2.10	Turbulent Premixed Combustion Modelling	34
3	Pseudo-Turbulent Probability Density Function and Scalar Dissipation Rate Models for Premixed Combustion	44
3.1	Introduction	44
3.2	Probability Density Function Models	46
3.3	Linear-Eddy Model	49
3.3.1	Evolutionary Equations	49
3.3.2	Stochastic Advection	50
3.4	LEM Simulation Parameters	51
3.4.1	LEM Methodology for Experimental Study	51
3.4.2	LEM Methodology for DNS Study	53
3.5	Construction of PDF Models	53
3.5.1	Convergence	56
3.6	Scalar Dissipation Rate	57
3.7	Results and Discussion	59
3.7.1	Probability Density Functions	59
3.7.2	Scalar Dissipation Rate	63
3.8	Concluding Remarks	64
4	Direct Comparison of PDF and SDR between LEM Simulations and Experiments for Turbulent, Premixed Methane-Air Flames	65
4.1	Introduction	65
4.2	Numerical Conditions: Premixed Combustion	69
4.2.1	Linear-Eddy Model	69
4.2.2	LEM Simulation Methods	71
4.3	Experimental Conditions: Stratified Swirl Burner	75
4.4	Construction of PDF Models	77
4.4.1	LEM Flame Profiles	77

Table of Contents

4.4.2	LEM PDF Construction	79
4.5	Results	82
4.5.1	Probability Density Function	82
4.5.2	Scalar Dissipation Rate	86
4.6	Discussion	91
4.6.1	Probability Density Function	91
4.6.2	Scalar Dissipation Rate	93
4.7	Conclusion	95
5	Conditional Source-term Estimation: Parallel Iterative Solution with Dynamic Ensembles	97
5.1	Introduction	97
5.2	Conditional Source-term Estimation	99
5.3	Optimisations	102
5.3.1	Dynamic Ensemble Selection	102
5.3.2	Matrix Inversion - Iterative Solver	110
5.4	Numerical Methods	113
5.4.1	Standalone Inversion Tests	113
5.4.2	LES of Axisymmetric Burner	114
5.5	Results and Discussion	116
5.5.1	Standalone Inversion Tests	116
5.5.2	LES of Axisymmetric Burner	121
5.6	Conclusion	125
6	LES of Turbulent Premixed Flames with Optimised CSE and LEM Submodels	126
6.1	Introduction	126
6.2	Theory	128
6.2.1	Governing Equations	128
6.2.2	Chemistry Reduction	128
6.2.3	Presumed PDF and SDR models	129
6.3	Numerical Methods / Simulation Setup	132
6.3.1	Pre-Process	132

Table of Contents

6.3.2	Non-Reactive Simulation	132
6.3.3	Reactive Simulation	134
6.4	Results and Discussion	135
6.5	Concluding Remarks	138
7	Conclusion and Future Work	139
7.1	Conclusions	139
7.2	Future Work	145
	Bibliography	147

List of Tables

Table 4.1	Relevant LEM simulation parameters: l_0 and η_K are the integral and Kolmogorov scales. Other constant parameters are invariant between the cases, including δ_f , s_L , C_λ and N_η , which are respectively, 588 μm , 0.214 m/s, 15.0 and 1.0. A minimum of 1,000 computational cells are used for each simulation.	74
Table 4.2	Operating Conditions: (1) SFR = ratio of split flow to swirlers to total flow, (2) SN = measured swirl number, ratio of tangential to axial momentum, (3) Maximum total u'/s_L at the midpoint of the flame brush at $z = 30$ mm [71, 194].	75
Table 5.1	Relevant LES simulation parameters. The values of a and b are computed according to the approach of Billson [15]: $a = \exp(-\Delta t/\tau_t)$ and $b = \sqrt{1 - a^2}$, where Δt and τ_t are respectively the timestep size and the integral time scale.	116

List of Tables

Table 5.2	Comparison of inversion times between LU-decomposition and LSQR at a number of different settings. The $ e $, $\vec{0}$ and \mathbf{x}_0 respectively denote the error tolerances in the solver, a zero initial solution and a non-zero initial solution. The associated L_2 -norms are printed below the convergence times. Typical matrix sizes for one-condition CSE are represented by the first two columns while typical matrix sizes for two-condition CSE are represented by the last three columns. Failed cases are denoted by ‘×’, where lower error tolerances are required to provide stable solutions. From the smallest to the largest matrix, 1,000, 500, 100, 20 and 2 samples are used to obtain the average inversion times.	118
Table 5.3	Comparison of inversion times for single ensembles of varying sizes with the MPI-LSQR implementation using a number of different processors. The numbers $x \times y$ in the first column represent the number of nodes (x) and the number of processors per node (y) used in the computation. The error tolerances are kept constant at 10^{-10} between all of the cases for consistency.	120
Table 5.4	Summary of simulation and execution times. The execution times for LSQR with Semi-Dynamic Ensemble Selection are typically 50% of the LU-decomposition method for larger systems. The executions times for MPI-LSQR with Dynamic Ensemble Selection are further decreased by 2–22% in comparison with the LSQR procedure. The ensemble sizes were decreased for the coarse grid (52,500) case in order to maintain locality for cell information, leading to more comparable execution times.	125
Table 6.1	Summary of experimental conditions. Symbols: U_B = bulk flow velocity in the central jet; u' = total turbulence intensity; Λ_L = longitudinal integral length scale; Re_t = turbulent Reynolds number; Ka = turbulent Karlovitz number; Da = turbulent Damköhler number. The values are retrieved from [173].	133

List of Tables

Table 6.2	Relevant synthetic turbulence parameters. Symbols: h_g = grid location; K = turbulent kinetic energy; Λ = integral length; a and b = temporal correlation parameters; Δt = timestep size. . .	134
Table 6.3	Summary of simulation times.	138

List of Figures

Figure 2.1	One-dimensional laminar flame structure for premixed, stoichiometric, methane-air combustion at standard conditions. . .	15
Figure 2.2	Laminar flame speeds and equilibrium temperatures of premixed methane-air flames at different equivalence ratios in standard ambient temperature (298.15 K) and pressure (100 kPa) conditions.	17
Figure 2.3	A schematic diagram of the turbulent kinetic energy spectrum at high Reynolds numbers.	21
Figure 2.4	Effect of large-scale turbulent eddies on a premixed flame front.	23
Figure 2.5	Modified Borghi diagram illustrating different turbulent premixed combustion regimes.	24
Figure 2.6	A schematic representation of the resolved and modelled energy spectrum in RANS.	29
Figure 2.7	Schematic view of the resolved and subgrid-scale structures in a grid with spacing of Δ : Resolved structures (light gray) and subgrid-scale structures (dark gray).	31
Figure 2.8	A symbolic representation of the resolved and modelled turbulent kinetic energy spectrum in LES.	32
Figure 2.9	CSE operational flowchart.	43
Figure 3.1	PDFs are exemplified at two axial locations of the DNS simulation domain ($L_d = 6$ mm). Solid: DNS; dash: modified laminar flamelet; dash dot: laminar flamelet; dot: β -PDFs [70].	49

List of Figures

Figure 3.2	<i>Triplet map</i> applied to a line. Cell values are transported within the eddy interval; interpolation is not required at any point [77]. The before (left) and after (right) domain values are illustrated.	51
Figure 3.3	Premixed laminar flame solution at an equivalence ratio of 0.73 for the six-step reduced mechanism (line) and Cantera [58] (symbols).	52
Figure 3.4	(a) (x_1, c_1) and (x_2, c_2) mark the truncation limits where only the cells within the interval are retained; (b) Modified laminar flamelet PDF with $\bar{c} = 0.5$ and $\overline{c'^2}_n = 0.33$ is constructed using the truncated flame profile.	54
Figure 3.5	Individual flame realisations from LEM calculations.	55
Figure 3.6	(a) The portions to be retained are within the intervals marked by 'o'. The truncation positions are selected such that the resultant PDF would have $\bar{c} = 0.5$ and $\overline{c'^2}_n = 0.33$. The truncation boundaries are different for each realisation because the profiles change with time. (b) The modified laminar flamelet (dash-circle) and LEM (solid) PDFs of similar mean and variance are shown.	56
Figure 3.7	LEM generated PDFs at $Re_t = 50$. Each line represents a different number of LEM profiles used for the construction of the PDF. Solid: 5000 profiles; dash: 2000 profiles; dash-dot: 1000 profiles.	57
Figure 3.8	Experimental PDFs conditioned at three different means of the reaction progress variable. Solid: experiment [167]; dot: modified laminar flamelet; dash: LEM. The vertical axis represents the probability at state c^* conditioned by \bar{c} , $P(c) _{\bar{c}}$.	61
Figure 3.9	PDF at $x/L_d = 0.391$ of the DNS simulation domain ($L_d = 6$ mm). Solid: DNS; dash: modified laminar flamelet; dash dot: LEM. The DNS and modified laminar flamelet PDFs are obtained from [70].	62
Figure 3.10	Scalar dissipation conditioned in c -space. Solid: experiment [167]; dot: laminar flame; dash: LEM at $Re_t = 38$; dash-dot: LEM at $Re_t = 200$.	63

Figure 4.1	Premixed laminar flame solution at an equivalence ratio of 0.73 for the six-step global mechanism (line) and Cantera [58] (symbols).	72
Figure 4.2	Borghi diagram showing locations of the ten prototype LEM flames (1a to 4b). The LEM test cases are represented by 'o'. The experimental flames are represented by triangles: SwB1 ('◁'), SwB2 ('▽'), and SwB3 ('▷'). It has been found that increasing the integral length beyond an order of magnitude above the laminar flame thickness while holding the turbulent fluctuations constant does not significantly alter the PDF profiles for the LEM simulations.	73
Figure 4.3	Characteristic LEM temperature profiles of two prototype flames at different Re_t . The individual profiles on each graph are separated by at least one large eddy turnover time.	78
Figure 4.4	(a) (x_1, c_1) and (x_2, c_2) mark the truncation limits; only the cells within the interval are retained. (b) The modified laminar flamelet PDF with $\bar{c} = 0.50$ and $\overline{c'^2}_n = 0.39$ constructed using the truncated flame profile is shown.	80
Figure 4.5	(a) The portions to be retained are within the intervals delimited by the circles. The four truncation positions are selected such that the resultant PDF would have $\bar{c} = 0.5$ and $\overline{c'^2}_n = 0.39$. The truncation boundaries are different for each temperature profile because of the transient effects. (b) The modified laminar flamelet (dash) and LEM (solid) PDFs of similar mean and variance are shown.	81
Figure 4.6	PDF models at various \bar{c} and $\overline{c'^2}_n$ are shown; each row represents one value of \bar{c} and three values of $\overline{c'^2}_n$ (ranging from left to right: 0.25, 0.49 and 0.75). Vertical and horizontal axis on each graph represent the probability and the progress variable, respectively. Solid: case 4b, dash: case 3b, dash dot: case 2b, dot: case 1b (notation is in accordance with Figure 4.2). . . .	83

List of Figures

Figure 4.7	PDFs measured from the SwB1 flame at various axial locations conditioned by different distribution means, \bar{c} (solid). The corresponding LEM PDFs of similar \bar{c} and $\overline{c'^2}_n$ are superimposed (dash).	84
Figure 4.8	PDFs measured from the SwB2 flame at various axial locations conditioned by different distribution means, \bar{c} (solid). The corresponding LEM PDFs of similar \bar{c} and $\overline{c'^2}_n$ are superimposed (dash).	85
Figure 4.9	PDFs measured from the SwB3 flame at various axial locations, conditioned by different distribution means, \bar{c} (solid). The corresponding LEM PDFs of similar \bar{c} and $\overline{c'^2}_n$ are superimposed (dash).	86
Figure 4.10	Non-dimensionalized conditional average of the SDR models from the ten prototype flames. The left and right columns illustrate changes in the turbulent fluctuations and integral length scales, respectively. The turbulent Reynolds number of each case is recorded in the parentheses. The laminar case is displayed in light gray as a reference.	88
Figure 4.11	Non-dimensionalized conditional average of the SDR from the three experimental flames. Each column represents a unique axial position from the flame stabilisation point while each row represents one swirling condition. The error bars indicate +/- one standard deviation from the mean. The laminar case is displayed in light gray as a reference.	89
Figure 4.12	Functional dependence of the scalar dissipation ($f(c^*)$) at three axial positions downstream from the flame stabilisation point (10, 30 and 60 mm) with the LEM results superimposed. The experimental flames are represented by symbols: SwB1 ('+'), SwB2 ('□'), and SwB3 ('◦'). The LEM result is represented by the solid line. The laminar case is displayed in light gray as a reference.	90

List of Figures

Figure 4.13	Values of χ_0 in relation to the distance downstream of the flame stabilisation point. The experimental flames are represented by symbols: SwB1 ('+'), SwB2 ('□'), and SwB3 ('o'). The LEM (black line) and laminar flame (gray line) results are not dependent on the axial position.	90
Figure 4.14	Non-dimensionalized unconditional mean scalar dissipation rates ($\bar{\chi}_c \times \delta_f / s_L$) as predicted by (a) LEM and (b) experiments with LEM solutions superimposed at various combinations of \bar{c} and \bar{c}^2_n . Solid: LEM; dash: SwB1; dash-dot: SwB2; dot: SwB3.	92
Figure 5.1	CSE operational flowchart.	101
Figure 5.2	Computational domain incorporating different number of static ensembles. The volume of each ensemble decreases with increasing number of processors.	104
Figure 5.3	The three possible scenarios that one would encounter during the construction of dynamic ensembles are illustrated: $N < M$, $N = M$ and $N > M$, where N and M are respectively the number of processors available and the number of ensembles required.	105
Figure 5.4	Semi-Dynamic Ensemble Selection processor map.	107
Figure 5.5	Dynamic Ensemble Selection schematic.	109
Figure 5.6	MPI-LSQR iteration. Each n participating processor performs $1/n$ of the total matrix-vector multiplications.	113
Figure 5.7	Comparison of inversion times between LU-decomposition and LSQR at a number of different error tolerance settings ($ e $). This is a graphical representation of the data in Table 5.2. Each of the five plots represents a different matrix size. For each plot, the bars represent inversion times of (from left to right): LU, LSQR ($\vec{0}$) ($ e = 10^{-10}$), LSQR (\mathbf{x}_0) ($ e = 10^{-10}$), LSQR (\mathbf{x}_0) ($ e = 10^{-8}$), LSQR (\mathbf{x}_0) ($ e = 10^{-6}$), and LSQR (\mathbf{x}_0) ($ e = 10^{-4}$). The '×' represents the failed cases.	117

Figure 5.8	Comparison of inversion times with the MPI-LSQR implementation using a number of different processors counts. This is a graphical representation of the data in Table 5.3. Each of the five plots represents a different matrix size. For each plot, the bars represent inversion times of (from left to right): LU, LSQR ($\vec{0}$), LSQR (\mathbf{x}_0), (1×2) , (2×1) , (1×3) , (3×1) , (1×4) , (4×1) , (1×8) , (8×1) , (1×12) , (12×1) , (1×24) , and (1×48) . The numbers $(x \times y)$ represent the number of nodes (x) and the number of processors per node (y) used in the MPI-LSQR computation. The inversion for the matrix of size $100,000 \times 2,500$ is not performed by the LU-decomposition due to an exceedingly exhaustive computational requirement; it is marked by ‘×’.	119
Figure 5.9	The radial temperature and species mass fraction profiles at six axial locations downstream of the inlet for the three grids are shown. The profiles at $z/D = 1$ are shown in black, then each successive jet diameter downstream is shown in a lighter shade of gray, up to $z/D = 6$. The jet diameter, D , is $1.12 \cdot 10^{-2}$ m. The symbols represent: LU–manual: ‘—’; LSQR–semi-dynamic: ‘△’; PLSQR–dynamic: ‘●’.	123
Figure 6.1	Burner cross-section.	133
Figure 6.2	Average binarised temperature-based progress variable profile and characteristic flame height from the LES result (left) and experimental measurements (right), where $D = 11.2$ mm. The solid line represents the half-burning surface, at which the average progress variable takes on a value of 0.5.	136

List of Symbols

Roman Symbols

$\bar{\bar{\mathbf{A}}}$	PDF matrix
A	Pre-exponential factor in Arrhenius equation
A_c	Cross-sectional area of streamtube
A_L	Laminar flame surface area
A_T	Turbulent flame surface area
a	Coefficient for Billson's temporal correlation (1)
$a_{i,\alpha}$	Viscosity fitting coefficients
\mathbf{b}	Mass fraction array
b	Coefficient for Billson's temporal correlation (2)
$b_{i,\alpha}$	Thermal conductivity fitting coefficients
C_{BML}	Bray Moss Libby parameter
C_{EBU}	Eddy Break-Up constant
C_{EDM}	Eddy Dissipation Model constant
C_κ	Universal constant for the energy spectrum in the inertial subrange
C_λ	LEM empirical parameter (1)
c	Progress variable
c^*	Discretised progress variable
\bar{c}	Mean of the progress variable

List of Symbols

\tilde{c}	Favre-averaged mean of the progress variable
$\overline{c'^2}$	Variance of the progress variable
$\widetilde{c'^2}$	Favre-averaged variance of the progress variable
$\overline{c'^2}_n$	Normalised variance of the progress variable
c_i	NASA polynomial coefficients
\bar{c}_p	Mixture-averaged heat capacity
$c_{p\alpha}$	Heat capacity of species α
D	Molecular diffusivity
Da	Damköhler number
D_α	Molecular diffusivity of species α
$D_{\alpha\beta}$	Binary diffusion coefficient
$d_{i,\alpha\beta}$	Binary diffusion fitting coefficients
E	Total Energy
E_a	Activation energy
\hat{e}_i	Unit vector in the i^{th} direction
\vec{g}	Gravity or body force in general
\bar{h}	Mixture-averaged enthalpy
h_α	Enthalpy of species α
$\Delta h_{f\alpha}^0$	Enthalpy of formation of species α
h_g	Grid location
$h_{s\alpha}$	Sensible enthalpy of species α
J_α	Mass diffusive flux of species α
K	Maximum number of processors in ensemble
Ka	First Karlovitz number
Ka_δ	Second Karlovitz number
k	Turbulent kinetic energy

List of Symbols

k_{Arr}	Arrhenious reaction rate constant
L	Characteristic flow length scale
l_0	Integral length scale
l_L	Longitudinal integral length scale
l_δ	Laminar flame thickness of inner layer
M	Number of ensembles
\mathcal{M}	Number of realisations
\dot{M}	Mass flow rate
\bar{M}	Mixture-averaged molar mass
M_α	Molar mass of species α
m_α	Mass of species α in mixture
N	Number of processors
\mathcal{N}	Number of species
N_η	LEM empirical parameter (2)
n_α	Mole number of species α in mixture
p	Pressure
\dot{Q}	Heat source term
\vec{q}	Total energy flux
\bar{R}	Mixture-averaged gas constant
R_{ij}	Two-point correlation, autocovariance
R_u	Universal gas constant
R_α	Gas constant of species α
Re	Reynolds number
Re_t	Turbulent Reynolds number
Sc_t	Turbulent Schmidt number
s_d	Displacement speed of iso-surface of the progress variable

List of Symbols

s_L	Laminar flame speed
s_T	Turbulent flame speed
s_α	Entropy of species α
T	Temperature
T_b	Burnt gas temperature
T_{eq}	Temperature at equilibrium
T_{ref}	Reference temperature for enthalpy of formation
T_u	Unburnt gas temperature
t	Time coordinate
Δt	Timestep size
U	Characteristic flow velocity
U_B	Bulk flow velocity
\vec{u}	Velocity
\bar{u}	Mean velocity
u'	Turbulent intensity
u''	Velocity fluctuation around the Favre average
u_K	Kolmogorov velocity
V_α	Diffusion velocity of species α
X_α	Mole fraction of species α
\mathbf{x}	Conditional mass fraction array
\vec{x}	Spatial coordinate
Y_α	Mass fraction of species α
$\overline{Y_P^2}$	Variance of the product mass fraction
z	Axial coordinate

Greek Symbols

α_{BML}	Bray Moss Libby PDF parameter (1)
----------------	-----------------------------------

List of Symbols

α_c	Thermal diffusivity
β_{BML}	Bray Moss Libby PDF parameter (2)
Δ	Grid spacing
δ_f	Laminar flame thickness
δ_{ij}	Kronecker delta
ε	Dissipation of turbulent kinetic energy
η_K	Kolmogorov length
κ	Wave number
Λ	Thermal conductivity
Λ_α	Thermal conductivity of species α
λ	Oxidizer-to-fuel equivalence ratio
λ_d	Damping coefficient in Tikhonov regularisation
λ_{LEM}	LEM eddy event frequency per unit length of the domain
μ	Dynamic viscosity
μ_α	Dynamic viscosity of species α
ν	Kinematic viscosity
ν_T	Coefficient of turbulent kinematic viscosity
Ξ	Flame wrinkling factor
Ξ_Δ	SGS flame wrinkling factor
ρ	Density
ρ_u	Unburnt mixture density
Σ	SGS flame surface density
σ_{ij}	Stress tensor
σ_{ij}^{sgs}	Subgrid-scale turbulent stress tensor
τ_{BML}	BML turbulent mixing time scale
τ_c	Chemical time scale

List of Symbols

τ_{ij}	Viscous stress tensor
τ_K	Kolmogorov time scale
τ_t	Turbulent integral time scale
ϕ	Fuel-to-oxidizer equivalence ratio
χ_c	Scalar dissipation rate
$\dot{\omega}_\alpha$	Chemical source term of species α

Superscripts

B	Temperature exponent in Arrhenius equation
T	Transpose of matrix

Subscripts

α, β	Species index
-----------------	---------------

List of Abbreviations

BML Bray Moss Libby

CFD Computational Fluid Dynamics

CMC Conditional Moment Closure

CSC Coherent Structure Capture

CSE Conditional Source-term Estimation

DNS Direct Numerical Simulation

EBU Eddy Break-Up

EDM Eddy Dissipation Model

FDf Filtered Density Function

FGM Flamelet-Generated Manifold

FPI Flame Prolongation of ILDM

FSD Flame Surface Density

LEM Linear-Eddy Model

LES Large Eddy Simulation

PCM Presumed Conditional Moment

List of Abbreviations

PDF Probability Density Function

RANS Reynolds-Averaged Navier-Stokes

SDR Scalar Dissipation Rate

SDS Semi-Deterministic Simulation

SFS Subfilter-Scale

SGS Subgrid-Scale

UHC Unburnt Hydrocarbons

URANS Unsteady Reynolds-Averaged Navier-Stokes

VLES Very Large Eddy Simulation

Acknowledgments

I would like to take this opportunity to express my appreciation to those who have helped and supported my journey in the pursuit of knowledge in the field of turbulent combustion over the past five years. First and foremost, I would like to extend my deepest gratitude to my supervisor, Professor W. Kendal Bushe, for his patience, trust and guidance over the past few years. Without his knowledge and advice on the myriad of problems, this research would not have been possible.

I would also like to thank Professors Carl Ollivier-Gooch, Dana Grecov and James Feng for dedicating their time to attend my proposal defense and to provide valuable feedback throughout different stages of this work. I am grateful to my friends and colleagues in the combustion research group for their positive company and valuable discussions; in particular, Dr. M. Mahdi Salehi invariably manages to spare time and share his insights and passion in the field of combustion research. I also thank Graham Hendra for the engaging discussions and thorough feedbacks, Girish Nivarti and Stefanie Asenbauer for enriching my graduate experience over the years.

I am sincerely thankful to my parents for the relentless encouragement and support, not only in the past five years of this work, but from the beginning. I would like to take a moment to extend my gratitude to all of my friends who have made this journey even more enjoyable and memorable, especially Roland Hui, and of course, my *juu*, Yvonne Law.

Last but not least, I would also like to acknowledge the financial support of the Natural Sciences and Engineering Research Council of Canada (NSERC) and Rolls-Royce Canada. The majority of the computations in this work were performed on the Mammouth Parallèle II supercomputer of Université of Sherbrooke.

Chapter 1

Introduction

1.1 Introduction

An abundance of energy is the prerequisite for the development of a flourishing civilisation. Clean, renewable and sustainable energies, such as solar, wind, hydro and geothermal sources currently constitute approximately 4% of the total global supply [2]. Despite the benefits of the clean energies, the upfront cost of the related infrastructures, relatively low power output and limited capacity for long-term storage hinder their general scalability. In consequence, fossil fuels will continue to be the dominant source of energy for all modes of transportation, industrial power generation and domestic heating in the foreseeable future. As the global population and standard of life continue to grow, it is evident that the accompanying demand for energy will increase as well.

The most convenient and effective approach to extract energy from fossil fuels is through combustion, a process in which the chemical energy stored within the molecular bonds of the fuel is released primarily as heat. The problem, however, is that combustion is not a perfect process. Pollutants such as carbon monoxide (CO), oxides of nitrogen (NO_x), oxides of sulfur (SO_x) and soot are inevitably formed. These substances are not only harmful to humans, but can also be detrimental to the environment [109]. Even in the most ideal scenario, the release of carbon dioxide (CO₂) to the atmosphere is unavoidable as this greenhouse gas is the product of the oxidation of hydrocarbons with oxygen in the air [162].

With these thoughts in mind, the world is confronted with two challenges related to the usage of fossil fuels in the immediate future: (i) the conservation of fossil fuel supplies, which can be extended by developing more efficient combustion devices; and (ii) the preservation of the atmosphere, which can be accomplished by the reduction of pollutants generated from anthropogenic combustion processes. Both of these challenges spark an increasing interest in the research of sustainable combustion technologies.

Turbulent, premixed combustion is an important part of current heat engine designs; examples of these devices include the common automobile engine and stationary gas turbines. The complex interactions between the turbulent flow, chemical kinetics and heat transfer at engine relevant pressures and temperatures establish a highly non-linear problem spanning orders of magnitude in both length and time scales. The difficulty here is to introduce improvements to these highly calibrated heat engines under such dynamic conditions.

Computational fluid dynamics (CFD) is now indispensable in the development of complex engines due to its low cost and time requirement compared to experiments. While direct numerical simulation (DNS) is the most accurate approach to understanding the interactions between turbulence and flame, performing DNS for engine relevant problems will remain impractical for the foreseeable future. With increasing computational capabilities, numerical simulation of these devices via the large eddy simulation (LES) paradigm is becoming more accessible to researchers. Nevertheless, due to the strong coupling between turbulence and chemistry in premixed flames, the prediction of chemical reaction source terms continue to be a modelling challenge.

The work in this thesis focuses on advancing two topics related to the simulation of turbulent, premixed combustion. First, new pseudo-turbulent probability density function (PDF) models of the reaction progress variable and scalar dissipation rate (SDR) models have been formulated. These PDF and SDR models are relevant to several approaches for premixed combustion; their applications are further explained in Chapters 3 and 4. Second, a modelling approach for the closure of chemical source terms in turbulent premixed combustion simulations, the Conditional Source-term Estimation (CSE), has been parallelised and improved for greater efficiency and autonomy in complex geometries.

1.2 Outline

This thesis is arranged such that the content of each chapter is self-contained; as such, the relevant background information and the key acronyms are repeated in the introduction of each chapter. In Chapter 2, a comprehensive introduction to the mathematical tools and theoretical models necessary for the discussion of the numerical simulation of turbulent, premixed combustion will be provided. Particularly, this chapter focuses on the relevant aspects of chemistry, turbulence and the interaction between these two highly coupled processes in premixed combustion. Furthermore, a number of modelling approaches for combustion simulation will be reviewed.

In Chapter 3, a new method of constructing probability density function models for premixed combustion is introduced. This method requires the use of a pseudo-turbulent, one-dimensional stochastic mixing model, referred to as the Linear-Eddy Model (LEM). The methodologies are fully documented in this chapter and the numerical results are compared to experimental data from a laboratory slot burner.

In Chapter 4, the LEM has been implemented to not only formulate PDF models for premixed combustion, but also to formulate the scalar dissipation rate models; the SDR is an important quantity to both premixed and non-premixed combustion simulation strategies. The LEM-formulated PDF and SDR models are compared to a number of turbulent, premixed, methane-air experimental flames, including swirling conditions. It is found that the PDF does not vary significantly with changing turbulent intensities or swirl numbers. However, the SDR tends to decrease in magnitude with increasing swirl and turbulent intensity. The LEM formulated models are well-matched with the experimental results in general.

In Chapter 5, the Conditional Source-term Estimation combustion model, a closure for chemical source terms in turbulent combustion simulations, has been fully parallelised for optimised computations on parallel clusters. Large eddy simulation of a turbulent, premixed methane-air burner has been performed with updated CSE modules. Comparison with the conventional CSE approach indicate that computing times can be significantly reduced, in most cases, by approximately 50%, while producing nearly identical results.

In Chapter 6, the LEM formulated PDF and SDR models are implemented

1.2 Outline

in conjunction with the parallelised CSE module in a large eddy simulation of a turbulent, premixed methane-air, laboratory-scale burner. The LES result appears to be well-matched with the experimental data, measured via Mie scattering, in terms of the flame height and overall temperature field distribution. This modelling approach is one step closer to being truly predictive as all of the *tunable* parameters are nearly eliminated.

In Chapter 7, a summary of conclusions from each of the preceding chapters will be highlighted, including brief discussions of the key accomplishments. In addition, future research topics related to the current work will be considered here.

Chapter 2

Literature Survey

2.1 Background

Turbulent combustion is a multi-scale phenomenon that requires a broad understanding of fluid mechanics, chemical kinetics, heat transfer, molecular diffusion and other physical principles [181]. To perform numerical simulations of this complex problem, further knowledge regarding combustion models and Computational Fluid Dynamics (CFD) is necessary. In this chapter, the fundamental mathematical description of turbulence and flame will first be provided, including the governing equations for Newtonian fluid motion, thermodynamic properties, chemical kinetics and a discussion of the scales of turbulence. This will be followed by a review of the different paradigms of fluid simulation and an overview of combustion models for turbulent, premixed combustion.

2.2 Conservation Equations

The derivation of the conservation equations for mass, momentum, energy and species mass fractions are well-established and can be found in many fluid mechanics textbooks [122, 132, 186]. These differential equations assume the continuum approximation, whereby fluid particles can be viewed as continuously distributed within the corresponding volume without spatial voids [136]. For methane-air combustion under atmospheric conditions, it is possible to further simplify the

2.2 Conservation Equations

equations and assume negligible Soret¹, Dufour² and radiation effects [192]. In tensor notation, the time-dependent conservation equations for a reactive multi-component gaseous mixture involving \mathcal{N} species are commonly written as [132]:

- Conservation of mass (commonly referred to as the continuity equation):

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

where ρ , t , \vec{u} and \vec{x} are respectively the mixture density, time coordinate, velocity and spatial coordinate. The Einstein notation is applied to the subscripts, i, j and k , in all of the conservation equations described in this section³.

- Conservation of species:

$$\frac{\partial(\rho Y_\alpha)}{\partial t} + \frac{\partial(\rho u_i Y_\alpha)}{\partial x_i} + \frac{\partial J_{i,\alpha}}{\partial x_i} = \dot{\omega}_\alpha, \quad \alpha = 1 \dots \mathcal{N}, \quad (2.2)$$

where Y_α and $\dot{\omega}_\alpha$ are the mass fraction and chemical source term of species α . The mass diffusive flux of species α is,

$$J_{i,\alpha} = -\rho D_\alpha \frac{\partial Y_\alpha}{\partial x_i}, \quad (2.3)$$

where D_α is the molecular diffusivity of species α . Observing the conservation of mass and species equations, there are a total of $\mathcal{N} + 1$ equations for the \mathcal{N} species, resulting in an over-determined system. This is usually resolved by removing one species equation, typically an inert species with a large mass fraction in the mixture. For methane-air oxidation, a fitting candidate is N_2 .

¹ Soret effect refers to the differential diffusion of particles with different molecular masses in the presence of a temperature gradient within the system.

² Dufour effect refers to an enthalpy flux induced by a concentration gradient.

³ k represents a directional component of the Cartesian coordinate system when it is applied to a subscript. In other sections of this thesis, k represents the turbulent kinetic energy.

2.2 Conservation Equations

- Conservation of momentum:

$$\frac{\partial(\rho u_j)}{\partial t} + \frac{\partial(\rho u_i u_j)}{\partial x_i} + \frac{\partial \sigma_{ij}}{\partial x_i} = \rho g_i, \quad (2.4)$$

where σ_{ij} and \vec{g} represent the stress tensor and the force of gravity, or body force in general. The stress tensor is defined as a combination of the pressure and viscous stress tensors, $\sigma_{ij} = p\delta_{ij} - \tau_{ij}$. The viscous stress tensor reads,

$$\tau_{ij} = \mu \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) - \frac{2}{3} \mu \delta_{ij} \frac{\partial u_k}{\partial x_k}, \quad (2.5)$$

where μ , p and δ_{ij} are the dynamic viscosity, pressure and Kronecker delta, respectively.

- Conservation of energy:

$$\frac{\partial(\rho E)}{\partial t} + \frac{\partial(\rho u_i E)}{\partial x_i} + \frac{\partial q_i}{\partial x_i} + \frac{\partial(\sigma_{ij} u_i)}{\partial x_j} = \rho g_i u_i + \dot{Q}, \quad (2.6)$$

where E , \vec{q} and \dot{Q} are respectively the total energy, the total energy flux and the heat source term. The heat source term is an external source, independent of the chemical energies. The energy flux is further defined as,

$$q_i = -\Lambda \frac{\partial T}{\partial x_i} - \rho \sum_{\alpha=1}^N h_\alpha D_\alpha \frac{\partial Y_\alpha}{\partial x_j}, \quad (2.7)$$

where T , Λ and h_α are respectively the temperature, thermal conductivity of the mixture and enthalpy of species α .

To close the system of equations and obtain the pressure, the usual ideal gas law is required,

$$p = \rho \bar{R} T = \rho R_u T / \bar{M}, \quad (2.8)$$

where \bar{R} , R_u and \bar{M} are the mixture-averaged gas constant, the universal gas constant and the mixture-averaged molar mass.

For completeness, the definition of the mass fraction of species α , Y_α , and the

mole fraction for species α , X_α , in a mixture of \mathcal{N} species are respectively,

$$Y_\alpha = \frac{m_\alpha}{\sum_{\beta=1}^{\mathcal{N}} m_\beta}, \quad (2.9)$$

$$X_\alpha = \frac{n_\alpha}{\sum_{\beta=1}^{\mathcal{N}} n_\beta}, \quad (2.10)$$

where m_α is the total mass content of species α and n_α is the total number of moles of species α within the mixture. A logical extension from the definitions of Y_α and X_α is that $\sum_{\alpha=1}^{\mathcal{N}} Y_\alpha = 1$ and $\sum_{\alpha=1}^{\mathcal{N}} X_\alpha = 1$. In general, the mass fraction will be written in the conservation equations while the mole fraction will be used in the numerical computations by chemical software packages, for example, CHEMKIN-II [74]. Through basic manipulations, the mass and mole fractions can be converted from one to the other from the following equations,

$$Y_\alpha = X_\alpha \frac{M_\alpha}{\bar{M}} = X_\alpha \frac{M_\alpha}{\sum_{\beta=1}^{\mathcal{N}} X_\beta M_\beta}, \quad (2.11)$$

$$X_\alpha = \bar{M} \frac{Y_\alpha}{M_\alpha} = \frac{Y_\alpha / M_\alpha}{\sum_{\beta=1}^{\mathcal{N}} Y_\beta / M_\beta}, \quad (2.12)$$

where M_α is the molar mass of species α .

2.3 Transport Properties

To correctly account for the transport of momentum, species and energy in a gaseous mixture, one must properly tabulate and evaluate the viscosities, thermal conductivities, diffusion coefficients and thermal diffusion coefficients; these four quantities constitute the properties of transport in a mixture. In general, there are two common practical approaches to calculating these transport characteristics [73], the mixture-averaged and the multi-component methods. For many applications, it is possible to use the basic mixture-averaged properties while maintaining a high degree of accuracy. In other cases, multi-component properties may be required.

While some of these properties can be computed by first principles, such as the standard kinetic theory expressions, vibrational modes, Lennard-Jones collision diameters and so forth [73], the quantities are typically tabulated as functions of

2.3 Transport Properties

the temperature and pressure using high-order polynomials. The viscosity (μ_α) and thermal conductivity (Λ_α) of each species are pressure-independent and are generally fitted with third order polynomials as functions of the temperature,

$$\ln \mu_\alpha = \sum_{i=1}^4 a_{i,\alpha} (\ln T)^{i-1}, \quad (2.13)$$

$$\ln \Lambda_\alpha = \sum_{i=1}^4 b_{i,\alpha} (\ln T)^{i-1}, \quad (2.14)$$

where $a_{i,\alpha}$ and $b_{i,\alpha}$ are the corresponding fitting coefficients. The binary diffusion coefficients ($D_{\alpha\beta}$) are similarly fitted,

$$\ln D_{\alpha\beta} = \sum_{i=1}^4 d_{i,\alpha\beta} (\ln T)^{i-1}, \quad (2.15)$$

where $d_{i,\alpha\beta}$ are the corresponding fitting coefficients. However, this property is pressure-dependent. As such, these coefficients are computed at unit pressure and the correct values are determined by dividing the evaluated coefficient from the fit by the actual pressure in the control volume.

To evaluate the mixture-averaged viscosity, one can invoke the semi-empirical formula of Wilke [188] and later modified by Bird *et al.* [16],

$$\mu = \sum_{\alpha=1}^{\mathcal{N}} \frac{X_\alpha \mu_\alpha}{\sum_{\beta=1}^{\mathcal{N}} X_\beta \Psi_{\alpha\beta}}, \quad (2.16)$$

where,

$$\Psi_{\alpha\beta} = \frac{1}{\sqrt{8}} \left(1 + \frac{M_\alpha}{M_\beta}\right)^{-\frac{1}{2}} \left(1 + \left(\frac{\mu_\alpha}{\mu_\beta}\right)^{\frac{1}{2}} \left(\frac{M_\beta}{M_\alpha}\right)^{\frac{1}{4}}\right)^2. \quad (2.17)$$

The mixture-averaged thermal conductivity can be computed with the combination averaging formula [108],

$$\Lambda = \frac{1}{2} \left(\sum_{\alpha=1}^{\mathcal{N}} X_\alpha \Lambda_\alpha + \frac{1}{\sum_{\alpha=1}^{\mathcal{N}} X_\alpha / \Lambda_\alpha} \right). \quad (2.18)$$

The mixture diffusion coefficient for species α , D_α , is,

$$D_\alpha = \frac{1 - Y_\alpha}{\sum_{\beta \neq \alpha}^{\mathcal{N}} X_\beta / D_{\beta\alpha}}. \quad (2.19)$$

Equations 2.16, 2.18 and 2.19 provide the three necessary quantities for the mixture-averaged formulation. The thermal diffusion coefficients are only included for the species with low molecular mass, H to He.

The detailed mathematical derivation of the multi-component formulation is beyond the scope of this section, but its defining characteristics will be highlighted. The multi-component formulation offers two advantages over the mixture-averaged transport properties. First, it does not have inherent limitations to the composition of the mixture qualities, such as the Fickian assumption, and therefore will be more accurate. Second, multi-component transport ensures the conservation of mass without the necessity of introducing correction factors. In exchange for the improved accuracy, this model requires a substantial increase in computational cost. An $\mathcal{N} \times \mathcal{N}$ matrix must be inverted to obtain the multi-component diffusion coefficients. Further, to compute the thermal conductivity and thermal diffusion coefficients, a $3\mathcal{N} \times 3\mathcal{N}$ system of equations must be solved [73]. Unless otherwise stated, mixture-averaged properties are used in the computation of the transport mechanics in this research.

2.4 Thermodynamics

The thermodynamic properties of chemical species are another important component to combustion modelling. In particular, these include the enthalpy (h_α), constant pressure heat capacity (c_{p_α}) and entropy (s_α) of species α . Typically the thermodynamic properties are fitted through high order polynomials as functions of temperature over two separate ranges for improved accuracy. These coefficients are then compiled into lookup tables; examples of these tables include the Chemkin Thermodynamic Data Base [75] and the database compiled by Gordon and McBride [59, 110]. For each species, there are seven tabulated coefficients at

2.4 Thermodynamics

each of the two temperature ranges:

$$\frac{c_{p\alpha}}{R_\alpha} = c_1 + c_2T + c_3T^2 + c_4T^3 + c_5T^4, \quad (2.20)$$

$$\frac{h_\alpha}{R_\alpha T} = c_1 + \frac{c_2}{2}T + \frac{c_3}{3}T^2 + \frac{c_4}{4}T^3 + \frac{c_5}{5}T^4 + \frac{c_6}{T}, \quad (2.21)$$

$$\frac{s_\alpha}{R_\alpha} = c_1 \ln(T) + c_2T + \frac{c_3}{2}T^2 + \frac{c_4}{3}T^3 + \frac{c_5}{4}T^4 + c_7. \quad (2.22)$$

This particular format is often referred to as the NASA polynomials [25].

