Combustion Modelling in Spark-Ignition Engines using Conditional Source-term Estimation

by

Girish Venkata Nivarti

B. Tech. (Honours), Indian Institute of Technology Kharagpur, 2011

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

MASTER OF APPLIED SCIENCE

in

The Faculty of Graduate and Postdoctoral Studies
(Mechanical Engineering)

THE UNIVERSITY OF BRITISH COLUMBIA
(Vancouver)

August 2013

Abstract

Conditional Source-term Estimation (CSE) is a chemical closure model for the simulation of turbulent combustion. In this work, CSE has been explored for modelling combustion phenomena in a spark-ignition (SI) engine. In the arbitrarily complex geometries imposed by industrial design, estimation of conditionally averaged scalars is challenging. The key underlying requirement of CSE is that conditionally averaged scalars be calculated within spatially localized sub-domains. A domain partitioning algorithm based on space-filling curves has been developed to construct localized ensembles of points necessary to retain the validity of CSE. Algorithms have been developed to evenly distribute points to the maximum extent possible while maintaining spatial locality. A metric has been defined to estimate relative inter-partition contact as an indicator of communication in parallel computing architectures. Domain partitioning tests conducted on relevant geometries highlight the performance of the method as an unsupervised and computationally inexpensive domain partitioning tool.

In addition to involving complex geometries, SI engines pose the challenge of accurately modelling the transient ignition process. Combustion in a homogeneous-charge natural gas fuelled SI engine with a relatively simple chamber geometry has been simulated using an empirical model for ignition. An oxygen based reaction progress variable is employed as the conditioning variable and its stochastic behaviour is approximated by a presumed probability density function (PDF). A trajectory generated low-dimensional manifold has been used to tabulate chemistry in a hyper-dimensional space described by the reaction progress variable, temperature and pressure. The estimates of pressure trace and pollutant emission trends obtained using CSE accurately match experimental measurements.

Preface

This thesis has been the result of collaboration with researchers from the University of British Columbia and Westport Innovations Inc. in Vancouver, Canada. The work has been funded by Mitacs Inc., Canada and Westport Innovations Inc., Canada. I am responsible for conducting all parts of the research and preparing the corresponding manuscripts.

- *Chapter Three:* contains work that was accomplished at the University of British Columbia under the supervision of Dr. Kendal Bushe. Part of the work has been published as conference proceedings under the title:

  - Girish V. Nivarti, M. Mahdi Salehi, W. Kendal Bushe, Conditional source-term estimation for large eddy simulation of turbulent premixed flames in complex geometries, *Proceedings of the Combustion Institute Canadian Section*, 03-001, 2013.

  Ms. Nasim Shahbazian and Dr. Jim Huang contributed geometry data for conducting validation tests of the developed method. A licensed copy of Tecplot360 software was used for plotting domain geometries. I am responsible for conducting all parts of the research and preparing the corresponding manuscripts. This work shall soon be submitted for journal publication.

- *Chapter Four:* involves research work conducted at Westport Innovations Inc. under the supervision of Dr. Jim Huang and Dr. W. Kendal Bushe under a non-confidentiality agreement signed through Mitacs Accelerate and Mitacs Globalink. The open-source software package, OpenFOAM was used for performing relevant simulations. A licensed MATLAB software copy was used for performing low-pass filter calculations. I am responsible for

writing the entire C++ code pertinent to the combustion model. I am responsible for conducting all parts of the research and preparing the corresponding manuscripts. This work shall soon be submitted for journal publication.

The entire manuscript was prepared using the LaTeX package.

Table of Contents

List of Tables

List of Figures

List of Algorithms

List of Symbols

α	Scalar conditional average
δ	Laminar flame thickness
δ_{ij}	Kronecker delta function
$\dot{\omega}$	Chemical production rate
ℓ_0	Integral length scale
ε	Dissipation rate of energy
η	Kolmogorov length scale
λ	Relative air-fuel ratio
$\mathscr{D}$	Molecular diffusion constant
$\mathscr{J}$	Molecular diffusion flux
ν	Coefficient of laminar kinematic viscosity
ν_T	Coefficient of turbulent kinematic viscosity
ρ	Fluid density
$\vec{x}$	Spatial coordinate
$\widetilde{c}$	Favre-averaged value of c
$\widetilde{c''^2}$	Favre-averaged variance of c
c	Reaction progress variable
Da	Damköhler number
k	Kinetic energy
Ka	Karlovitz number
P	Pressure
Re	Reynolds number
Re_t	Turbulent Reynolds number
S_L	Laminar flame speed
Sc	Schmidt number
T	Temperature
u'	Turbulence intensity
u_i	i^{th} component of velocity
Y_k	k^{th} species mass fraction