The mixture heat capacity and enthalpy can subsequently be calculated through linear combinations of the species specific properties, either by mass or mole fractions,

$$\bar{c}_p = \sum_{\alpha=1}^N c_{p\alpha} Y_\alpha, \quad (2.23)$$

$$\bar{h} = \sum_{\alpha=1}^N h_\alpha Y_\alpha. \quad (2.24)$$

The Chemkin Thermodynamic Data Base [75] has been used for all of the thermodynamic computations performed by this work unless specified otherwise.

Aside from tabulating the thermodynamic properties for the purposes of numerical simulation, it is often useful to analytically relate the properties, particularly,

$$h_\alpha = \int_{T_{ref}}^{T^*} c_{p\alpha}(T) dT + \Delta h_{f\alpha}^0 = h_{s\alpha} + \Delta h_{f\alpha}^0, \quad (2.25)$$

where $\Delta h_{f\alpha}^0$ is the species standard enthalpy of formation taken at a reference temperature, T_{ref} , typically 298 K. The integration of the species constant pressure heat capacity from T_{ref} to the current temperature, T^* , yields the sensible enthalpy of the species, $h_{s\alpha}$. Further, the total energy of the mixture (Equation 2.6) can be computed from the mixture enthalpy as follows,


$$E = \bar{h} - \frac{p}{\rho} + \frac{1}{2} u_i u_i. \quad (2.26)$$

2.5 Chemistry

One can obtain valuable information related to the combustion process from thermodynamics alone; however, this limits the calculations to global values such as enthalpy changes and adiabatic flame temperatures. For modern engine design, where the prediction of thermal efficiency and pollutant formation are equally critical, thermodynamics alone is not sufficient to characterise the system. This is where a detailed understanding of the underlying chemical phenomenon allows one to compute a crucial aspect of combustion — the time scales or rates at which the reactions occur. In the most accurate representation, the chemical reactions are fully coupled to the fluid flow, leading to the finite rate chemistry representation.

2.5.1 Chemical Kinetics

For methane-air combustion, finite rate chemistry was first introduced as a one-step, global reaction,




This representation, however, is not complete as a large number of minor species and radicals are formed and consumed during the combustion process via elementary reactions. The GRI-Mech 3.0 [156] is a well-known example of a detailed finite rate mechanism that represents methane-air oxidation. For this basic hydrocarbon, the GRI-Mech 3.0 chemistry mechanism prescribes 325 reactions and 53 species to model the underlying phenomenon.

The rate at which each of these elementary reactions occurs is dependent on the activation energy required to change the reactants to products, the local concentration of the reactants and the local temperature. In general, it has been established that the Arrhenius law can adequately capture the reaction rate constant, k_{Arr} , for most elementary chemical reactions,

$$k_{Arr} = A \cdot T^B e^{-E_a/(R_u T)}, \quad (2.28)$$

where A , B and E_a are respectively the pre-exponential factor, temperature exponent and activation energy; these three parameters are determined experimentally.

To obtain the overall reaction rate, one can multiply the reaction rate constant, k_{Arr} , with the concentrations of the reactants raised to the power of their stoichiometric coefficient. Using a simple two-reactant reaction as an example,



the overall consumption rate of Reactant(A) would be,

$$\frac{d[\text{Reactant(A)}]}{dt} = -k_{Arr} \cdot [\text{Reactant(A)}]^1 [\text{Reactant(B)}]^1, \quad (2.30)$$

where $[\]$ represents the concentration of the species. For more complex mechanisms, the overall rate of consumption or production of a species would be a linear combination of all elementary reactions in which the species participates; this follows the basic principles of the law of mass action. Indeed, there are a number of specialised reactions that deviate from the standard Arrhenius law for certain chemical mechanisms, including three-body reactions, fall-off reactions, chemically-activated reactions and pressure-dependent reactions [74].

2.5.2 Reduced Chemistry

While a detailed mechanism can accurately compute the reaction rates and describe the underlying physics of the combustion process, it is much too expensive for numerical simulations using the current generation of computational hardware for industrially relevant problems. As a result, a number of different modelling approaches have been developed to overcome this obstacle.

Detailed mechanisms can be reduced by the use of quasi steady-state and partial-equilibrium assumptions [32, 125]. The resultant skeletal mechanisms would include fewer species and elementary reactions compared to the full chemistry, which would then decrease the demand on the computational hardware. In exchange, however, the skeletal mechanisms are typically accurate only for a restricted range of equivalence ratios and flame conditions. Moreover, it is still possible for the time scales to differ greatly between reactions, leading to a stiff system [102]. Nevertheless, a large number of reduced mechanisms are available for premixed, methane-air combustion due to their low computational require-

ment [10, 11, 29, 114, 118, 146].

A more sophisticated class of techniques for chemistry reduction are low-dimensional manifold approaches [104]. In these methods, the underlying concept is to eliminate the slow time scales such that the system can be generalised by a lower-dimensional manifold in the composition space. Low-dimensional manifolds are well-suited for simulation purposes as the chemistry is fully tabulated prior to the flow computation. This tabulation only needs to be done once for each combination of thermodynamic and chemical composition. Both the chemical source terms and species mass fractions are stored as a function of a few controlling variables for which transport equations are usually solved. Examples of these variables for methane-air combustion include H_2O and CO_2 mass fractions. The chemical properties are then retrieved from the table in real time, directing computing resources to the flow calculation instead of the reaction kinetics.

The Flamelet-Generated Manifold (FGM) [121] is one such approach in which a prototype flame is used to build a lower-dimensional manifold. In premixed combustion, a one-dimensional laminar flame can be used to generate these trajectories. A combination of skeletal mechanism and FGM approaches has been used to complete the work in this thesis; the skeletal mechanisms will be primarily used to construct the Probability Density Function (PDF) models for premixed combustion in a pre-processing manner, whereas the FGM will be coupled directly to the three-dimensional reactive flow calculations. The reduction techniques will be discussed in greater detail in Chapters 4 and 5, respectively.

2.6 Laminar Premixed Flame

Before the discussion of turbulent premixed flames, it is important to briefly review the general structure of a laminar premixed flame. A premixed flame can be viewed as a moving front in which the exothermic reaction is sustained by the continuous consumption of the nearby reactants. In the ideal planar laminar premixed flame, the flame front propagates steadily from the burnt mixtures towards the unburnt reactants; the speed at which this front propagates divides the premixed flame into two categories, the subsonic deflagration and the supersonic detonation waves [91, 96]. In this thesis, the discussion will be limited to the subsonic deflagration waves.

2.6 Laminar Premixed Flame

From a classical standpoint, as first proposed by Mallard and Le Chatelier and later developed by Peters [128], the structure of a laminar premixed flame contains three distinct layers: the preheat layer, the inner layer and the oxidation layer. In the preheat layer, the reaction rates are relatively low and the temperature increase is caused primarily by conduction from the inner layer. This layer tends to be thicker than the other two due to the lack of chemical reactions in the governing mechanism. Next, the inner layer is where the fuel is consumed, radicals are depleted and the bulk of the heat of combustion is released. This layer tends to be the thinnest as the chemical time scales are relatively short. Last, the final products, accompanied by minor heat release, are formed through the slow reactions in the oxidation layer. Figure 2.1 illustrates the structure of a one-dimensional laminar premixed flame. The three layers combine to form the flame front, which separates the unburnt fresh gas mixture from the burnt product mixture. The temperature of the mixtures increases across the flame front while the mixture density decreases such that the pressure changes are minimal across the flame.




Figure 2.1: One-dimensional laminar flame structure for premixed, stoichiometric, methane-air combustion at standard conditions.

Two important parameters are typically used to characterise a laminar premixed flame: the flame front propagation speed, commonly referred to as the laminar flame speed or the laminar burning velocity, s_L , and the laminar flame thickness, δ_f . The laminar flame speed describes the flame front propagation normal to itself. This property is dependent primarily on the equivalence ratio and temperature of the premixed reactants and secondarily on the overall pressure of the system. The laminar flame thickness is usually defined by either the diffusive or thermal thicknesses. In this thesis, the thermal thickness is used exclusively unless otherwise specified. The thermal thickness takes on the following form,

$$\delta_f = \frac{T_b - T_u}{\left. \frac{dT}{dx} \right|_{max}}, \quad (2.31)$$

where T_b and T_u are respectively the burnt and unburnt gas temperatures; $\left. \frac{dT}{dx} \right|_{max}$ refers to the maximum gradient of the temperature across the flame front. Another parameter that characterises a premixed flame, which is commonly found in the literature, is the thickness of the inner layer, l_δ . This parameter is typically one order of magnitude smaller than the thermal thickness [128].

For a laminar premixed flame at standard ambient temperature (298.15 K) and pressure (100 kPa), the ratio of fuel-to-oxidizer notably alters the laminar flame speed, the overall flame thickness, and the equilibrium flame temperature. This ratio is typically expressed as a non-dimensional quantity, referred to as the equivalence ratio (ϕ), which compares the current fuel-to-oxidizer ratio to the stoichiometric fuel-to-oxidizer ratio; stoichiometric combustion is achieved if the fuel and oxidizer are both perfectly consumed and converted to stable products in the reaction. Mathematically, the equivalence ratio is defined as,

$$\phi = \frac{m_{fuel}/m_{oxidizer}}{(m_{fuel}/m_{oxidizer})_{stoich}} = \frac{n_{fuel}/n_{oxidizer}}{(n_{fuel}/n_{oxidizer})_{stoich}}, \quad (2.32)$$

where m_x and n_x respectively represent the mass and number of moles of the quantity designated by x . Three possible scenarios can occur: (i) fuel-rich combustion, $\phi > 1$; (ii) stoichiometric combustion, $\phi = 1$; and (iii) fuel-lean combustion, $\phi < 1$. Another commonly found dimensionless quantity is the oxidizer-to-fuel equivalence ratio, λ . This is simply the inverse of ϕ , where $\lambda = \phi^{-1}$. Figure 2.2 illustrates

2.6 Laminar Premixed Flame

the laminar flame speeds and equilibrium temperatures of premixed methane-air flames at different equivalence ratios in standard ambient temperature and pressure conditions. The values are calculated by Cantera [58] using the detailed GRI 3.0 mechanism [156] and mixture-averaged transport.




Figure 2.2: Laminar flame speeds and equilibrium temperatures of premixed methane-air flames at different equivalence ratios in standard ambient temperature (298.15 K) and pressure (100 kPa) conditions.

To further expand on the equivalence ratios, lean premixed combustion is generally the most efficient as the excess oxidizer serves to exhaust the fuel supply while driving the reaction to completion. If the process is controlled appropriately, it can also reduce pollutant formation, such as NO_x , due to the reduced combustion temperatures. However, the shortcoming is a reduced flame speed, which can lead to a lower power output for the same engine dimensions and intake conditions. Stoichiometric combustion is typically found in spark-ignited automobile engines, as it yields a suitable balance between power, efficiency and pollutant formation. Rich combustion is not frequently desired as it results in Unburnt Hydrocarbons (UHC) due to the excess fuel. This leads to a reduction in thermal efficiency and the UHC would become an additional source of pollutants. It is important to note that the flow conditions would be turbulent for engine-relevant combustions rather than laminar, though the rich, stoichiometric and lean combustion generalisations still stand.

2.7 Scales of Turbulence

In practical combustion applications, the fluid motion predominately exists in a turbulent rather than a laminar regime. A flow is considered turbulent if it can be characterised as unsteady, three-dimensional, chaotic, vortical and dissipative [136, 186]. This highly non-linear, multi-scale phenomena is, in principle, fully described by the equations shown in Section 2.2. The conservation of momentum equations are referred to as the Navier-Stokes equations, which were first derived in the early 19th century by G. Stokes and M. Navier, independently. They elegantly describe the coupling between the density, velocity, temperature and pressure within the fluid. Coincidentally, due to the strong non-linearity of this system of equations, analytical solutions currently do not exist; such analytical solutions can only be derived for laminar conditions or highly simplified systems. Indeed, *simply* providing the mathematical proof for the existence and uniqueness of a solution to the Navier-Stokes equations is one of the most challenging theoretical problems of the 21st century [48].

In the absence of a robust analytical framework, turbulence research has been dependent on experiments and numerical simulations. Through observations in the late 19th century by Reynolds [138], a fundamental non-dimensional parameter that characterises the intensity of turbulent flows emerged, the Reynolds number, Re . It is the ratio of inertial forces to viscous forces and is defined as,

$$Re \equiv \frac{U \cdot L}{\nu}, \quad (2.33)$$

where U , L and ν are respectively the characteristic velocity, the characteristic length and kinematic viscosity of the fluid. The bulk flow velocity would typically be used as U , while a suitable length scale determined by the geometry of the flow field would be used as L . In the internal pipe flow experiments performed by Reynolds [138], the transition between laminar and turbulent states occurred at a Reynolds number between 2,300 to 4,000. For external boundary layer flows over flat plates, however, the transition between laminar and turbulent regimes typically occur at a Reynolds number of 300,000 to 500,000. Nevertheless, despite having rather extensive analytical understanding of laminar flows and empirical relations for turbulent flows, the transition between the two regimes is still not well charac-

terised [5].

Two parameters that are typically used in the literature to specify the characteristics of a turbulent flow field are the turbulent kinetic energy, k , and the integral length scale, l_0 . By the Reynolds decomposition, which will be discussed in greater detail in Section 2.9, the instantaneous velocity, u_i , can be separated into a mean velocity, \bar{u}_i , and a fluctuating component, u'_i . The turbulent kinetic energy can then be defined as,

$$k = \frac{1}{2} \left(\overline{(u'_x)^2} + \overline{(u'_y)^2} + \overline{(u'_z)^2} \right), \quad (2.34)$$

where u'_x, u'_y, u'_z are respectively the fluctuating components of the velocity in each direction of the Cartesian coordinate system. For isotropic⁴, homogeneous⁵ turbulent fields, it is possible to further simplify the turbulent kinetic energy to,

$$k = \frac{3}{2} u'^2, \quad (2.35)$$

where u' is the root-mean-square of the velocity fluctuations, also commonly known as the turbulent intensity. The integral length scale indicates the size of the turbulent motions, or turbulent *eddies*, in which the majority of the turbulent kinetic energy manifests [136]. It is usually considered the large-scale motion of the turbulent field and is typically on the same order of magnitude as the characteristic length, L . To formally characterise the spatial structure of a random field, one can introduce a two-point correlation, or autocovariance, defined as [136],

$$R_{ij}(\vec{r}, \vec{x}, t) \equiv \langle u_i(\vec{x}, t) u_j(\vec{x} + \vec{r}, t) \rangle, \quad (2.36)$$

where \vec{r} is some distance away from \vec{x} . The integral length scale related to the x_1 -coordinate direction can then be written as,

$$L_{11}(\vec{x}, t) \equiv \frac{1}{R_{11}(0, \vec{x}, t)} \int_0^\infty R_{11}(\hat{e}_1 r, \vec{x}, t) dr, \quad (2.37)$$

where \hat{e}_1 is the unit vector in the x_1 -coordinate direction. Indeed, one can define an integral length scale for each direction of the coordinate system. For the special

⁴Statistically invariant under rotations and reflections of the coordinate system.

⁵Statistically invariant under translations of the coordinate system.

case of isotropic homogeneous turbulence, one can associate the integral length scale, l_0 , with any one of the L_{ii} . From these parameters, one is able to define another non-dimensional number, the turbulent Reynolds number, Re_t ,

$$Re_t = \frac{u' l_0}{\nu}. \quad (2.38)$$

As a first approximation, the turbulent Reynolds number is generally an order of magnitude lower than the Reynolds number for the same flow field.

Aside from the non-dimensional numbers, it has also been observed that the production mechanism driving the turbulent motions, or turbulent *eddies*, transfers its kinetic energy from the largest characteristic scales to the smallest scales until it is finally dissipated by viscous action; this concept of energy transfer is referred to as the *energy cascade* [139]. Kolmogorov later formulated two similarity hypotheses pertaining to the different ranges of the energy cascade. The first of these postulated that the statistics of the small-scale motions can be ubiquitously determined by the kinematic viscosity of the fluid, ν , and the rate of dissipation of the turbulent kinetic energy, ε , provided the turbulent flow is at a sufficiently high Reynolds number [89]. Through dimensional analysis, three parameters can be derived from the postulation describing the smallest scales of turbulence; these are respectively named the Kolmogorov length (η_K), time (τ_K) and velocity (u_K),

$$\eta_K \equiv (\nu^3/\varepsilon)^{1/4}, \quad (2.39)$$

$$\tau_K \equiv (\nu/\varepsilon)^{1/2}, \quad (2.40)$$

$$u_K \equiv (\varepsilon\nu)^{1/4}. \quad (2.41)$$

The second of Kolmogorov's hypotheses postulated that the statistics of the motions of scale l in the range $\eta_K \ll l \ll l_0$ can be uniquely determined by the rate of dissipation of the turbulent kinetic energy, ε , and independent of the kinematic viscosity of the fluid, ν , provided the turbulent flow is at a sufficiently high Reynolds number [89]. This range of length scales is referred to as the inertial subrange. The inertial subrange exhibits a unique quality where the energy carried by a turbulent eddy of size l decreases with increasing wavenumber, κ ($\kappa = 2\pi/l$), with a known

functional form. Particularly, the energy spectrum in the inertial subrange is,

$$E(\kappa) = C_\kappa \varepsilon^{2/3} \kappa^{-5/3}, \quad (2.42)$$

where C_κ is a universal constant [136]. Figure 2.3 illustrates a typical energy spectrum observed in turbulent motion at high Reynolds numbers.

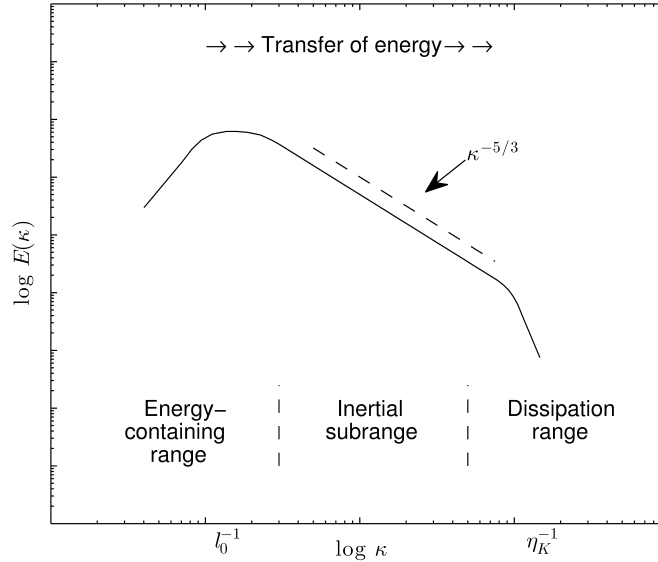


Figure 2.3: A schematic diagram of the turbulent kinetic energy spectrum at high Reynolds numbers.

A very informative relationship can be established between the integral scales and the Kolmogorov scales through an estimation of the dissipation rate, ε , in terms of the large-scale features of the flow. The ε is approximately equal to the kinetic energy production rate because of the transfer of energy through the energy cascade [136]. The kinetic energy is proportional to U^2 while the time scale of the energy transfer is estimated to be on the order of the large-scale motion, l_0/U . The dissipation rate can thus be scaled as,

$$\varepsilon \approx \frac{U^2}{l_0/U} \approx \frac{U^3}{l_0}. \quad (2.43)$$

Substituting this relation into Equation 2.39, one obtains,

$$\eta_K = \left(\frac{\nu^3}{\varepsilon} \right)^{1/4} \approx \left(\frac{\nu^3 l_0}{U^3} \right)^{1/4}. \quad (2.44)$$

Finally, the ratio between the integral length scale and Kolmogorov length scale can be approximated,

$$\frac{l_0}{\eta_K} \approx \frac{l_0 U^{3/4}}{\nu^{3/4} l_0^{1/4}} \approx \left(\frac{U l_0}{\nu} \right)^{3/4} \approx Re^{3/4}. \quad (2.45)$$

Following a similar derivation, one can additionally obtain the relationship between the integral time scale, τ_i , and the Kolmogorov time scale, τ_K ,

$$\frac{\tau_i}{\tau_K} = \frac{l_0/U}{(\nu/\varepsilon)^{1/2}} \approx \frac{l_0/U}{(\nu l_0/U^3)^{1/2}} \approx \left(\frac{U l_0}{\nu} \right)^{1/2} \approx Re^{1/2}. \quad (2.46)$$

Equations 2.45 and 2.46 suggest that the disparity between the largest and smallest turbulent scales of the flow field increases rapidly as the Reynolds number increases.

2.8 Regimes of Premixed Turbulent Combustion

The overall behaviour of the premixed flame would differ under various turbulent conditions as a result of the changing turbulence-chemistry interactions. It has been recognised that these interactions can be categorised into large and small scales [36]. Figure 2.4 provides a conceptual overview to the behaviour of a premixed flame front exposed to large-scale turbulence. At very low turbulent intensities, the eddies typically do not alter the overall shape and thickness of the premixed flame. The flame front retains its planar geometry and propagates towards the reactants at approximately the laminar flame speed, s_L . As the turbulent intensity increases, the premixed flame would become corrugated due to the interaction with the vortical motion of the eddies. The total surface area of the flame front would increase due to the corrugation, resulting in a higher rate of consumption of the reactants. This increased effective flame speed is referred to as the turbulent

flame speed, s_T . As proposed by Damköhler [36],

$$\frac{s_T}{A_T} = \frac{s_L}{A_L}, \quad (2.47)$$

where A_T and A_L are the turbulent flame surface area and the laminar flame surface area. The ratio of A_T/A_L is referred to as the flame wrinkling factor, Ξ . As the turbulent Reynolds number increases, the kinetic energy from the large-scale motions would be transported to smaller and smaller eddies; such eddies are then able to interact with the flame front more effectively. Past some threshold, it is even possible for the smallest eddies to not just wrinkle the flame front, but the Kolmogorov-scale turbulent motion could also affect the transport mechanisms within the flame — this is the regime of small-scale turbulent interactions.




Figure 2.4: Effect of large-scale turbulent eddies on a premixed flame front.

This complex coupling between the chemistry and turbulence leads to a separation of combustion characteristics dependent on the physical parameters that are representative of the system. The length and time scales of the turbulent flow field and the chemical reactions then provide a basis for the identification of different regimes for turbulent premixed combustion; the analysis of which is best summarised by regime diagrams, as originally proposed by Borghi [19]. A number of such diagrams have been proposed to characterise the turbulence-chemistry interactions for turbulent premixed combustion, including the work of Abdel-Gayed and Bradley [3], Borghi [19], Bray [23], Düsing *et al.* [46], Peters [126, 127], Pitsch and Duchamp de Lageneste [131] and Poinso *et al.* [133]. Figure 2.5 illustrates a commonly referenced regime diagram in the literature, the modified Borghi diagram, as proposed by Peters [127]. The ratio of the turbulence intensity to the

laminar flame speed is plotted against the ratio of the integral length scale over the laminar flame thickness on logarithmic scales.




Figure 2.5: Modified Borghi diagram illustrating different turbulent premixed combustion regimes.

Prior to the detailed discussion of the combustion regimes illustrated by the modified Borghi diagram, it is important to acknowledge a few additional dimensionless parameters aside from the turbulent Reynolds number. Scaling arguments between the ratios of turbulence and chemistry lead to the definition of the following entities:

- The Damköhler number, Da , which is the ratio between the turbulent integral time scale to the chemical time scale,

$$Da \equiv \frac{\tau_t}{\tau_c} = \frac{l_0/u'}{\delta_f/s_L}. \quad (2.48)$$

- The First Karlovitz number, Ka , which is the ratio between the chemical

time scale to the Kolmogorov time scale,

$$Ka \equiv \frac{\tau_c}{\tau_K} = \frac{\delta_f/s_L}{\eta_K/u_K} = \frac{\delta_f^2}{\eta_K^2}, \quad (2.49)$$

where by dimensional analysis,

$$\eta_K = \left(\frac{\nu^3}{\varepsilon}\right)^{1/4}, \quad s_L \approx \frac{\nu}{\delta_f}, \quad u_K = (\varepsilon\nu)^{1/4}. \quad (2.50)$$

- The Second Karlovitz number, Ka_δ , which is the ratio between the chemical time scale, based on the inner layer, to the Kolmogorov time scale,

$$Ka_\delta \equiv \frac{l_\delta^2}{\eta_K^2}. \quad (2.51)$$

Assuming the thickness of the inner layer is approximately 10% of the laminar flame thickness [128], then $Ka_\delta \approx 0.01 \cdot Ka$.

- The turbulent Reynolds number, Re_t , which has been previously defined in Equation 2.38, can be rewritten in terms of the Damköhler and Karlovitz numbers,

$$Re_t = Da^2 Ka^2. \quad (2.52)$$

Armed with the preceding dimensionless parameters, one can observe five distinct regimes of turbulent premixed combustion within the modified Borghi diagram:

- The laminar flame regime: $Re_t < 1$. The turbulent intensity is at a negligible level. The flame behaviour can be fully encapsulated by the description of the laminar premixed flame in Section 2.6.
- The wrinkled flamelet regime: $Re_t > 1$, $Ka < 1$ and $u' < s_L$. Since the chemical time scale is shorter than the integral turbulent time scale, the internal flame structure remains unaffected by the turbulent motion. More precisely, the thickness of flame front is smaller than the smallest eddies, such that

the laminar structure remains unchanged. The turbulent motion would only serve to transport the flame on a large scale.

- The corrugated flamelet regime: $Re_t > 1$, $Ka < 1$ and $u' > s_L$. This is similar to the wrinkled flamelet regime, except the turbulent intensity is at a higher level. The increased turbulent motion may lead to the formation of isolated pockets of fresh and burnt gases, hence the name. The smallest eddies are still larger than the flame front thickness, allowing the internal flame structure to remain mostly laminar.
- The thin reaction zone: $Re_t > 1$ and $1 < Ka < 100$. The Kolmogorov length scale is shorter than the overall flame thickness such that the smallest eddies are able to penetrate into the preheat layer, though not into the inner layer [132]. Part of the internal flame structure can be modified by the eddies and therefore should no longer be modelled with locally laminar assumptions.
- The broken reaction zone: $Ka > 100$. The Kolmogorov length scale is shorter than the thickness of the inner layer such that the smallest eddies are able to penetrate into the inner layer and modify the flame structure. The turbulent mixing is extremely fast and the reaction rate is limited by chemistry [132].

While the turbulence-chemistry interaction is often described from the perspective of eddies penetrating into the various layers of the flame, it is important to realise that the flame correspondingly modifies the turbulent field. For example, the temperature across the premixed flame front can increase by a factor of eight or so (Figure 2.2), whereas the kinematic viscosity of air varies with temperature approximately as $\nu \sim T^{1.7}$. This results in a significant decrease in the Reynolds number in the product mixture, which may lead to relaminarisation of the locally turbulent field [132]. Moreover, a rapid increase in gas temperatures in the thin front leads to a rapid increase in gas velocities in that region. This acceleration in turn modifies the turbulent field, which results in an effect referred to as the flame-generated turbulence [132]. The intricate feedback between turbulence and

chemistry is one of the fundamental reasons why turbulent combustion modelling will continue to be a challenge in the foreseeable future.

Practical combustion devices, such as gas turbines and automotive internal combustion engines, often operate within the corrugated flamelet regime and thin reaction zone, resulting in an increased interest for models that are representative of these regimes. A brief review to the turbulent combustion models available today will be discussed in Section 2.10.

2.9 Computational Fluid Dynamics (CFD)

The complexity of the governing conservation equations suggests that the analytical solutions to any practical fluid dynamics problems involving turbulence may not be derived for some time [48]. This limits researchers to two possible approaches in the quest of understanding the nature of turbulence: experimental observations and computation of the governing equations using numerical solvers, which is the study of computational fluid dynamics. Though the results from a well-designed experiment are highly accurate and reliable, these are often costly and the desired metrics may not be readily measured in combustion devices. Moreover, simply introducing a measuring tool to the system may alter the characteristics of the device. Computational fluid dynamics, thus, emerges as an important tool for understanding the effects of turbulence and flame from an theoretical perspective and for predicting the power and pollutant outputs in practical combustion devices. Three general paradigms of CFD currently exist: Direct Numerical Simulation (DNS), Large Eddy Simulation (LES) and Reynolds-Averaged Navier-Stokes (RANS). Each of these strategies has certain benefits over one another [34, 141] and their characteristics will be briefly summarised in this section.

DNS, as the name implies, directly resolves all of the spatial and temporal scales of the turbulent field without subjecting the simulation to any modelling approaches. This necessarily requires the computational grid to resolve the three dimensions of space and one dimension of time at the Kolmogorov scales. More precisely, the grid spacing in each direction of the simulation domain, Δ^6 , must be on the order of η_K , while the size of the timesteps must be on the order of τ_K .

⁶ Assuming the grid spacing is uniform in x , y and z directions.

Referring to Equations 2.45 and 2.46, one can roughly estimate the computational requirement of DNS to scale approximately as $Re^{11/4}$ [136]. The tremendous computational effort renders DNS unfeasible for practical applications as the turbulent Reynolds number can exceed well beyond 10^5 [112]. Another difficulty is the implementation of the correct initial and boundary conditions; if these are inaccurate, then even a perfect grid would be unable to provide an accurate solution. DNS is presently a useful research tool for studying the nature of turbulence, for both reacting and non-reacting flows, in relatively low Reynolds number environments. Moreover, the detailed data obtained from the DNS can also be used to validate other simulation models [99, 112].

The restricted applicability of DNS motivates researchers to introduce methods that can decrease the computational cost of fluid simulations. This can often be achieved by reducing the degrees of freedom of the problem by solving for averaged quantities. The most intuitive approach is to compute the time-averaged solution of the Navier-Stokes — this method is the well-known Reynolds-Averaged Navier-Stokes simulation paradigm. By the Reynolds decomposition, the instantaneous velocity, u_i , can be separated into a mean velocity, \bar{u}_i , and a fluctuating component, u'_i ,

$$u_i = \bar{u}_i + u'_i. \quad (2.53)$$

The mean velocity is defined as either,

$$\bar{u}_i = \frac{1}{t^*} \int_0^{t^*} u_i dt, \quad (2.54)$$

or,

$$\bar{u}_i = \frac{\sum_{j=1}^{\mathcal{M}} u_{i,j}}{\mathcal{M}}. \quad (2.55)$$

In the former equation, t^* is the period over which the average should be performed, where t^* must be larger than the time scale of the large-scale motion, l_0/U^7 . In the latter equation, $u_{i,j}$ is the velocity in the j^{th} realisation in an ensemble of \mathcal{M} realisations, where \mathcal{M} is taken to be sufficiently large. This suggests that while the mean field can be obtained from RANS, all of the turbulent features are unresolved and

⁷This time scale is also commonly referred to as the large-eddy turnover time.

require modelling. Figure 2.6 illustrates a symbolic representation of the resolved and modelled energy spectrum of the solution in RANS. The statistical characteristics of RANS prevent the simulation from isolating transient turbulent structures and therefore it is not a useful tool for studying the fundamental behaviour of turbulence. However, the general thermodynamic performance of combustion devices can be obtained if the turbulent fluctuation is properly modelled. As the computational cost of RANS is significantly lower than that of DNS and LES, it is still a preferred method for industrial applications.




Figure 2.6: A schematic representation of the resolved and modelled energy spectrum in RANS.

Prior to deriving the RANS governing equations, it is useful to introduce a density-weighted averaging method referred to as Favre averaging, which serves the purpose of simplifying the governing equations by eliminating the fluctuating terms related to the density in compressible flows. It is defined as,

$$\tilde{u}_i \equiv \frac{\overline{\rho u_i}}{\bar{\rho}}. \quad (2.56)$$

The Reynolds decomposition can then be restated in terms of the Favre notation,

$$u_i = \tilde{u}_i + u_i'', \quad (2.57)$$

with u_i'' being the fluctuations around the Favre average. Substituting into Equ-

tion 2.4, one can obtain the Favre-averaged conservation of momentum equation⁸,

$$\frac{\partial \bar{\rho} \tilde{u}_j}{\partial t} + \frac{\partial}{\partial x_i} \left[\bar{\rho} \tilde{u}_i \tilde{u}_j + \underbrace{\overline{\rho u_i'' u_j''}}_{\text{unclosed}} + \bar{p} \delta_{ij} - (\tilde{\tau}_{ij} + \underbrace{\overline{\tau_{ij}''}}_{\text{unclosed}}) \right] = \bar{\rho} g_i. \quad (2.58)$$

In general, $\overline{\tau_{ij}''}$ can be safely neglected for most flows as $|\tilde{\tau}_{ij}| \gg |\overline{\tau_{ij}''}|$. The $\overline{\rho u_i'' u_j''}$ represents the Reynolds stress tensor, which remains unclosed. Indeed, providing suitable models for this quantity is the fundamental study of RANS. Many closure models for the Reynolds stress tensor have been proposed in the past; each has excellent features in certain applications in comparison with the others. These models include: (i) the one-equation formulation of Spalart-Allmaras [159]; (ii) the two-equation formulations, such as the k - ε model of Launder [94] and the k - ω model of Wilcox [187]; (iii) the stress-transport models of Launder [95]; and (iv) the algebraic models of Baldwin-Lomax [6].

As computational hardware became more powerful in the past decades, a modification to the traditional RANS approach, referred to as the Unsteady Reynolds-Averaged Navier-Stokes (URANS), had been proposed. A change in formulation in this method allows for some of the low-frequency modes, usually in the range of a few hundred hertz, to be resolved [141]. It has a number of different names, including, Semi-Deterministic Simulation (SDS), Very Large Eddy Simulation (VLES) and Coherent Structure Capture (CSC) [9, 179]. Though this unsteady formulation contains more information than the usual RANS, the majority of the wave modes are still modelled; as a result, this approach shares many of the underlying deficiencies as RANS.

It appears, then, a shift in the fundamental averaging process is necessary to capture the turbulent motion in fluid flows. A solution is to resolve the large-scale structures and to model the small-scale structures in the flow field by the application of a spatial filter — this is the underlying principle of large eddy simulation [141]. In turbulent flows, features such as mixing, instabilities and certain transient effects that are governed by the large-scale structures can be computed exactly by LES. Furthermore, modelling the fine-scale structures is justified by the first hypothesis of Kolmogorov, where the small turbulent motions tend to have

⁸ This process can be repeated similarly for the continuity and energy equations.

universal characteristics dependent only on ν and ϵ of the fluid, especially at higher Reynolds numbers [89]. The unresolved flow features that are too small for the grid are commonly described as Subgrid-Scale (SGS) or Subfilter-Scale (SFS) structures. A two-dimensional schematic view of the resolved and subgrid-scale flow structures in a grid with spacing of Δ is shown in Figure 2.7.

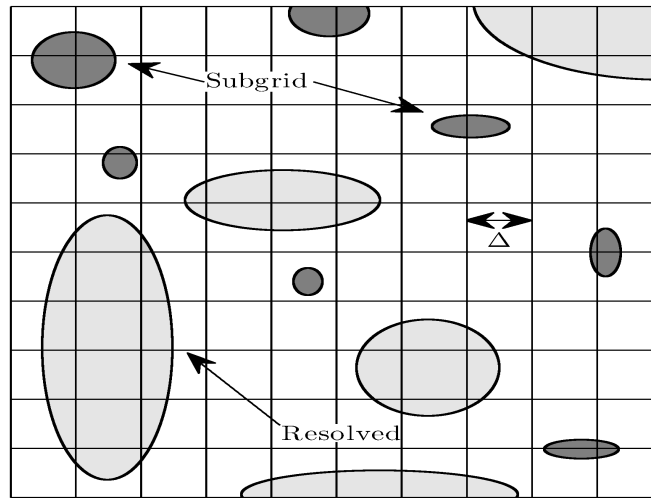


Figure 2.7: Schematic view of the resolved and subgrid-scale structures in a grid with spacing of Δ : Resolved structures (light gray) and subgrid-scale structures (dark gray).

The grid spacing, Δ , acts as an implicit filter that eliminates all of the high-frequency wave modes from direct computation. For anisotropic grids, the filter width is calculated as [37],

$$\Delta = (\Delta_1 \Delta_2 \Delta_3)^{1/3}. \quad (2.59)$$

In practice, the filter width is typically assumed to be proportional to the local cell size and does not need to be exactly equal. Figure 2.8 symbolically demonstrates the separation between the resolved and modelled scales in the energy spectrum for LES. The boundary is solely established by the filter width; as one decreases Δ , one can resolve more high-frequency modes at the expense of increasing the computational cost. Further discussion on the topic of implicit filtering can be found in [61, 141, 150, 151].




Figure 2.8: A symbolic representation of the resolved and modelled turbulent kinetic energy spectrum in LES.

While the implicit filter is commonly found in LES implementations because of its convenience and its ability to take complete advantage of the grid resolution, it includes the unavoidable drawbacks of truncation errors and a filter function that cannot be exactly determined. This may be of issue as wave modes are not identically filtered in different regions of the grid. For explicit LES, the spatial filtering is defined as [97],

$$\langle u_i(\vec{x}, t) \rangle = \int_V G(\vec{x} - \vec{x}') u_i(\vec{x}', t) d\vec{x}', \quad (2.60)$$

where $G(\vec{x})$ and $\langle u_i \rangle$ are respectively the low-pass filter and the filtered velocity component. The three classical filters commonly found in the literature are: (i) the spectral filter, where all of the wave numbers $\kappa > \pi/\Delta$ are eliminated; (ii) the top-hat filter, where a volume average over a sphere of radius $\Delta/2$ is performed [136]; and (iii) the Gaussian filter, where it represents the joint normal distribution with a mean of zero and a covariance of $\Delta^2/12$. Each of these filter functions satisfies the following normalisation condition,

$$\int G(\vec{x}) d\vec{x} = 1. \quad (2.61)$$

An additional complexity arises in the spatial filtering introduced by LES that does not appear in RANS. The differentiation and filtering operations do not necessarily commute if a non-uniform filter is applied to the governing equations. It has been found that the commutation errors are of second order, $\mathcal{O}(\Delta^2)$ [54]. As such, it may be unnecessary to employ discretisation schemes higher than order

two because the commutation errors would dominate.

One can also derive the momentum equation in the context of spatial filtering and Favre averaging for LES by first redefining \tilde{u}_i , which now represents the Favre-filtered velocity component,

$$\tilde{u}_i \equiv \frac{\langle \rho u_i \rangle}{\langle \rho \rangle}. \quad (2.62)$$

Substituting into Equation 2.4, one can obtain the Favre-filtered conservation of momentum equation⁹,

$$\langle \rho \rangle \frac{\partial \tilde{u}_j}{\partial t} + \frac{\partial}{\partial x_i} \left[\langle \rho \rangle \tilde{u}_i \tilde{u}_j + \langle p \rangle \delta_{ij} - \tilde{\tau}_{ij} + \underbrace{\sigma_{ij}^{sgs}}_{\text{unclosed}} \right] = \langle \rho \rangle g_i, \quad (2.63)$$

where $\sigma_{ij}^{sgs} \equiv \langle \rho \rangle (\widetilde{u_i u_j} - \tilde{u}_i \tilde{u}_j)$ is the subgrid-scale turbulent stress tensor and $\tilde{\tau}_{ij}$ is the viscous stress tensor defined as,

$$\tilde{\tau}_{ij} = \langle \mu \rangle \left(\frac{\partial \tilde{u}_i}{\partial x_j} + \frac{\partial \tilde{u}_j}{\partial x_i} \right) - \frac{2}{3} \langle \mu \rangle \delta_{ij} \frac{\partial \tilde{u}_k}{\partial x_k}. \quad (2.64)$$

To provide closure for the SGS turbulent stress tensor, two main classes of modelling approaches are available, the functional modelling and the structural modelling. In the functional approach, the unclosed, subgrid scales are modelled according to the resolved scales. Some examples include the one-equation eddy-viscosity model [149]¹⁰, the Smagorinsky model [155], the dynamic Smagorinsky model [53] and a number of variants of the preceding models [54, 100, 111]. These models are based on solving equations for a turbulent viscosity, ν_T , which can then be related to σ_{ij}^{sgs} . In the structural approach, the subgrid scales are computed without directly relating the resolved fields to the SGS quantities. Examples include the scale-similarity model [7], the differential stress equation model [38], the non-linear anisotropic model [103] and so forth [141].

⁹ This process can be repeated similarly for the continuity and energy equations. The closure of the additional fluctuating terms can be found in [141].

¹⁰ The one-equation eddy-viscosity model is used for all of the LES calculations in this thesis.

2.10 Turbulent Premixed Combustion Modelling

First moment closure, whereby the filtered or averaged reaction rates are directly computed by the filtered or averaged values of the scalar fields via the Arrhenius rate equations, can lead to considerable discrepancies compared to the fully resolved values, due to the highly non-linear interactions between chemical source terms and the fluctuating scalar fields. The chemically reactive regions are confined to the inner layer, which exists at a smaller length scale than the average LES grid resolution. In consequence, combustion models are required for LES, which, by their nature, cannot fully resolve the turbulence-chemistry interactions. A number of different combustion models have been proposed for LES; many of these are derived from concepts developed for RANS approaches. While the remainder of this section will present brief descriptions of some of the most commonly cited combustion models, it is important to mention that the modelling of the filtered reaction rates remains a topic of significant research in the combustion community.

Eddy Break-Up Model

One of the simplest closures for chemical source terms is an algebraic approach referred to as the Eddy Break-Up (EBU) model, first introduced by Spalding [160]. This model assumes fast chemistry and thus the rate determining process for the chemical source terms would be completely governed by the rate of mixing. The chemical source term under such turbulent conditions for the products ($\bar{\omega}_P$) are estimated to be proportional to the turbulent time scale and the variance of mass fraction of the product ($\overline{Y_P'^2}$),

$$\bar{\omega}_P = C_{\text{EBU}} \bar{\rho} \frac{\varepsilon}{k} \overline{Y_P'^2}, \quad (2.65)$$

where C_{EBU} is the Eddy Break-Up constant.

Although the EBU model can be successfully applied to some RANS simulations, more accurate closures are desired for LES [181]. The problem here is that the EBU model does not take into account any aspect of chemistry; it can neither distinguish between different equivalence ratios nor identify different types of fuels. Indeed, the reaction rate is linked directly to the tunable Eddy Break-Up con-

stant, which must be adjusted on a case-by-case basis. As a result, the chemistry effects are completely *ad hoc* and non-physical results may ensue [181].

An improvement to the EBU model was made subsequently by Magnussen and Hjertager [105], termed the Eddy Dissipation Model (EDM). The modification replaces the variance of mass fraction of the product by the mean mass fraction of the deficient species, for example, fuel in lean mixtures or oxygen in rich mixtures. This allows the model to distinguish between rich and lean mixtures; however, it shares the same drawbacks as the EBU in that the chemistry effects are completely *ad hoc* and the constant, C_{EDM} for this model, would require adjustment for each new application.

Bray Moss Libby (BML) Model

The Bray Moss Libby (BML) model, similar to the EBU and EDM models, is an algebraic approach first introduced by Bray, Moss and Libby [20, 98]. While it has been improved a number of times since its introduction [21, 24], the BML model fundamentally assumes the flame to be infinitely thin, where the progress variable would intermittently take on a value of 0 or 1. This model generally uses a temperature-based progress variable,

$$c = \frac{T - T_u}{T_b - T_u}, \quad (2.66)$$

where T , T_b and T_u are respectively the local temperature, the burnt gas temperature and the unburnt gas temperature; this formulation is also referred to as the reduced temperature. Under the thin-flame assumption, the PDF of the progress variable can be reduced to the following form with two Dirac delta functions,

$$P(c) = \alpha_{\text{BML}} \delta(c) + \beta_{\text{BML}} \delta(1 - c), \quad (2.67)$$

where $\alpha_{\text{BML}} + \beta_{\text{BML}} = 1$. Combining this PDF and the scalar dissipation rate formulation derived by Bray and Moss [20], it is possible to obtain the average reac-

tion rate of the progress variable,

$$\bar{\omega}_c = \frac{\bar{\rho} \tilde{\chi}_c}{2C_{\text{BML}} - 1}, \quad (2.68)$$

where $\tilde{\chi}_c$ is the scalar dissipation rate of the progress variable. Furthermore, the parameter, C_{BML} , is defined as [20],

$$C_{\text{BML}} = \frac{c \bar{\omega}_c}{\tilde{\omega}_c}. \quad (2.69)$$

One method to obtain an approximation to the scalar dissipation rate would be a simple balance equation [106],

$$\bar{\rho} \tilde{\chi}_c = \bar{\rho} \tilde{\chi}_c \approx \frac{\bar{\rho} \tilde{c}^2}{\tau_{\text{BML}}}, \quad (2.70)$$

where \tilde{c}^2 and τ_{BML} are the variance of the progress variable and the turbulent mixing time scale in the BML model, respectively. Finally, if one assumes the τ_{BML} to be related to the turbulent kinetic energy, k , and the dissipation rate, ε , as $\tau_{\text{BML}} = k/\varepsilon$, one can recover the following equation for the production rate of the progress variable,

$$\bar{\omega}_c = \frac{1}{2C_{\text{BML}} - 1} \bar{\rho} \frac{\varepsilon}{k} \tilde{c}^2. \quad (2.71)$$

Indeed, a brief comparison of the preceding equation with Equation 2.65 would reveal that the BML model is closely related to the EBU model; some view the BML model as a theoretical derivation of the EBU model, except it is superior in that the C_{BML} parameter partially captures the effects of chemistry. However, much like its counterpart, the BML model is not able to account for the effects of chemistry accurately.

Flame Surface Density (FSD) Model

The Flame Surface Density (FSD) model is a flamelet approach that assumes the flame front is a surface that can be convected, diffused and strained by the turbulent interactions [64]. Assuming c is the usual progress variable, where it can

be based on the temperature or a mass fraction:

$$c = \frac{T - T_u}{T_b - T_u}, \quad (2.72)$$

$$c = \frac{Y_{\text{CO}_2} - Y_{\text{CO}_2}^u}{Y_{\text{CO}_2}^b - Y_{\text{CO}_2}^u}, \quad (2.73)$$

where Y_{CO_2} , $Y_{\text{CO}_2}^u$ and $Y_{\text{CO}_2}^b$ are respectively the local, unburnt and burnt CO_2 mass fractions¹¹. By definition, $c = 0$ in reactants and $c = 1$ in products, and the transport equation for c can be established as,

$$\frac{\partial(\rho c)}{\partial t} + \frac{\partial(\rho u_i c)}{\partial x_i} = \frac{\partial}{\partial x_i} \left(\rho \mathbf{D}_c \frac{\partial c}{\partial x_i} \right) + \dot{\omega}_c = \rho s_d |\nabla c|. \quad (2.74)$$

Here, s_d is the local displacement speed of the iso-surface of the progress variable. One can filter the preceding equation and obtain a term related to the displacement of the iso-surface, $\overline{\rho s_d |\nabla c|}$, which requires closure. The concept behind the FSD model is to relate this displacement term to the flame surface density,

$$\overline{\rho s_d |\nabla c|} \approx \rho_u s_L \Sigma = \rho_u s_L \Xi_\Delta |\nabla c|, \quad (2.75)$$

where ρ_u , Σ and Ξ_Δ are respectively the unburnt mixture density, the SGS flame surface density and the SGS wrinkling factor. The closure of Σ has been proposed in the form of transport equations [17, 64], algebraic formulations [17] and a scale-similarity model [87]. Furthermore, the closure of Ξ_Δ can be found in the form of transport equations [50, 184] and algebraic expressions [30, 31, 51]. Understandably, the accurate closure of the SGS terms and the modelling of the effects of strain on the flame surface constitute some of the challenges in applying the FSD model in LES calculations [47, 130, 132].

¹¹ The mass fraction CO_2 can be used as the progress variable for lean to stoichiometric methane-air combustion under atmospheric conditions.

G-Equation Model

The G -Equation model was introduced by Markstein [107] and is based on a level set approach in the calculation of a G -field. G typically takes on a value of less than 0 in the unburnt mixture, 0 at the flame front and greater than 0 in the burnt mixture. The transport equation for G is,

$$\frac{\partial G}{\partial t} + \tilde{u}_i \frac{\partial G}{\partial x_i} = s_T |\nabla G|, \quad (2.76)$$

where s_T is the turbulent flame speed. While there exists several empirical correlations for s_T [44] in the RANS context, this parameter is associated with some uncertainties in experimental measurements [132]. A detailed discussion of the G -equation model in the RANS context can be found in [128].

In the LES context, a number of G -equation approaches have been proposed [69, 83, 131]. Interestingly, the G -field, as originally proposed, exhibits a special symmetry which makes it incompatible with LES [119]. Only recently has Pitsch [129] achieved consistency in the LES filtering of the G -field by the addition of a term which includes the effects of curvature on the flame propagation. Despite the simple formulation, the G -equation approach has been used in a number of simulations, including LES of a premixed swirl-stabilised combustor [68], low-swirl premixed flame [117] and other flames [69, 83, 88]. As with the problem found in the RANS formulation, closure for the turbulent flame speed is required; there does not appear to be a universal consensus on a suitable model or correlation for this parameter [14].

Artificially Thickened Flame Model

The artificially thickened flame model, proposed by Butler and O'Rourke [27], was introduced to overcome the problem of unresolved flame fronts on coarse computational grids. The premise of the model, as the name implies, is to artificially increase the thickness of the flame while preserving the flame speed. This is achieved by manipulating the basic scaling laws from the theory of laminar premixed flames. In general, the laminar flame thickness, δ_L , can be related to the molecular diffu-

sivity, \mathbf{D} , and the laminar flame speed, s_L ,

$$\delta_L \propto \frac{\mathbf{D}}{s_L}. \quad (2.77)$$

Furthermore, the laminar flame speed is proportional to the molecular diffusivity and mean reaction rate, $\bar{\omega}$, as follows,

$$s_L \propto \sqrt{\mathbf{D}\bar{\omega}}. \quad (2.78)$$

By increasing the flame thickness by a factor of Q , from δ_L to $Q \cdot \delta_L$, one must also increase the diffusivity from \mathbf{D} to $Q \cdot \mathbf{D}$ and decrease the mean reaction rate from $\bar{\omega}$ to $\bar{\omega}/Q$. An appropriate value of Q would be in the range of 5–30 to properly resolve the artificially thickened flame front on typical LES grids. This model was first applied to LES by Veynante and Poinot [1] and has been developed further thereafter [35]. Since its introduction, the thickened flame model has been used to simulate combustors with complex geometry [152], to study flashback and blow-off effects in a swirl burner [158], and to simulate the ignition sequence of a gas turbine engine [18].

While the artificially thickened flame model demonstrates merit in accounting for ignition and near-wall interaction, the factor, Q , is not complete and the physics of the turbulence-chemistry interaction is not represented precisely. Indeed, an additional adjustable parameter, termed the efficiency factor, is required in the species transport equations to negate the changes introduced to the Damköhler number by Q .

Presumed Conditional Moment (PCM) Model

The Presumed Conditional Moment (PCM) Model was originally proposed by Vervisch *et al.* [180] for non-premixed combustion in the RANS paradigm. This model was later extended to LES as a SGS model by Domingo *et al.* for the simulation of a premixed flame [39]. The model assumes a locally laminar structure for a progress variable in the turbulent flames; the mean reaction rates are simply

calculated by a convolution with the most appropriate PDF model,

$$\tilde{\omega}_c = \int_0^1 \overline{\omega_c|c^*} \tilde{P}(c^*; \vec{x}, t) dc^*, \quad (2.79)$$

where ω_c represents the source term for the selected progress variable. The presumed PDF model can be the usual β -PDF for non-premixed combustion or the modified laminar flamelet PDF for premixed combustion [70, 153]. The $\overline{\omega_c|c^*}$ is typically obtained from canonical, one-dimensional laminar flame calculations and tabulated prior to the LES or RANS simulations.

One of the primary advantages of PCM is that it can be conveniently combined with tabulated chemistry reduction techniques, such as the Flame Prolongation of ILDM (FPI) model [55] described in Section 2.5.2. In this way, the PCM-FPI simulations can emulate the effects of detailed chemistry at a fraction of the computational cost of strategies requiring reaction mechanisms. The PCM-FPI approach has demonstrated good agreement with experimental results in the simulations of a lifted jet flame [40] and a lean premixed turbulent swirl flame [52]. Furthermore, this strategy has also been successfully applied to premixed Bunsen flames for a laboratory-scale burner [143, 145, 153]. The PCM-FPI appears to be a promising technique for the simulations of premixed and non-premixed combustion, allowing for detailed chemistry at a reasonable computational cost. Its main disadvantage is that it is only formally valid in the flamelet regime, which could restrict its value for predictions of flames in industrial applications.

Transported PDF Model

The transported PDF model, proposed by Pope [134], solves a transport equation for the joint PDF of the velocity and composition. Unlike the previous models, this one does not assume a locally laminar structure for the turbulent flames and is applicable to combustion problems outside of the flamelet regime. The chemical source terms are perfectly closed at the expense of unclosed diffusion terms in the transport equation for the joint PDF. This tradeoff is of particular concern for premixed combustion where the turbulent mixing and chemical kinetics are intricately linked. Nevertheless, the transported PDF model has been em-

ployed in a number of RANS simulations for turbulent premixed laboratory-scale flames [28, 101, 115, 163, 182].

The transported PDF model has also been extended to LES by the introduction of a Filtered Density Function (FDF) [56, 135] and successfully implemented for the LES of turbulent non-premixed [154] and premixed flames [189]. However, this model is not ideally suited for industrial problems as the Lagrangian Monte Carlo methods are required for the solution of the transport equation of the joint PDF. The additional computational complexity and cost hinder the transported PDF model from practical applications.

Conditional Moment Closure (CMC)

Conditional Moment Closure (CMC) was first proposed by Klimenko [85] and Bilger [13] independently in the context of non-premixed combustion for RANS. CMC has a theoretical advantage over flamelet models in that it does not assume the turbulent flame to be spatially thin. The RANS-CMC approach has demonstrated reliability in the chemical closure for non-premixed combustion [47, 132] and for diesel spray combustion in engines [124]. The successes led to the extension of CMC to premixed combustion [84, 166]. More recently, CMC has been applied in the simulation of stoichiometric pilot-stabilised Bunsen flames [4]. Furthermore, the LES formulations of CMC have also been developed for both non-premixed [82, 116] and premixed [174] combustion, though the LES-CMC implementation for industrial applications appears to be less frequent.

The underlying hypothesis of CMC assumes that the conditionally-averaged chemical source terms ($\overline{\dot{\omega}_\alpha|c^*}$) are closed by evaluating the reaction rate expressions with the conditionally-averaged scalars,

$$\overline{\dot{\omega}_\alpha|c^*} \approx \dot{\omega}_\alpha(\overline{T|c^*}, \overline{Y_\alpha|c^*}, \overline{\rho|c^*}), \quad (2.80)$$

where c^* is the discretised conditioning variable. $\overline{T|c^*}$, $\overline{Y_\alpha|c^*}$ and $\overline{\rho|c^*}$ are the conditional temperature, mass fraction of species α and density, respectively. The unconditionally-filtered chemical source term is subsequently calculated by integrating the conditional source term with the PDF of the conditioning variable,

$$\tilde{P}(c^*; \vec{x}, t),$$

$$\tilde{\omega}_\alpha = \int_0^1 \overline{\omega_\alpha | c^*} \tilde{P}(c^*; \vec{x}, t) dc^*. \quad (2.81)$$

There are three main challenges in applying CMC to practical problems for turbulent premixed combustion. First, the conditionally-averaged scalars are calculated from transport equations that include the conditioning variable as an additional independent variable; this added dimension translates to significantly higher computational cost. Second, these transport equations contain a number of unclosed terms that are not well understood and require modelling [12]. As a result, the closure for these terms are typically not accurately represented. Finally, for premixed combustion, the solution is extremely sensitive to the model for the conditional scalar dissipation. Models for this term, which permit CMC to be used as a chemical closure strategy for premixed combustion, have only recently been developed.

Conditional Source-term Estimation (CSE)

Conditional Source-term Estimation, first proposed by Bushe and Steiner [26], is a combustion model that provides closure for the mean chemical source terms for turbulent flames. This model has been shown to produce promising results for diffusion flames [67, 93, 161, 183]. Further, Salehi *et al.* has demonstrated that CSE can also be applied to premixed flames and produce meaningful results [42, 143–145, 153]. Recent effort has extended CSE to lifted turbulent flames [43] and flames operating in the Moderate and Intense Low Oxygen Dilution (MILD) combustion mode [92].

CSE invokes the first-order Conditional Moment Closure hypothesis such that the conditionally-averaged chemical source terms are determined by evaluating the reaction rate expressions with the conditionally-averaged scalars (Equation 2.80). The unconditional mean chemical source term is calculated in the same manner, by integrating the conditional source term with the PDF of the conditioning variable (Equation 2.81). The Favre-averaged presumed PDF models adopt the functional form:

$$\tilde{P}(c^*; \vec{x}, t) \approx \tilde{P}(c^*; \tilde{c}, \tilde{c}^{\prime 2}), \quad (2.82)$$

where \tilde{c} and \tilde{c}^2 are the Favre-averaged mean and variance of the progress variable. The usual spatial and temporal coordinates are represented by \vec{x} and t . Several PDF models are tabulated in this format; a PDF model commonly implemented in conjunction with CSE is the modified laminar flamelet PDF [70, 143–145, 153].

The central concept behind CSE is that there is more information available if the filtered field is considered in its entirety rather than as unrelated data points. Each filtered volume, governed by the same underlying physical processes, then contributes additional information about the localised reactive flow. Specifically, the relevant conditional scalar averages are obtained by the inversion of the following integral equation [26],

$$\tilde{Y}_\alpha(\vec{x}, t) = \int_0^1 \tilde{P}(c^*; \vec{x}, t) \overline{Y_\alpha | c^*}(\vec{x}, t) dc^*, \quad (2.83)$$

where $\tilde{Y}_\alpha(\vec{x}, t)$ is the Favre-averaged mass fraction. Figure 2.9 illustrates the interaction between the CFD module, the CSE module, the chemistry module (FGM in this example) and the PDF. The mass fractions, means and variances of the progress variable are calculated by the CFD module via transport equations. These means and variances of the progress variable are then used to obtain the PDFs, which in conjunction with the mass fractions, serve as the input to the CSE module via Equation 2.83. The computed conditional mass fractions are subsequently used to determine the mean reaction rates through the FGM and the PDF model.

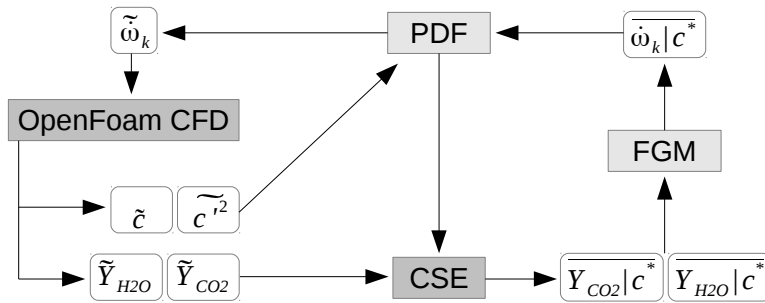


Figure 2.9: CSE operational flowchart.

Chapter 3

Pseudo-Turbulent Probability Density Function and Scalar Dissipation Rate Models for Premixed Combustion

3.1 Introduction

There are several different models for turbulent premixed combustion that rely on having an accurate probability density function (PDF) of a reaction progress variable [39, 55, 70, 121]; these PDFs are conventionally expressed as a function of the mean and variance of that progress variable. Such presumed PDF approaches are used to account for finite rate chemistry, with tabulated data used to include detailed chemistry into simulations of turbulent combustion. Presumed PDF models are being developed for both the Reynolds-Averaged Navier-Stokes (RANS) and LES paradigms. However, the accuracy of these methods inevitably depends to some considerable degree on the accuracy of the function presumed for the PDF of the progress variable.

The oft-used β -PDF is able to recover the extreme properties expected of the true PDF, such as having δ s appear at 0 and 1 for maximal variance or a single δ

appear at the mean for a zero variance; however, it fails to reproduce the shape of the true PDF [70]. Here, the issue is the shape of the PDF is strongly dependent on how the chemical reaction rates vary as a function of the progress variable. Hence, different chemical kinetics will lead to a different shape of the PDF of progress variable and the distribution of the β -PDF will, for most cases, lead to significant inaccuracies. Most critically, the inaccuracies can lead to a bias, as one typically finds that the error is consistently always at the same values of progress variable in any particular flame.

In an attempt to account for the effects of chemistry on the shape of the PDF, Bray *et al.* [22] proposed using a premixed laminar flame to model the functional dependence of the PDF on progress variable. Their original formulation only provided coverage for flames with very high variances. Jin *et al.* [70] proposed a modification of the Bray PDF that extends the original formulation to cover all possible mean and variance combinations. This is accomplished by truncating the function as needed to match the required mean and variance. It was found to significantly improve the fit of the PDF to that extracted from DNS results. However there remains the issue that, at the point of truncation, the model PDF has a sharp drop to zero; the true PDF tends to be more rounded.

In this chapter, a one-dimensional pseudo-turbulent method is proposed to replace the laminar flame that is conventionally used to generate PDF models for the reaction progress variable for turbulent reacting fields; the hope is that this strategy will smooth out these unphysical peaks in the distributions. Furthermore, for flames that are not in the flamelet regime, the pseudo-turbulent calculation might capture whatever effects small-scale eddies penetrating the reaction zone might have on the shape of the distribution.

Hence, the objective of this study is to verify the ability of the Linear-Eddy Model (LEM) to capture the overall PDF distributions for two test scenarios in the premixed combustion mode. Particularly, PDFs generated by LEM simulations are compared to both experimental and DNS results. For the experimental comparison, LEM is used to simulate a premixed methane-air v-flame exposed to low intensity turbulence. Temperature-based PDF models for the reaction progress variable are subsequently constructed. In addition, the scalar dissipation rates are also investigated for this v-flame. For the DNS comparison, LEM is used to sim-

ulate a premixed methane-air flame in an isotropic, homogeneous turbulent field; the product-based PDF models are then constructed.

3.2 Probability Density Function Models

Product-based and temperature-based progress variables can be formulated respectively as follows,

$$c(x, t) = \frac{Y_p(x, t)}{Y_{peq}}, \quad (3.1)$$

$$c(x, t) = \frac{T(x, t) - T_u}{T_b - T_u}, \quad (3.2)$$

where Y_{peq} , T_b and T_u are respectively the mass fraction of the products at complete combustion, burnt gas mixture temperature and unburnt gas temperature. In either case, one can formulate a model PDF with the following form,

$$P(c^*; x, t) \approx P(c^*; \bar{c}, \overline{c'^2}), \quad (3.3)$$

where \bar{c} and $\overline{c'^2}$ are the mean and variance of the progress variable. Such a form is compatible with several current models for the turbulence-chemistry interactions in turbulent premixed combustion.

Four different PDF models will be considered here, three of which have been discussed previously: the β -PDF, the Bray PDF and a modified laminar flamelet PDF. The β -PDF is typically a bimodal distribution widely used in non-premixed combustion.

The Bray, or laminar flamelet PDF, was originally formulated to model flames in the thin flamelet regime. The Bray PDF is defined as [22]:

$$P(c^*; \bar{c}, \overline{c'^2}) = A\delta(c^*) + Bf(c^*) + C\delta(1 - c^*), \quad (3.4)$$

where $f(c^*)$ is calculated using the solution of an unstrained laminar flame:

$$f(c^*) = \frac{1}{|\nabla c^*|}. \quad (3.5)$$

The coefficients A , B and C are obtained based on the first three moment equations

of c :

$$\int_0^1 P(c^*; \bar{c}, \bar{c}^2) dc^* = 1 \quad (3.6)$$

$$\int_0^1 c^* P(c^*; \bar{c}, \bar{c}^2) dc^* = \bar{c} \quad (3.7)$$

$$\int_0^1 c^{*2} P(c^*; \bar{c}, \bar{c}^2) dc^* = \bar{c}^2. \quad (3.8)$$

The Bray PDF is problematic in that it cannot produce PDFs for all values of mean and variance [143].

Jin *et al.* [70] proposed a modification to this PDF whereby the function $f(c^*)$ is truncated to zero over a range of c^* and one or both of the δ s is eliminated from the PDF in order to expand the applicable range of the PDF to all possible values of mean and variance of c . In doing this, it was first recognized that the PDF will ultimately be stored in a look-up table for discrete intervals in c^* . That discretisation is used to provide bounds for using the laminar flame solution to assemble the PDF. A lower bound at c_{\min} is defined which would be the maximum c^* in the first discrete interval (that is to say that this discrete interval in c^* would be bounded by 0 and c_{\min}) and an upper bound at c_{\max} is defined to be the minimum c^* in the last discrete interval (this discrete interval in c^* would be bounded by c_{\max} and 1).

A region bounded by $[x_1, x_2]$ — with corresponding values of the progress variable $[c_1, c_2]$ — is located in the laminar flame such that the mean and variance of c sampled within $[x_1, x_2]$ match a given combination of \bar{c} and \bar{c}^2 ; x_1 and x_2 then are compared to x_{\min} and x_{\max} , which are the points in the laminar flame solution that correspond to c_{\min} and c_{\max} . The functional form of the modified PDF changes depending on the locations of x_1 and x_2 relative to x_{\min} and x_{\max} . Four cases are then possible:

- Case 1: $x_{\min} < x_1 < x_2 < x_{\max}$

$$P(c^*; \vec{x}, t) = \begin{cases} 0 & \text{if } c^* < c_1 \text{ or } c^* > c_2 \\ Bf(c^*) & \text{if } c_1 \leq c^* \leq c_2 \end{cases}$$

- Case 2: $x_1 < x_{\min} < x_2 < x_{\max}$

$$P(c^*; \vec{x}, t) = \begin{cases} A\delta(c^*) + Bf(c^*) & \text{if } c^* \leq c_2 \\ 0 & \text{if } c^* > c_2 \end{cases}$$

- Case 3: $x_{\min} < x_1 < x_{\max} < x_2$

$$P(c^*; \vec{x}, t) = \begin{cases} 0 & \text{if } c^* < c_1 \\ Bf(c^*) + C\delta(1 - c^*) & \text{if } c^* \geq c_1 \end{cases}$$

- Case 4: $x_1 < x_{\min} < x_{\max} < x_2$

$$P(c^*; \vec{x}, t) = A\delta(c^*) + Bf(c^*) + C\delta(1 - c^*). \quad (3.9)$$

As in the Bray PDF, each of these cases has three unknown coefficients (among A , B , C , x_1 and x_2). The first three moment equations of c (Eqs. 3.6 through 3.8) are used to evaluate these. In practical implementation [144], the PDF is calculated in a pre-processing operation where the laminar premixed flame is first computed, then it is used with the above rules to fill in the complete range of mean and variance of c and store the PDF in a table for later retrieval.

Figure 3.1 shows PDFs collected from DNS data compared to the three different model PDFs described above. The β -PDF overestimates the probability over a wide range of progress variable for $c < 0.6$, then utterly fails to capture the peak around $c = 0.85$ and instead rises to a δ at $c = 1$. The original laminar flame PDF of Bray underestimates the PDF virtually over the entire range of progress variable and also includes a δ at $c = 1$ that is not exhibited by the DNS. The modified laminar flamelet PDF clearly does the best job of representing the DNS PDF but it has too high a peak towards the product side of the distribution and exhibits a sharp cut-off not evident in the DNS. To try to smooth off the sharp cut-off, a one-dimensional pseudo-turbulent flame calculation obtained using the Linear-Eddy Model of Kerstein [77] is proposed to replace the usual laminar flame calculation for the function f in the formulation above. The four PDF cases described for the modified laminar flamelet PDF will be the same, including the need to identify locations in space, x_1 and x_2 , at which to truncate the solution so as to fit the first




Figure 3.1: PDFs are exemplified at two axial locations of the DNS simulation domain ($L_d = 6$ mm). Solid: DNS; dash: modified laminar flamelet; dash dot: laminar flamelet; dot: β -PDFs [70].

three moments of c .

3.3 Linear-Eddy Model

The Linear-Eddy Model is selected due to its minimal empiricism and its ability to replicate turbulent flow conditions inexpensively [77–81]. It can be fundamentally divided into two modules. The deterministic component consists of one-dimensional gas dynamics evolutionary equations. The stochastic component consists of random *eddy events*. The two processes are implemented simultaneously during the simulation to achieve a pseudo-turbulent effect. The coupling between the stochastic advective process and deterministic evolutionary equations on a one-dimensional computational domain permits the LEM to simulate turbulent flows with higher Reynolds numbers than multi-dimensional models. The current LEM variant is adapted from the original formulation to accommodate turbulent premixed reacting flows.

3.3.1 Evolutionary Equations

The following set of evolutionary equations is derived for an unsteady, isobaric, quasi-one-dimensional flame [76]. Small changes in pressure across the premixed

flame reaction zone permit the assumption of an isobaric system. The continuity, energy, species transport equations and equation of state are respectively,

$$\dot{M} = \rho u_x A_c, \quad (3.10)$$

$$\rho A_c \frac{\partial T}{\partial t} + \dot{M} \frac{\partial T}{\partial x} - \frac{1}{c_p} \frac{\partial}{\partial x} \Lambda A_c \frac{\partial T}{\partial x} + \frac{A_c}{c_p} \sum_{\alpha=1}^{\mathcal{N}} \rho c_{p\alpha} Y_\alpha V_\alpha \frac{\partial T}{\partial x} + \frac{A_c}{c_p} \sum_{\alpha=1}^{\mathcal{N}} \dot{\omega}_\alpha h_\alpha M_\alpha = 0, \quad (3.11)$$

$$\rho A_c \frac{\partial Y_\alpha}{\partial t} + \dot{M} \frac{\partial Y_\alpha}{\partial x} + \frac{\partial}{\partial x} (\rho A_c Y_\alpha V_\alpha) - A_c \dot{\omega}_\alpha M_\alpha = 0, \quad (3.12)$$

$$\rho = \frac{p \bar{M}}{R_u T}. \quad (3.13)$$

The \dot{M} , u_x , A_c and V_α respectively represent the mass flow rate, streamwise velocity (assumed to be in the x direction), the cross-sectional area of the streamtube¹ and the diffusion velocity of species α . The other symbols are previously defined in Section 2.2.

3.3.2 Stochastic Advection

LEM postulates a random process that rearranges fluid elements along a line in order to simulate the chaotic vortices that appear in turbulent flows. The inversion generates discontinuous fluid motions, which lead to random walks of fluid elements. Each inversion requires three random variables: epoch, position and length. These variables are determined from the turbulent Reynolds number (Re_t) governing the flow via Kolmogorov scaling [77]. For the purposes of the LEM simulations to be performed here, the turbulent Reynolds number should correspond to the turbulent Reynolds number of the flame to be modeled. The eddy inversions, illustrated by Figure 3.2, are now referred to as *triplet maps*.

¹ A_c is assumed to be unity for all of the LEM simulations in this thesis.




Figure 3.2: *Triplet map* applied to a line. Cell values are transported within the eddy interval; interpolation is not required at any point [77]. The before (left) and after (right) domain values are illustrated.

This mapping strategy is able to conserve the scalar fields because the content of each control volume within the selected interval is only rearranged to a neighbouring position. Interpolation is not required at any point. In this way, inadvertent changes in the solution vectors due to continuous implementation of the stochastic advection will not occur.

3.4 LEM Simulation Parameters

3.4.1 LEM Methodology for Experimental Study

A reduced, six-step mechanism designed by Chen *et al.* [29] to simulate premixed methane-air combustion is implemented to replicate the experimental flame. The propagation speed and overall temperature profile for the unstrained laminar flame has been calibrated to the reference solution from Cantera [58] by modifying the reaction rates of the mechanism. This mechanism is selected because it is able to reproduce the temperature profile sufficiently well compared to the full chemistry; other reduced mechanisms with fewer reaction steps tend to be less satisfactory in this regard. For the purposes of this chapter, the temperature gradient must be fairly accurate as the functional form of the PDF model depends on the normalized inverse of this parameter.

All of the species properties are calculated through CHEMKIN-II [74], including specific heats, diffusion coefficients, thermal conductivities and enthalpies.

3.4 LEM Simulation Parameters

Mixture-averaged transport is adopted to reduce computational time. Figure 3.3 illustrates the flame solutions from Cantera and the six-step mechanism. The overall flame thickness, propagation speed, equilibrium temperature and mass fractions of major species are well matched.




Figure 3.3: Premixed laminar flame solution at an equivalence ratio of 0.73 for the six-step reduced mechanism (line) and Cantera [58] (symbols).

The inflow mixture is set to atmospheric pressure and 294 K at an equivalence ratio of 0.73. The integral scale and turbulent Reynolds number are respectively 1.9 mm and 38. The LEM simulations are performed with a mixed first order upwind and second order centered spatial scheme. Explicit timesteps are taken in order to cope with the discontinuities introduced by the stochastic component of the model. Each realisation is a snapshot of the instantaneous temperature profile recorded at $1/100^{\text{th}}$ large eddy turnover time intervals. 10,000 realisations are used to generate the PDFs and scalar dissipation curves at each of the specified turbulent Reynolds numbers.

3.4.2 LEM Methodology for DNS Study

A simplified, two-step, five-species mechanism is used to simulate a methane-air premixed flame in a one-dimensional, homogeneous turbulent field. The two reactions are chosen to be as economical as possible while retaining crucial characteristics of premixed methane-air flames. Further information regarding the mechanism and inlet conditions can be found in [62]. The integral scale and turbulent Reynolds number are respectively 4.5 mm and 51. All species have equal and constant heat capacities of 1250 J/kg·K. The binary diffusion coefficients and thermal conductivities are equal for all species and vary only with temperature. The mechanism and thermochemical parameters are selected to emulate the DNS flame as closely as possible.

Similar numerical methods prescribed in the previous subsection are applied; however, each realisation is a snapshot of the instantaneous product mass fraction recorded at $1/10^{\text{th}}$ large eddy turnover time intervals. 10,000 realisations are averaged to generate the final PDF models.

3.5 Construction of PDF Models

The method used to construct the LEM PDF models will first be discussed in the context of the modified laminar flamelet PDF then extended to the LEM formulation. As per Section 3.2, the relevant parameters governing the distributions are the mean (\bar{c}) and variance ($\overline{c'^2}$) of the PDFs. Truncations (in space) applied to the one-dimensional flame profile would lead to changes in both \bar{c} and $\overline{c'^2}$. The mean of the distribution will increase for every point removed from the unburnt mixture boundary while the contrary is true for points removed from the burnt mixture boundary. The variance can be adjusted in a similar manner. Figure 3.4(a) shows an example of such truncations. The symbols correspond to the truncation limits, (x_1, c_1) and (x_2, c_2) , where only the cells within the interval are retained. The arbitrary values of mean and normalised variance for illustrative purposes are set to $\bar{c} = 0.5$ and $\overline{c'^2}_n = 0.33$; the normalised variance is determined by the limiting variance of a

perfectly segregated mixture of hot and cold gases:

$$\overline{c'^2}_n = \frac{\overline{c'^2}}{\bar{c}(1-\bar{c})}. \quad (3.14)$$

The modified laminar flamelet PDFs with desired means and variances can then be constructed with the truncated flame profiles according to methods outlined in Section 3.2. An example is shown in Figure 3.4(b). Crucially, it can be deduced that there exists one unique modified laminar flamelet PDF for every mean and variance combination.

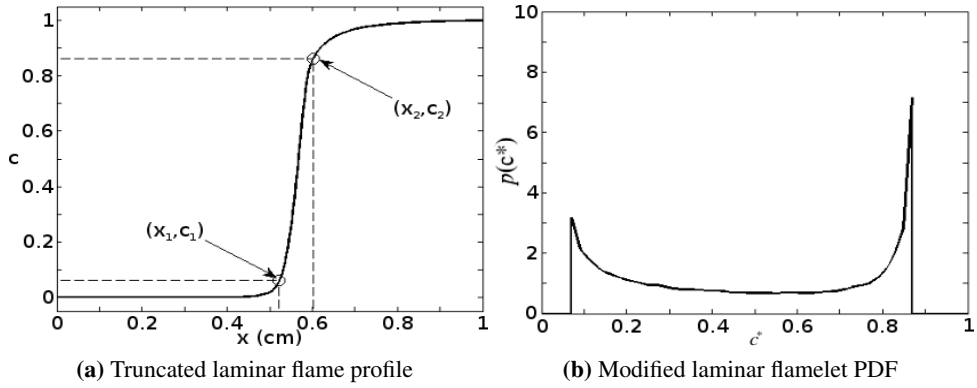


Figure 3.4: (a) (x_1, c_1) and (x_2, c_2) mark the truncation limits where only the cells within the interval are retained; (b) Modified laminar flamelet PDF with $\bar{c} = 0.5$ and $\overline{c'^2}_n = 0.33$ is constructed using the truncated flame profile.

The LEM PDFs are constructed analogously; however, the LEM flame profiles are not steady with time. Figure 3.5 illustrates five sample realisations obtained from LEM calculations where each profile is slightly different but equally valid. Thus, it is necessary to average the contributions from *each* flame realisation in order to arrive at the final model. The truncation strategy prescribed for the modified laminar flamelet profiles to obtain the desired mean and variance is applied to each LEM result.

Figure 3.6(a) illustrates two truncated LEM profiles with $\bar{c} = 0.5$ and $\overline{c'^2}_n = 0.33$, similar to the laminar flame shown in Figure 3.4(a). The truncation positions are shifted for individual realisations because the distributions are inherently dif-

3.5 Construction of PDF Models




Figure 3.5: Individual flame realisations from LEM calculations.

ferent. The final PDF model is acquired by averaging numerous PDFs built from such truncated LEM flame profiles. Figure 3.6(b) illustrates an example of the converged solution constructed by averaging 10,000 LEM PDFs at a relatively low turbulent Reynolds number. It is apparent the LEM PDF displays peaks of lower magnitude at each boundary, thereby redistributing the probability of finding c . This *smoothing* effect is typical of PDFs constructed with LEM.

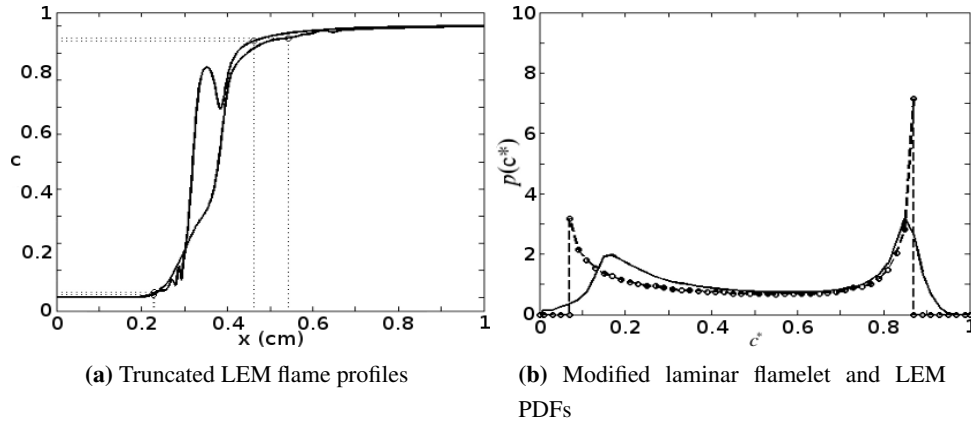


Figure 3.6: (a) The portions to be retained are within the intervals marked by 'o'. The truncation positions are selected such that the resultant PDF would have $\bar{c} = 0.5$ and $\overline{c'^2}_n = 0.33$. The truncation boundaries are different for each realisation because the profiles change with time. (b) The modified laminar flamelet (dash-circle) and LEM (solid) PDFs of similar mean and variance are shown.