List of Acronyms

BML Bray-Moss-Libby

CMC Conditional Moment Closure

CSE Conditional Source-term Estimation

DNS Direct Numerical Simulation

EBU Eddy Break-up

FSD Flame Surface Density

LES Large-eddy Simulation

LEM Linear-eddy Model

MMC Multiple Mapping Conditioning

ODT One-dimensional Turbulence

PCM Presumed Conditional Moment

PDF Probability Density Function

RANS Reynolds Averaged Navier Stokes

TGLDM Trajectory Generated Low-dimensional Manifold

AMR Adaptive Mesh Refinement

NP Non-deterministic Polynomial-time

RSB Recursive Spectrum Bisection

RCB Recursive Coordinate Bisection

SFC Space-filling Curve

ATDC After Top Dead Center

BTDC Before Top Dead Center

CA Crank Angle

EVO Exhaust Valve Opens

GIW Gross Indicated Work

GRI Gas Research Institute

HRR Heat Release Rate

IVC Intake Valve Closes

MFB Mass Fraction Burned

SI Spark-Ignition

TDC Top Dead Center

Acknowledgements

Turbulent combustion is a formidable subject of inquiry. Over the past few years of research, I have evolved as a person through learning from my mistakes. Throughout my errors, doubts, and fears, my supervisor, Prof. W. Kendal Bushe, has shown me courage and calm. In his sheltering aura of knowledge and sheer cool, I have risen to a stage where I can comprehend the basic questions within our discipline. I owe this work and the related successes in my life entirely to him.

My academic growth has been accelerated by the impeccable mentorship of my professors and collaborating researchers. Foremost amongst them, Dr. Mahdi Salehi has been a friend, philosopher and guide to me. I humbly thank him for his company and his benefactions, academic and otherwise. I sincerely thank my professors, Prof. Neil Balmforth, Prof. Mona Berciu and Prof. Carl Ollivier-Gooch, whose passion towards their disciplines has been inspirational. I am grateful to Dr. Jim Huang for introducing me to the world of industrial problems which was once bizarre to me. I also thank my colleagues Hong Tsui and Graham Hendra for their positive company.

I owe my relentlessly optimistic attitude to my mother, my father and my sister who, through their love and support, allowed me to grow healthily and live passionately in this continent half the world away from home. I am deeply indebted to Lorenzo Lane, Dr. Artan Sheshmani, Amanda Grochowich, Razan Hamdi and Tim Höllering for their true friendship, inspiration, love, guidance and *science*! I dedicate this thesis to the land where I first delved into research, where some of the most memorable moments of my life were experienced and where, today, I stand beckoned by the vast ocean of learning...

To Canada.

1 | Introduction

Energy has remained vital to the growth of our civilisation for centuries. But over the past few decades, the rising population and rapid urban development have caused an acute inflation of energy demand, encouraging a globally pervasive debate on energy policy. A consensus has emerged mandating the alleviation of domestic and industrial carbon footprints on the environment. Amidst the ongoing polemic on current power generation technologies, the Stern review [1] categorically decries fossil fuel combustion as a source of harmful pollutants and greenhouse gases. Today, low-polluting renewable sources of energy form a topic of significant research. Moreover, a growing conscience for our future generations has initiated a shift towards sustainable lifestyles.

The ground reality, however, is that building an infrastructure of alternative energy hinges on existing resources. Burgeoning demands of our society cannot be fulfilled by renewable energy, particularly with the current pace of development. The significance of combustion has been clearly illustrated by Tollefson and Monastersky [2] among others. While fossil fuel reserves are said to be declining, new repositories of combustible fuels are being discovered and invented in forms hitherto unknown. All roads to the immediate future of our civilisation depend inevitably on combustion – the challenge is to use every drop of fuel judiciously, with diminishing environmental effects.

Scientific interest in combustion research has been fuelled further by the long advancing strides of technology with each finding. In contrast to their counterparts from early twentieth century, combustion systems today are outcomes of prolific theoretical and experimental research. Contemporary aircraft engines are much more efficient, cleaner and safer than their immediate precursors. Although ex-

periments have been crucial to these drastic developments in engine design, they have limited capability to probe the intricate geometries of modern devices. The advent of high performance computing technology has not only aided experimental research in this regard, but it has become a necessary predictive tool in the analysis of alternative fuels. In fact, Moin and Kim [3] describe supercomputers as powerful assets which might uncover the secrets long held behind natural phenomena that characterise the behaviour of flames.

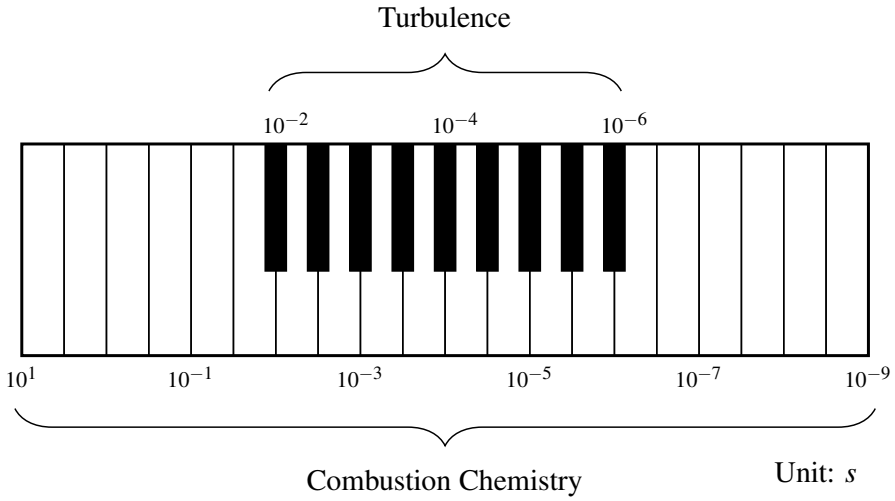

Figure 1.1: Diversity of time scales in a turbulent flame.

In combustion devices of engineering interest, fuel mixtures burn as turbulent reactive flows. Chaotic properties of turbulence impart a disparity of scales to the flow physics, while the chemistry in flames comprises an even broader spectrum of time and length scales. Nitrogen oxides (collectively termed, NO_x) are products of gradual transformation at high activation temperatures, whereas myriads of radicals such as hydroxyl ions (OH^-) survive fleeting nanoseconds. As discussed by Pope [4], small scales and many species are but among the manifold other formidable challenges that render turbulent combustion to be mathematically and computationally prohibitive for decades to come. Even the existence of an analytical solution to the non-linear governing equations of turbulent flows is an open Millenium Prize Problem, formally stated by Fefferman [5].

In the absence of an exact solution, computational fluid dynamics (CFD) has emerged as an approach to numerically estimate solutions of time sensitive industry scale problems. Fundamental aspects of flame behaviour, including ignition, extinction and flame propagation, are modelled based on experimental knowledge, asymptotic analysis, theoretical insight and often, scientific faith. Combustion modelling seeks to resolve the interaction between chemistry and turbulence at small scales; this is infamously known as the chemical closure problem. When the fuel and oxidiser are mixed prior to combustion, the chemical reactions and turbulent flow exhibit a stronger coupling, particularly at the high turbulence intensity observed in engines. Premixed flames, however, can propagate in lean fuel mixtures and demand relatively low ignition temperatures. These conditions provide an unsuitable environment for the emission of pollutants such as unburned hydrocarbons, carbon monoxide (CO) and NO_x. As modern engines increasingly exploit a greater fraction of the premixed burn of fuel, there is an urgent need for developing precise modelling tools.

Conditional source-term estimation (CSE) is a combustion model formulated by Bushe and Steiner [6] in 1999 at the Center for Turbulence Research, Stanford. In contrast to previously developed models for chemical closure, CSE is not restricted to any specific regime of flame behaviour. CSE is poised to accurately describe turbulence-chemistry interactions at high turbulence intensities due to the lack of common binding assumptions. During the past decade, this model has been successful in calculating quantities of practical interest for canonical non-premixed flame problems; the application to premixed flame simulations has been relatively recent. The subject of this thesis is to extend the applicability of CSE to problems of industrial context. The work has been presented as a part on the theoretical development of CSE and another on its application to problems in contemporary natural gas engines.