3.5.1 Convergence

The number of LEM profiles required to construct a statistically converged PDF solution appears to depend on two criteria; the first of these is the turbulent Reynolds number governing the stochastic advection. A flow of higher turbulent intensity requires additional LEM profiles for a smooth probability distribution. The second criterion is the desired resolution in the progress variable (c) space. The number of LEM profiles needed to construct a converged PDF will increase with the resolution; typically 5000 realisations is sufficient for $Re_t < 200$ combined with 0.02 resolution in c -space. This behaviour is illustrated by Figure 3.7. It can be seen that the PDFs constructed from 2000 and 5000 profiles are nearly identical at $Re_t = 50$.

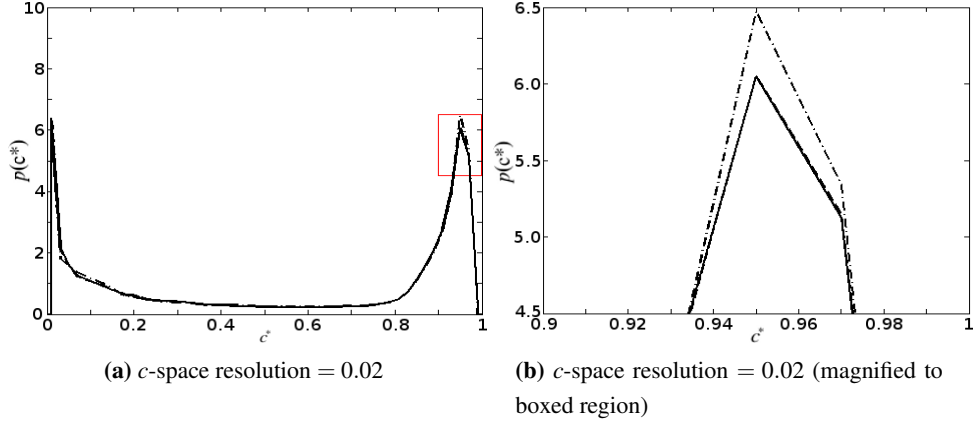


Figure 3.7: LEM generated PDFs at $Re_t = 50$. Each line represents a different number of LEM profiles used for the construction of the PDF. Solid: 5000 profiles; dash: 2000 profiles; dash-dot: 1000 profiles.

The computational time required to fully construct a LEM PDF model for premixed combustion is directly proportional to the number of LEM profiles needed for a converged solution. A typical quad-core computer can produce on the order of thousands of flame profiles per 24-hour period, depending on the complexity of the chemistry and the large eddy turnover time of the system. However, the computational time can be reduced linearly with the number of processors as the LEM simulation on each processor runs independently. This allows for quick and straightforward implementation on large computing clusters. If sufficient processors are available, one can effectively amass the flame profiles on the order of hours. After collecting the required number of realisations, the PDF model can be built in a matter of minutes.

3.6 Scalar Dissipation Rate

For premixed combustion, the Scalar Dissipation Rate (SDR), χ_c , of a temperature-based reaction progress variable is defined as,

$$\chi_c(T, \phi) = \alpha_c(T, \phi) |\nabla c| |\nabla c|, \quad (3.15)$$

where $\alpha_c(T, \phi)$ and ∇c are the thermal diffusivity at the local temperature and equivalence ratio and the gradient of the progress variable, respectively. The scalar dissipation rate is an important quantity to both non-premixed and premixed combustion modeling [45, 181]. It is an unclosed term that appears in the transport equation for the variance of progress variable; χ_c directly measures the decaying speed of fluctuations through turbulent micromixing [181]. Since the burn rate of many combustion processes depends on the contact area and local gradient between reactants, it is inevitable that the mean burning rate of the flame either explicitly and implicitly depends on the scalar dissipation in most combustion models. For example, Conditional Moment Closure (CMC) uses scalar dissipation rate conditioned by the progress variable to calculate the micromixing term [181]. Coincidentally, modeling $\chi_c|c^*$ emerges as one of the main difficulties in applying CMC to turbulent premixed combustion problems [4].

A relationship can be established between the conditional scalar dissipation rate ($\chi_c|c^*$) and the PDF of the reaction progress variable ($P(c^*)$) [181],

$$\chi_c = \int_0^1 (\chi_c|c^*) \cdot P(c^*) dc^*. \quad (3.16)$$

Equation 3.16 is particularly interesting to the analysis because LEM realisations could be used to construct both $\chi_c|c^*$ and $P(c^*)$. The methods applied to obtain the PDF model has been described in the previous section; the conditional scalar dissipation rates can be retrieved similarly. First, the thermal diffusivity at the local temperature and equivalence ratio, $\alpha_c(T, \phi)$ is generated from laminar flame simulations by Cantera [58] and stored prior to calculations. Second, the progress variable gradients (∇c) at the corresponding progress variables (c) are obtained from LEM flame realisations. This is then combined with the stored values of α_c according to Equation 3.15 at each discrete interval of c to form the conditional scalar dissipation rate. Third, the temperature gradients are arranged into 50 bins of size 0.02, varying between 0 to 1 in c -space. The values within each bin are averaged and additional LEM realisations are introduced until statistical convergence is achieved. It is worth noting that, for a given chemistry and set of LEM parameters, there would be one pseudo-invariant conditional average of scalar dissipation rate predicted. Practically, it appears that the addition of an entry for the uncon-

ditional mean of scalar dissipation in the PDF look-up table that would contain that pseudo-invariant conditional scalar dissipation rate already combined with the companion PDF using Equation 3.16 would be appropriate; this would be done for every combination of mean and variance of progress variable stored in the table for easy look-up when needing the closure for this term in the progress variable variance transport equation.

3.7 Results and Discussion

3.7.1 Probability Density Functions

LEM – Experiment Comparison

The premixed methane-air experiment was performed by a multi-slot burner developed at the University of Cambridge. The equivalence ratio (ϕ), inlet temperature and pressure were respectively 0.73, 294 K and one atmosphere. The flame structure, temperature and relevant species concentrations were measured via a combination of line imaging of spontaneous Raman scattering, Rayleigh scattering, two-photon laser-induced fluorescence (LIF) of CO and crossed planar imaging of OH-LIF. Further information on the Cambridge slot burner and multiscale laser diagnostics used for the experiment could be found in [8].

PDF models are built with LEM using the temperature-based reaction progress variable defined by Equation 3.2 at three preconditioned means. Experimental, modified laminar flamelet and LEM PDFs are shown in Figure 3.8. For a preconditioned mean of $\bar{c} = 0.05$, it is evident from Figure 3.8(a) that the LEM formulated PDF captures both the position and magnitude of the experimental result in a superior manner than the modified laminar flamelet PDF towards the left boundary. The modified laminar flamelet PDF well under-predicted the distribution for $0.02 < c^* < 0.2$ and over-predicted the result close to $c^* = 0$. However, it is able to capture the slight increase in probability near $c^* = 0.9$ whereas LEM cannot. The small compromise in the LEM PDF within this vicinity is compensated by the significant improvement for small c values.

The more curious observation is shown by Figure 3.8(b), where both the mod-

ified laminar and LEM PDFs differ from the experimental result. This could be caused by at least two different possible sources of error. First, the underlying reaction properties from the reduced mechanism may be incorrect; it may not be able to fully capture the reaction rates, thus causing the temperature gradients which govern the PDF to be inaccurate. Second, the temperature gradients for the experiment are measured 25-30 mm *downstream* of the slot burner [168]. This could potentially skew the distribution such that low c values are inadequately represented comparing to high c values. Typically one would expect to find slight to moderate peaks towards the boundaries of the PDF (for $\bar{c} \approx 0.5$) because reactants and products tend to dominate the total mass fraction (excluding inert species) at any given time. This means there should be a higher probability to find the reaction either near the beginning or the finish. The modified laminar flamelet and LEM PDFs both predict this result, opposite of the experimental data. One could speculate measurements at an axial location closer to the slot exit may reveal a different PDF even at the same preconditioned mean. Nevertheless, it can be seen that the LEM PDF achieves better overall resemblance to the experimental PDF than the previous formulation.

3.7 Results and Discussion




Figure 3.8: Experimental PDFs conditioned at three different means of the reaction progress variable. Solid: experiment [167]; dot: modified laminar flamelet; dash: LEM. The vertical axis represents the probability at state c^* conditioned by \bar{c} , $P(c)|_{\bar{c}}$.

For a preconditioned mean of $\bar{c} = 0.95$, it can be seen from Figure 3.8(c) that LEM closely captures the experimental distribution in peak magnitude and overall shape. The modified laminar flamelet PDF tends to mildly over-predict the probabilities for $0.7 < c^* < 0.9$ and is not able to acquire either the peak location or the magnitude for $c^* > 0.9$. Considering the three preconditioned means, it can be concluded that LEM PDFs tend to be more comparable to the experimental results than the modified laminar flamelet PDFs.

LEM – DNS Comparison

The product-based PDF model constructed with LEM is superimposed with the DNS and modified laminar flamelet PDFs in Figure 3.9. It can be seen that the LEM formulated PDF is able to capture the gradual decay of probabilities near $0.85 < c^* < 0.95$ exhibited by the DNS result; the peak values near complete combustion are more closely approximated than the modified laminar flamelet PDF. LEM is able to achieve this smoothing effect because of the stochastic nature of the model. Each realisation is slightly different; when averaged, they lead to a more realistic distribution of flame speeds and temperature gradients. Midrange values of c^* reveal excellent agreement for all three PDF models. However, it is evident that LEM under-predicts the probability for $0.02 < c^* < 0.2$. This could be an indication of over-mixing in the preheat layer, leading to a higher rate of product formation. Such a process inadvertently increases the gradient of the progress variable, thus lowering the probability of finding it at those states. The colder region of the preheat layer is particularly susceptible to this problem because the flame generally does not have enough time to recover before successive *triple maps*. In general, it can be seen that LEM can acquire the positions, distributions and peak values of the DNS PDF in a superior manner than the modified laminar flamelet method.




Figure 3.9: PDF at $x/L_d = 0.391$ of the DNS simulation domain ($L_d = 6$ mm). Solid: DNS; dash: modified laminar flamelet; dash dot: LEM. The DNS and modified laminar flamelet PDFs are obtained from [70].

3.7.2 Scalar Dissipation Rate

Each scalar dissipation result is constructed using 10,000 LEM realisations. Figure 3.10 illustrates the scalar dissipation rate conditioned in c -space from experimental data [167] and LEM at two turbulent Reynolds numbers.




Figure 3.10: Scalar dissipation conditioned in c -space. Solid: experiment [167]; dot: laminar flame; dash: LEM at $Re_t = 38$; dash-dot: LEM at $Re_t = 200$.

It is apparent that the model is able to reflect the decrease in peak amplitudes of the distributions relative to the laminar flamelet result. This amplitude decrease is likely attributed to thickening of the flame, also seen in previous DNS studies [147]. Another scalar dissipation rate generated from a higher turbulent Reynolds number is also investigated; the tendency for the peak temperature gradients to decrease with Re_t coincides with current understanding of flame broadening as turbulent intensity increases. This investigation also reveals that LEM may not be sufficiently mixing the flow field as $Re_t = 200$ still does not appear to lower the temperature gradients to the level of the experimental result.

Moreover, it may indicate the reduced mechanism is reacting too rapidly, causing an increased temperature gradient. Though the peak amplitudes are generally higher than the experimental values for $0.5 < c^* < 0.8$, it serves to validate the general physics of turbulent flames is appropriately represented by the stochastic advection introduced by LEM. The model has demonstrated the ability to acquire the qualitative behaviour of the temperature gradients and scalar dissipation rates

satisfactorily even with a six-step, reduced mechanism.

3.8 Concluding Remarks

One objective of this chapter has been to examine an alternate PDF model for premixed combustion which uses the LEM paradigm to mimic the effect of turbulent mixing on the PDF. LEM has the advantages of low computational cost and minimal empiricism. PDFs generated by LEM have demonstrated the ability to qualitatively match DNS and experimental results with greater success than previously suggested PDF models. LEM has also been used to predict the decrease in conditional scalar dissipation rate caused by flame thickening due to turbulent effects in premixed combustion, albeit over-predicting the experimental results in peak magnitudes. The stochastic nature of LEM enables turbulent reactions in a simple geometry to be captured inexpensively and qualitatively. It can be concluded that LEM is suitable for data pre-processing of simulation strategies that require the use of PDFs and scalar dissipation rates.

Chapter 4

Direct Comparison of PDF and SDR between LEM Simulations and Experiments for Turbulent, Premixed Methane-Air Flames

The work in this chapter is an extension of the LEM study in Chapter 3. The range of turbulent conditions simulated by the LEM has been increased and the comparison with experimental results is more complete. In addition, a number of improvements to the software has been introduced, including the parallelisation of the code that constructs the probability density function (PDF) and scalar dissipation rate (SDR) models. As this chapter is an extension of the previous, some introductory material and acronyms are inevitably repeated with the thought of improving readability.

4.1 Introduction

Several practical models for turbulent premixed combustion rely on an accurate representation of the probability density function of a reaction progress variable, which is often parameterised by the mean and variance of that progress variable [39, 55, 70, 121, 143, 145]. These presumed PDF approaches are often implemented in

conjunction with tabulated chemical variables to achieve detailed chemistry calculations in turbulent combustion simulations. Such models have been developed for both the Reynolds-Averaged Navier-Stokes (RANS) and Large Eddy Simulation (LES) paradigms. Previous work has shown that the accuracy of these methods depends to a considerable extent on the accuracy of the function presumed for the PDF of the progress variable [70].

A number of different presumed PDF models have been previously investigated for premixed combustion. The often used β -PDF does recover the extreme properties expected of the true PDF, such as δ functions at the zero and unity extremes of reaction progress for maximal variance, and single δ functions at the mean for zero variance. However, it fails to reproduce the shape of the true PDF in more general cases [70]. The issue is related to the fact that the shape of the true PDF appears to be a function of how the chemical reaction rates vary as a function of the progress variable; hence, different chemical kinetics lead to different shapes of the PDF of progress variable. The form of the β -PDF — which is of course entirely independent of the chemical kinetics — can lead to significant inaccuracies. Most critically, there can be a biasing error, as the discrepancies tend to occur at the same values of the progress variable for any particular flame.

One of the primary concerns in the design of modern engines is the reduction of harmful pollutants; specifically, the current generation of numerical models must be able to predict both the thermodynamic properties of the reacting mixture and the formation of these minor species with sufficient accuracy to resolve the parts per million produced. A prominent example is the prediction of prompt flame NO_x via the Fenimore pathway [49]. The predicted NO_x values from this mechanism are strongly affected by the choice of the PDF model, as this pathway is sensitive to the predicted temperatures and flame profile in the reactive regions.

In an attempt to account for the effects of chemistry on the shape of the PDF, Bray *et al.* [22] proposed using a premixed laminar flame to model the functional dependence of the PDF on progress variable. The proposed probability dependence of the flame existing at any given state is inversely proportional to the magnitude of the gradient of the temperature. Their original formulation only provided coverage for flames with very high variance. Jin *et al.* [70] then proposed a modification to the Bray PDF that extends the original formulation to cover all possible mean and

variance combinations. This is accomplished by truncating the PDF shape function as needed to match the mean and variance parameters. It was found to significantly improve the fit of the PDF to that extracted from Direct Numerical Simulation (DNS) results. A shortcoming of this method is that, at the point of truncation, the model PDF has a sharp drop to zero, whilst the true PDF tends to be more rounded.

To address this issue, a one-dimensional turbulent method was proposed to take the place of the typical laminar flame calculation to tabulate pseudo-turbulent PDF models for RANS and LES closures. The Linear-Eddy Model (LEM), an inexpensive one-dimensional stochastic mixing model, has demonstrated the ability to capture important effects from the interaction between chemistry and turbulence on the PDF distributions sufficiently well [148, 157, 175]. This model provides us with a mechanism to investigate the general flame characteristics at very high turbulence intensities, much beyond the capability of current DNS strategies. In turn, it permits us to analyze the behavior of the PDF constructed at these highly turbulent states.

While the LEM has been implemented to investigate the shape of the PDF distributions [148, 157], such simulations were performed to analyze the PDF for specific flames. The current study is primarily interested in the tabulation of a PDF lookup table useful for subsequent RANS and LES flame computations (a pre-processing operation not unlike pre-calculating the β -function and storing that in a lookup table). There was an initial suggestion that one ought to try to match the turbulence statistics in the LEM calculations to those that one expects to find in the later turbulent flame calculation. Indeed, a possibility is that one might need to add a dimension to the lookup table (something like the local turbulent Reynolds number) to account for variations in local turbulence properties in the turbulent flame calculation and their effect on the shape of the PDF. A large part of the motivation for the LEM work presented here is to answer the following questions: Do the turbulence properties affect the shape of the PDF? If so, how? How important is it to match the LEM turbulence properties used in generating the PDF lookup table to those that will be found in the turbulent flow to be calculated later?

A related question can be asked about the local gradient of progress variable in a premixed turbulent flame, which is closely related to the SDR. For premixed

combustion, the SDR, χ_c , of a temperature-based reaction progress variable is,

$$\chi_c(T, \phi) = \alpha_c(T, \phi) |\nabla c| \cdot |\nabla c|, \quad (4.1)$$

where $\alpha_c(T, \phi)$ and ∇c are the thermal diffusivity at the local temperature and equivalence ratio and the gradient of the progress variable, respectively.

The SDR is an important quantity to both non-premixed and premixed combustion modeling [45, 181]. It is an unclosed term that appears in the transport equation for the variance of progress variable, where χ_c directly measures the decay rate of fluctuations through turbulent micromixing [181]. Since the burn rate of many combustion processes depends on the contact area and local gradient between the reactants, it is reasonable for most combustion models to assume that the mean burning rate of the flames either explicitly or implicitly depends on the scalar dissipation rate. For example, Conditional Moment Closure (CMC) uses the scalar dissipation rate conditioned by the progress variable to calculate micromixing [181]; not surprisingly, modeling the conditional scalar dissipation term $\chi_c|c^*$ conditioned on the local and global progress of reaction emerges as one of the main difficulties in applying CMC to turbulent premixed flames [4].

The same LEM method can be used to generate a model for the conditional scalar dissipation, where one simply conditionally averages the scalar dissipation in the LEM temperature profiles that are used to construct the PDF model. The unconditional mean SDR can be obtained by taking the inner product of the conditional SDR with the model PDF. This allows for the construction of pseudo-turbulent PDF and SDR models that are perfectly consistent with one another. While it is possible to obtain PDF and SDR distributions for premixed combustion from DNS, the associated cost is generally prohibitive for flows with relevant turbulent conditions [70]. More importantly, studies typically tend to focus on the analysis of the PDFs and SDRs at specific points within the domain. This leads to the problem that although DNS and experiments can provide valuable insight to the behavior of the PDF and SDR [168, 170, 193], they cannot provide usable input models for subsequent RANS and LES combustion calculations, which is the primary motivation behind the current work. The incorporation of turbulence characteristics in the PDF distributions could provide a solution to the deficiencies seen in a number of

current PDF models, such as the *ad hoc* β -pdf or the laminar approaches. Having more accurate PDF and SDR models would be beneficial to a number of RANS and LES strategies as closures for turbulent premixed combustion typically rely on some variant of the PDF or SDR model for the reaction progress variable as the input [45, 181].

Recent detailed measurements of species and temperature have been made by the Cambridge-Sandia swirl burner [169–171, 194]. This swirl burner was designed specifically to explore the influence of stratification on the flame. However, the very detailed nature of the scalar and velocity measurements have made the data set attractive as a target for premixed flame model validation as well [72, 137]. In particular, the comprehensive database allows conditioning on a number of different variables, including equivalence ratio (for the stratified flames), temperature (or progress of reaction) or any other suitable scalar.

In this chapter, the experimental dataset is used to obtain detailed PDFs of the progress of reaction. The temperature is used to characterise the extent of reaction, for direct comparison with the PDFs generated from LEM simulations for three different swirl (and turbulence) levels. The measured and computed variances are considered, as well as the detailed shape of the PDFs, in the comparison. In addition, the unconditional mean SDR values obtained from experiment and LEM simulations are also directly evaluated. In the following sections, the numerical approach will first be discussed, followed by a summary of the experiments and the data treatment used.

4.2 Numerical Conditions: Premixed Combustion

4.2.1 Linear-Eddy Model

The Linear-Eddy Model has been demonstrated to replicate the flow statistics for simple turbulent conditions with acceptable accuracy [77–81]. Given the one-dimensional nature of the model, the computational costs remain relatively low for most practical cases. Here, the LEM is used in a pre-processing manner for the tabulation of discrete PDF and SDR models, which can be implemented in subsequent RANS and LES applications.

The LEM can be divided into two modules. The deterministic component consists of the usual one-dimensional gas dynamics evolution equations, whereas the stochastic component consists of random *eddy events*. The turbulence concept of LEM postulates a random process that rearranges fluid elements along a line in order simulate the chaotic vortices that appear in turbulent fields. These one-dimensional vortices, known as *triplet maps*, generate discontinuous fluid motions, which lead to a random walk of fluid elements. Three quantities control the behavior of each eddy event: the frequency, the size and the position. Firstly, the eddy event frequency per unit length of the domain, λ_{LEM} , is governed by [79],

$$\lambda_{LEM} = \frac{54 \nu Re_t (l_0/\eta_K)^{5/3} - 1}{5 C_\lambda l_0^3 (1 - (\eta_K/l_0)^{4/3})}, \quad (4.2)$$

where ν , Re_t , l_0 and η_K are the kinematic viscosity, turbulent Reynolds number, the integral length scale and the Kolmogorov scale, respectively. The empirical parameter C_λ should typically be tuned to the flame in question for LEM studies [148]; however, as the primary interest here is to investigate the changes in the PDF models associated with variations in the turbulent fluctuations and integral length scales, C_λ is held constant for consistency between test cases. The value of 15.0 for C_λ is adopted from [148]. Similarly, the parameter used to scale the Kolmogorov length, N_η , which typically requires tuning to case specific conditions, is also held at a constant value of 1.0 for all of the LEM cases. This parameter is applied to the inertial scaling law in the following form,

$$\eta_K = N_\eta l_0 Re_t^{-3/4}. \quad (4.3)$$

Secondly, the size of the turbulent eddies are sampled from a PDF¹,

$$f(l) = \frac{5}{3} \frac{l^{-8/3}}{\eta_K^{-5/3} - l_0^{-5/3}}, \quad (4.4)$$

that varies within the interval of $\eta_K < l < l_0$, derived from inertial scaling [79]. Thirdly, the position of the eddies are determined from a random uniform distribu-

¹ For clarity, it should be explicitly mentioned that this PDF of eddy sizes is different than the PDF of the reaction progress variable.

tion within the domain.

The deterministic and stochastic modules are implemented simultaneously during the simulation to achieve a pseudo-turbulent effect; this coupling between the stochastic advective process and deterministic evolution equations in a one-dimensional computational domain permits the LEM to simulate turbulent flows with higher Reynolds numbers than multi-dimensional models. The current LEM variant follows the formulation of [120] to accommodate premixed turbulent reacting flows.

To a large extent, the applicability of LEM depends on the homogeneity of the turbulence and the degree of which the problem can be characterised in one dimension within the simulation domain. This suggests that the method may fail where there is significant coupling between turbulence non-homogeneity at the small scales and the flow field. For many problems in premixed combustion, the spatial scales where reactions take place tend to be of the order of the laminar flame thickness, and in many cases, experiments show that at such small scales of turbulence, its characteristics have decayed to conditions that are reasonably isotropic and uniform.

4.2.2 LEM Simulation Methods

A global, six-step mechanism by Chang *et al.* [29] designed to simulate premixed methane-air combustion is implemented to generate the flame profiles. The six global reaction rates have been reduced by approximately 8% to achieve the correct unstrained laminar flame propagation speed observed in the reference solution from Cantera's GRI MECH 3.0 [156] calculation. This extra calibration procedure ensures increased accuracy for this value of equivalence ratio of 0.73².

All of the species properties are calculated using CHEMKIN-II [74], including specific heats, diffusion coefficients, thermal conductivities and enthalpies. The thermodynamic coefficients are based on the CHEMKIN Thermodynamic Data Base [75]. Figure 4.1 illustrates the flame solutions from Cantera and the Chang

²Reduced mechanisms are typically designed to operate over a range of thermodynamic conditions. However, due to a more restricted degree of freedom from the reduced number of steps, the level of accuracy may not be identical over the range of operations. For this numerical study, the primary interest is with premixed flames at $\phi = 0.73$; thus, the reduced mechanism is optimised for this particular equivalence ratio.

4.2 Numerical Conditions: Premixed Combustion

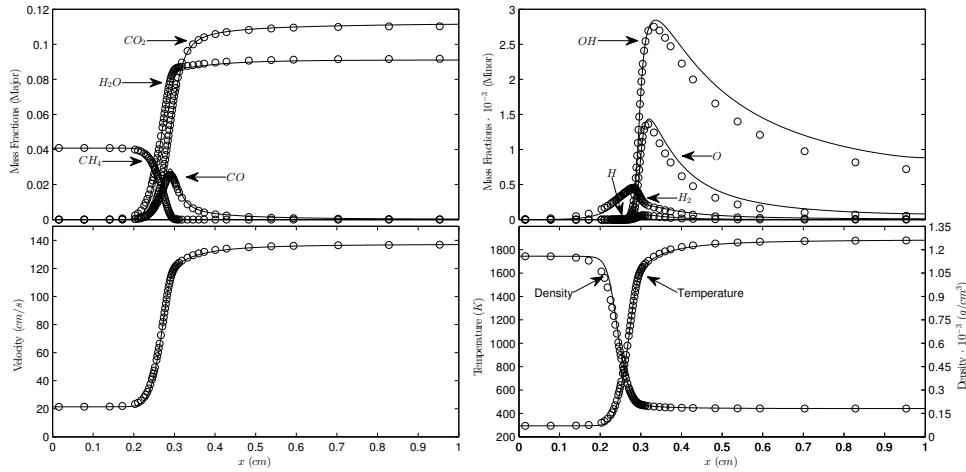


Figure 4.1: Premixed laminar flame solution at an equivalence ratio of 0.73 for the six-step global mechanism (line) and Cantera [58] (symbols).

mechanism. Mixture-averaged transport is adapted to reduce computational time. The inflow mixture was set to atmospheric pressure and 294 K at an equivalence ratio of 0.73. The calculated laminar flame speed, s_L , under these conditions is 0.214 m/s, and the calculated laminar flame thickness, δ_f , is 588 μm . The overall flame thickness, propagation speed, equilibrium temperature and mass fractions of major and minor species are sufficiently well-matched between the reduced Chang and full GRI mechanisms.

The LEM simulations are performed with a mixed first order upwind and second order centered spatial scheme. Explicit time steps are taken in order to cope with the stochastic nature of the model. The instantaneous temperature profiles are recorded at regular time intervals on the order of 1% of a large eddy turnover time for the freely propagating flames. Some 3,500 to 16,000 temperature profiles are stored for post-simulation construction of the LEM PDF models. The exact number of profiles required to create a statistically converged PDF model is dependent primarily on the turbulence intensity of the flame. Ten prototype turbulent premixed flames with various integral length scales and turbulent Reynolds numbers are tested. The flow parameters are varied from case to case such that the prototype flames adhere to predetermined locations on the Borghi diagram, as illustrated in




Figure 4.2: Borghi diagram showing locations of the ten prototype LEM flames (1a to 4b). The LEM test cases are represented by 'o'. The experimental flames are represented by triangles: SwB1 ('◁'), SwB2 ('▽'), and SwB3 ('▷'). It has been found that increasing the integral length beyond an order of magnitude above the laminar flame thickness while holding the turbulent fluctuations constant does not significantly alter the PDF profiles for the LEM simulations.

Figure 4.2. The LEM parameters used to simulate the freely propagating flames at the prescribed turbulent conditions are summarised in Table 4.1.

Throughout the chapter, a temperature-based reaction progress variable is used:

$$c = \frac{T - T_u}{T_{eq}(\phi) - T_u}, \quad (4.5)$$

where $T_{eq}(\phi)$ is the equilibrium temperature at the equivalence ratio ϕ and T_u is the unburnt gas temperature. From the simulations or experiments, it is possible to

4.2 Numerical Conditions: Premixed Combustion

Case	Re_t	l_0 (cm)	η_K (cm)	u'/s_L	Domain (cm)	Min. cells/cm	Cells used
1a	5	0.02992	$8.95 \cdot 10^{-3}$	9.826	1.0	670	1,000
1b	50	0.2992	$1.59 \cdot 10^{-2}$	9.826	1.0	337	1,000
1c	500	2.992	$2.83 \cdot 10^{-2}$	9.826	4.0	318	1,272
2a	50	0.02992	$1.59 \cdot 10^{-3}$	98.26	1.0	3,771	3,772
2b	500	0.2992	$2.83 \cdot 10^{-3}$	98.26	1.0	2,121	2,122
2c	5000	2.992	$5.03 \cdot 10^{-3}$	98.26	4.0	1,193	4,772
3a	500	0.02992	$2.83 \cdot 10^{-4}$	982.6	1.0	21,204	21,204
3b	5000	0.2992	$5.03 \cdot 10^{-4}$	982.6	1.0	11,924	11,924
3c	50000	2.992	$8.95 \cdot 10^{-4}$	982.6	4.0	6,706	26,824
4b	50000	0.2992	$8.95 \cdot 10^{-5}$	9826	1.0	67,053	67,054

Table 4.1: Relevant LEM simulation parameters: l_0 and η_K are the integral and Kolmogorov scales. Other constant parameters are invariant between the cases, including δ_f , s_L , C_λ and N_η , which are respectively, $588 \mu\text{m}$, 0.214 m/s , 15.0 and 1.0 . A minimum of 1,000 computational cells are used for each simulation.

construct a PDF with the following form,

$$P(c^*; x, t) \approx P(c^*; \bar{c}, \overline{c'^2}), \quad (4.6)$$

where \bar{c} and $\overline{c'^2}$ are the mean and variance of the progress variable, and c^* is the discretized variable representing the continuous c -space. In another words, the PDF at any space and time within the simulation domain can be approximated by the mean and variance of the distribution. Such a formulation is compatible with several current models for the turbulence and chemistry interactions for turbulent premixed combustion [39, 55, 70, 121]. A detailed description of the LEM PDF formulation can be found in [175], while a brief overview of the construction of the LEM PDF is provided in Section 4.4. Furthermore, a normalised variance is defined, $\overline{c'^2}_n$, determined by the limiting variance of a perfectly segregated mixture of hot and cold gases as:

$$\overline{c'^2}_n = \frac{\overline{c'^2}}{\bar{c}(1 - \bar{c})}. \quad (4.7)$$

This formulation is more suitable for the discrete tabulation of the PDF table, as the boundary values conveniently vary between zero and unity.

4.3 Experimental Conditions: Stratified Swirl Burner

The flames and experimental techniques are summarised briefly in this section³ as they have been described in previous papers [170–172]. A turbulent flame is stabilised on a bluff body, with reactants fed through two concentric streams at specified inner and outer equivalence ratios (ϕ). In the present chapter, only pre-mixed cases are considered. Swirl can be added to the flame by splitting the outer stream through a tangential inlet.

The global equivalence ratio for the flame is nominally 0.75 (measured via species is 0.73 +/- 0.018), and the range of local equivalence ratios spans 0.375–1.125. Scalar data obtained from Rayleigh/Raman/CO-LIF line measurements at 103 μm resolution allows the behaviour of key combustion species — CH_4 , CO_2 , CO , H_2 , H_2O and O_2 — to be probed within the instantaneous flame front. Simultaneous cross-planar OH-PLIF is used to determine the orientation of the instantaneous flame normal in the scalar measurement window, allowing real gradients of the temperature and hence the progress variable to be obtained.

The operating conditions considered in the present analysis are listed in Table 4.2. Measurements were taken at six heights along the centerline of the flame, by collecting data in 6 mm linear segments. The notation SwBN_z is used to denote cases at different locations, where N is the case number and z is the axial coordinate. Three premixed flames under different swirl conditions are considered.

Case	SFR ⁽¹⁾	SN ⁽²⁾	u'/s_L ⁽³⁾
SwB1	0	0	10.2
SwB2	0.25	0.26	12.2
SwB3	0.33	0.45	14.3

Table 4.2: Operating Conditions: (1) SFR = ratio of split flow to swirlers to total flow, (2) SN = measured swirl number, ratio of tangential to axial momentum, (3) Maximum total u'/s_L at the midpoint of the flame brush at $z = 30$ mm [71, 194].

The bulk velocity in the outer annulus, $U_o = 18.7$ m/s, is set at more than

³It is important to explicitly mention that the experimental expertise and analysis are provided by Mr. M. Kamal and Professor S. Hochgreb of the University of Cambridge. They are the co-authors of the work documented in this chapter.

4.3 Experimental Conditions: Stratified Swirl Burner

twice the value of the velocity in the inner annulus, $U_i = 8.3$ m/s, so as to generate substantial levels of shear, and thus turbulence, between the two flows. Co-flow air was supplied around the outer annulus with a bulk velocity $U_{co} = 0.4$ m/s to provide well-characterised boundary conditions. The Reynolds numbers derived from the bulk velocities at the exit geometry are $Re_i = 5,960$ for the inner flow and $Re_o = 11,500$ for the outer flow.


Multi-scalar laser diagnostics were applied at the Turbulent Combustion Laboratory in Sandia National Laboratories, and extensively described in [170–172]. The diagnostics setup allows for the line measurement of temperature (Rayleigh scattering) and major species (Raman scattering and CO-LIF) at $103 \mu\text{m}$ projected pixel resolution with simultaneous cross planar OH-PLIF at $48 \mu\text{m}$ projected pixel resolution. Signal to noise ratios are of the order of 150 for temperatures, and about 60 for equivalence ratio, with estimated accuracies of 2% and 5%, respectively. As the optical resolution of the Raman-Rayleigh-LIF measurements is smaller than the spatial sampling rate, the resolution of the temperature and major species measurements is limited by the sampling resolution ($103 \mu\text{m}$) and the laser beam diameter (0.22 mm , $1/e^2$). The optical resolution ($1/e^2$) of the OH-PLIF measurement is between 98 and $144 \mu\text{m}$ and therefore the spatial resolution of the OH-PLIF measurements is limited by the optical resolution rather than the sampling resolution. The OH-PLIF cross planar technique allows the flame normal in 3D to be assessed relatively to the measurement line, thus allowing the real gradients to be obtained by projection. Radial profiles were obtained by moving the burner horizontally in 4 mm steps, producing overlapping steps in the relative position of the 6 mm wide measurement window, with 300 laser shots taken at each step. Radial profiles were taken at axial increments of 10 mm above the burner exit to capture changes in flame structure with axial distance. Substantially larger datasets ($5,000$ - $30,000$ shots) were taken at the intersection of the mean flame brush and the mixing layer for stratified flames only. Here, only the premixed cases confined to within 2.5% of the nominal value of $\phi = 0.75$ are considered.

4.4 Construction of PDF Models


4.4.1 LEM Flame Profiles

Figure 4.3 illustrates a few selected LEM temperature profiles for two prototype flames of varying turbulence intensities. It can be seen that the influence of the triplet maps becomes more pronounced and modify a larger portion of the flame as the turbulent Reynolds number increases. Consequently, the turbulent temperature profiles become more distinguishable from the laminar counterpart, with larger effective flame thicknesses. These perturbations propagate through the flame and reduce in strength until they either blend into the flame profile or get transported out of the domain.

4.4 Construction of PDF Models



(a) Case 2b ($Re_t = 500$).



(b) Case 4b ($Re_t = 50,000$).

Figure 4.3: Characteristic LEM temperature profiles of two prototype flames at different Re_t . The individual profiles on each graph are separated by at least one large eddy turnover time.

At first inspection, it appears that the average temperature gradients should increase with the number of triplet maps implemented per unit time; however, the diffusion mechanism from the one-dimensional evolution equations quickly diminishes the sharp discontinuities introduced in the slope of the temperature field. As

a consequence, the flame begins to broaden as it recovers from this perturbed state. The net effect is to decrease the overall conditional averages of the gradients when considered across the entire width of the flame. The temperature profiles depicted in Figure 4.3 can then be transformed into the progress variable space via Equation 4.5.

4.4.2 LEM PDF Construction

The method used to construct the LEM PDF models is first discussed in the context of the modified laminar flamelet PDF, which is then extended to the LEM formulation. The relevant parameters governing the behavior of the distributions are the mean (\bar{c}) and variance ($\overline{c'^2}$), as previously mentioned by Eq. 4.6. Truncations of the profiles in the spatial domain are applied to the one-dimensional flame profiles, leading to changes in both \bar{c} and $\overline{c'^2}$. The mean of the distribution increases for every point removed from the unburnt mixture ($c = 0$) boundary, whilst the opposite is true for points removed from the burnt mixture ($c = 1$) boundary. Moreover, the variance decreases for every point removed from either boundary.

Figure 4.4(a) demonstrates an example of these truncations. The symbols correspond to the truncation limits, (x_1, c_1) and (x_2, c_2) , where only the cells within the interval are retained. Arbitrary values of mean and variance are set to $\bar{c} = 0.50$ and $\overline{c'^2}_n = 0.39$ for illustrative purposes. The modified laminar flamelet PDFs with the prescribed means and variances can then be constructed with the truncated flame profiles according to methods described by previous work [70]. An example of the modified laminar flamelet PDF of $\bar{c} = 0.50$ and $\overline{c'^2}_n = 0.39$ is shown in Figure 4.4(b). Crucially, it can be deduced that there exists one unique modified laminar flamelet PDF for every mean and variance combination.

4.4 Construction of PDF Models

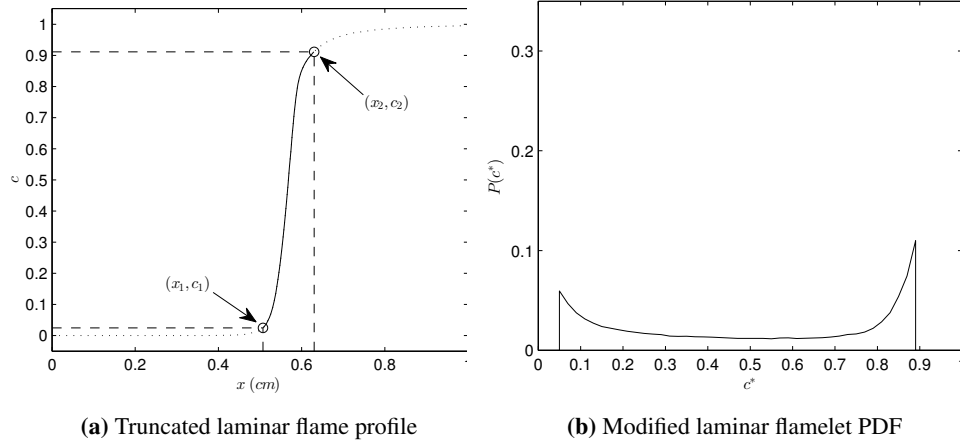


Figure 4.4: (a) (x_1, c_1) and (x_2, c_2) mark the truncation limits; only the cells within the interval are retained. (b) The modified laminar flamelet PDF with $\bar{c} = 0.50$ and $\overline{c'^2}_n = 0.39$ constructed using the truncated flame profile is shown.

The LEM PDFs are constructed in a similar manner compared to the modified laminar flamelet PDFs. Instead of one unique, steady, laminar temperature profile, the LEM generates transient temperature profiles; however, each LEM temperature profile is equally valid and must be equally weighted in an ensemble average to arrive at the final PDF model. Consequently, the truncation strategy prescribed for the modified laminar flamelet profiles to obtain the desired mean and variance is applied to each LEM temperature profile. Here, when a (x_1, x_2) pair is fixed in real space, averaging together the many LEM realizations leads to a particular mean and variance of progress variable. In principle then, one could vary the positions of x_1 and x_2 for each temperature profile until the desired mean and variance for the PDF is achieved. In practice, all possible discrete (x_1, x_2) combinations are used to populate the PDF lookup table, which is stored as a function of the mean and variance. During subsequent RANS and LES calculations, this table is called and interpolated as needed to obtain the PDF for any required mean and variance combination.