2 | Background

Accurate modelling of turbulent reactive flows hinges on a sound understanding of the complex constitutive processes. Knowing the fluid mechanical properties and their variation with the temperatures exhibited by a combustion system is absolutely essential for describing transport phenomena and other mechanisms such as heat transfer, molecular diffusion and convective turbulent transport [7]. Upon mixing, reacting species undergo chemical transformations that follow non-linear, inter-dependent mechanisms. Pollutant formation and fuel consumption are governed by these mechanisms; in particular, nitrogen undergoes transformations to produce environmentally dangerous oxides, NOx [8]. Even methane, the simplest of simplest of hydrocarbons, may undergo over 340 elementary reactions involving up to 50 intermediate species upon combustion in the presence of air [9]. Detailed mechanisms are often simplified and tabulated prior to simulation in order to improve computational efficiency [10]. Tabulation of reaction constants and chemical reaction mechanisms has not been accomplished for higher hydrocarbons [11, 12]. In gas phase combustion, the energy released from reactions affects the transport properties of the system and enhances turbulent micro-mixing, whilst the increased micro-mixing helps escalate reaction rates by bringing reactants together faster. Resolution of these non-linear turbulence-chemistry interactions is a persistent challenge in gas-phase combustion. This chapter provides a conceptual background of mathematical investigative tools necessary for developing models to resolve this coupling.

2.1 Turbulence

Turbulent flows present a non-linear, multi-scale problem described in complete detail by the inherently coupled system of Navier-Stokes equations. Originally written as integral conservation laws by White [13], these differential equations define conservation of mass and momentum assuming the continuum hypothesis, wherein flow variables may be represented as smooth functions in space and time [14]. Conservation of mass principle (continuity equation) is written as

$$\frac{\partial \rho}{\partial t} + \frac{\partial \rho u_j}{\partial x_j} = 0 \tag{2.1}$$

where ρ is the density of the fluid and u_j is the velocity in the j^{th} principal direction. Conservation of momentum is described by the following transport equation

$$\frac{\partial \rho u_i}{\partial t} + \frac{\partial \rho u_i u_j}{\partial x_j} = -\frac{\partial p}{\partial x_i} + \frac{\partial \tau_{ij}}{\partial x_j} + \mathscr{F}_i. \tag{2.2}$$

where τ_{ij} represents the stress tensor and p, the pressure at each point in the fluid, while $\mathscr{F}_i$ accounts for external forces. For the constant-property Newtonian fluids, this can be written as

$$\tau_{ij} = -p\delta_{ij} + \mu \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) \tag{2.3}$$

wherein δ_{ij} is the Kronecker delta and μ is the viscosity of considered fluid whereas $-p$ is the normal static pressure observed at any point. In addition with an equation for energy used in one of its various forms, the system is closed using an appropriate equation of state [15].

Although this system of equations is readily solved for laminar flows, no known analytical solutions exist for turbulent flows even in physical simple domains such as pipes. Features that characterise turbulent flows have been known for centuries, such as high rates of diffusivity, 3-dimensional irregularities and vorticity, but the transition to turbulence has not been well described [16]. In the absence of such critical knowledge, turbulence research relies on physical observations and numerical analysis for valuable insight into this ubiquitous phenomenon. As established

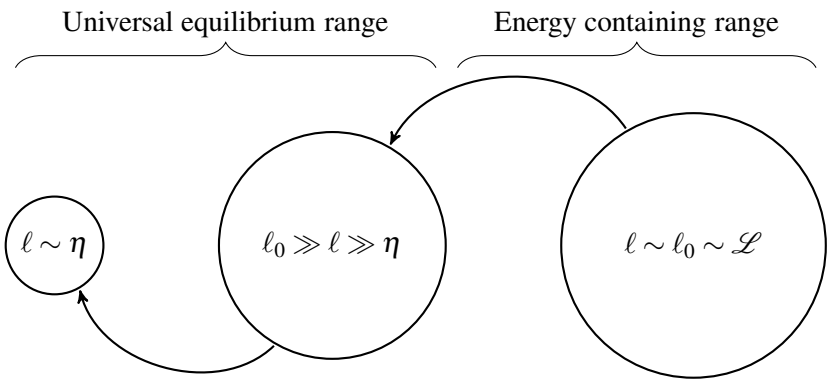

Figure 2.1: Kolmogorov's energy cascade hypothesis.