Figure 4.5(a) illustrates two truncated LEM profiles that correspond to PDFs with $\bar{c} = 0.5$ and $\overline{c'^2}_n = 0.33$, similar to the laminar flame shown in Figure 4.4(a). The truncation positions are shifted for individual temperature profiles because the

4.4 Construction of PDF Models

flame changes with time. The final PDF model is acquired by averaging a number of PDFs built from such truncated LEM flames. Figure 4.5(b) illustrates an example of the converged solution constructed by averaging 10,000 LEM PDFs at a relatively low turbulent Reynolds number. It is apparent the LEM PDF displays peaks of lower magnitude at each boundary than the modified laminar flamelet PDF model. This smoothing effect is a direct consequence of the application of triplet maps to the temperature profiles. The mapping introduces variations in the temperature gradients at a given temperature (Figure 4.3(b)); as a result, the probability of the flame existing at any given state varies with different temperature profiles. When averaged, the PDF decreases to zero at each boundary more gradually [175].

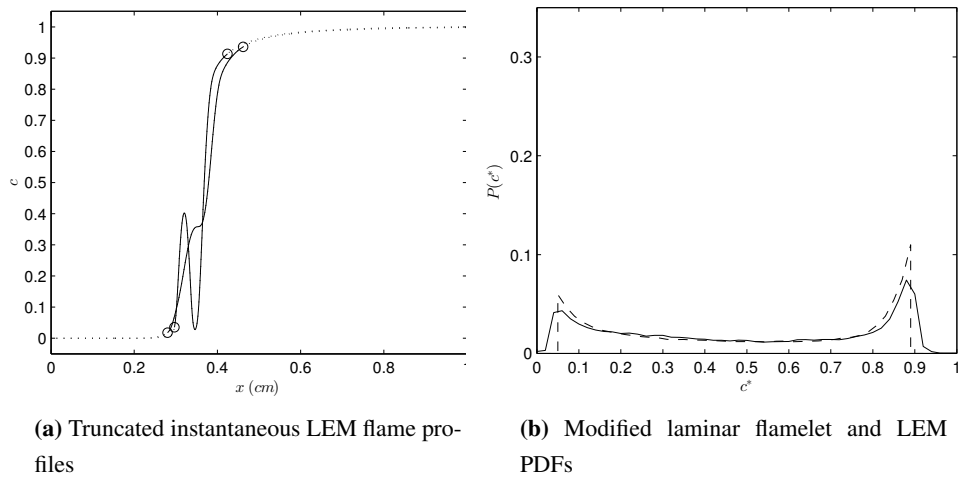


Figure 4.5: (a) The portions to be retained are within the intervals delimited by the circles. The four truncation positions are selected such that the resultant PDF would have $\bar{c} = 0.5$ and $\overline{c'^2} = 0.39$. The truncation boundaries are different for each temperature profile because of the transient effects. (b) The modified laminar flamelet (dash) and LEM (solid) PDFs of similar mean and variance are shown.

4.5 Results

4.5.1 Probability Density Function

The LEM PDF models spanning a range of distribution means and variances are illustrated by Figure 4.6. This study reveals that the characteristics of the PDF models for all ten prototype flames display remarkable similarities across the scope of test scenarios despite a three-order change of magnitude in turbulent fluctuation intensity and length scale on the Borghi diagram. The PDFs tend to become slightly less bivariate as the turbulence intensities increase. This effect is particularly noticeable for cases with means and normalised variances having values close to 0.5.

The PDF distributions captured at various axial positions for the three experimental flames (SwB1-3) operating at different swirling conditions are illustrated by Figures 4.7-4.9. The LEM PDFs tabulated as a function of c^* from one particular prototype flame has been superimposed for direct comparison between the model and the swirl burner data⁴. In general, a high degree of similarity can be observed between the experimental and LEM PDFs in overall features, magnitudes and peak positions. The 3×5 panels correspond to different distances from the stabilisation point of the flame and radial positions across the flame brush, which are characterised by various combinations of \bar{c} and $\overline{c'^2}_n$.

Starting with SwB1 (Figure 4.7), the model PDFs are well-matched for all normalised variance values below a distribution mean of approximately 0.7. The reactant edge of the flame brush can be found at low \bar{c} (left columns); and the PDF peaks at $c^* = 0$. This agrees with the calculated LEM PDFs, but the measured PDFs are somewhat wider than the calculations suggest. At the intermediate values of \bar{c} , two low peaks are found, one at zero and one around 0.8-0.9, which are well-captured by the LEM model, even at low variance levels. Closer to the product end of the flame brush (right columns), the PDF peak appears not at unity, but at slightly lower values, between 0.8 and 0.9, much like the experimental values. The peak at lower values than full reaction progress is attributed to the limited time for completion of reaction due to turbulent mixing of cold and hot gases. However,

⁴LEM prototype flame 1b is selected for this demonstration. However, it is evident from Figure 4.6 that the location on the Borghi diagram does not significantly influence the resulting LEM PDF for a given variance and mean combination.

4.5 Results

heat transfer to the base of the flame can also contribute to the lowering of the extent of reaction.




Figure 4.6: PDF models at various \overline{c} and $\overline{c_n^2}$ are shown; each row represents one value of \overline{c} and three values of $\overline{c_n^2}$ (ranging from left to right: 0.25, 0.49 and 0.75). Vertical and horizontal axis on each graph represent the probability and the progress variable, respectively. Solid: case 4b, dash: case 3b, dash dot: case 2b, dot: case 1b (notation is in accordance with Figure 4.2).

4.5 Results




Figure 4.7: PDFs measured from the SwB1 flame at various axial locations conditioned by different distribution means, \bar{c} (solid). The corresponding LEM PDFs of similar \bar{c} and $\overline{c'^2}_n$ are superimposed (dash).

Proceeding to the first swirl flame (Figure 4.8), it can be seen that the LEM and experimental PDFs share as many similarities as with the no swirl case. The model and experiment agree fairly well for most values of variances at a distribution mean of less than approximately 0.7. The experimental PDFs tend to be slightly less bivariate in comparison to the no swirl case, reflected by the decreasing distribution variance. An interesting difference can be observed for the results having $\bar{c} \approx 0.7$ and $\overline{c'^2}_n < 0.4$. The probability of c^* continues to be non-zero towards the 0 boundary for the experimental PDFs whereas the LEM PDFs tend to go to zero. This could suggest that the swirling flow may prolong the time during which the flame spends in the preheat layer, effectively increasing the probability of finding the flame at a low c^* state and decreasing the gradient of the progress variable around this region of the flame.

4.5 Results




Figure 4.8: PDFs measured from the SwB2 flame at various axial locations conditioned by different distribution means, \bar{c} (solid). The corresponding LEM PDFs of similar \bar{c} and \bar{c}_n^2 are superimposed (dash).

For the SwB3 flame (Figure 4.9), the LEM and experimental PDFs are typically well-matched for distributions with $0.3 < \bar{c} < 0.7$. For PDFs with low values of \bar{c} ($\bar{c} \approx 0.11$), it can be seen that the experimental PDF peaks have lower magnitudes and in some cases, are shifted towards a non-zero c^* value ($c^* = 0.05$). More so, the distributions appear to be slightly wider than the LEM counterparts. This suggests that increasing the swirl number of the flame may have greater impact on the colder regions of the flame, where the conditional mean is low — such a behavior is not apparent in the SwB1 and SwB2 cases. Overall, the PDFs for all cases, with and without swirl, are still captured for the corresponding combination of means and variances. However, the agreement is generally better for the intermediate values of \bar{c} than for the extremes.

4.5 Results




Figure 4.9: PDFs measured from the SwB3 flame at various axial locations, conditioned by different distribution means, \bar{c} (solid). The corresponding LEM PDFs of similar \bar{c} and $\overline{c'^2}_n$ are superimposed (dash).

4.5.2 Scalar Dissipation Rate

A relationship can be established between the conditional scalar dissipation rate ($\overline{\chi_c|c^*}$) and the PDF of the reaction progress variable ($P(c^*)$) [4],

$$\bar{\chi}_c = \int_0^1 (\overline{\chi_c|c^*}) P(c^*) dc^*, \quad (4.8)$$

where $\bar{\chi}_c$ is the unconditional mean scalar dissipation rate, or simply the mean scalar dissipation rate. Equation 4.8 is particularly interesting to the analysis because the LEM temperature profiles can be used to construct both $\overline{\chi_c|c^*}$ and $P(c^*)$. A more detailed description of the methodology used in constructing the PDF and SDR models can be found in Chapter 3.

There are two observable behaviors from the conditionally averaged SDR models from the LEM. First, as the characteristic flame properties are shifted upwards

in the Borghi diagram (increasing u'/s_L), the peak magnitudes of the dissipation rates tend to decrease. Second, if the characteristic flame properties are shifted rightward in the Borghi diagram (increasing l_0/δ_f), the peak magnitudes of the SDRs tend to increase towards the limit of a laminar flame. It is difficult to determine from the current simulation results whether the flame broadening (increasing u'/s_L) or the integral scale increment (increasing l_0/δ_f) effect is dominant within a specific combustion regime. It is clear, however, that the processes will compete with one another in general, leading to relatively unchanged conditionally averaged SDR distributions for the special case of isotropic, homogeneous flames.

For a given chemical mechanism and set of LEM parameters, there should be one pseudo-invariant conditional scalar dissipation rate (cSDR) averaged from the temperature profile datasets. This cSDR can then be multiplied with the companion PDF through an inner product to arrive at the unconditional mean SDR via Equation 4.8. To reduce simulation time during practical implementation for complex combustion models, the values of the unconditional mean SDR can be pre-tabulated at desired combinations of means and variances of the reaction progress variable for efficient retrievals. The same LEM datasets from the PDF study can be used to generate one conditionally averaged SDR model for each of the ten prototype flames. The results are presented in Figure 4.10.

Figure 4.11 illustrates the pseudo-invariant conditional SDR distributions from the experiments. The conditional SDR typically decreases with increasing swirl and appears to be most similar to the laminar limit at $z = 30$ mm. This axial location roughly coincides with the intersection of the mixing and shear layers, promoting the development of a highly turbulent and dissipative region, after which the SDR quickly diminish. The comparison with the LEM model is striking: in the model, there is little sensitivity to the local turbulence characteristics, whereas the experiments indicate that there are significant changes in the SDR values; the magnitude tends to be more sensitive to the effects of swirl. Nevertheless, the shape of the distributions are well-captured, with the maximum in the vicinity of $c^* = 0.65$ to 0.75 . The maximum generally moves rightward, towards $c^* = 0.75$, as the SDR magnitude decreases.

4.5 Results




Figure 4.10: Non-dimensionalized conditional average of the SDR models from the ten prototype flames. The left and right columns illustrate changes in the turbulent fluctuations and integral length scales, respectively. The turbulent Reynolds number of each case is recorded in the parentheses. The laminar case is displayed in light gray as a reference.

As mentioned in the Introduction, several models for turbulent combustion relate the local rate of micromixing to the conditional scalar dissipation. A common approach to modelling scalar dissipation [113, 164, 165] is to extract the functional dependence on the conditioning variable and model the conditional scalar dissipation as,

$$\overline{\chi|c^*} = \chi_0 \times f(c^*), \quad (4.9)$$

where χ_0 is then calculated by some other means and is not a function of the conditioning variable. As usual, the integration of $f(c^*)$ over the progress variable space results in the value of unity. Here, the same decomposition is applied to the results to see if the underlying functional dependence of scalar dissipation on the conditioning variable changes — that is, to test whether or not such a decomposition might be sensible.

4.5 Results

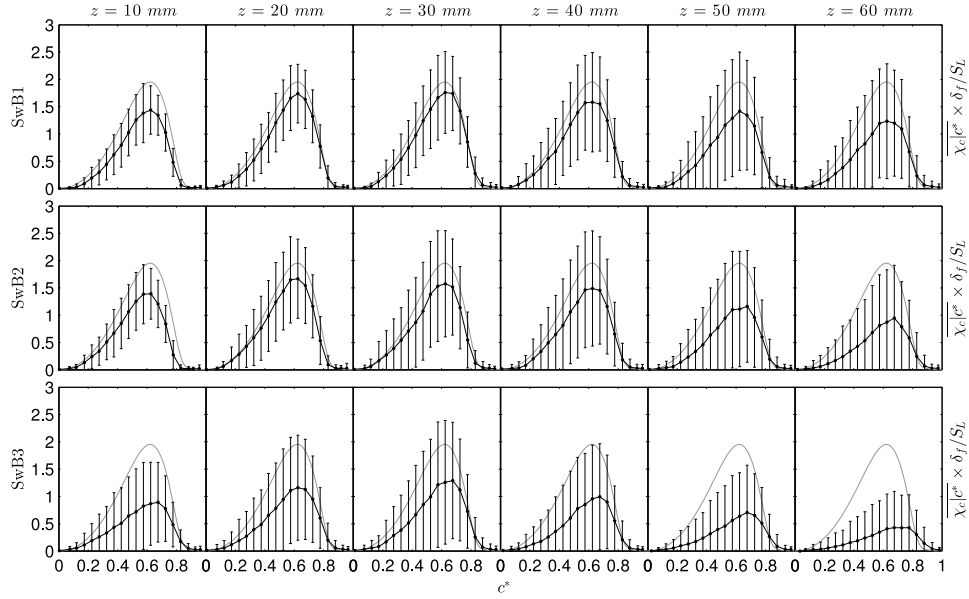


Figure 4.11: Non-dimensionalized conditional average of the SDR from the three experimental flames. Each column represents a unique axial position from the flame stabilisation point while each row represents one swirling condition. The error bars indicate \pm one standard deviation from the mean. The laminar case is displayed in light gray as a reference.

Figure 4.12 depicts the functional dependence of scalar dissipation ($f(c^*)$) at three axial positions from the flame stabilisation point. This figure confirms the observation that the cSDRs tend to shift rightward with increasing swirl intensity, particularly at the beginning and end of the flame brush, corresponding to the axial positions of 10 mm and 60 mm. The reason for such changes in the shape of the normalised cSDR is not immediately clear from the current results; rather than speculate on the possible causes, this question can be an opportunity for future work. In general, the LEM results tend to remain close to that of the laminar flame as the model is unable to emulate the effects of swirl. Furthermore, Figure 4.13 illustrates the values of the normalization constant, χ_0 , in relation to the axial position for the experimental flames. It is interesting that the maximum values of χ_0 for all of the experimental flames occur at the intersection of the mixing and shear layers near $z = 30$ mm.

4.5 Results




Figure 4.12: Functional dependence of the scalar dissipation ($f(c^*)$) at three axial positions downstream from the flame stabilisation point (10, 30 and 60 mm) with the LEM results superimposed. The experimental flames are represented by symbols: SwB1 ('+'), SwB2 ('□'), and SwB3 ('○'). The LEM result is represented by the solid line. The laminar case is displayed in light gray as a reference.




Figure 4.13: Values of χ_0 in relation to the distance downstream of the flame stabilisation point. The experimental flames are represented by symbols: SwB1 ('+'), SwB2 ('□'), and SwB3 ('○'). The LEM (black line) and laminar flame (gray line) results are not dependent on the axial position.

The inner product of the pseudo-invariant conditional SDR with the appropriate PDF distribution via Equation 4.8 provides the solution for the unconditional mean SDR, $\overline{\chi}_c$, or simply the mean SDR. The modelling of this term remains as one of the final challenges in applying some combustion models to premixed flame calculations [181]. Figure 4.14(a) illustrates the behavior of the mean SDR as a

function of the mean and variance of the progress variable as calculated by LEM. The mean SDR peaks around $\bar{c} = 0.5$ and $\overline{c'^2}_n \approx 0$ as the PDF and cSDR distributions have coinciding maxima for these values of \bar{c} and $\overline{c'^2}_n$. The experimental values can be seen in Figure 4.14(b), where the magnitudes of the mean SDR are evidently lower than the predicted values. The LEM predictions are best when the swirl intensity of the burner is at its lowest. This suggests that swirling flows induce a decrease in the magnitude of the SDR, both conditional and unconditional, which cannot be replicated by a one-dimensional turbulence model. Moreover, recalling the PDF results from the previous section, one can conclude that swirl plays a more prominent role in the determination of the SDR than the PDF distributions.

4.6 Discussion

4.6.1 Probability Density Function

The LEM formulated PDF models display minimal variations in distribution shape and magnitude with changing turbulent fluctuations and integral lengths. Typically, the PDF distributions become slightly less bivariate with increasing turbulent fluctuations. Towards the reactant boundary ($c^* = 0$), the turbulent fluctuations appear to increase the diffusive mixing in the preheat layer, leading to slightly faster flame development in this region. In turn, this causes a decrease in the duration of the flame residing at low values of c^* . Towards the products boundary ($c^* = 1$), the model suggests that the eddy interactions impede the flame from fully oxidizing towards the end of the reaction. This does not imply a reduction in the flame speed, but rather a barrier to reaching full chemical equilibrium. This effect can be observed in Figure 4.6, where the PDF models constructed from flames with higher turbulent fluctuations have the rightward boundaries shifted slightly towards the left. The combination of these phenomena causes the LEM PDF to deviate from the bivariate distribution seen in laminar PDF models as turbulence intensity increases, in agreement with the experimental results. The model of course neglects any potential heat transfer issues, which are present in the experiments, and may influence the approach to equilibrium at the base of the flame.

The swirl burner results (Figures 4.7-4.9) reveal that the shape of the PDFs





Figure 4.14: Non-dimensionalized unconditional mean scalar dissipation rates ($\bar{\chi}_c \times \delta_f/s_L$) as predicted by (a) LEM and (b) experiments with LEM solutions superimposed at various combinations of \bar{c} and \bar{c}_n^2 . Solid: LEM; dash: SwB1; dash-dot: SwB2; dot: SwB3.

is dependent on the distance from the base of the flame and secondarily on the amount of swirl in the flame. The dependence on distance from the base arises from the different turbulence characteristics across the flame brush as the flow evolves downstream. The variance increases with distance from the base, and with swirl, and this directly affects the corresponding PDFs. The experimental PDFs are perhaps surprisingly well-captured by the two input parameters of \bar{c} and $\overline{c'^2}_n$. Nevertheless, once the variance and mean are determined, there are only so many sensible ways for the PDF to be arranged within the LEM context, so perhaps this justifies the good agreement.

The data suggests that variations in the turbulent Reynolds number and integral lengths induce a rather modest effect on the shape of the PDF; it appears that it may not be necessary to add an additional dimension related to the turbulence intensity to the lookup table. Indeed, while it would be ideal to construct a table using representative conditions, using a model generated from a flame with lower turbulence intensities (which is the less computationally intensive option) may suffice. The number of temperature profiles required to construct converged LEM PDF models can be markedly reduced while the additional error contribution remains small. This is especially true for premixed flames with turbulent Reynolds numbers and integral lengths within an order of magnitude of the targeted flame on the Borghi diagram. Further discretion may be required when applying different PDF models for flames exposed to more intense swirling conditions. The experimental results reveal that the overall flame gradients and hence the flame surface density function reduce in magnitude under swirl. However, such changes may not greatly alter the shape of the PDF of the progress variable because of the normalization procedure required in the construction of the LEM PDF models.

4.6.2 Scalar Dissipation Rate

As previously mentioned, there are two observable behaviors from the conditionally averaged SDRs computed by the LEM. First, as the intensity of turbulence increases, the peak magnitudes of the mean dissipation rates tend to decrease. Such an effect has been previously observed in DNS studies and is attributed to flame broadening caused by the decrease in peak temperature gradients [147]. This ef-

fect is emulated by the LEM via the interactions between the turbulent eddies and the scalar fields. These small eddies penetrate into the reaction zone, increasing the local gradients of the scalar fields. This is followed by the dissipation of the sharp discontinuities introduced in the slope of the scalar fields due to the coupled diffusion mechanism. Effectively, the flame is broadened, increasing its thickness and decreasing the gradients. Second, as the turbulent length scale increases at constant turbulent fluctuation intensity, the peak magnitudes of the SDRs tend to increase towards the limit of a laminar flame⁵. Such an effect could be connected to the increasing Kolmogorov length associated with the integral scale increment. An increase in the Kolmogorov length means the smallest eddies will become too large to effectively change the local structure of the flame. In other words, the eddies have more difficulty penetrating the preheat layer and the reaction zone, thus reducing the interaction between the vortices and the flame. With a smaller number of interactions, the pseudo-turbulent SDR distributions will become more similar to the laminar SDR distribution.

It is important to emphasise that the current numerical results do not include the effect of strain. Additional modelling parameters would need to be introduced to the transport equations for such calculations. Previous studies have shown that regions under highly positive strain are correlated with regions of increased scalar dissipation, though this is not an exclusive association [90, 178]. For turbulent flames exposed to high strain rates, neglecting this effect may lead to non-negligible changes to both the PDF and SDR distributions.

The pseudo-invariant conditional SDR distributions and mean SDR values from the experimental flames show a more pronounced decrease in magnitude with increasing swirl effects. It appears that the intersection of the mixing and shear layers near $z = 30$ mm leads to quickly diminishing gradients of c for positions downstream of this location. This suggests that additional parameters which quantify the dissipative effects of mixing and swirling may be necessary in order to accurately model the conditional SDR distributions. Moreover, the entrainment of cold air into the premixed flame caused by the increase in fluid motion may

⁵ For unstrained flows, the maximum conditionally averaged SDR distribution corresponds to the laminar flame with the same chemistry; any turbulence interactions decrease the magnitude of the SDR distribution with respect to the laminar limit.

induce changes in the temperature based reaction progress variable that the LEM cannot currently simulate. It is, however, curious that the PDF distributions from both the LEM and experiments reveal far less susceptibility to changes when exposed to physical effects, such as entrainment of air and swirl, in comparison to the conditional SDR distributions.

Nevertheless, the quantity that requires closure in combustion models is the unconditional mean SDR, defined by Equation 4.8. In this regard, it appears that the LEM can reproduce values of $\bar{\chi}_c$ well within 50% of the experimental equivalent, but the accuracy depends primarily on the swirl number of the flow. For most points, the LEM over-estimates the value of $\bar{\chi}_c$, which is understandable considering the conditional SDR distributions from the model indicate higher magnitudes compared to the experiments. Further research could explore the possibility of including additional modelling parameters into the LEM formulated conditional SDR model according to physical effects governing the flow field. In this case, perhaps a post-processing correction can account for the entrainment of air induced by the swirl. This could amount to applying a simple correction factor to the magnitude of the cSDR, which is calculated as a function of the local anisotropy.

4.7 Conclusion

This chapter investigates the effects of variation in swirl and turbulence intensity on the probability distribution of the reaction progress variable for a series of globally lean, turbulent premixed flames, both experimentally and numerically.

The PDF models constructed from the LEM simulations indicate that increasing turbulence intensity has a rather modest impact on the distributions. The minor changes are only observed at medium values (approximately 0.5) of means and normalised variances. The PDFs from the experimental flames show a bivariate distribution for all cases, with peaks in the unburned and burned regions. The thermal flame thickness is generally larger than the unstrained (and strained) laminar flame thickness. The data therefore suggests that the thermal gradient is smeared by turbulent diffusion at scales on the order of the flame thickness. The LEM PDF model demonstrates good agreement with the experimental results in terms of nominal changes with turbulence. Particularly, it is able to capture the overall

4.7 Conclusion

shape and the effect of smoothing towards the left and right boundaries of the PDF distributions.

The LEM results suggest that the conditional SDR would decrease towards the upper regions of the Borghi diagram and would increase in peak magnitudes towards the unstrained laminar limit with increasing integral scales. The experimental conditional SDR for all axial locations of the flame brush shows a greater decrease in magnitude than predicted by the numerical model, though this could be partially caused by an interaction between shearing and mixing layers within a specific region of the swirl flames, leading to greater dissipation downstream of this location. The values for the unconditional mean SDR predicted by LEM are typically within 50% of the experimental values. Moreover, the accuracy of the numerical model seems to be dependent on the swirl number of the flow.

Overall, the results point towards the idea that the PDF for premixed flames remains relatively steady under a variety of turbulent conditions, including changes in the integral length, velocity fluctuation and swirl number. As a consequence, it appears to be practical to use a representative pseudo-turbulent PDF model for a range of turbulent conditions. Future studies could focus on understanding the physical interpretation behind the relative invariance in the PDF model with variations in both the turbulent fluctuations and integral length scales. Another possible direction is to determine if this invariance also extends to other fuels in the premixed combustion regime. Swirl appears to have greater influence on the local flame gradients; this, in turn, leads to more substantial changes in the SDR, which the model presented here cannot capture in its current form. Perhaps a practical solution in the near future would be to implement a correction factor to the LEM formulated SDR models based on the local anisotropy of the flow field.

Chapter 5

Conditional Source-term Estimation: Parallel Iterative Solution with Dynamic Ensembles

5.1 Introduction

Turbulent combustion is an important part of current heat engine designs; examples of these devices include the common automobile engine and gas turbines. A fundamental challenge today is to reduce the pollutant formation of such engines as emission regulations become increasingly stringent. However, the complex interactions between the turbulent flow, chemical kinetics and heat transfer at engine-relevant pressures and temperatures establish a highly non-linear problem spanning orders of magnitude in both length and time scales. The difficulty here is to introduce improvements to these highly calibrated heat engines under such dynamic conditions.

Computational fluid dynamics (CFD) is indispensable in the development of complex engines due to its low cost and time requirement compared to experiments. While direct numerical simulation (DNS) is the most accurate approach

to understanding the interactions between turbulence and flame, performing DNS for engine relevant problems will remain impractical for the foreseeable future. With increasing computational capabilities, numerical simulation of these devices via the large eddy simulation (LES) paradigm is becoming more accessible to researchers. Nevertheless, due to the strong coupling between turbulence and chemistry in premixed flames, the prediction of chemical reaction source terms continue to be a modelling challenge.

A number of simulation strategies has been developed to model the interaction between turbulence and flame [84, 86, 132, 134]. For Conditional Moment Closure (CMC) [84], the transport equations for the reactive species are conditionally averaged on a conserved scalar. The primary advantages of this strategy are the reduced spatial dependence of the conditional averages and negligible fluctuations around the conditional averages, which considerably simplify the reaction rate closure. In addition, detailed chemistry can be included at relatively low computational cost without the thin flame restriction encountered in flamelet models. However, one extra transport equation for the conditional mass fraction would need to be solved for *each* reactive species, which increases the dimensionality of the problem. Furthermore, there are unclosed terms associated with these CMC transport equations that require modelling. Bushe and Steiner [26] later proposed a technique based on the CMC known as the Conditional Source-term Estimation (CSE). Rather than solving extra transport equations, the conditional mass fractions are recovered through an inverse calculation. A more detailed description of CSE is provided in Section 5.2. This model has been shown to produce promising results for diffusion flames [67, 93, 161, 183]. Further, Salehi *et al.* has demonstrated that CSE can also be applied to premixed flames and produce meaningful results [42, 143–145, 153]. Recent effort has extended CSE to lifted turbulent flames [43] and flames operating in the Moderate and Intense Low Oxygen Dilution (MILD) combustion mode [92].

In this chapter, two optimisations to enhance the capability of CSE have been introduced. The first optimisation is to increase the flexibility of the CSE model by dynamically allocating the physical space required for the individual ensemble averages. The second optimisation is to integrate a well established iterative solver, Least-Squares with QR-factorisation (LSQR) [123], to tackle the inverse problem. These optimisations can be applied to both Reynolds-averaged Navier-

Stokes (RANS) and LES paradigms in the premixed, partially-premixed and non-premixed combustion modes. A comparison between the existing CSE methodologies and the new optimisations is made using LES calculations with OpenFOAM 2.1.1 [185] simulating the Gülder burner [190].

5.2 Conditional Source-term Estimation

Combustion models are required for LES or RANS simulations, which, by their nature, cannot fully resolve the turbulence-chemistry interactions spatially and temporally. First moment closure, whereby the filtered or averaged reaction rates are directly computed by the filtered or averaged values of the scalar fields via the Arrhenius rate equations, can lead to considerable discrepancies compared to the fully resolved values, due to the highly non-linear interactions between chemical source terms and the fluctuating scalar fields. Conditional Source-term Estimation is one such combustion model that provides closure for the mean chemical source terms for turbulent flames. CSE invokes the first-order Conditional Moment Closure hypothesis such that the conditionally-averaged chemical source terms are closed by evaluating the reaction rate expressions with the conditionally-averaged scalars,

$$\overline{\dot{\omega}_\alpha|c^*} \approx \dot{\omega}_\alpha(\overline{T|c^*}, \overline{Y_\alpha|c^*}, \overline{\rho|c^*}), \quad (5.1)$$

where c^* is the discretised conditioning variable. $\overline{T|c^*}$, $\overline{Y_\alpha|c^*}$ and $\overline{\rho|c^*}$ are the conditional temperature, mass fraction of species α and density, respectively. The unconditional mean chemical source term is calculated by integrating the conditional source term with the probability density function (PDF) of the conditioning variable, $P(c^*; \vec{x}, t)$,

$$\tilde{\omega}_\alpha = \int_0^1 \overline{\dot{\omega}_\alpha|c^*} \tilde{P}(c^*; \vec{x}, t) dc^*. \quad (5.2)$$

The Favre-averaged presumed PDF models adopt the functional form:

$$\tilde{P}(c^*; \vec{x}, t) \approx \tilde{P}(c^*; \tilde{c}, \tilde{c}'^2), \quad (5.3)$$

where \tilde{c} and \tilde{c}'^2 are the Favre-averaged mean and variance of the progress variable. The usual spatial and temporal coordinates are represented by \vec{x} and t . Several

PDF models are tabulated in this format; a PDF model commonly implemented in conjunction with CSE is the modified laminar flamelet PDF [70, 143–145, 153].

The central concept behind CSE is that there is more information available if the filtered field is considered in its entirety rather than as unrelated data points. Each filtered volume, governed by the same underlying physical processes, then contributes additional information about the localised reactive flow. Specifically, the relevant conditional scalar averages are obtained by the inversion of the following integral equation [26],

$$\tilde{Y}_\alpha(\vec{x}, t) = \int_0^1 \tilde{P}(c^*; \vec{x}, t) \overline{Y_\alpha|c^*}(\vec{x}, t) dc^*, \quad (5.4)$$

where $\tilde{Y}_\alpha(\vec{x}, t)$ is the Favre-averaged mass fraction¹. Figure 5.1² illustrates the interaction between the CFD module, the CSE module, the chemistry module — which, in this implementation presented here uses a low-dimensional manifold chemical mechanism reduction scheme called the Flamelet-Generated Manifold (FGM) [121, 142], which reduces the chemical kinetic mechanism down to a simple two-dimensional table lookup based on two reaction progress variables (the mass fractions of vapour water and carbon dioxide) — and the PDF. The mass fractions, means and variances of the progress variable are calculated by the CFD module via transport equations. These means and variances of the progress variable are then used to obtain the PDFs, which in conjunction with the mass fractions, serve as the input to the CSE module via Equation 5.4. The computed conditional mass fractions are subsequently used to determine the mean reaction rates through the FGM and the PDF model. In this chapter, the discussion is focused on the optimisation of two steps within the CSE framework: (1) to dynamically allocate the physical boundary that governs the cells used as the input to Equation 5.4; and (2) to compute the conditional mass fractions with greater efficiency.

In practice, Equation 5.4 can be expressed in a matrix form for each species,

$$\mathbf{b} = \overline{\mathbf{A}\mathbf{x}}, \quad (5.5)$$

¹Both $\overline{Y_\alpha|c^*}(\vec{x}, t)$ and $\overline{Y_\alpha|c^*}$ represent the conditional mass fraction of species α . The specification, (\vec{x}, t) , expresses the notion that the preceding quantity is a function of space and time. This simplified notation is applicable to all of the scalar fields and conditional scalar fields in this section.

² This figure is the same as Figure 2.9. It is repeated here for convenience.




Figure 5.1: CSE operational flowchart.

where \mathbf{b} , $\bar{\mathbf{A}}$ and \mathbf{x} respectively represent the mass fraction, the PDF of the conditioning variable and the conditional mass fraction. The length of array \mathbf{b} corresponds to the number of data points within the ensemble and the length of array \mathbf{x} corresponds to the number of bins in which the PDF is discretised. From past CSE experience [143–145, 153], it appears that a suitable minimum size of \mathbf{b} is around $\mathcal{O}(10,000)$ while the size of \mathbf{x} is typically held constant at 50. These two numbers consequently determine the size of matrix $\bar{\mathbf{A}}$, which would be in the vicinity of $\mathcal{O}(10,000)$ by 50. As a result, Equation 5.5 represents an over-constrained inversion problem; the solution will therefore be sensitive to small perturbations in \mathbf{b} and regularisation must be imposed for stability and uniqueness of \mathbf{x} [63].

The Tikhonov regularisation approach has been implemented with CSE extensively in the past for both premixed and non-premixed combustion [42, 67, 92, 93, 143–145, 153, 161, 183]. For premixed combustion, it was first proposed by Jin *et al.* [70] to use the solution of an unstrained one-dimensional laminar premixed flame as the regularisation array. The motivation is that this regularisation method would implicitly allow the CSE combustion model to switch to the flamelet solution in case of ill-posedness. In this way, one can obtain a robust combustion model that is suitable for both flamelet and non-flamelet combustion regimes.

As an aside, another feasible *a priori* solution to implement as the constraint with the Tikhonov regularisation would be the conditional mass fraction from the previous timestep. The advantage of this approach is that perhaps the inverted conditional mass fractions would be more reflective of the current values within the computation domain; however, the disadvantage is that the errors from the previously inverted solutions may slowly accumulate. As a result of this, it is possible for this regularisation method to decrease stability instead of providing

stability as the simulation progresses.

In general though, there is a tradeoff between stability and accuracy in the inversion problem [142]. One can always obtain a stable solution to the conditional mass fraction by increasing the regularisation parameter in the system; however, this is not desired for turbulent combustion as the solution will tend towards the regularisation array: an unstrained one-dimensional laminar premixed flame. The preferred approach for achieving stability while minimising regularisation is to increase the size of the system. This translates to ensuring a minimum number of data points is included in each instance of inversion of Equation 5.4. As mentioned previously, this minimum number of points is roughly $\mathcal{O}(10,000)$.

5.3 Optimisations

5.3.1 Dynamic Ensemble Selection

As discussed in the preceding section, the central concept behind CSE is that there is more information available if the filtered field is considered in its entirety rather than as unrelated data points. This leads to the grouping of localised cells referred to as *ensembles*. It is important to emphasise the distinction between the CFD subdomains used by OpenFOAM 2.1.1 [185] and the ensembles used by CSE before describing the new algorithm in detail. The number of subdomains is directly linked to the number of processors available to the user for the simulation. These subdomains are typically divided in such a way as to minimise boundary surface areas (hence, communication overhead) between the processors. Moreover, the number of cells within each subdomain is made roughly equal for maximum load balance. These subdomains are predetermined by the CFD module and remain constant throughout the simulation. With the CSE ensembles, the physical boundaries of these partitions must change during the simulation if one wishes to incorporate a roughly equal number of *reactive* cells in each ensemble. This would ensure proper load balance and more importantly, sufficient number of reactive cells in all of the ensembles for the inversion of Equation 5.4.

The accepted method of dividing up the domain for CSE is to manually split up the flow field and assign each ensemble to a processor prior to the simulation;

this is a close analogue to the partitioning of the computational domain for parallel processing in the CFD module. This manual ensemble division ultimately depends on the physical insight of the user for the particular flow field [143]. For example, a reasonable set of ensembles for a jet flame would be axisymmetric slices along the axial direction. The two major advantages with this method are the ease of implementation and good performance for flames with clear symmetries [143–145, 153]. Manual partition of the domain into geometrically equal ensembles, however, has several associated disadvantages. First, this technique may not be suitable for complex geometries or flames with more complicated structures. In such cases, the optimal ensemble assignment might not be immediately evident to the user prior to the initial simulation. Second, the reactive regions (that is, regions in which chemical reactions are taking place) for turbulent flames will constantly shift within the simulation field and will be distributed non-uniformly throughout the domain. This could lead to load balance issues where equal amounts of computational resources are assigned to regions of different reactivity. Figure 5.2 illustrates an axisymmetric jet flame in a rectangular simulation domain divided axially into different number of ensembles; for the illustrated jet flame, the ensembles near the base of the flame ($z = 0.0$ m) will be less “reactive” than the top of the flame. In this case, a cell is defined to be reactive if it satisfies the following condition,

$$\tilde{c} \cdot (1 - \tilde{c}) > 10^{-2}. \quad (5.6)$$

This simple condition eliminates the computational cells that are mostly reactants and mostly products. Depending on the inlet conditions and definition of the progress variable, the value of \tilde{c} in the inflow boundary condition is typically set to 0 or 1; thus the cells within close proximity to the inflow boundary at the base of the flame would be considered non-reactive according to Equation 5.6. Third, due to the nature of the CSE inverse problem, one must provide a certain number of data points in order to invert Equation 5.4 appropriately. This issue becomes of concern when a large number of processors is to be used for the simulation. From Figure 5.2, it is evident that increasing the number of processors effectively decreases the thickness of each ensemble; when the thickness becomes too small, the number of available reactive cells to perform the inversion would become in-

5.3 Optimisations

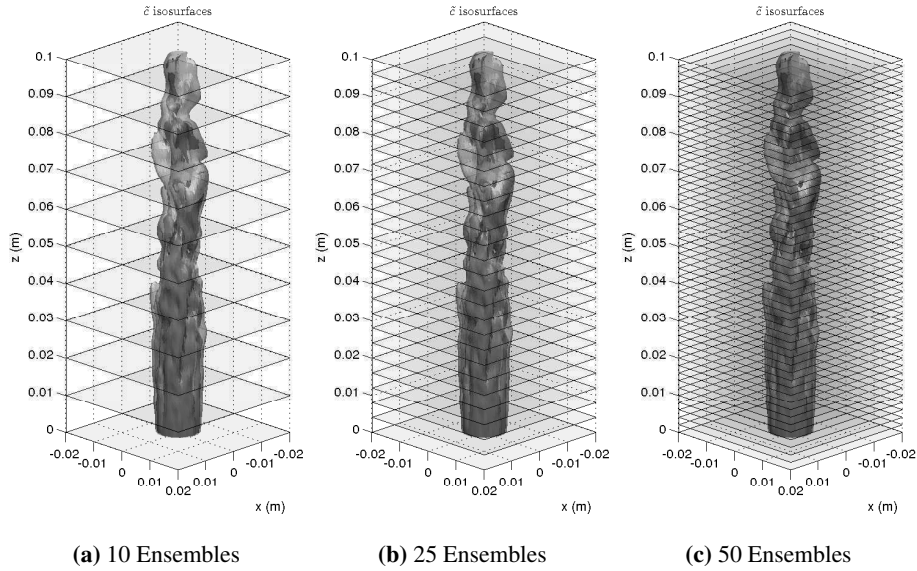


Figure 5.2: Computational domain incorporating different number of static ensembles. The volume of each ensemble decreases with increasing number of processors.

sufficient. A practical method to alleviate part of this problem is to introduce an overlap in the physical spaces between neighbouring ensembles [143–145, 153]. This increases the size of the ensembles and includes an added benefit of a more gradual transition of the conditional mass fraction between ensembles as some data is now shared — albeit at the expense of increased communication overhead between the processors sharing data. Clearly, however, this does not provide a long term solution to the problem.

These disadvantages suggest that a more optimal method to decompose the flow field would be to allocate the ensembles and processors in real time according to the location of the flame for maximum load balance. As a consequence, the physical boundaries of the ensembles will vary and thus likely lead to a more complicated implementation, in addition to increased communication of cell information between processors. The main difficulty here is determining a robust algorithm to construct the ensembles.

When constructing dynamic ensembles, the number of processors, N , available




Figure 5.3: The three possible scenarios that one would encounter during the construction of dynamic ensembles are illustrated: $N < M$, $N = M$ and $N > M$, where N and M are respectively the number of processors available and the number of ensembles required.

may be less than, equal to, or greater than the number of ensembles, M . For any given simulation, N will be fixed, whereas the number of ensembles (M) and the physical boundaries of each ensemble will continuously change within the domain. The goal is to divide the domain such that $\mathcal{O}(10,000)$ reactive cells are incorporated into each ensemble. It is then conceivable for the algorithm to experience all three scenarios over the course of one simulation. Thus, it is imperative to devise strategies to manage each case as it arises. Figure 5.3 illustrates the N and M divisions in a rectangular simulation domain for the three possible cases.

For the first case, where the number of processors is less than the number of ensembles ($N < M$), the implementation could be quite straightforward: one could simply divide the number of ensembles as evenly as possible between the processors. If N is much smaller than M , the remainders in division can be almost uniformly distributed and load balancing would not be a primary issue. If N is slightly smaller than M , then the difference in computational time between the quickest and slowest processors would be approximately one ensemble inversion.

The second case, where the number of processors equals the number of ensembles ($N = M$), is ideal: each processor can be assigned to one ensemble. This is what often occurs in conventional CSE simulations, where the user would carefully select the number of processors appropriate to the specific problem.

The third case, where the number of processors is greater than the number of ensembles ($N > M$), represents a new frontier in CSE. Here, one could conceive of two different approaches, depending on the latency associated with message pass-

ing on the particular machine on which the calculation is to be done. The first would be to simply have one local master processor collect the information needed for the CSE calculation for its neighbours, perform the CSE calculation for the ensemble, then distribute the calculated conditionally averaged reaction rates back to the neighbouring processors. In this way, the conventional approach to CSE would still be applicable. However, computational resources are not efficiently allocated as a large number of processors will be at an idle state while only the local masters perform the inversions for the ensembles: the resulting code would not scale. The other approach would be to use a different, parallel algorithm to take advantage of the parallel computing environment by having the CSE calculation occur in parallel for ensembles spread over a cluster of processors. During practical implementation, this would amount to solving each ensemble in parallel to maximise efficiency and hopefully lead to a code that would scale. As largely-parallel computation becomes a standard in LES combustion codes, the $N > M$ case is likely to emerge as the only practical scenario for turbulent flame simulations with CSE in the future.

In this chapter, two dynamic ensemble selection strategies that automatically adjust the physical boundaries of the ensembles as required to maintain the correct number of reactive cells in each ensemble are devised.

Semi-Dynamic Ensemble Selection Algorithm

The premise of this new algorithm is relatively straightforward. To begin, one must define what it is for a cell to be “reactive” — and, hence, to be meaningfully contributing to the CSE inversion process by virtue of having a PDF for the conditioning variable that is not merely a delta function at either zero or unity. For premixed combustion, a reactive cell is one where the quantity $\tilde{c} \cdot (1 - \tilde{c})$ is greater than 10^{-2} , as previously defined by Equation 5.6.

Prior to the first timestep of the simulation, the domain is decomposed by the OpenFOAM CFD module; the centroid of each subdomain is computed and passed to the master processor. The master then constructs an $N \times K$ array, where N is the number of processors and K is the maximum number of processors with which to form an ensemble, that indicates the nearest neighbours to each of the proces-

Processor	Nearest Neighbours			
	1 st	2 nd	3 rd	→ K th
1	2	3	4	→
2	1	3	4	→
3	2	4	1	→
↓	↓	↓	↓	→
N	N-1	N-2	N-3	→

Processor	Reactive Cells
1	3,541
2	4,334
3	5,872
4	3,546
↓	↓
N	1,235

Figure 5.4: Semi-Dynamic Ensemble Selection processor map.

sors (Figure 5.4). During runtime, an integer containing the number of reactive cells within each subdomain is sent to the master processor and sorted according to the $N \times K$ array. For example, the construction of the ensemble for processor 3 is demonstrated by Figure 5.4. The algorithm sums the number of reactive cells in order of proximity, starting with itself, then processors 2, 4, 1 and so on, until the total number of reactive cells exceeds the minimum predefined number; the data from these processors constitutes the ensemble of processor 3. This summation procedure is repeated for each subdomain until all of the ensembles are constructed. Indeed, the number of reactive cells in each subdomain changes at every timestep; the integers shown in the “Reactive Cells” column in Figure 5.4 are for illustrative purposes only. Unless otherwise specified, all of the inter-processor communications are performed using the standard OpenFOAM class, UPstream [60], which is based on the Message Passing Interface (MPI) system.

Effectively, each CFD subdomain becomes the centre of a CSE ensemble and its nearest neighbours will contribute information to this central core. For particularly reactive regions, perhaps one or two neighbours will contain sufficient reactive cells for a proper inversion; for less reactive regions, cell information is gathered from the surrounding subdomains in increasing radius from the core until a sufficient number of reactive cells is achieved or the maximum number of neighbours, K , is exceeded. From past experience, it appears that a sufficient number of reactive cells for the CSE inversion is on the order of $\mathcal{O}(10,000)$. For the current study, this parameter is varied between 5,000 and 15,000 cells depending on the grid density. This implementation represents a special case of $N = M$, where the number of ensembles is constant in quantity but varies in physical size. This algo-

rithm is referred to as the Semi-Dynamic Ensemble Selection because the number of processors always corresponds to the number of ensembles.

A possible limitation of this algorithm is that a number of a processors may repeat the same CSE calculations in the extreme scenario where only one or two subdomains contain reactive cells. For this case, it would be possible for $\mathcal{O}(K)$ processors to repeat the same inversion. The repeated inversions, however, would not slow down the overall CSE calculation because all of the processors must be synchronised after every timestep; the otherwise idling processors would need to wait for the one to two highly reactive subdomains to finish their CSE calculation before proceeding. In this regard, the Semi-Dynamic Ensemble Selection algorithm simply repeats the calculation on the idling processors without additional penalty. This slight inefficiency is one of the motivations to introduce the Dynamic Ensemble Selection algorithm, where the load would be equally distributed to *all* of the processors in this limiting case. Another consideration is that the preceding method of ensemble selection is ideally suited for geometries that contain no spatial voids and for subdomains that are arranged along the principle flow direction, which is typical of experimental burners. For complex geometries, however, more judicious decomposition of the domain may be required.

Dynamic Ensemble Selection Algorithm

There are two primary advantages associated with the Dynamic Ensemble Selection algorithm in comparison to the Semi-Dynamic Ensemble Selection algorithm and static ensembles. First, the number of processors used in the simulation can be completely decoupled from the number of ensembles. This allows for a more flexible allocation of resources and reduces complexity for the user during the initial setup of the simulation. In addition, this allows for a high number of processors to be used in parallel without introducing exceedingly small ensembles or large overlaps between the ensembles. Second, a parallel solver can be implemented to solve the inverse problem as multiple processors are now assigned to each matrix.

The Dynamic Ensemble Selection algorithm consists of six steps with each step summarised in Figure 5.5. Similar to the Semi-Dynamic Ensemble Selection algorithm, the reactive cells within each CFD subdomain are first determined by the

5.3 Optimisations

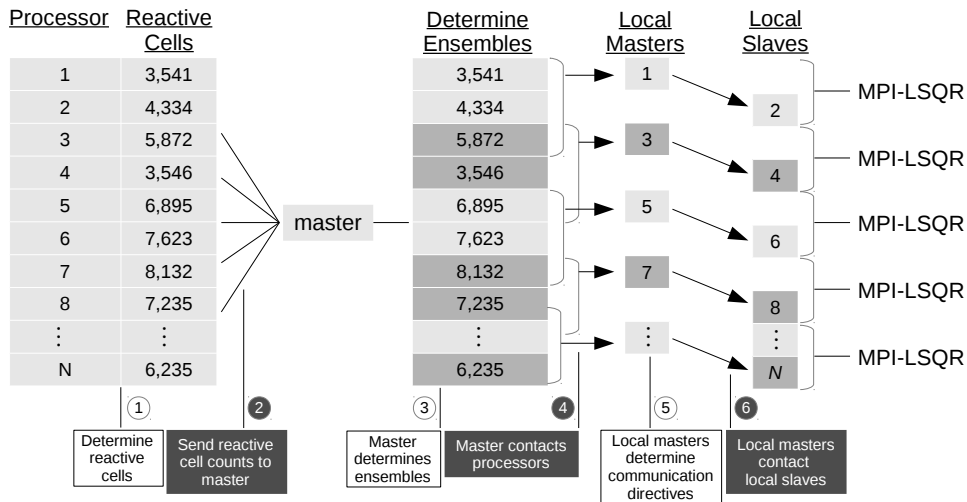


Figure 5.5: Dynamic Ensemble Selection schematic.

criterion set in Equation 5.6. Second, once the number of reactive cells is identified on each processor, the N integers are sent to the master processor. Third, the master allocates the ensembles according to the user defined quantity that is related to the minimum number of reactive cells per ensemble; as before, this number is of $\mathcal{O}(10,000)$. Fourth, the master communicates with all of the processors and assigns their role, local master or local slave, in each ensemble. Fifth, the local masters compute the communication directives for the local slaves belonging to the respective ensembles according to the classic round robin algorithm. Sixth, the local masters distribute the directives to the appropriate local slaves. This process of ensemble allocation occurs at the beginning of each timestep, prior to any other calculations.

To further consider the layout of the ensembles in this algorithm, the first and last processors of each ensemble are deliberately overlapped with the previous and next ensembles. This sharing of reactive cell data between ensembles allows for a smoother transition of the resulting conditional mass fractions from one part of the domain to the next. The ensembles illustrated in Figure 5.5 can be examined. For this example, each ensemble consists of data from three processors: Ensemble 1 includes data from processors 1–3; ensemble 2 includes data from processors 3–5; and ensemble 3 includes data from processors 5–7. This pattern continues

until the N processors are fully divided. To properly invert the matrix in ensemble 1, processor 1 must receive the reactive cell data from both processors 2 and 3, while processor 2 must receive the reactive cell data from processors 1 and 3. Only processors 1 and 2 will be used to compute the inverse of this ensemble in parallel; processor 3 is not directly used in this inversion calculation as it will be used in ensemble 2. Similarly, for ensemble 2, processor 3 must receive the reactive cell data from both processors 4 and 5, while processor 4 must receive the reactive cell data from processors 3 and 5. The two processors, 3 and 4, will then compute the inverse of this ensemble in parallel. This pattern continues until the M ensembles are fully divided. In the event where the last ensemble in the domain contains less than the user-defined minimum number of reactive cells, it will be combined with the nearest ensemble.

While this procedure appears to be complex, the additional computational cost of the Dynamic Ensemble Selection is relatively low. In the first step, the reactive cells must be identified regardless of the ensemble selection method, thus, imposing no additional cost. Second, only *one* integer is sent to the master from each processor to identify the number of reactive cells. Third, the master is only required to sort N integers, where N is the total number of processors in the simulation. Fourth, the master sends three integers back to each processor; these integers identify the positions of the first, the last and the overlapping processors belonging to each ensemble. Fifth, to further reduce computational time, the round robin algorithm used to determine the communication directives by the local masters is pre-computed and tabulated into lookup tables prior to the simulation. Sixth, the local masters send the communication directives to the proper local slaves, with each ensemble typically incorporating one to three local slaves. These directives consist of approximately $2 \cdot n$ integers, where n is the number of processors in each ensemble. Once the ensembles have been assigned, the next step is to compute the inverse solution with the parallel iterative solver described in Section 5.3.2.

5.3.2 Matrix Inversion - Iterative Solver

Inverting the PDF matrix, denoted as $\overline{\mathbf{A}}$, is the most time-consuming step in a CSE calculation. As mentioned, the typical inversion matrix is $\mathcal{O}(10,000)$ by 50 and the

accepted method uses LU-decomposition with a zeroth order Tikhonov regularisation [143–145, 153]. Practical implementation of the LU-decomposition usually requires a mathematical manipulation to first reduce the size of the inversion, involving a matrix multiplication of $\overline{\overline{\mathbf{A}}}^T \overline{\overline{\mathbf{A}}}$. The benefits of the LU method are that the solution is exact and the inversion time is deterministic. However, in the context of CSE, an exact solution is not required for the ensemble average, as there are at least two sources of error implicit to the inversion of Equation 5.4. First, the PDF of the reaction progress variable that is used as the kernel of inversion will inherently contain some modelling error; and second, the zeroth order Tikhonov regularisation using the laminar flamelet solution contributes additional error. As a consequence, provided the error generated from the iterative solver is relatively small compared to the aforementioned sources, it is tenable to obtain a pseudo-inverse of the matrix instead of an exact inverse of the matrix for the conditional scalars. This relatively modest reduction in accuracy is a small compromise to the substantial reduction in computational effort for the matrix inversions.

To address this, the LU solver is proposed to be replaced with an established iterative approach, Least-Squares with QR-factorisation, or LSQR [123]. This iterative method is designed for the computation of $\overline{\overline{\mathbf{A}}}\mathbf{x} = \mathbf{b}$, where $\overline{\overline{\mathbf{A}}}$ and \mathbf{b} are known. The LSQR implements the Lanczos method that leads to a factorization of the tridiagonal reduced matrix for $\overline{\overline{\mathbf{A}}}^T \overline{\overline{\mathbf{A}}}$. This method is often the method of choice for over-determined or under-determined systems [140]. Formally, it is equivalent to the following minimisation problems:

$$\min || \overline{\overline{\mathbf{A}}}\mathbf{x} - \mathbf{b} ||^2, \quad (5.7)$$

$$\min || \overline{\overline{\mathbf{A}}}\mathbf{x} - \mathbf{b} ||^2 + \lambda_d^2 || \mathbf{x} ||^2, \quad (5.8)$$

where the matrix $\overline{\overline{\mathbf{A}}}$ is supposed to be large and sparse³ and λ_d is the damping coefficient. The method is based on the bidiagonalization procedure of Golub and Kahan [57]. It is analytically equivalent to the standard method of conjugate gradients, but possesses more favorable numerical properties [123]. The cost per it-

³The PDF matrix is not sparse; however, the LSQR solver has proven to be robust and efficient for all of the performed tests. These include standalone matrix inversions and CSE inversions in the context of LES.

eration for this solver is primarily dependent on the matrix-vector multiplications of $\overline{\mathbf{A}} \cdot \mathbf{x}$ and $\overline{\mathbf{A}}^T \cdot \mathbf{b}$ and typically converges in less than 25 iterations for correctly regularised one-condition CSE matrices.

Another considerable advantage of LSQR over LU-decomposition is that the iterative solver can be parallelised for the computation of large-scale inversion problems. The implementation of a parallel LSQR solver would allow the user to efficiently solve each CSE inversion problem when the number of processors exceeds the number of ensembles. A Scalable Parallel LSQR (SPLSQR) algorithm has been developed recently by Huang *et al.* [66] in the field of seismic tomography; SPLSQR has demonstrated scalability beyond $\mathcal{O}(10,000)$ processors for large-scale inversion problems. In this study, a modified version of their scalable parallel algorithm adapted for CSE problems has been developed; this parallel solver will be referred to as MPI-LSQR in the subsequent discussions to reduce ambiguity. In general, the code structure of the MPI-LSQR and LSQR solvers are very comparable, though two unique distinctions can be identified. First, each processor participating in the parallel computation of MPI-LSQR is responsible only for a portion of the solution vector. Second, the partially updated solution vectors would need to be distributed between all of the participating processors at each iteration.

Figure 5.6 illustrates the steps required by the MPI-LSQR algorithm to obtain a converged solution for the case of three participating processors. The most current working arrays of length \mathbf{x} and \mathbf{b} , respectively representing $\overline{Y_\alpha|c^*}$ and \tilde{Y}_α , are first gathered and distributed between the processors. These arrays, shown in white, are then checked for convergence according to the user defined error tolerances. If the arrays satisfy the error tolerances, the current solution is accepted. If not, then an iteration of the solver will be performed; each iteration consists of two matrix-vector multiplications on the order of $\overline{\mathbf{A}} \cdot \mathbf{x}$ and $\overline{\mathbf{A}}^T \cdot \mathbf{b}$, followed by two small normalisation calculations. Crucially, each processor only needs to compute one third of the total multiplications for this particular illustration. Each processor then obtains a local portion of the new \mathbf{x} and \mathbf{b} arrays. Next, these local arrays are gathered and distributed between the processors. This cycle continues until convergence is achieved. In applying the MPI-LSQR solver to the inverse problem, a computational scheme that is parallel in both the CFD and the CSE modules has

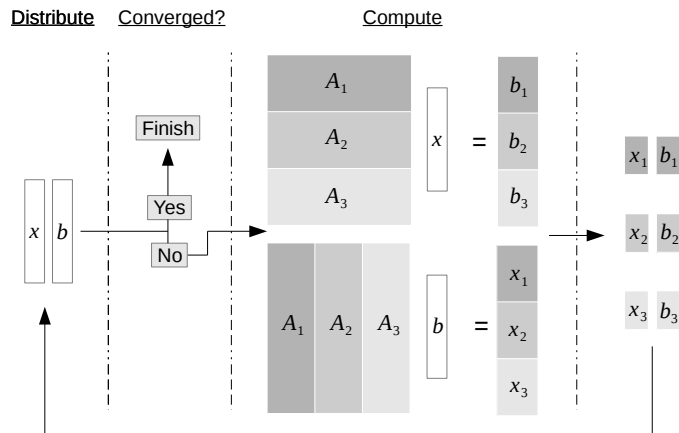


Figure 5.6: MPI-LSQR iteration. Each n participating processor performs $1/n$ of the total matrix-vector multiplications.

been obtained.

5.4 Numerical Methods

5.4.1 Standalone Inversion Tests

The LSQR and MPI-LSQR solvers have been thoroughly tested before implementation with the CSE calculations. Two separate tests have been performed. First, the original formulation of LSQR by Paige and Saunders [123] has been modified to accommodate the input of an initial solution, \mathbf{x}_0 . The tests have shown that this improves the convergence rate by 10–15%, reducing the number of iterations to 20–25 for the average one-condition CSE problem. During practical implementation, \mathbf{x}_0 is set to the conditional mass fraction of the previous timestep; the usual zeroth order Tikhonov regularisation with the laminar flamelet solution is used [142]. Second, the LSQR test is repeated with the MPI-LSQR solver subjected to the same initial condition and Tikhonov regularisation protocols. During combustion simulation with CSE, in the case where the number of processors exceeds that of the ensembles ($N > M$), it is possible for each of the processors to carry a different conditional mass fraction from the previous timestep into the ensemble. In this case, the conditional mass fraction used as the input of \mathbf{x}_0 would

have to be a weighted average of the participating processors.

For the standalone LSQR inversion test, five matrices of different dimensions are examined. The PDF matrix, $\overline{\mathbf{A}}$, is first randomly generated and multiplied with a known vector, \mathbf{x} , to create the sample problem of the form $\overline{\mathbf{A}}\mathbf{x} = \mathbf{b}$. The $\overline{\mathbf{A}}$ and \mathbf{b} are then inverted individually by the LU-decomposition (as the control) and the LSQR solver at various settings. In order to better represent the conditions that the LSQR solver would encounter during practical CSE implementation, the initial condition input, \mathbf{x}_0 , is generated by adding random +/- 5% fluctuations to the known \mathbf{x} . These random fluctuations emulate the conditional mass fraction change per timestep. Though this is likely an overestimate of the change, it serves to provide an upper limit of the convergence time. Moreover, the error bounds controlling the termination condition of the LSQR solver have been adjusted to determine their effects on accuracy and convergence time. The +/- 5% fluctuations are not necessary for the LU-decomposition method as it does not accept an input solution and the computational times are strictly dependent on the size of the system. The inversions are performed on a single i7-2600K processor with a maximum frequency of 3.8 GHz.

For the standalone MPI-LSQR inversion test, quantifying the scalability of the new MPI-LSQR algorithm is the primary concern. The inversion problems are generated identically as the LSQR inversion test and the parallel solver is subjected to the same initial condition and Tikhonov regularisation protocols. The inversions have been computed using a number of different computational node (1–12) and processor (1–48) combinations. All of the parallel computations in this study are performed on the Mammouth Parallèle II supercomputer of Université of Sherbrooke. The participating nodes consisted of processors with frequencies of 2.2 GHz.

5.4.2 LES of Axisymmetric Burner

The simulation setup of the stoichiometric flame conditions described in a previous study [153] has been followed. The CFD code is based on OpenFOAM 2.1.1 [185] and a FGM model is used for the chemistry reduction [121, 142]. Second order implicit temporal and second order spatial schemes are employed. The LES model

is the standard *oneEqEddy* from OpenFOAM, which is equivalent to the k -equation eddy-viscosity model. A cylindrical simulation domain with a radius of 0.025 m and a height of 0.10 m is used along with three grids of different hexahedral cell counts and processor numbers. The three grids consisting of 52,500, 284,000 and 880,000 cells are respectively assigned to 4, 24 and 48 processors. While the grid density is lower than required for accurate LES simulations, demonstrating the characteristics of the new optimisation algorithms takes precedence. The goal here is to compare the combined performance of the three inversion strategies under identical simulation conditions: MPI-LSQR solver with Dynamic Ensemble Selection, LSQR solver with Semi-Dynamic Ensemble Selection and conventional LU-decomposition with static (manual) ensemble selection.

The inflow condition is isotropic homogeneous turbulence generated from a spectral method using the energy spectrum of Haworth and Poinso [65]. The turbulent kinetic energy (k), peak energy wavenumber and maximum wavenumber are respectively $33.3 \text{ m}^2/\text{s}^2$, $2\pi/l_0$ and $2\pi/2\Delta$, where l_0 and Δ are the integral length and grid spacing. The integral length is set to $4.8 \cdot 10^{-3}$ m and the grid spacing varied from $8.96 \cdot 10^{-4}$ m to $2.24 \cdot 10^{-4}$ m. The time correlation for the velocity fluctuation is based on the method developed by Billson [15]. The relevant simulation parameters are summarised in Table 5.1. The inlet diameter is $1.12 \cdot 10^{-2}$ m with stoichiometric premixed methane-air reactants at 300 K and 101.3 kPa. These parameters are designed to simulate the experimental conditions as closely as possible [153]. The inlet bulk flow, which includes the spectral generated turbulence, is 15.58 m/s. A laminar, hot pilot surrounds the inlet with an inflow of combustion products moving at 16.81 m/s.

The simulations are carried out in two stages. The first stage is performed by the very coarse grid (52,500 cells) to quickly discharge the initial condition and to develop the flame in the domain. This is allowed to carry on for approximately eight flow-through times, from $t = 0.0$ s to $t = 0.05$ s. The second stage is initiated by interpolating the result of the very coarse grid to the two finer grids. Data is then collected from $t = 0.05$ s to $t = 0.25$ s for the three grids. Radial temperature profiles are collected every 0.001 s at axial positions of $z/D = 1-6$, where D is the diameter of the inlet ($D = 1.12 \cdot 10^{-2}$ m). The simulations are performed using the MPI-LSQR solver with Dynamic Ensemble Selection, LSQR solver with

5.5 Results and Discussion

Grid	52,500 cells	284,000 cells	880,000 cells
k	$33.3 \text{ m}^2/\text{s}^2$	$33.3 \text{ m}^2/\text{s}^2$	$33.3 \text{ m}^2/\text{s}^2$
Peak energy wavenumber	$1.31 \cdot 10^3 \text{ m}^{-1}$	$1.31 \cdot 10^3 \text{ m}^{-1}$	$1.31 \cdot 10^3 \text{ m}^{-1}$
Maximum wavenumber	$3.51 \cdot 10^3 \text{ m}^{-1}$	$7.01 \cdot 10^3 \text{ m}^{-1}$	$1.40 \cdot 10^4 \text{ m}^{-1}$
l_0	$4.80 \cdot 10^{-3} \text{ m}$	$4.80 \cdot 10^{-3} \text{ m}$	$4.80 \cdot 10^{-3} \text{ m}$
Δ	$8.96 \cdot 10^{-4} \text{ m}$	$4.48 \cdot 10^{-4} \text{ m}$	$2.24 \cdot 10^{-4} \text{ m}$
Δt	$8.00 \cdot 10^{-6} \text{ s}$	$4.00 \cdot 10^{-6} \text{ s}$	$2.50 \cdot 10^{-6} \text{ s}$
τ_t	$9.96 \cdot 10^{-4} \text{ s}$	$9.96 \cdot 10^{-4} \text{ s}$	$9.96 \cdot 10^{-4} \text{ s}$
a	0.992	0.996	0.998
b	0.125	0.0884	0.0700

Table 5.1: Relevant LES simulation parameters. The values of a and b are computed according to the approach of Billson [15]: $a = \exp(-\Delta t/\tau_t)$ and $b = \sqrt{1 - a^2}$, where Δt and τ_t are respectively the timestep size and the integral time scale.

Semi-Dynamic Ensemble Selection and LU-decomposition with static ensembles for the three grids. It is important to emphasise that the only changes between the simulations are the inversion and ensemble selection algorithms; the grids, CFD module, chemistry reduction, timestep sizes and computing facilities are otherwise identical. The performance and results are documented in Section 5.5.2.

5.5 Results and Discussion

5.5.1 Standalone Inversion Tests

Table 5.2 compares the computational cost of inverting matrices of various sizes and L_2 -norms⁴ of the solutions that are relevant to one and two-condition⁵ CSE problems using the LU-decomposition method and the LSQR solver. Figure 5.7 graphically compares the computational cost of inversion from Table 5.2. It is apparent that the LSQR solver demonstrates noticeably shorter calculation times than LU-decomposition for one-condition CSE and is orders of magnitude faster

⁴ L_2 -norm is defined as: $L_2 = \sqrt{\sum_{i=1}^n (x_{i,LSQR} - x_{i,exact})^2}$, where n is the number of discrete bins in the tabulated PDF model, which is typically 50 for one-condition CSE problems.

⁵Two-condition CSE is also referred to as Doubly Conditional Source-term Estimation (DCSE) in the literature [43].

5.5 Results and Discussion




Figure 5.7: Comparison of inversion times between LU-decomposition and LSQR at a number of different error tolerance settings ($|e|$). This is a graphical representation of the data in Table 5.2. Each of the five plots represents a different matrix size. For each plot, the bars represent inversion times of (from left to right): LU, LSQR ($\vec{0}$) ($|e| = 10^{-10}$), LSQR (\mathbf{x}_0) ($|e| = 10^{-10}$), LSQR (\mathbf{x}_0) ($|e| = 10^{-8}$), LSQR (\mathbf{x}_0) ($|e| = 10^{-6}$), and LSQR (\mathbf{x}_0) ($|e| = 10^{-4}$). The ‘ \times ’ represents the failed cases.

for two-condition CSE problems⁶. These results are anticipated as it has been documented that the LSQR solver is well-suited and exhibits exceptional convergence behaviour for over-determined inversion problems [41, 140]. Another favourable feature of LSQR is permitting the user to determine the relative and absolute error bounds according to the problem. For typical one-condition CSE problems, it is conceivable to increase the error tolerances to produce solutions with L_2 -norms on the order of 10^{-5} , further decreasing the convergence time. However, it should be mentioned that as the error tolerances are increased, some cases begin to reveal convergence issues as expected. These cases are marked by ‘ \times ’.

Table 5.3 compares the computational cost of inverting matrices of various sizes and L_2 -norms of the solutions that are relevant to one and two-condition CSE problems using the MPI-LSQR implementation. Figure 5.8 graphically compares the computational cost of inversion from Table 5.3. The inversions in the first

⁶ The focus is on one-condition CSE problems in this study. However, the LSQR solver is independent of the inversion problem and it is meaningful to illustrate its potential to solve two-condition CSE matrices for future work.

5.5 Results and Discussion

Dimensions	$1 \cdot 10^4 \times 50$	$3 \cdot 10^4 \times 50$	$2 \cdot 10^3 \times 1 \cdot 10^3$	$6 \cdot 10^3 \times 1 \cdot 10^3$	$1 \cdot 10^5 \times 2.5 \cdot 10^3$
LU	0.0319 s	0.301 s	9.97 s	30.7 s	4 hours+
LSQR ($\vec{0}$) $ e = 10^{-10}$ L_2 -norm \rightarrow	0.0258 s $4.5 \cdot 10^{-10}$	0.0912 s $4.4 \cdot 10^{-10}$	1.23 s $8.4 \cdot 10^{-6}$	1.95 s $1.3 \cdot 10^{-6}$	98.2 s $1.3 \cdot 10^{-6}$
LSQR (\mathbf{x}_0) $ e = 10^{-10}$ L_2 -norm \rightarrow	0.0253 s $1.5 \cdot 10^{-11}$	0.0875 s $3.0 \cdot 10^{-11}$	1.22 s $5.0 \cdot 10^{-7}$	1.90 s $4.7 \cdot 10^{-8}$	93.6 s $8.9 \cdot 10^{-8}$
LSQR (\mathbf{x}_0) $ e = 10^{-8}$ L_2 -norm \rightarrow	0.0216 s $3.6 \cdot 10^{-8}$	0.0753 s $8.9 \cdot 10^{-10}$	0.975 s $4.3 \cdot 10^{-5}$	1.54 s $4.6 \cdot 10^{-6}$	75.3 s $4.2 \cdot 10^{-6}$
LSQR (\mathbf{x}_0) $ e = 10^{-6}$ L_2 -norm \rightarrow	0.0168 s $8.6 \cdot 10^{-6}$	0.0639 s $7.4 \cdot 10^{-8}$	0.699 s $3.9 \cdot 10^{-3}$	1.26 s $3.7 \cdot 10^{-4}$	×
LSQR (\mathbf{x}_0) $ e = 10^{-4}$ L_2 -norm \rightarrow	0.0142 s $1.4 \cdot 10^{-4}$	0.0491 s $1.6 \cdot 10^{-5}$	×	×	×

Table 5.2: Comparison of inversion times between LU-decomposition and LSQR at a number of different settings. The $|e|$, $\vec{0}$ and \mathbf{x}_0 respectively denote the error tolerances in the solver, a zero initial solution and a non-zero initial solution. The associated L_2 -norms are printed below the convergence times. Typical matrix sizes for one-condition CSE are represented by the first two columns while typical matrix sizes for two-condition CSE are represented by the last three columns. Failed cases are denoted by ‘×’, where lower error tolerances are required to provide stable solutions. From the smallest to the largest matrix, 1,000, 500, 100, 20 and 2 samples are used to obtain the average inversion times.




Figure 5.8: Comparison of inversion times with the MPI-LSQR implementation using a number of different processors counts. This is a graphical representation of the data in Table 5.3. Each of the five plots represents a different matrix size. For each plot, the bars represent inversion times of (from left to right): LU, LSQR ($\vec{0}$), LSQR (\mathbf{x}_0), (1×2) , (2×1) , (1×3) , (3×1) , (1×4) , (4×1) , (1×8) , (8×1) , (1×12) , (12×1) , (1×24) , and (1×48) . The numbers $(x \times y)$ represent the number of nodes (x) and the number of processors per node (y) used in the MPI-LSQR computation. The inversion for the matrix of size $100,000 \times 2,500$ is not performed by the LU-decomposition due to an exceedingly exhaustive computational requirement; it is marked by ‘x’.

three rows are performed with one processor via the LU-decomposition and LSQR solver; these are considered the reference states. The inversions from rows 4–15 are performed by the MPI-LSQR solver with different numbers of nodes and processors. Moreover, the error tolerance ($|e|$) is held constant at $1 \cdot 10^{-10}$ for all of the LSQR and MPI-LSQR cases for consistency. The L_2 -norms produced by this value of error tolerance would be much lower than required by CSE; as such, the listed times illustrate an upper bound for the inversion calculations. During practical implementation, one would increase the $|e|$ -tolerance to the vicinity of $1 \cdot 10^{-8}$ – $1 \cdot 10^{-6}$ to increase the convergence rate by approximately 30%. The method used to generate the sample problems and initial condition inputs for rows 3–15 are as previously described in Section 5.4.1.

In general, it can be observed from Table 5.3 and Figure 5.8 that the inversion

5.5 Results and Discussion

Dimensions	$1 \cdot 10^4 \times 50$	$3 \cdot 10^4 \times 50$	$2 \cdot 10^3 \times 1 \cdot 10^3$	$6 \cdot 10^3 \times 1 \cdot 10^3$	$1 \cdot 10^5 \times 2.5 \cdot 10^3$
LU	0.228 s	0.694 s	12.3 s	59.0 s	N/A
LSQR ($\vec{0}$)	0.0877 s	0.249 s	2.03 s	5.39 s	313 s
L_2 -norm \rightarrow	$4.1 \cdot 10^{-10}$	$4.3 \cdot 10^{-10}$	$8.6 \cdot 10^{-6}$	$1.3 \cdot 10^{-6}$	$1.3 \cdot 10^{-6}$
LSQR (\mathbf{x}_0)	0.0881 s	0.337 s	2.80 s	5.41 s	361 s
L_2 -norm \rightarrow	$1.6 \cdot 10^{-11}$	$2.9 \cdot 10^{-11}$	$4.9 \cdot 10^{-7}$	$5.5 \cdot 10^{-8}$	$1.6 \cdot 10^{-8}$
1×2	0.0717 s	0.205 s	1.85 s	2.98 s	126 s
L_2 -norm \rightarrow	$1.3 \cdot 10^{-11}$	$3.1 \cdot 10^{-11}$	$5.0 \cdot 10^{-7}$	$5.7 \cdot 10^{-8}$	$1.6 \cdot 10^{-8}$
2×1	0.0739 s	0.197 s	1.88 s	2.72 s	137 s
L_2 -norm \rightarrow	$1.3 \cdot 10^{-11}$	$2.9 \cdot 10^{-11}$	$5.0 \cdot 10^{-7}$	$5.6 \cdot 10^{-8}$	$1.5 \cdot 10^{-8}$
1×3	0.0532 s	0.156 s	1.01 s	1.60 s	86.1 s
L_2 -norm \rightarrow	$1.3 \cdot 10^{-11}$	$2.9 \cdot 10^{-11}$	$4.9 \cdot 10^{-7}$	$6.1 \cdot 10^{-8}$	$1.5 \cdot 10^{-8}$
3×1	0.0573 s	0.168 s	0.978 s	1.50 s	86.5 s
L_2 -norm \rightarrow	$1.3 \cdot 10^{-11}$	$2.9 \cdot 10^{-11}$	$4.9 \cdot 10^{-7}$	$5.0 \cdot 10^{-8}$	$1.6 \cdot 10^{-8}$
1×4	0.0503 s	0.137 s	0.574 s	1.08 s	67.0 s
L_2 -norm \rightarrow	$1.2 \cdot 10^{-11}$	$3.0 \cdot 10^{-11}$	$4.9 \cdot 10^{-7}$	$5.8 \cdot 10^{-8}$	$1.8 \cdot 10^{-8}$
4×1	0.0460 s	0.134 s	0.574 s	0.991 s	63.7 s
L_2 -norm \rightarrow	$1.1 \cdot 10^{-11}$	$3.0 \cdot 10^{-11}$	$4.9 \cdot 10^{-7}$	$5.7 \cdot 10^{-8}$	$1.6 \cdot 10^{-8}$
1×8	0.0386 s	0.107 s	0.242 s	0.573 s	47.2 s
L_2 -norm \rightarrow	$1.4 \cdot 10^{-11}$	$3.0 \cdot 10^{-11}$	$4.8 \cdot 10^{-7}$	$5.6 \cdot 10^{-8}$	$1.7 \cdot 10^{-8}$
8×1	0.0398 s	0.0998 s	0.229 s	0.571 s	50.5 s
L_2 -norm \rightarrow	$1.2 \cdot 10^{-11}$	$2.9 \cdot 10^{-11}$	$4.9 \cdot 10^{-7}$	$5.1 \cdot 10^{-8}$	$1.6 \cdot 10^{-8}$
1×12	0.0347 s	0.106 s	0.192 s	0.512 s	44.3 s
L_2 -norm \rightarrow	$1.1 \cdot 10^{-11}$	$2.9 \cdot 10^{-11}$	$4.8 \cdot 10^{-7}$	$6.0 \cdot 10^{-8}$	$1.7 \cdot 10^{-8}$
12×1	0.0337 s	0.0978 s	0.202 s	0.559 s	38.3 s
L_2 -norm \rightarrow	$1.7 \cdot 10^{-11}$	$2.9 \cdot 10^{-11}$	$4.9 \cdot 10^{-7}$	$5.4 \cdot 10^{-8}$	$1.5 \cdot 10^{-8}$
1×24	0.0311 s	0.0927 s	0.161 s	0.392 s	24.3 s
L_2 -norm \rightarrow	$1.3 \cdot 10^{-11}$	$2.9 \cdot 10^{-11}$	$5.0 \cdot 10^{-7}$	$6.1 \cdot 10^{-8}$	$1.7 \cdot 10^{-8}$
1×48	0.0295 s	0.0983 s	0.152 s	0.390 s	23.6 s
L_2 -norm \rightarrow	$1.9 \cdot 10^{-11}$	$3.0 \cdot 10^{-11}$	$4.9 \cdot 10^{-7}$	$5.8 \cdot 10^{-8}$	$1.7 \cdot 10^{-8}$

Table 5.3: Comparison of inversion times for single ensembles of varying sizes with the MPI-LSQR implementation using a number of different processors. The numbers $x \times y$ in the first column represent the number of nodes (x) and the number of processors per node (y) used in the computation. The error tolerances are kept constant at 10^{-10} between all of the cases for consistency.

times decrease steadily, to a certain limit, with increasing number of processors. The limit appears to be around 24–48 processors, where gains tend to be minimal with additional processors. For some cases, the inversion time of LSQR (\mathbf{x}_0) is slightly longer than LSQR ($\vec{0}$). However, this is simply a consequence of the method by which the terminating condition is calculated from the relative error tolerance in the LSQR solver. An inspection of Table 5.3 reveals that the L_2 -norm from LSQR (\mathbf{x}_0) is typically an order of magnitude smaller than the L_2 -norm of LSQR ($\vec{0}$), leading to longer inversion times. With the inclusion of an initial condition (\mathbf{x}_0), it is possible to decrease the relative error tolerance in the solver to achieve the same L_2 -norm as cases with a zero initial condition ($\vec{0}$). Nonetheless, for consistency between cases, the relative and absolute error tolerances are held at the same values.

Larger systems of equations (two-condition CSE matrices) tend to demonstrate more benefit from increasing the processor count, though this is anticipated as communication overhead differentially affects smaller systems. A rather surprising result is that the internodal and intranodal communication times appear to be comparable for the machine used. For instance, examining the 1×12 and the 12×1 inversion times, where the first number represents the number of nodes while the second number represents the number of processors used per node, the values are typically within 15% of one another. More importantly, inversions performed by processors centralised on one node do not consistently yield greater performance than processors scattered over a number of nodes. This implies that the ensemble assignment could be very flexible for future implementation of this solver with CSE, as the penalty for computing a parallel inversion with processors from different nodes appears to be small.

5.5.2 LES of Axisymmetric Burner

The temperature and species mass fraction profiles at six axial locations for the nine simulations are depicted by Figure 5.9. For each grid density, the results averaged over 200 ms for the LU-decomposition with static ensembles and LSQR with Semi-Dynamic Ensemble Selection are nearly indistinguishable from one another; this implies that the conditional mass fractions produced by both of the methods

are very comparable. The slight deviations could be caused by a difference in the ensemble sizes between the static ensemble and Semi-Dynamic Ensemble Selection methods; additionally, the spectral turbulent inlet condition could play a minor role. It is interesting, however, to observe the transition exhibited by the CO mass fraction profiles with increasing grid resolution; particularly, the CO appears to increase around $r/D = 0.6$ farther downstream from the burner nozzle. This behaviour of the CO mass fraction could be associated with the currently implemented chemistry reduction technique, where the FGM may not fully capture the formation of species with slow time scales. If that is the case, it would not be possible for the algorithm, in its present form, to accurately represent the CO profiles. As the main purpose of this study is to compare the three ensemble selection algorithms and the three inversion strategies, the LES results from the LU-decomposition with static ensembles and the LSQR with Semi-Dynamic Ensemble Selection are remarkably similar.

The temperature and species mass fraction profiles produced by the MPI-LSQR and Dynamic Ensemble Selection are slightly different in comparison with the other two algorithms. This can be attributed to at least two possible factors. First, the total number of ensembles no longer corresponds to the number of processors in the MPI-LSQR cases, as explained in Section 5.3.1. The number of ensembles for the 284,000 grid simulation is typically around 10–15, instead of 24 for the LU and LSQR simulations. Similarly, the number of ensembles for the 880,000 grid simulation is typically around 20–25, instead of 48 for the LU and LSQR simulations. This change in the number of ensembles in the domain would have an impact on the conditional mass fractions, which would then alter the overall reaction rates in each CFD subdomain, leading to different temperature profiles. Second, the conditional mass fraction of the previous timestep that is used as the initial condition input is now a weighted average of the participating CFD subdomains in each ensemble instead of one CFD subdomain. While this change in the initial condition will not lead to a noticeable difference in the solution over one timestep, the difference would accumulate over thousands of timesteps. Despite these two relatively important modifications, the MPI-LSQR results are very comparable to the LU and LSQR results. In general, it appears that the number of ensembles and the exact layout of the ensembles do not significantly alter the converged solution. Similar




Figure 5.9: The radial temperature and species mass fraction profiles at six axial locations downstream of the inlet for the three grids are shown. The profiles at $z/D = 1$ are shown in black, then each successive jet diameter downstream is shown in a lighter shade of gray, up to $z/D = 6$. The jet diameter, D , is $1.12 \cdot 10^{-2}$ m. The symbols represent: LU-manual: ‘—’; LSQR-semi-dynamic: ‘ \triangle ’; PLSQR-dynamic: ‘ \bullet ’.

to the LU and LSQR results, it is observed that the CO mass fraction exhibits a local maximum near $r/D = 0.6$ with increasing grid resolution farther downstream from the burner nozzle.

The computational time for the nine simulations are tabulated in Table 5.4. The execution times indicate that the LSQR solver in conjunction with the Semi-Dynamic Ensemble Selection method can reduce the total computational cost by up to 50% in comparison with the conventional LU-decomposition method. Moreover, the MPI-LSQR solver with Dynamic Ensemble Selection can further reduce the computational times by another 2–22% in comparison with LSQR⁷. While the performance increase is relatively small by changing the implementation from the LSQR solver with the Semi-Dynamic Ensemble Selection to the MPI-LSQR solver with Dynamic Ensemble Selection, the primary advantage is to enable the CSE combustion module to continuously form more suitable ensembles as the flame evolves in shape and position within the domain. From these simulations and the standalone tests, it is conceivable that a parallelised version of LSQR along with the Dynamic Ensemble Selection algorithm could serve as a good candidate to replace the LU-decomposition with static ensembles for future versions of CSE for all modes of combustion.

⁷ The execution times for the 880,000 cell cases are averaged over two runs to eliminate possible network instability errors.

5.6 Conclusion

Grid	52,500	52,500	284,000	880,000
Sim. time (s)	0.0 → 0.05	0.05 → 0.25	0.05 → 0.25	0.05 → 0.25
Δt (s)	$8 \cdot 10^{-6}$	$8 \cdot 10^{-6}$	$4 \cdot 10^{-6}$	$2.5 \cdot 10^{-6}$
LU exec. (hrs)	1.7	11.9	88.8	302.0
Processors	4×3.8 GHz	4×2.1 GHz	24×2.1 GHz	48×2.2 GHz
LSQR exec. (hrs)	1.2	8.6	37.5	177.2
Processors	4×3.8 GHz	4×2.1 GHz	24×2.1 GHz	48×2.2 GHz
MPI-LSQR exec. (hrs)	0.98	7.4	35.2	173.6
Processors	4×3.8 GHz	4×2.1 GHz	24×2.1 GHz	48×2.2 GHz

Table 5.4: Summary of simulation and execution times. The execution times for LSQR with Semi-Dynamic Ensemble Selection are typically 50% of the LU-decomposition method for larger systems. The execution times for MPI-LSQR with Dynamic Ensemble Selection are further decreased by 2–22% in comparison with the LSQR procedure. The ensemble sizes were decreased for the coarse grid (52,500) case in order to maintain locality for cell information, leading to more comparable execution times.

5.6 Conclusion

A first study of Conditional Source-term Estimation using a modified version of the LSQR iterative solver along with Dynamic Ensemble Selection for a premixed, turbulent methane-air Bunsen burner has been performed. The LSQR solver has been modified to accept an initial condition as the input and parallelised for scalability. The standalone tests and LES results indicate that the iterative solver is fully capable of reproducing similar solutions to the LU-decomposition method while substantially reducing computational time. The data sharing algorithm between ensembles has also proven to be robust in achieving the proper communication directives for grids of three sizes, three different processor counts and different computational hardware. The integration of these two algorithms allows one to establish the framework for a highly autonomous and parallel version of CSE that can adapt to the changing flame distribution in complex geometries.

Chapter 6

LES of Turbulent Premixed Flames with Optimised CSE and LEM Submodels

The work in this chapter serves to unify the submodels and optimisations developed in Chapters 4 and 5. A turbulent, premixed, stoichiometric methane-air flame of a laboratory-scale Bunsen burner operating in the thin reaction zone has been simulated via large eddy simulation (LES). The optimised and parallelised version of Conditional Source-term Estimation (CSE) developed in Chapter 5 has been implemented; furthermore, the new pseudo-turbulent probability density function (PDF) and scalar dissipation rate (SDR) models formulated from the Linear-Eddy Model (LEM), as discussed in Chapter 4, are used to complete the CSE inversion and to close the transport equation for the variance of the progress variable, respectively.

6.1 Introduction

As described in the previous chapter, a number of simulation strategies has been developed to model the interaction between turbulence and flame [84, 86, 132, 134]. For Conditional Moment Closure (CMC) [84], the transport equations for the reactive species are conditionally averaged on a conserved scalar. The primary advantages of this strategy are the reduced spatial dependence of the conditional averages

and negligible fluctuations around the conditional averages, which considerably simplifies the reaction rate closure. In addition, detailed chemistry can be included at relatively low computational cost without the thin flame restriction encountered in flamelet models. However, one extra transport equation for the conditional mass fraction would need to be solved for *each* reactive species, increasing the dimensionality of the problem. Furthermore, there are unclosed terms associated with the CMC transport equations that require modelling. Bushe and Steiner [26] later proposed a technique based on the CMC known as the Conditional Source-term Estimation (CSE). Rather than solving extra transport equations, the conditional mass fractions are recovered through an inverse calculation. A detailed description of the CSE combustion model has been provided in Section 5.2. This model has been shown to produce promising results for diffusion flames [67, 93, 161, 183]. Further, Salehi *et al.* has demonstrated that CSE can also be applied to premixed flames and produce meaningful results [42, 143–145, 153]. Recent effort has extended CSE to lifted turbulent flames [43] and flames operating in the Moderate and Intense Low Oxygen Dilution (MILD) combustion mode [92].

In this chapter, the performance of an optimised version of the CSE combustion model designed for parallel computations over a cluster of processors via a large eddy simulation of a turbulent, premixed methane-air laboratory burner operating in the thin reaction zone [173] is evaluated. In addition, recently developed probability density function (PDF) and scalar dissipation rate (SDR) submodels required to complete the CSE model and the variance transport equation for the progress of reaction are implemented concurrently in a large eddy simulation combustion code for the first time. The pseudo-turbulent PDF and SDR submodels are formulated through the Linear-Eddy Model (LEM) [148, 157] and benefit from the inclusion of turbulence and correct chemistry in the framework compared to previous submodels; as a result, the LEM-based PDF model has demonstrated improved accuracy over previous modelling approaches [175]. In the following section, a detailed description of the underlying theory behind each numerical model used in this work will be provided. Next, the simulation setup for both the non-reactive and reactive flows will be described. This will be followed by the LES result and concluding remarks.

6.2 Theory

6.2.1 Governing Equations

The governing equations for the gaseous flow field is based on the standard OpenFOAM 2.1.1 package [185], which would be similar to the ones outlined in Section 2.2. The energy equation, however, is solved in the form of the total enthalpy. To integrate the CFD code with the CSE combustion module, two additional transport equations for the mean and variance of the reaction progress variable are required. The transport equation of the normalised mean of the progress variable takes the following form [132],

$$\frac{\partial(\bar{\rho}\bar{c})}{\partial t} + \frac{\partial(\bar{\rho}\tilde{u}_i\bar{c})}{\partial x_i} = \frac{\partial}{\partial x_i} \left(\bar{\rho} \frac{v_T}{Sc_t} \frac{\partial \bar{c}}{\partial x_i} \right) + \bar{\omega}_c, \quad (6.1)$$

where Sc_t is the turbulent Schmidt number¹. The mass fraction of CO_2 is used as the progress variable in this study and the normalisation is calculated in accordance with Equation 2.73. The transport equation of the variance of the progress variable with some simplification reads [132],

$$\frac{\partial(\bar{\rho}\widetilde{c'^2})}{\partial t} + \frac{\partial(\bar{\rho}\tilde{u}_i\widetilde{c'^2})}{\partial x_i} = \frac{\partial}{\partial x_i} \left(\bar{\rho} \frac{v_T}{Sc_t} \frac{\partial \widetilde{c'^2}}{\partial x_i} \right) + 2\bar{\rho} \frac{v_T}{Sc_t} \frac{\partial \bar{c}}{\partial x_i} \frac{\partial \bar{c}}{\partial x_i} - 2\bar{\rho}D_c \frac{\partial \bar{c}}{\partial x_i} \frac{\partial \bar{c}}{\partial x_i} + 2\bar{c}'\bar{\omega}_c, \quad (6.2)$$

where $\bar{\rho}D_c \frac{\partial \bar{c}}{\partial x_i} \frac{\partial \bar{c}}{\partial x_i} \approx \bar{\rho}\tilde{\chi}_c$ ²; $\tilde{\chi}_c$ is the scalar dissipation rate of the progress variable. An implicit, second order temporal scheme from the standard OpenFOAM package [185], designated as *backward*, is implemented in the calculation. The LES subgrid model is known as the *oneEqEddy* in the framework of OpenFOAM, which is equivalent to the *k*-equation eddy-viscosity model.

6.2.2 Chemistry Reduction

To maintain reasonable computational times for combustion simulations, several notable techniques have been developed in the past decades to simplify the reac-

¹ Sc_t is set to 0.7 unless otherwise specified.

²In some literature, the factor of 2 could be included in the definition of the scalar dissipation rate.

tion kinetics. These methods typically involve the use of quasi steady-state and partial-equilibrium assumptions [32, 125] to reduce the complexity of the systems. However, low-dimensional manifold approaches, first proposed by Maas and Pope [104], are particularly suited for complex simulations as they retain the features of detailed mechanisms while significantly reducing the overall computational cost. For this study, we have used a manifold approach referred to as the Flamelet-Generated Manifold (FGM) [121]. The particular manifold we are using is tabulated as functions of two progress variables, the normalised CO_2 and H_2O mass fractions; it is thus referred to as a two-dimensional FGM, or 2D-FGM. The 2D-FGM approach has been successfully coupled with CSE in the past [153].

For premixed combustion involving short hydrocarbon chains, it is generally acceptable to use a one-dimensional, unstrained laminar flame to construct the model manifold. Each one-dimensional flame is provided with a realisable unburnt fuel-air mixture; the solution to the flame is tabulated in the manifold with the end point represented by the combustion products. The FGM in this study is tabulated as functions of two progress variables, the normalised CO_2 and H_2O mass fractions. During the LES calculation, the mass fractions and production rates of all major species can then be retrieved from the FGM if the two progress variables are provided. It should be mentioned that FGM would typically encounter difficulties in predicting slow chemical species, such as NO_x and, to a lesser degree, CO . This is caused by the observation that the true, instantaneous reaction rates of these slow species are not always close to the tabulated manifold. As the primary objective in this study is to determine the overall performance and differences between the new ensemble selection and inversion algorithms, the FGM model is suitable for the task.

6.2.3 Presumed PDF and SDR models

A number of different presumed PDF models have been previously investigated for premixed combustion. The often used β -PDF does recover the extreme properties expected of the true PDF, such as δ -functions at the zero and unity extremes of reaction progress for maximal variance and single δ -functions at the mean for zero variance. However, it fails to reproduce the shape of the true PDF in more general

cases of premixed flames [70]. The issue is related to the fact that the shape of the true PDF appears to be a function of how the chemical reaction rates vary as a function of the progress variable; hence, different chemical kinetics lead to different shapes of the PDF of progress variable. To account for the effects of chemistry on the shape of the PDF, Bray *et al.* [22] proposed using a premixed laminar flame to model the functional dependence of the PDF on progress variable. Their original formulation only provided coverage for flames with very high variance. Jin *et al.* [70] later proposed a modification of the Bray PDF that extends the original formulation to cover all possible mean and variance combinations. This is accomplished by truncating the PDF shape function as needed to match the mean and variance parameters. It was found to significantly improve the fit of the PDF to that extracted from DNS results. This PDF model by Jin *et al.*, referred to as the modified laminar flamelet PDF (MLF-PDF), is commonly implemented in conjunction with the CSE combustion model [70, 143–145, 153]. A shortcoming of this method is that, at the point of truncation, the model PDF has a sharp drop to zero, whilst the true PDF tends to be more rounded.

To address this issue, a one-dimensional turbulent method was proposed to take the place of the typical laminar flame calculation to tabulate pseudo-turbulent PDF models for RANS and LES closures [175]. The Linear-Eddy Model [148, 157], an inexpensive one dimensional stochastic mixing model, has demonstrated the ability to capture important effects from the interaction between chemistry and turbulence on the PDF distributions sufficiently well [175]. This PDF model, formulated by the pseudo-turbulent LEM calculation, is implemented in place of the conventional modified laminar flamelet PDF for the LES calculations in the current chapter.

A related quantity is the scalar dissipation rate, χ_c , which is an important parameter to both premixed and non-premixed combustion modelling [45, 181]. For premixed combustion, the SDR, χ_c , of a species-based reaction progress variable is,

$$\chi_c(T, \phi) = D_c(T, \phi) |\nabla c| \cdot |\nabla c|, \quad (6.3)$$

where $D_c(T, \phi)$ and ∇c are the diffusivity of the selected species at the local temperature and equivalence ratio and the gradient of the progress variable, respectively. The SDR appears as an unclosed term in the transport equation for the vari-

ance of progress variable, where χ_c directly measures the decay rate of fluctuations through turbulent micromixing [181]. Since the burning rate of many combustion processes depends on the contact area and local gradient between the reactants, it is reasonable for most combustion models to assume that the mean burning rate of the flames either explicitly or implicitly depends on the scalar dissipation rate. For example, Conditional Moment Closure uses the scalar dissipation rate conditioned by the progress variable to calculate micromixing [181]; not surprisingly, modelling the conditional scalar dissipation term conditioned on the local and global progress of reaction, $\chi_c|c^*$, emerges as one of the main difficulties in applying CMC to turbulent premixed flames [4].

In common practice, the SDR can be estimated through additional algebraic modelling [40] or by including an extra transport equation related to the flame surface density as proposed in the Coherent Flame Model (CFM) [132]. However, a fundamental relationship can be established between the conditional scalar dissipation rate ($\overline{\chi_c|c^*}$) and the PDF of the reaction progress variable ($P(c^*)$) [4],

$$\overline{\chi_c} = \int_0^1 \overline{(\chi_c|c^*)} P(c^*) dc^*, \quad (6.4)$$

where $\overline{\chi_c}$ is the unconditional mean scalar dissipation rate, or simply the mean scalar dissipation rate. Through the Linear-Eddy Model, it is possible to obtain pseudo-turbulent flame data that can directly compute both the $\overline{\chi_c|c^*}$ and the $P(c^*)$, leading to the quantity required to close the variance transport equation, $\overline{\chi_c}$. This single method can create both the PDF and SDR submodels, which are consistent with each other. Furthermore, this formulation of the SDR model has the added benefit of pre-tabulation, which would reduce the computational time for the LES or RANS calculations. A detailed description of the methodology used in constructing the PDF and SDR submodels can be found in [175]. This LEM SDR model is implemented in place of the conventional algebraic and transport equation methods in the current chapter.

6.3 Numerical Methods / Simulation Setup

6.3.1 Pre-Process

The 2D-FGM, PDF and SDR submodels require pre-tabulation prior to the large eddy simulation of the reactive flow. The temperature, pressure and equivalence ratio corresponding to the experimental condition for the methane-air burner are respectively, 300 K, 101.325 kPa and 1.0. All of the submodels are specific to the chemistry and thermodynamics conditions; thus, the temperature, pressure and equivalence ratio adhere to the quantities used in the experiments. The 2D-FGM is constructed with the open source chemical kinetics and thermodynamics software, Cantera [58], using the GRI-MECH 3.0 mechanism [156]. A detailed description of the 2D-FGM chemistry reduction technique can be found here [142]. Approximately five to ten hours are required to construct each 2D-FGM table with a single i7-2600K processor; however, the procedure can be readily parallelised if shorter convergence time is desired as the unstrained laminar flame solutions with different realisable initial conditions are independent of one another.

In addition to the thermo-chemical composition, the current PDF and SDR submodels are dependent on the turbulent intensity. As such, the overall turbulent fluctuation is approximately matched to the experimental condition (1.89 m/s) in the LEM calculations. Following the procedure from previous work [175], approximately 10,000 pseudo-turbulent flame profiles are collected at 1/100 large eddy turnover time intervals at the specified equivalence ratio. These profiles are subsequently used to construct the converged PDF and SDR submodels implemented as the input to the CSE combustion model and the variance transport equation. The computational time required to accumulate the 10,000 flame profiles and constructing the PDF and SDR submodels is approximately 48 hours on a single i7-2600K processor using a six-step global mechanism.

6.3.2 Non-Reactive Simulation

A premixed, stoichiometric methane-air flame in the thin reaction zone has been simulated. Similar to the experiments, the turbulent fluctuation (u') is first determined via the cold flow. This is possible as it has been found that the total

turbulence intensity value along the axial direction on the centerline of this burner remains relatively constant between reacting and non-reacting conditions [33].

For the experimental burner, a turbulence grid is fitted a distance, h_g , upstream of the nozzle. The value of h_g for the current flame is 44.5 mm. Figure 6.1 illustrates the schematic of the axisymmetric burner nozzle, while a detailed description of the burner assembly can be found in [191]. It should be mentioned that the pilot has been turned off during the experimental cold flow measurements; the simulations also adhere to this implementation for consistency. The relevant parameters characterising the experimental conditions in this study are summarised in Table 6.1.




Figure 6.1: Burner cross-section.

U_B (m/s)	u' (m/s)	Λ_L (mm)	Re_t	Ka	Da
21.0	1.89	2.64	316	2.6	9.6

Table 6.1: Summary of experimental conditions. Symbols: U_B = bulk flow velocity in the central jet; u' = total turbulence intensity; Λ_L = longitudinal integral length scale; Re_t = turbulent Reynolds number; Ka = turbulent Karlovitz number; Da = turbulent Damköhler number. The values are retrieved from [173].

To simulate the grid generated turbulence achieved in the experiments, a spectral inlet condition assuming isotropic homogeneous turbulence using the energy

spectrum of Haworth and Poinso [65] has been implemented. The time correlation for the velocity fluctuation is based on the method developed by Billson [15]. The time correlation parameters, denoted a and b , represent the fraction of the current velocity field, a , to be added to the newly generated synthetic fluctuations, b . They are computed by the following relations: $a = \exp(-\Delta t / \tau_t)$ and $b = \sqrt{1 - a^2}$, where Δt and τ_t are respectively the timestep size and the integral time scale. The parameters for the spectral inlet condition required to generate the turbulent intensity for the cold flow observed in the experiment are determined on the basis of trial and error. The values are summarised in Table 6.2.

Following the procedure of the experiments, the reported value of u' (1.89 m/s) is averaged within the cylindrical volume of $-0.25 < r/D < 0.25$ and $0.2 < h/D < 0.5$, where r and h are respectively the radius and the axial distance downstream of the burner nozzle. The overall simulation domain length, diameter and hexahedral cell count for the cold flow are 100 mm, 60 mm and 2.25 million cells, respectively. The lower boundary of the simulation domain is aligned with location of the turbulence grid.

h_g (mm)	K (m ² /s ²)	Max. κ (m ⁻¹)	Λ (mm)	a	b	Δt (s)
44.5	570	$2.10 \cdot 10^4$	2.64	0.985	0.171	$2 \cdot 10^{-6}$

Table 6.2: Relevant synthetic turbulence parameters. Symbols: h_g = grid location; K = turbulent kinetic energy; Λ = integral length; a and b = temporal correlation parameters; Δt = timestep size.

6.3.3 Reactive Simulation

With the appropriate turbulent parameters determined from the cold flow simulation, the reactive flow is initialised with premixed reactants in the central jet and stoichiometric methane-air combustion products in the pilot stream. The central jet bulk flow velocity is as described in Table 6.1, while the pilot bulk flow velocity is prescribed a magnitude of 1.0 m/s at the lower boundary of the concentric pilot pipe (Figure 6.1). The overall simulation domain length, diameter and hexahedral cell count for the reactive flow are 178.9 mm, 60 mm and 3.16 million cells, respectively. The simulation is performed with 96 2.2 GHz processors. Instantaneous

temperature profiles are collected after ten characteristic flame height flow-through times from initialisation for 30 ms at 0.2 ms intervals.

6.4 Results and Discussion

The temperature field for the current flame is measured via Mie scattering³, whereby the progress of reaction from the experiments is computed via binarised images depending on light intensity [173]. The light intensity binarisation threshold due to the evaporation of the oil particles corresponds to approximately 700 K; the progress of reaction is assigned a value of 0 for temperatures below 700 K and 1 otherwise. These binarised images are subsequently averaged and the iso-contour of 0.5 is used to determine the characteristic flame height; this approach has also been adopted for the analysis of the LES temperature field.

From Figure 6.2, it can be observed that the overall flame height and temperature distribution are well matched between the simulation and the experiment. The slight difference in the width of the flame could be caused by the synthetic turbulence inlet condition at the location of the turbulence grid. Improvements to the central jet velocity profile at the synthetic turbulence inlet can be made, as the velocity profile for the experimental burner exhibits a parabolic shape which is not fully captured by the current implementation of the boundary condition. The parabolic shape leads to a slightly slower transition of the progress variable from zero to one along the centerline of the jet and faster transition of the progress variable from zero to one along the outer rim of the jet, observed in the right half of Figure 6.2. The difference in overall behavior of the flame, however, may not be significant as the turbulence intensity has been carefully calibrated from the cold flow simulation.

The accurate prediction of the overall flame distribution and height can be attributed to a number of improvements made in the current LES calculation in comparison with previous CSE implementations for this laboratory burner, which operated in slightly different turbulent conditions [153]. In terms of the CFD module, the synthetic turbulent boundary condition is applied at the base of the premixed

³The experimental data is provided by Mr. P. Tamadonfar and Dr. Ö. Gülder of the University of Toronto [173].




Figure 6.2: Average binarised temperature-based progress variable profile and characteristic flame height from the LES result (left) and experimental measurements (right), where $D = 11.2$ mm. The solid line represents the half-burning surface, at which the average progress variable takes on a value of 0.5.

gas pipe, which is located 44.5 mm from the burner nozzle (Figure 6.1); the inclusion of the premixed gas pipe in the simulation allows the turbulent flow field to develop under more realistic conditions before interacting with the premixed flame front. Additionally, the synthetic turbulent intensity is carefully calibrated in the non-reactive simulation to match the experimental measurement prior to the reactive calculation.

For the CSE module, we have implemented a recently developed dynamic ensemble selection algorithm [177] and a parallel iterative solver in a detailed LES calculation for the first time. The dynamic ensemble selection algorithm increases the flexibility of the CSE module as the physical spaces in which the conditional averages are computed are automatically adjusted to the changing flame positions throughout the simulation. Such a feature ensures that an optimal number of reactive cells is provided for the inversion of Equation 5.4 for all of the ensembles, leading to increased load balance and more accurately inverted conditional mass fractions. Furthermore, this automatic procedure reduces the complexity of implementation from the perspective of the user in two notable ways: i) substantial fore-knowledge of the flame distribution and flow is no longer required as the ensembles can adapt to these changes as the simulation progresses; and ii) the number of processors used in the LES computation can now be optimised only according to the CFD module, without having to account for the manual division of physical spaces associated with the ensembles in the CSE module. In addition, the implementation of the parallel iterative solver, parallel-LSQR, contributes to a reduction in computational time by approximately 50%, while maintaining accuracy in the inverse solution as compared to conventional CSE approaches [177].

The LEM-formulated PDF model, which has been rigorously compared to experimentally measured PDFs in a rudimentary sense [176], is implemented for the first time in a LES of a turbulent premixed flame. This LEM PDF model has demonstrated improved accuracy compared to its predecessor, the MLF-PDF [175], which could account for part of the improvement in the prediction of the flame distribution and height observed in this study. Similarly, the LEM-formulated SDR model is implemented for the first time in a LES calculation; it has been shown that the LEM SDR matches experimental results satisfactorily for non-swirling flames [176]. The SDR term is embedded in the transport equation of the variance of the progress variable; improvements to this parameter can lead to a more accurate representation of the PDF model in each control volume.

The computational time required to achieve the current result is documented in Table 6.3. However, the times are expected to vary depending on the machine and communication characteristics between nodes. In general, one can anticipate a converged solution for similar Bunsen-type burners using the simulation strategies

outlined in this work within 200–300 hours of computation using 96 2.2 GHz processors. Indeed, increasing the number of processors would also increase the rate of convergence, though scalability has not been thoroughly studied.

Processors	Cell count	Simulation time (s)	Run time (hr)
$96 \times 2.2\text{-GHz}$	$3.16 \cdot 10^6$	0.030	225

Table 6.3: Summary of simulation times.

6.5 Concluding Remarks

A turbulent, premixed, stoichiometric methane-air laboratory Bunsen flame operating in the thin reaction zone has been simulated with a recently developed version of the Conditional Source-term Estimation combustion model designed for parallel computational environments. This version of CSE, in conjunction with newly formulated pseudo-turbulent PDF and SDR models, are implemented concurrently in a LES calculation for the first time. Comparison with experimental data suggests this combination of models can accurately predict the flame height and overall temperature distribution for our test case. Moreover, this implementation of CSE can produce converged solutions with rather modest computational effort.

Each of the implemented models, the CSE, 2D-FGM, LEM PDF and LEM SDR, is self-contained; in principle, it is possible to simulate any other turbulent premixed flame with the same thermo-chemical composition and similar turbulent intensity without the need to readjust any model parameters. The approach outlined in this work allows for progression towards a fully dynamic, predictive CFD solver for turbulent premixed combustion.

Chapter 7

Conclusion and Future Work

7.1 Conclusions

The objective of this thesis is to advance numerical simulation strategies for turbulent premixed combustion in terms of accuracy, autonomy and efficiency. The first half of this work is focused on the development of new presumed PDF and SDR models formulated by the Linear-Eddy Model for the purposes of subsequent LES and RANS simulations. The LEM-based PDF and SDR models, tabulated as a function of the mean and variance of the reaction progress variable are able to account for the basic features of turbulence-chemistry interactions, a characteristic which all of the previous models lack. These new models are then compared to experiments in detail, through a range of turbulent intensities and swirl conditions for validation. The second half of this work is focused on the optimisation and parallelisation of a combustion model for LES and RANS, the Conditional Source-term Estimation. The methodologies developed herein can be straightforwardly extended to partially-premixed and non-premixed combustion. As the final study of this thesis, the new LEM-based PDF and SDR models are combined with the redesigned CSE module in a large eddy simulation of a premixed, methane-air laboratory flame operating in the thin reaction zone.

This work began with the observation that a single PDF model, the β -PDF, can be applied to non-premixed combustion model based on conditional moment approaches; however, there does not appear to be a strong consensus on the most

suitable PDF model for turbulent premixed combustion simulation strategies. The β -PDF, Bray PDF and modified laminar flamelet PDF have been previously studied in the context of premixed, methane-air flames; each of these models has demonstrated some form of deficiency in capturing the correct PDF shape. For the purpose of CSE calculations, it is critical that the PDF distribution is modelled as accurately as computationally allowable, since the PDF is implemented twice for each reaction rate calculation — once to obtain the conditional mass fractions and once to obtain the source terms — at every timestep.

Presented with this challenge, it would be appropriate to develop a computationally inexpensive methodology for the construction of a more accurate PDF model for premixed combustion. It has been found that the Linear-Eddy Model, a one-dimensional stochastic mixing model, can reproduce turbulent statistics with relatively low computational cost. As a result, this model is selected as a suitable candidate to construct a PDF model that includes turbulence in its formulation. The one-dimensional premixed flame simulation code, including a suitable Linear-Eddy Model stochastic mixing module, is developed in this work. The resultant LEM-based PDFs show greater resemblance to the DNS and experimental measurements in comparison with existing presumed PDF models. Concurrently, the LEM data used to construct the PDF model can be applied to formulate a pseudo-turbulent conditional scalar dissipation rate model. The combination of these two models would allow for the implementation of consistent LEM-based PDF and SDR models in the subsequent LES calculations. It is shown that the LEM SDR is able to predict the decrease in the magnitude of these distributions caused by flame thickening due to turbulent effects in premixed combustion, albeit over-predicting the experimental results in peak magnitudes.

The PDF and SDR models formulated in this work are suitable for a number of simulation strategies for turbulent premixed combustion, including Presumed Conditional Moment, Conditional Moment Closure and Conditional Source-term Estimation. The LEM-based SDR model is particularly useful for the CMC approach as the lack of accurate SDR models is limiting the general applicability of CMC in this combustion mode. The successful development of a one-dimensional pseudo-turbulent flame code, LEM PDF model and LEM SDR model motivated the subsequent topic of research, which is to extend the current study from a purely

practical basis — that is, to simply provide better submodels for advanced combustion models — to a theoretical investigation, where perhaps it would be possible to understand the underlying characteristics of the PDF and SDR distributions as a function of turbulent fluctuations and swirling conditions.

The Cambridge swirl burner, an experimental bluff body stabilised burner, is selected as the basis of validation of the new LEM-based PDF and SDR models. A broad range of turbulent conditions, including changes in intensities and swirl numbers, are simulated and compared with the experimental results. The algorithms for the construction of the LEM PDF and LEM SDR models are also parallelised and optimised at this stage.

The LEM PDF model demonstrates good agreement with the experimental results in terms of nominal changes with turbulence. Particularly, it is able to capture the overall shape and the effect of smoothing towards the left and right boundaries of the PDF distributions — a feature that cannot be reproduced by any of the previous PDF models. Further, the LEM results suggest that the conditional SDR would decrease towards the upper regions of the Borghi diagram and would increase in peak magnitudes towards the unstrained laminar limit with increasing integral scales. The experimental conditional SDRs for all axial locations of the flame brush show a greater decrease in magnitude than predicted by the numerical model, though this could be partially caused by an interaction between shearing and mixing layers within a specific region of the swirl flames, leading to greater dissipation downstream of this location. The values for the unconditional mean SDR predicted by LEM are typically within 50% of the experimental values. Moreover, the accuracy of the numerical model seem to be dependent on the swirl number of the flow.

Overall, the results suggest that the PDF tabulated as a function of the mean and variance of the reaction progress variable for premixed flames remains relatively steady under a variety of turbulent conditions, including changes in the integral length, velocity fluctuation and swirl number. From a practical perspective, this finding is remarkable as it implies that one can use a representative pseudo-turbulent PDF model for a range of turbulent conditions in subsequent RANS or LES calculations, which greatly simplifies the complexity of implementation with the corresponding combustion models. The validation with the Cambridge swirl

burner data firmly establishes that LEM can be used to formulate a high-fidelity PDF model for turbulent premixed combustion.

Swirl appears to have greater influence on the local flame gradients; this, in turn, leads to more substantial changes in the SDR, which the LEM cannot fully capture in its current form. However, the magnitude of the SDR as calculated by LEM is, in the most extreme cases, over-predicted by a factor of two in comparison with the experimental flames in this study. Moreover, the shape of the conditional SDR distribution is well-captured with the LEM methodology. These characteristics imply: (i) that it could be possible to apply an empirical correction for the conditional SDR as a simple function of the swirl number and turbulent fluctuation; and (ii) the LEM-formulated SDR is a competitive candidate amongst the available SDR models for turbulent premixed combustion modelling.

As a result of this study, an accurate, pseudo-turbulent PDF model suitable for subsequent RANS and LES calculations has been formulated. Further, a companion SDR model that is consistent with the LEM PDF has been also been established. It is hoped that these new models would not only serve to demonstrate the theoretical findings described above, but would also be useful to a number of the simulation strategies for turbulent premixed combustion. For the reason that these models are pre-tabulated prior to the detailed flame simulation, they are ideally suited for large-scale, complex problems by focusing the computational resources to solving the fluid dynamics calculation rather than the submodel parameters.

The next part of this thesis work is dedicated to the parallelisation and optimisation of the Conditional Source-term Estimation combustion model for computations in large clusters. Prior to this research, the CSE model had several implementational obstacles that limited its performance in largely parallel computational environments; these include the pre-allocation of static ensembles and single-processor inversion calculations. In reference to the former, it would be difficult for the user to achieve load balance during the simulation as the flame distribution evolves dynamically within the domain. Further, the exact division of the ensembles is entirely dependent on the foreknowledge of the flame morphology by the user; this leads to the question of completeness of the combustion model as different ensemble divisions could potentially lead to different solutions. Another problem is that the number of reactive cells may not be fully satisfied if the

ensembles are not carefully divided according to the available processors.

To address these shortcomings, two optimisations to enhance the capability of CSE are introduced. The first optimisation is to increase the flexibility of the CSE model by dynamically allocating the physical space required for the individual ensemble averages. The second optimisation is to integrate a well-established iterative solver, Least-Squares with QR-factorisation, to tackle the CSE inversion problem. As the purpose of this work is to prepare the CSE module for cluster operation, the LSQR solver is also parallelised to ensure scalability of the combustion code. These optimisations can be applied to both RANS and LES paradigms in the premixed, partially-premixed and non-premixed combustion modes.

The new implementations are subsequently validated with several generic LES calculations of a turbulent, premixed methane-air Bunsen burner. The standalone inversion tests and LES results indicate that the parallelised LSQR solver is fully capable of reproducing similar solutions to the conventional CSE methods while substantially reducing computational time. The data sharing algorithm between ensembles also proves to be robust in achieving the proper communication directives for three sets of grids, three different processor counts and three different sets of computational hardware. The integration of the dynamic ensemble selection and parallel LSQR algorithms allows one to establish the framework for a highly autonomous and parallel version of CSE that can adapt to the changing flame distribution in complex geometries.

As combustion modelling for turbulent premixed flames is currently a subject of significant research because of the intricate interactions between turbulence and chemistry, the newly optimised version of CSE demonstrates itself to be a highly practical approach in several notable ways. First, the computational times have been substantially reduced in comparison with previous implementations, by approximately 50% in most LES cases. Second, this variant can be considered more *complete* as the need for manual adjustment of the ensembles has been eliminated. Third, the CSE combustion model is more accessible to the average user as the ensembles are automatically constructed to best suit the flame position and reactivity within the simulation domain at any given time step — substantial foreknowledge of the flame morphology is no longer required for the division of the ensembles. Another benefit to the optimisations described herein is that the CSE model is

7.1 Conclusions

designed for RANS or LES calculations under all ranges of air-fuel ratios and turbulent intensities; this implies that the algorithms developed is applicable to all of the combustion modes within the framework of both RANS and LES. In consequence, the current variant of CSE is not only more suitable for academic purposes of validation, but is also more suitable for industrial problems from a practical perspective.

In the final study of this thesis work, a detailed LES calculation of a premixed methane-air laboratory-scale burner operating in the thin reaction zone is performed; this simulation neatly connects the LEM-formulated PDF and SDR models with the redesigned CSE combustion module. Comparison with experimental data, measured by Mie scattering, suggests this combination of models can accurately predict the flame height and overall distribution. Moreover, this implementation of CSE can produce converged solutions with rather modest computational effort. Each of the implemented models, the CSE, FGM, LEM PDF and LEM SDR, is self-contained; in principle, it is possible to simulate any other turbulent premixed flame with the same thermo-chemical composition and similar turbulent intensity without the need to adjust any modelling parameters.

This LES of the laboratory-scale burner is an appropriate terminus for the current thesis as it serves to demonstrate the accomplishments of all of the previous chapters. The pseudo-turbulent PDF and SDR models presented in Chapters 3 and 4 are fully integrated into the optimised CSE combustion model of Chapter 5 for a detailed simulation. Each chapter represents an important advancement in the modelling of turbulent premixed combustion, in terms of accuracy and efficiency. The summation of approaches outlined herein allows for progression towards a fully dynamic, predictive CFD solver for turbulent premixed combustion devoid of any tunable parameters. With continued research, it is hoped that these methods can eventually be implemented in industrially relevant LES calculations to accurately predict the formation of minor species and the thermodynamic characteristics of the next generation of combustion devices. Furthermore, perhaps the outlined approaches can be refined in the future to a level where they can provide a more fundamental understanding of the nature of the turbulence-chemistry interaction with high-fidelity large eddy simulations.

7.2 Future Work

The first half of this research is concerned with the development of more accurate PDF and SDR models for turbulent premixed combustion. While it appears that the LEM PDF can replicate the measured PDF to a high degree of precision for a broad range of turbulence conditions, the same cannot be said of the LEM SDR model. It would be of importance to understand the underlying mechanisms that decrease the magnitude of the SDR with increasing swirl intensity. Further, it may be possible to formulate an empirical correlation between the magnitude of the SDR and the various turbulent characteristics, such as swirl and turbulent fluctuations.

The second half of this research pertains to further developing the Conditional Source-term Estimation combustion model for turbulent premixed flames. Having optimised a version of CSE that is capable of predicting basic flame attributes from a Bunsen-type laboratory burner with very modest computational resources, it would be interesting to examine other burners, fuels and combustion modes.

A particularly attractive set of experiments for simulation purposes is produced by the Cambridge swirl burner [167]. The temperature and species mass fractions are measured accurately and well-documented [168, 170–172, 194]. Further, the burner geometry and boundary conditions are well-characterised and suitable for simulations. It would be interesting to use the new formulation of CSE to simulate this burner and to study the swirling effects. The ability to simulate swirl burners is fundamental to combustion models designed for industrial applications as modelling gas turbine combustors is one of their most important functions.

Another direction of research would be to test different fuels with the CSE combustion model. This study has focused on methane-air combustion as it is the simplest of hydrocarbon fuels. More complex hydrocarbons and alternative fuel sources, such as biofuels, could be interesting candidates due to their increasing popularity as fuels for a new generation of engines and slower chemical time scales. The extended chemical times allows for greater turbulent mixing within the reactive flow, leading to more substantial chemistry-turbulence interactions. This extended period allows for more thorough validation of the combustion model, provided detailed experimental data is available. In addition, perhaps it would be feasible to implement different chemistry reduction techniques in conjunction with CSE, such

7.2 Future Work

as skeletal mechanisms or other reduced manifold approaches for such complex fuels.

In reality, combustion rarely occurs in a purely premixed or non-premixed mode; it is, thus, imperative for combustion models to perform equally well in the partially-premixed mode. The primary challenge for CSE here is that the inversion calculations are substantially more expensive. However, the newly developed dynamic ensemble selection algorithms and the newly integrated parallel iterative solver from this thesis work should alleviate a considerable part of the computational load. Thus, another possible research direction is the optimisation of the PLSQR and ensemble selection algorithms for the partially-premixed combustion mode in the context of CSE.

Bibliography

- [1] *Large eddy simulation of combustion instabilities in turbulent premixed burners*, 1997. SAND86-8841. → pages 39
- [2] Key World Energy Statistics 2015. <https://www.iea.org/publications/freepublications/publication/key-world-energy-statistics-2015.html>, 2015. Accessed: 2016-01-14. → pages 1
- [3] R. Abdel-Gayed, D. Bradley, and F.-K. Lung. Combustion regimes and the straining of turbulent premixed flames. *Combustion and Flame*, 76(2): 213–218, 1989. → pages 23
- [4] S. Amzin, N. Swaminathan, J. W. Rogerson, and J. H. Kent. Conditional Moment Closure for turbulent premixed flames. *Combustion Science and Technology*, 184(10-11):1743–1767, 2012. → pages 41, 58, 68, 86, 131
- [5] G. K. Batchelor. *The Theory of Homogeneous Turbulence*. Cambridge University Press, 1982. → pages 19
- [6] B. S. Baldwin and H. Lomax. Thin-layer approximation and algebraic model for separated turbulent flows. In *Paper 91-0610, AIAA*, 1978. → pages 30
- [7] J. Bardina, J. Ferziger, and W. Reynolds. chapter Improved subgrid-scale models for large eddy simulation. Fluid Dynamics and Co-located Conferences. American Institute of Aeronautics and Astronautics, Jul 1980. → pages 33
- [8] R. Barlow, G.-H. Wang, P. Anselmo-Filho, M. Sweeney, and S. Hochgreb. Application of Raman/Rayleigh/LIF diagnostics in turbulent stratified flames. *Proceedings of the Combustion Institute*, 32(1):945–953, 2009. → pages 59

- [9] F. Bastin, P. Lafon, and S. Candel. Computation of jet mixing noise due to coherent structures: the plane jet case. *Journal of Fluid Mechanics*, 335: 261–304, 3 1997. → pages 30
- [10] A. Belcadi, M. Assou, E. Affad, and E. Chatri. Construction of a reduced mechanism for modelling premixed combustion of methane-air. *Combustion Theory and Modelling*, 11(4):603–613, 2007. → pages 14
- [11] A. Belcadi, M. Assou, E. Affad, and E. Chatri. CH₄/NO_x reduced mechanisms used for modeling premixed combustion. *Energy and Power Engineering*, 4(4):264–273, 2012. → pages 14
- [12] R. Bilger, S. Pope, K. Bray, and J. Driscoll. Paradigms in turbulent combustion research. *Proceedings of the Combustion Institute*, 30(1): 21–42, 2005. → pages 42
- [13] R. W. Bilger. Conditional Moment Closure for turbulent reacting flow. *Physics of Fluids A*, 5(2):436–444, 1993. → pages 41
- [14] R. W. Bilger, S. B. Pope, K. N. C. Bray, and J. F. Driscoll. Paradigms in turbulent combustion research. *Proceedings of the Combustion Institute*, 30:21–42, 2005. → pages 38
- [15] M. Billson, L.-E. Eriksson, and L. Davidson. Jet noise prediction using stochastic turbulence modeling. In *AIAA paper 2003-3282, 9th AIAA/CEAS Aeroacoustics Conference*, 2003. → pages x, 115, 116, 134
- [16] R. B. Bird, W. E. Stewart, and E. N. Lightfoot. Transport phenomena. *AIChE Journal*, 7(2):5J–6J, 1961. → pages 9
- [17] M. Boger, D. Veynante, H. Boughanem, and A. Trouvé. Direct numerical simulation analysis of flame surface density concept for large eddy simulation of turbulent premixed combustion. *Symposium (International) on Combustion*, 27(1):917–925, 1998. → pages 37
- [18] M. Boileau, G. Staffelbach, B. Cuenot, T. Poinsot, and C. Bérat. LES of an ignition sequence in a gas turbine engine. *Combustion and Flame*, 154(12): 2–22, 2008. → pages 39
- [19] R. Borghi. *Recent Advances in the Aerospace Sciences: In Honor of Luigi Crocco on His Seventy-fifth Birthday*, chapter On the Structure and Morphology of Turbulent Premixed Flames, pages 117–138. Springer US, Boston, MA, 1985. ISBN 978-1-4684-4298-4. → pages 23

- [20] K. Bray and J. Moss. A unified statistical model of the premixed turbulent flame. *Acta Astronautica*, 4(34):291–319, 1977. → pages 35, 36
- [21] K. Bray, P. A. Libby, and J. Moss. Unified modeling approach for premixed turbulent combustion—part I: General formulation. *Combustion and Flame*, 61(1):87–102, 1985. → pages 35
- [22] K. Bray, M. Champion, P. Libby, and N. Swaminathan. Finite rate chemistry and presumed PDF models for premixed turbulent combustion. *Combustion and Flame*, 146(4):665–673, 2006. → pages 45, 46, 66, 130
- [23] K. N. C. Bray. *Turbulent Reacting Flows*, chapter Turbulent flows with premixed reactants, pages 115–183. Springer Berlin Heidelberg, Berlin, Heidelberg, 1980. ISBN 978-3-540-38273-7. → pages 23
- [24] K. N. C. Bray, M. Champion, and P. A. Libby. *Turbulent Reactive Flows*, chapter The Interaction Between Turbulence and Chemistry in Premixed Turbulent Flames, pages 541–563. Springer US, New York, NY, 1989. ISBN 978-1-4613-9631-4. → pages 35
- [25] A. Burcat. Thermochemical data for combustion calculations. In *Combustion Chemistry*, pages 455–473. Springer US, 1984. → pages 11
- [26] W. K. Bushe and H. Steiner. Conditional Moment Closure for large eddy simulation of non-premixed turbulent reacting flows. *Physics of Fluids*, 11(7):1896–1906, 1999. → pages 42, 43, 98, 100, 127
- [27] T. Butler and P. O’Rourke. A numerical method for two dimensional unsteady reacting flows. *Symposium (International) on Combustion*, 16(1): 1503–1515, 1977. → pages 38
- [28] S. Cannon, B. Brewster, and L. Smoot. PDF modeling of lean premixed combustion using in situ tabulated chemistry. *Combustion and Flame*, 119(3):233–252, 1999. → pages 41
- [29] W. Chang and J. Chen. Reduced mechanisms for premixed and non-premixed combustion, 1999. URL <http://firebrand.me.berkeley.edu/griredu.html>. Accessed: 2016-01-14. → pages 14, 51, 71
- [30] F. Charlette, C. Meneveau, and D. Veynante. A power-law flame wrinkling model for LES of premixed turbulent combustion part I: non-dynamic formulation and initial tests. *Combustion and Flame*, 131(12):159–180, 2002. → pages 37

- [31] F. Charlette, C. Meneveau, and D. Veynante. A power-law flame wrinkling model for LES of premixed turbulent combustion part II: dynamic formulation. *Combustion and Flame*, 131(12):181–197, 2002. → pages 37
- [32] J. Y. Chen. A general procedure for constructing reduced reaction mechanisms with given independent relations. *Combustion Science and Technology*, 57(1-3):89–94, 1988. → pages 13, 129
- [33] Y.-C. Chen, N. Peters, G. Schneemann, N. Wruck, U. Renz, and M. S. Mansour. The detailed flame structure of highly stretched turbulent premixed methane-air flames. *Combustion and Flame*, 107(3):223–244, 1996. → pages 133
- [34] T. J. Chung. *Computational Fluid Dynamics*. Cambridge University Press, 2002. ISBN 0521594162. → pages 27
- [35] O. Colin, F. Ducros, D. Veynante, and T. Poinsot. A thickened flame model for large eddy simulations of turbulent premixed combustion. *Physics of Fluids*, 12(7):1843–1863, 2000. → pages 39
- [36] G. Damköhler. Der Einfluss der Turbulenz auf die Flammgeschwindigkeit in Gasgemischen. *Z. Electrochem*, 46:601–652, 1940. → pages 22, 23
- [37] J. W. Deardorff. A numerical study of three-dimensional turbulent channel flow at large Reynolds numbers. *Journal of Fluid Mechanics*, 41:453–480, 4 1970. → pages 31
- [38] J. W. Deardorff. The use of subgrid transport equations in a three-dimensional model of atmospheric turbulence. *ASME Transactions, Journal of Fluids Engineering*, 95:429–438, 1973. → pages 33
- [39] P. Domingo, L. Vervisch, S. Payet, and R. Hauguel. DNS of a premixed turbulent V flame and LES of a ducted flame using a FSD-PDF subgrid scale closure with FPI-tabulated chemistry. *Combustion and Flame*, 143(4):566–586, 2005. → pages 39, 44, 65, 74
- [40] P. Domingo, L. Vervisch, and D. Veynante. Large eddy simulation of a lifted methane jet flame in a vitiated coflow. *Combustion and Flame*, 152(3):415–432, 2008. → pages 40, 131
- [41] J. J. Dongarra, I. S. Duff, D. C. Sorensen, and H. A. van der Vorst. *Numerical Linear Algebra on High-Performance Computers*. Society for Industrial and Applied Mathematics, 1998. → pages 117

- [42] D. Dovizio, M. M. Salehi, and C. B. Devaud. RANS simulation of a turbulent premixed bluff body flame using Conditional Source-term Estimation. *Combustion Theory and Modelling*, 17(5):935–959, 2013. → pages 42, 98, 101, 127
- [43] D. Dovizio, J. W. Labahn, and C. B. Devaud. Doubly Conditional Source-term Estimation (DCSE) applied to a series of lifted turbulent jet flames in cold air. *Combustion and Flame*, 162(5):1976–1986, 2015. → pages 42, 98, 116, 127
- [44] J. F. Driscoll. Turbulent premixed combustion: Flamelet structure and its effect on turbulent burning velocities. *Progress in Energy and Combustion Science*, 34(1):91–134, 2008. → pages 38
- [45] J. Duclos, D. Veynante, and T. Poinso. A comparison of flamelet models for premixed turbulent combustion. *Combustion and Flame*, 95(1 - 2): 101–117, 1993. → pages 58, 68, 69, 130
- [46] M. Düsing, A. Sadiki, and J. Janicka. Towards a classification of models for the numerical simulation of premixed combustion based on a generalized regime diagram. *Combustion Theory and Modelling*, 10(1): 105–132, 2006. → pages 23
- [47] T. Echekki and E. Mastorakos. *Turbulent Combustion Modeling: Advances, New Trends and Perspectives*, volume 95. Springer Science, 2011. → pages 37, 41
- [48] C. L. Fefferman. Existence and smoothness of the Navier-Stokes equation. URL <http://www.claymath.org/millennium-problems/navier%E2%80%9393stokes-equation>. Accessed: 2016-01-15. → pages 18, 27
- [49] C. Fenimore. Formation of nitric oxide in premixed hydrocarbon flames. *Symposium (International) on Combustion*, 13(1):373–380, 1971. → pages 66
- [50] C. FUREBY. Large eddy simulation of combustion instabilities in a jet engine afterburner model. *Combustion Science and Technology*, 161(1): 213–243, 2000. → pages 37
- [51] C. Fureby. A fractal flame-wrinkling large eddy simulation model for premixed turbulent combustion. *Proceedings of the Combustion Institute*, 30(1):593–601, 2005. → pages 37

- [52] J. Galpin, A. Naudin, L. Vervisch, C. Angelberger, O. Colin, and P. Domingo. Large eddy simulation of a fuel-lean premixed turbulent swirl-burner. *Combustion and Flame*, 155(12):247–266, 2008. → pages 40
- [53] M. Germano, U. Piomelli, P. Moin, and W. H. Cabot. A dynamic subgrid-scale eddy viscosity model. *Physics of Fluids A*, 3(7):1760–1765, 1991. → pages 33
- [54] S. Ghosal and P. Moin. The basic equations for the large eddy simulation of turbulent flows in complex geometry. *Journal of Computational Physics*, 118(1):24–37, 1995. → pages 32, 33
- [55] O. Gicquel, N. Darabiha, and D. Thévenin. Laminar premixed hydrogen/air counterflow flame simulations using flame prolongation of ILDM with differential diffusion. *Proceedings of the Combustion Institute*, 28(2):1901–1908, 2000. → pages 40, 44, 65, 74
- [56] P. Givi. Model-free simulations of turbulent reactive flows. *Progress in Energy and Combustion Science*, 15(1):1–107, 1989. → pages 41
- [57] G. Golub and W. Kahan. Calculating the singular values and pseudo-inverse of a matrix. *Journal of the Society for Industrial and Applied Mathematics Series B Numerical Analysis*, 2(2):205–224, 1965. → pages 111
- [58] D. G. Goodwin, H. K. Moffat, and R. L. Speth. Cantera: An object-oriented software toolkit for chemical kinetics, thermodynamics, and transport processes. <http://www.cantera.org>, 2015. Version 2.2.0. → pages xiv, xv, 17, 51, 52, 58, 72, 132
- [59] S. Gordon and B. J. McBride. Computer program for calculation of complex chemical equilibrium compositions and applications: I. analysis. Technical report, NASA Reference Publication 1311, 1994. → pages 10
- [60] C. J. Greenshields. Upstream Class Reference, 2011. URL <http://foam.sourceforge.net/docs/cpp/a02750.html>. Accessed: 2016-01-20. → pages 107
- [61] F. Grinstein, L. Margolin, and W. Rider. *Implicit Large Eddy Simulation: Computing Turbulent Fluid Dynamics*. Cambridge University Press, 2007. ISBN 9780521869829. → pages 31
- [62] R. W. Grout. An age extended progress variable for conditioning reaction rates. *Physics of Fluids*, 19(10):105–107, 2007. → pages 53

- [63] P. C. Hansen. Numerical tools for analysis and solution of Fredholm integral equations of the first kind. *Inverse Problems*, 8(6):849, 1992. → pages 101
- [64] E. Hawkes and R. Cant. A flame surface density approach to large eddy simulation of premixed turbulent combustion. *Proceedings of the Combustion Institute*, 28(1):51–58, 2000. → pages 36, 37
- [65] D. C. Haworth and T. J. Poinsot. Numerical simulations of Lewis number effects in turbulent premixed flames. *Journal of Fluid Mechanics*, 244: 405–436, 11 1992. → pages 115, 134
- [66] H. Huang, J. M. Dennis, L. Wang, and P. Chen. A scalable parallel LSQR algorithm for solving large-scale linear system for tomographic problems: A case study in seismic tomography. *Procedia Computer Science*, 18(0): 581–590, 2013. → pages 112
- [67] J. Huang and W. K. Bushe. Simulation of transient turbulent methane jet ignition and combustion under engine-relevant conditions using Conditional Source-term Estimation with detailed chemistry. *Combustion Theory and Modelling*, 11(6):977–1008, 2007. → pages 42, 98, 101, 127
- [68] Y. Huang, H.-G. Sung, S.-Y. Hsieh, and V. Yang. Large eddy simulation of combustion dynamics of lean-premixed swirl-stabilized combustor. *Journal of Propulsion and Power*, 19(5):782–794, Sep 2003. → pages 38
- [69] H. G. Im, T. S. Lund, and J. H. Ferziger. Large eddy simulation of turbulent front propagation with dynamic subgrid models. *Physics of Fluids*, 9(12): 3826–3833, 1997. → pages 38
- [70] B. Jin, R. Grout, and W. Bushe. Conditional Source-term Estimation as a method for chemical closure in premixed turbulent reacting flow. *Flow, Turbulence and Combustion*, 81:563–582, 2008. → pages xiii, xiv, 40, 43, 44, 45, 47, 49, 62, 65, 66, 68, 74, 79, 100, 101, 130
- [71] M. M. Kamal, R. Zhou, S. Balusamy, and S. Hochgreb. Favre- and Reynolds-averaged velocity measurements: Interpreting PIV and LDA measurements in combustion. *Proceedings of the Combustion Institute*, 35 (3):3803–3811, 2015. → pages x, 75
- [72] V. Katta and W. M. Roquemore. C/H atom ratio in recirculation-zone-supported premixed and non-premixed flames. *Proceedings of the Combustion Institute*, 34(1):1101–1108, 2013. → pages 69

- [73] R. J. Kee, G. Dixon-lewis, J. Warnatz, M. E. Coltrin, and J. A. Miller. A fortran computer code package for the evaluation of gas-phase, multicomponent transport properties. Technical report, 1986. → pages 8, 10
- [74] R. J. Kee, F. M. Rupley, and J. A. Miller. CHEMKIN-II: A FORTRAN chemical kinetics package for the analysis of Gas-Phase chemical kinetics. 1989. → pages 8, 13, 51, 71
- [75] R. J. Kee, F. M. Rupley, and J. A. Miller. The CHEMKIN Thermodynamic Data Base, 1991. → pages 10, 11, 71
- [76] R. J. Kee, F. M. Rupley, J. A. Miller, M. E. Coltrin, J. F. Grcar, E. Meeks, H. K. Moffat, A. E. Lutz, G. Dixon-Lewis, M. D. Smooke, J. Warnatz, G. H. Evans, R. S. Larson, R. E. Mitchell, L. R. Petzold, W. C. Reynolds, M. Caracotsios, W. E. Stewart, P. Glarborg, C. Wang, and O. Adigun. PREMIX Program, CHEMKIN Collection, Release 3.6, 2000. → pages 49
- [77] A. R. Kerstein. Linear-eddy modeling of turbulent transport. Part 2: Application to shear layer mixing. *Combustion and Flame*, 75(3-4): 397–413, 1989. → pages xiv, 48, 49, 50, 51, 69
- [78] A. R. Kerstein. Linear-eddy modeling of turbulent transport. Part 3: Mixing and differential molecular diffusion in round jets. *Journal of Fluid Mechanics*, 216:411–435, 1990. → pages
- [79] A. R. Kerstein. Linear-eddy modeling of turbulent transport. Part 6: Microstructure of diffusive scalar mixing fields. *Journal of Fluid Mechanics*, 231:361–394, 1991. → pages 70
- [80] A. R. Kerstein. Linear-eddy modeling of turbulent transport. Part 4: Structure of diffusion flames. *Combustion Science and Technology*, 81 (1-3):75–96, 1992. → pages
- [81] A. R. Kerstein. Linear-eddy modeling of turbulent transport. Part 7: Finite-rate chemistry and multi-stream mixing. *Journal of Fluid Mechanics*, 240:289–313, 1992. → pages 49, 69
- [82] S. H. Kim and H. Pitsch. Conditional filtering method for large eddy simulation of turbulent non-premixed combustion. *Physics of Fluids*, 17 (10), 2005. → pages 41

- [83] W.-W. Kim, S. Menon, and H. C. Mongia. Large eddy simulation of a gas turbine combustor flow. *Combustion Science and Technology*, 143(1-6): 25–62, 1999. → pages 38
- [84] A. Klimenko and R. Bilger. Conditional Moment Closure for turbulent combustion. *Progress in Energy and Combustion Science*, 25(6):595–687, 1999. → pages 41, 98, 126
- [85] A. Y. Klimenko. Multicomponent diffusion of various admixtures in turbulent flow. *Fluid Dynamics*, 25(3):327–334, 1990. → pages 41
- [86] A. Y. Klimenko and S. B. Pope. The modeling of turbulent reactive flows based on multiple mapping conditioning. *Physics of Fluids (1994-present)*, 15(7):1907–1925, 2003. → pages 98, 126
- [87] R. Knikker, D. Veynante, and C. Meneveau. A priori testing of a similarity model for large eddy simulations of turbulent premixed combustion. *Proceedings of the Combustion Institute*, 29:2105–2111, 2002. → pages 37
- [88] E. Knudsen and H. Pitsch. A dynamic model for the turbulent burning velocity for large eddy simulation of premixed combustion. *Combustion and Flame*, 154(4):740–760, 2008. → pages 38
- [89] A. N. Kolmogorov. The local structure of turbulence in incompressible viscous fluid for very large reynolds numbers. *Dokl. Akad. Nauk SSSR*, 30: 299–303, 1921. → pages 20, 31
- [90] P. S. Kothnur and N. T. Clemens. Effects of unsteady strain rate on scalar dissipation structures in turbulent planar jets. *Physics of Fluids*, 17(12):–, 2005. → pages 94
- [91] K. K. Kuo. *Principles of Combustion*. John Wiley & Sons, 1986. → pages 14
- [92] J. Labahn, D. Dovizio, and C. Devaud. Numerical simulation of the delft-jet-in-hot-coflow (DJHC) flame using Conditional Source-term Estimation. *Proceedings of the Combustion Institute*, 35(3):3547–3555, 2015. → pages 42, 98, 101, 127
- [93] J. W. Labahn and C. B. Devaud. Investigation of Conditional Source-term Estimation applied to a non-premixed turbulent flame. *Combustion Theory and Modelling*, 17(5):960–982, 2013. → pages 42, 98, 101, 127

- [94] B. E. Launder and D. B. Spalding. *Mathematical Models of Turbulence*. Academic Press, 2002. → pages 30
- [95] B. E. Launder, G. J. Reece, and W. Rodi. Progress in the development of a Reynolds-stress turbulence closure. *Journal of Fluid Mechanics*, 68: 537–566, 4 1975. → pages 30
- [96] C. K. Law. *Combustion Physics*. Cambridge University Press, 2006. → pages 14
- [97] A. Leonard. Energy cascade in large eddy simulations of turbulent fluid flows. In F. Frenkiel and R. Munn, editors, *Turbulent Diffusion in Environmental Pollution Proceedings of a Symposium held at Charlottesville*, volume 18, Part A of *Advances in Geophysics*, pages 237–248. Elsevier, 1975. → pages 32
- [98] P. A. Libby and K. N. C. Bray. Variable density effects in premixed turbulent flames. *AIAA Journal*, 15(8):1186–1193, Aug 1977. → pages 35
- [99] P. A. Libby and F. A. Williams. *Turbulent Reacting Flows*. London : Academic, 1994. ISBN 0124479456. → pages 28
- [100] D. K. Lilly. A proposed modification of the Germano subgrid-scale closure method. *Physics of Fluids A*, 4(3), 1992. → pages 33
- [101] R. Lindstedt and E. Vaos. Transported PDF modeling of high-Reynolds-number premixed turbulent flames. *Combustion and Flame*, 145(3):495–511, 2006. → pages 41
- [102] T. Lu and C. K. Law. Toward accommodating realistic fuel chemistry in large-scale computations. *Progress in Energy and Combustion Science*, 35 (2):192–215, 2009. → pages 13
- [103] T. Lund and E. A. Novikov. Parameterization of subgrid-scale stress by the velocity gradient tensor. *Center for Turbulence Research (Stanford University and NASA)*, 1992. → pages 33
- [104] U. Maas and S. Pope. Simplifying chemical kinetics: Intrinsic low-dimensional manifolds in composition space. *Combustion and Flame*, 88(34):239–264, 1992. → pages 14, 129
- [105] B. Magnussen and B. Hjertager. On mathematical modeling of turbulent combustion with special emphasis on soot formation and combustion. *Symposium (International) on Combustion*, 16(1):719–729, 1977. → pages 35

- [106] T. Mantel and R. Borghi. A new model of premixed wrinkled flame propagation based on a scalar dissipation equation. *Combustion and Flame*, 96(4):443–457, 1994. → pages 36
- [107] G. H. Markstein. *Non-steady Flame Propagation*. Pergamon Press, 1964. → pages 38
- [108] S. Mathur, P. K. Tondon, and S. C. Saxena. Thermal conductivity of binary, ternary and quaternary mixtures of rare gases. *Molecular Physics*, 12: 569–579, 1967. → pages 9
- [109] S. McAllister, J. Y. Chen, and A. C. Fernandez-Pello. *Fundamentals of Combustion Processes*. Springer, 2011. → pages 1
- [110] B. J. McBride and S. Gordon. Computer program for calculation of complex chemical equilibrium compositions and applications: II. users manual and program description. Technical report, NASA Reference Publication 1311, 1996. → pages 10
- [111] C. Meneveau, T. S. Lund, and W. H. Cabot. A Lagrangian dynamic subgrid-scale model of turbulence. *Journal of Fluid Mechanics*, 319: 353–385, 7 1996. → pages 33
- [112] P. Moin and K. Mahesh. Direct Numerical Simulation: A tool in turbulence research. *Annual Review of Fluid Mechanics*, 30(1):539–578, 1998. → pages 28
- [113] C. Montgomery, G. Kosály, and J. Riley. Direct numerical solution of turbulent non-premixed combustion with multistep hydrogen-oxygen kinetics. *Combustion and Flame*, 109(1-2):113–144, 1997. → pages 88
- [114] A. Msaad, A. Belcadi, M. Mahdaoui, E. Aaffad, and M. Mouqallid. Reduced detailed mechanism for methane combustion. *Energy and Power Engineering*, 4(1):28–33, 2012. → pages 14
- [115] J. R. Nanduri, D. R. Parsons, S. L. Yilmaz, I. B. Celik, and P. A. Strakey. Assessment of RANS-based turbulent combustion models for prediction of emissions from lean premixed combustion of methane. *Combustion Science and Technology*, 182(7):794–821, 2010. → pages 41
- [116] S. Navarro-Martinez, A. Kronenburg, and F. D. Mare. Conditional Moment Closure for large eddy simulations. *Flow, Turbulence and Combustion*, 75 (1):245–274, 2005. → pages 41

- [117] K.-J. Nogenmyr, C. Fureby, X. Bai, P. Petersson, R. Collin, and M. Linne. Large eddy simulation and laser diagnostic studies on a low swirl stratified premixed flame. *Combustion and Flame*, 156(1):25–36, 2009. → pages 38
- [118] I. V. Novosselov and P. C. Malte. Development and application of an eight-step global mechanism for CFD and CRN simulations of lean-premixed combustors. *Journal of Engineering for Gas Turbines and Power*, 130(2):021502–021502–9, 2008. → pages 14
- [119] M. Oberlack, H. Wenzel, and N. Peters. On symmetries and averaging of the G-equation for premixed combustion. *Combustion Theory and Modelling*, 5(3):363–383, 2001. → pages 38
- [120] M. Oevermann, H. Schmidt, and A. Kerstein. Investigation of autoignition under thermal stratification using linear eddy modeling. *Combustion and Flame*, 155(3):370–379, 2008. → pages 71
- [121] V. Oijen and D. Goey. Modelling of premixed laminar flames using flamelet-generated manifolds. *Combustion Science and Technology*, 161(1):113–137, 2000. → pages 14, 44, 65, 74, 100, 114, 129
- [122] E. S. Oran and J. P. Boris. *Numerical Simulation of Reactive Flow*. Elsevier Science Publishing Co., Inc., 1987. → pages 5
- [123] C. C. Paige and M. A. Saunders. LSQR: An algorithm for sparse linear equations and sparse least squares. *ACM Trans. Math. Softw.*, 8(1):43–71, Mar. 1982. → pages 98, 111, 113
- [124] G. D. Paola, E. Mastorakos, Y. M. Wright, and K. Boulouchos. Diesel engine simulations with multi-dimensional Conditional Moment Closure. *Combustion Science and Technology*, 180(5):883–899, 2008. → pages 41
- [125] N. Peters. Numerical and asymptotic analysis of systematically reduced reaction schemes for hydrocarbon flames. In R. Glowinski, B. Larrouturou, and R. Temam, editors, *Numerical Simulation of Combustion Phenomena*, volume 241 of *Lecture Notes in Physics*, pages 90–109. Springer Berlin Heidelberg, 1985. ISBN 978-3-540-16073-1. → pages 13, 129
- [126] N. Peters. Laminar flamelet concepts in turbulent combustion. *Symposium (International) on Combustion*, 21(1):1231–1250, 1988. → pages 23
- [127] N. Peters. The turbulent burning velocity for large-scale and small-scale turbulence. *Journal of Fluid Mechanics*, 384:107–132, 1999. → pages 23

- [128] N. Peters. *Turbulent Combustion*. Cambridge University Press, 2000. → pages 15, 16, 25, 38
- [129] H. Pitsch. A consistent level set formulation for large eddy simulation of premixed turbulent combustion. *Combustion and Flame*, 143(4):587–598, 2005. → pages 38
- [130] H. Pitsch. Large eddy simulation of turbulent combustion. *Annual Review of Fluid Mechanics*, 38(1):453–482, 2006. → pages 37
- [131] H. Pitsch and L. D. de Lageneste. Large eddy simulation of premixed turbulent combustion using a level-set approach. *Proceedings of the Combustion Institute*, 29(2):2001–2008, 2002. → pages 23, 38
- [132] T. Poinso and D. Veynante. *Theoretical and Numerical Combustion*. Edwards, Philadelphia, 2nd edition, 2005. → pages 5, 6, 26, 37, 38, 41, 98, 126, 128, 131
- [133] T. Poinso, D. Veynante, and S. Candel. Quenching processes and premixed turbulent combustion diagrams. *Journal of Fluid Mechanics*, 228:561–606, 7 1991. → pages 23
- [134] S. Pope. PDF methods for turbulent reactive flows. *Progress in Energy and Combustion Science*, 11(2):119–192, 1985. → pages 40, 98, 126
- [135] S. Pope. Computations of turbulent combustion: Progress and challenges. *Symposium (International) on Combustion*, 23(1):591–612, 1991. Twenty-Third Symposium (International) on Combustion. → pages 41
- [136] S. B. Pope. *Turbulent Flows*. Cambridge University Press, 2000. → pages 5, 18, 19, 21, 28, 32
- [137] F. Proch and A. M. Kempf. Numerical analysis of the cambridge stratified flame series using artificial thickened flame LES with tabulated premixed flame chemistry. *Combustion and Flame*, 161(10):2627–2646, 2014. → pages 69
- [138] O. Reynolds. On the dynamical theory of incompressible viscous fluids and the determination of the criterion. *Philosophical Transactions of the Royal Society of London. A*, 186:123–164, 1895. → pages 18
- [139] L. F. Richardson. *Weather Prediction by Numerical Process*. Cambridge University Press, 1922. → pages 20

- [140] Y. Saad and H. A. van der Vorst. Iterative solution of linear systems in the 20th century. *Journal of Computational and Applied Mathematics*, 123(12):1–33, 2000. → pages 111, 117
- [141] P. Sagaut. *Large Eddy Simulation for Incompressible Flows*. Springer-Verlag Berlin Heidelberg, 2006. ISBN 3540263446. → pages 27, 30, 31, 33
- [142] M. Salehi. *Numerical Simulation of Turbulent Premixed Flames with Conditional Source-term Estimation*. PhD thesis, University of British Columbia, 2012. → pages 100, 102, 113, 114, 132
- [143] M. M. Salehi and W. K. Bushe. Presumed PDF modeling for RANS simulation of turbulent premixed flames. *Combustion Theory and Modelling*, 14(3):381–403, 2010. → pages 40, 42, 43, 47, 65, 98, 100, 101, 103, 104, 111, 127, 130
- [144] M. M. Salehi, W. K. Bushe, and K. J. Daun. Application of the Conditional Source-term Estimation model for turbulence-chemistry interactions in a premixed flame. *Combustion Theory and Modelling*, 16(2):301–320, 2012. → pages 48
- [145] M. M. Salehi, W. K. Bushe, N. Shahbazian, and C. P. Groth. Modified laminar flamelet presumed probability density function for LES of premixed turbulent combustion. *Proceedings of the Combustion Institute*, 34(1):1203–1211, 2013. → pages 40, 42, 43, 65, 98, 100, 101, 103, 104, 111, 127, 130
- [146] A. Sánchez, A. Lépinette, M. Bollig, A. Liñán, and B. Lázaro. The reduced kinetic description of lean premixed combustion. *Combustion and Flame*, 123(4):436–464, 2000. → pages 14
- [147] R. Sankaran, E. R. Hawkes, J. H. Chen, T. Lu, and C. K. Law. Structure of a spatially developing turbulent lean methane-air bunsen flame. *Proceedings of the Combustion Institute*, 31(1):1291–1298, 2007. → pages 63, 93
- [148] V. Sankaran and S. Menon. Structure of premixed turbulent flames in the thin-reaction-zones regime. *Proceedings of the Combustion Institute*, 28(1):203–209, 2000. → pages 67, 70, 127, 130
- [149] U. Schumann. Subgrid-scale model for finite difference simulations of turbulent flows in plane channels and annuli. *Journal of Computational Physics*, 18(4):376–404, 1975. → pages 33

- [150] A. Scotti, C. Meneveau, and D. K. Lilly. Generalized Smagorinsky model for anisotropic grids. *Physics of Fluids A*, 5(9):2306–2308, 1993. → pages 31
- [151] A. Scotti, C. Meneveau, and M. Fatica. Dynamic Smagorinsky model on anisotropic grids. *Physics of Fluids*, 9(6):1856–1858, 1997. → pages 31
- [152] L. Selle, G. Lartigue, T. Poinsot, R. Koch, K.-U. Schildmacher, W. Krebs, B. Prade, P. Kaufmann, and D. Veynante. Compressible large eddy simulation of turbulent combustion in complex geometry on unstructured meshes. *Combustion and Flame*, 137(4):489–505, 2004. → pages 39
- [153] N. Shahbazian, M. M. Salehi, C. P. Groth, Ö. L. Gülder, and W. K. Bushe. Performance of Conditional Source-term Estimation model for LES of turbulent premixed flames in thin reaction zones regime. *Proceedings of the Combustion Institute*, 35(2):1367–1375, 2015. → pages 40, 42, 43, 98, 100, 101, 103, 104, 111, 114, 115, 127, 129, 130, 135
- [154] M. Sheikhi, T. Drozda, P. Givi, F. Jaber, and S. Pope. Large eddy simulation of a turbulent non-premixed piloted methane jet flame (Sandia Flame D). *Proceedings of the Combustion Institute*, 30(1):549 – 556, 2005. → pages 41
- [155] J. Smagorinsky. General Circulation Experiments with the Primitive Equations. *Monthly Weather Review*, 91:99, 1963. → pages 33
- [156] G. P. Smith, D. M. Golden, M. Frenklach, N. W. Moriarty, B. Eiteneer, M. Goldenberg, C. T. Bowman, R. K. Hanson, S. Song, W. C. Gardiner, V. V. Lissianski, and Z. Qin. GRI-Mech 3.0. http://www.me.berkeley.edu/gri_mech/, 1997. Accessed: 2016-01-14. → pages 12, 17, 71, 132
- [157] T. M. Smith and S. Menon. One-dimensional simulations of freely propagating turbulent premixed flames. *Combustion Science and Technology*, 128(1-6):99–130, 1997. → pages 67, 127, 130
- [158] Y. Sommerer, D. Galley, T. Poinsot, S. Ducruix, F. Lacas, and D. Veynante. Large eddy simulation and experimental study of flashback and blow-off in a lean partially premixed swirled burner. *Journal of Turbulence*, 5:N37, 2004. → pages 39
- [159] P. Spalart and S. Allmaras. A one-equation turbulence model for aerodynamic flows. In *Paper 92-0439, AIAA*, 1992. → pages 30

- [160] D. Spalding. Mixing and chemical reaction in steady confined turbulent flames. *Symposium (International) on Combustion*, 13(1):649–657, 1971. → pages 34
- [161] H. Steiner and W. K. Bushe. Large eddy simulation of a turbulent reacting jet with Conditional Source-term Estimation. *Physics of Fluids*, 13(3): 754–769, 2001. → pages 42, 98, 101, 127
- [162] N. H. Stern. *Stern Review: on the Economics of Climate Change*. London HM Treasury, 2006. → pages 1
- [163] M. Stöllinger and S. Heinz. PDF modeling and simulation of premixed turbulent combustion. *Monte Carlo Methods and Applications*, 14(4): 343–377, 2008. → pages 41
- [164] N. Swaminathan and R. Bilger. Assessment of combustion submodels for turbulent non-premixed hydrocarbon flames. *Combustion and Flame*, 116 (4):519–545, 1999. → pages 88
- [165] N. Swaminathan and R. W. Bilger. Scalar dissipation, diffusion and dilatation in turbulent H₂-air premixed flames with complex chemistry. *Combustion Theory and Modelling*, 5(3):429–446, 2001. → pages 88
- [166] N. Swaminathan and R. W. Bilger. Analyses of Conditional Moment Closure for turbulent premixed flames. *Combustion Theory and Modelling*, 5(2):241–260, 2001. → pages 41
- [167] M. Sweeney, S. Hochgreb, and R. Barlow. Cambridge stratified slot burner data. 2009. URL <http://www.dspace.cam.ac.uk/handle/1810/226470>. → pages xiv, 61, 63, 145
- [168] M. Sweeney, S. Hochgreb, and R. Barlow. The structure of premixed and stratified low turbulence flames. *Combustion and Flame*, 158(5):935–948, 2011. → pages 60, 68, 145
- [169] M. Sweeney, S. Hochgreb, M. Dunn, and R. Barlow. A comparative analysis of flame surface density metrics in premixed and stratified flames. *Proceedings of the Combustion Institute*, 33(1):1419–1427, 2011. → pages 69
- [170] M. S. Sweeney, S. Hochgreb, M. J. Dunn, and R. S. Barlow. The structure of turbulent stratified and premixed methane/air flames I: Non-swirling flows. *Combustion and Flame*, 159(9):2896–2911, 2012. → pages 68, 75, 76, 145

- [171] M. S. Sweeney, S. Hochgreb, M. J. Dunn, and R. S. Barlow. The structure of turbulent stratified and premixed methane/air flames II: Swirling flows. *Combustion and Flame*, 159(9):2912–2929, 2012. → pages 69
- [172] M. S. Sweeney, S. Hochgreb, M. J. Dunn, and R. S. Barlow. Multiply conditioned analyses of stratification in highly swirling methane/air flames. *Combustion and Flame*, 160(2):322–334, 2013. → pages 75, 76, 145
- [173] P. Tamadonfar and Ö. L. Gülder. Effects of mixture composition and turbulence intensity on flame front structure and burning velocities of premixed turbulent hydrocarbon/air bunsen flames. *Combustion and Flame*, pages –, 2015. → pages xi, 127, 133, 135
- [174] B. Thornber, R. Bilger, A. Masri, and E. Hawkes. An algorithm for LES of premixed compressible flows using the Conditional Moment Closure model. *Journal of Computational Physics*, 230(20):7687–7705, 2011. → pages 41
- [175] H. Tsui and W. Bushe. Linear-eddy model formulated probability density function and scalar dissipation rate models for premixed combustion. *Flow, Turbulence and Combustion*, 93(3):487–503, 2014. → pages 67, 74, 81, 127, 130, 131, 132, 137
- [176] H. Tsui, M. Kamal, S. Hochgreb, and W. Bushe. Direct comparison of PDF and scalar dissipation rates between LEM simulations and experiments for turbulent, premixed methane air flames. *Combustion and Flame*, 165: 208–222, 2016. → pages 137
- [177] H. P. Tsui and W. K. Bushe. Conditional Source-term Estimation using dynamic ensemble selection and parallel iterative solution. *Combustion Theory and Modelling*, 2016. (in press). → pages 137
- [178] M. Tsurikov. *Experimental investigation of the fine-scale structure in turbulent gas-phase jet flows*. PhD thesis, University of Texas at Austin, 2002. → pages 94
- [179] D. Vandromme and H. Haminh. *Turbulence and Coherent Structures: Selected Papers from “Turbulence 89: Organized Structures and Turbulence in Fluid Mechanics”, Grenoble, 18–21 September 1989*, pages 507–523. Springer Netherlands, Dordrecht, 1991. ISBN 978-94-015-7904-9. → pages 30

- [180] L. Vervisch, R. Hauguel, P. Domingo, and M. Rullaud. Three facets of turbulent combustion modelling: DNS of premixed V-flame, LES of lifted non-premixed flame and RANS of jet-flame. *Journal of Turbulence*, 5:N4, 2004. → pages 39
- [181] D. Veynante and L. Vervisch. Turbulent combustion modeling. *Progress in Energy and Combustion Science*, 28(3):193–266, 2002. → pages 5, 34, 35, 58, 68, 69, 90, 130, 131
- [182] W. Vicente, M. Salinas, E. Barrios, and C. Dopazo. PDF modeling of CO and NO formation in lean premixed methane flames. *Combustion Science and Technology*, 176(4):585–601, 2004. → pages 41
- [183] M. Wang, J. Huang, and W. Bushe. Simulation of a turbulent non-premixed flame using Conditional Source-term Estimation with trajectory generated low-dimensional manifold. *Proceedings of the Combustion Institute*, 31(2):1701–1709, 2007. → pages 42, 98, 101, 127
- [184] H. Weller, G. Tabor, A. Gosman, and C. Fureby. Application of a flame-wrinkling LES combustion model to a turbulent mixing layer. *Symposium (International) on Combustion*, 27(1):899–907, 1998. → pages 37
- [185] H. G. Weller. OpenFOAM 2.1.1. <http://www.openfoam.org/version2.1.1/>, 2012. Accessed: 2016-01-14. → pages 99, 102, 114, 128
- [186] F. M. White. *Fluid Mechanics*. Boston: McGraw-Hill Book Company, 5 edition, 2003. → pages 5, 18
- [187] D. C. Wilcox. Reassessment of the scale-determining equation for advanced turbulence models. *AIAA Journal*, 26(11):1299–1310, 1988. → pages 30
- [188] C. R. Wilke. A viscosity equation for gas mixtures. *The Journal of Chemical Physics*, 18(4):517–519, 1950. → pages 9
- [189] S. L. Yilmaz, M. B. Nik, P. Givi, and P. A. Strakey. Scalar filtered density function for large eddy simulation of a bunsen burner. *Journal of Propulsion and Power*, 26(1):84–93, Jan 2010. → pages 41
- [190] F. T. Yuen and Ö. L. Gülder. Premixed turbulent flame front structure investigation by Rayleigh scattering in the thin reaction zone regime. *Proceedings of the Combustion Institute*, 32(2):1747–1754, 2009. → pages 99

Bibliography

- [191] F. T. C. Yuen. *Experimental Investigation of the Dynamics and Structure of Lean-Premixed Turbulent Combustion*. PhD thesis, University of Toronto, 2009. → pages 133
- [192] K. K. yun Kuo and R. Acharya. *Fundamentals of Turbulent and Multi-Phase Combustion*. Wiley, 1st edition, 2012. → pages 6
- [193] J. Zhang, F. Gao, G. Jin, and G. He. Conditionally statistical description of turbulent scalar mixing at subgrid-scales. *Flow, Turbulence and Combustion*, 93(1):125–140, 2014. → pages 68
- [194] R. Zhou, S. Balusamy, M. S. Sweeney, R. S. Barlow, and S. Hochgreb. Flow field measurements of a series of turbulent premixed and stratified methane/air flames. *Combustion and Flame*, 160(10):2017–2028, 2013. → pages x, 69, 75, 145