by Reynolds [17], turbulent flows are characterized by a single non-dimensional parameter based on the characteristic velocity ($\mathscr{U}$) and length scales ($\mathscr{L}$) of the flow, defined as

$$\text{Reynolds number, Re} \equiv \mathscr{U} \cdot \mathscr{L}/\nu \tag{2.4}$$

where ν is the kinematic viscosity of the fluid. At high Reynolds numbers, typical for most turbulent flows, a disparity of scales is observed [14]. The *energy cascade* hypothesis illustrated in Fig. 2.1 describes the transfer of kinetic energy (k) from the largest scales (ℓ_0), which are determined by the scale of the domain geometry ($\mathscr{L}$), to the smallest scales where all the energy dissipates by viscous action. Kolmogorov hypothesised that in a turbulent flow with high enough Reynolds number, the statistics of motion at the smallest scales are isotropic and are universal to all turbulent flows. The smallest scales can be determined solely by ν and the dissipation rate, ε, and thus all information of geometry is lost at these scales. The unique scaling of various parameters at play results in a universal form for the ratio of largest scales, ℓ_0, and the smallest, Kolmogorov scales defined by

$$\eta \equiv (\nu^3/\varepsilon)^{1/4} \tag{2.5a}$$

for the length scales, and by

$$\tau_\eta \equiv (\nu/\varepsilon)^{1/2} \tag{2.5b}$$

for the time scales [16, 18]. The relative magnitude of the smallest length and time scales, η and τ_η respectively, referred to as the Kolmogorov scales [14], compared to the largest scales, l_0 and τ_0, is also approximated in a simple form by using Kolmogorov's hypotheses as

$$\eta/\ell_0 \sim Re^{-3/4} \tag{2.6a}$$

$$\tau_\eta/\tau_0 \sim Re^{-1/2} \tag{2.6b}$$

where τ_0 and ℓ_0 represent integral time and length scales respectively. An exact numerical solution of Navier-Stokes equations is necessary to obtain the flow variables, such as velocity and pressure, as known functions of space and time. The computational cost of a numerical simulation are determined by the resolution requirements – the solution domain must be large enough to contain the energy-containing motions, and the grid spacing Δx must be small enough to resolve the dissipative scales [14]. However, it is evident from Eq. 2.6 that the instantaneous range of scales in a turbulent flow increases rapidly with the Reynolds number. Majority of engineering problems have too wide a range of scales to be directly computed – given the necessary grid spacing, the number of grid points required vary as $N_g \sim Re^{9/4}$ [14], which when combined with the time resolution requirements lends the alarming result given by

$$\text{Computational effort} \sim Re^{11/4}. \tag{2.7}$$

It is evident that resolving all scales of motion using Direct Numerical Simulation (DNS) for practical flows is infeasible [19]. The vast information obtained using DNS of a practical combustion device often has little value since quantities such as mean pollutant formation rates, average power, and average fuel consumption are of greater interest to the industry. Significant progress has been made towards a faster estimate of these parameters by describing the flow as a chaotic process using random variables. Simpler descriptions seek to reduce computational expense: instead of solving for the instantaneous flow-field, a statistical time evolution of the flow is sought. Not only are the computational abilities inadequate, knowledge of initial conditions is seldom precise, accurate or sufficient. Be it ex-

perimental measurements or numerical estimates of initial and boundary values, the difference between realisation and intention is often disparate. The mathematical treatment of turbulent flows using random variables is a useful exercise that helps achieve reasonably accurate predictions. However, these methods suffer from the drawback of introducing terms which cannot be obtained from the governing equations and thus require modelling.

2.1.1 Reynolds and Favre Averaging

The starting point of statistical analysis of turbulent flows is the decomposition of each quantity into a mean and a fluctuating component as follows:

$$Q = \overline{Q} + Q', \tag{2.8a}$$

where the components follow the basic moment rules:

$$\overline{\overline{Q}} = \overline{Q}, \quad \overline{Q'} = 0. \tag{2.8b}$$

Reynolds Averaged Navier Stokes (RANS) equations describe the transport of averaged quantities through the Navier-Stokes equations using this decomposition. The immense computational expense involved in DNS is, therefore, avoided by calculating only the mean flow field. In compressible flows such as turbulent reactive flows, the terms are instead Favre-averaged by considering a density weighted average to simplify the resulting equations while still including the effect of density fluctuations. Decomposition into Favre-terms is written as

$$Q = \widetilde{Q} + Q'' \tag{2.9a}$$

where the Favre-average is defined by the principles

$$\widetilde{Q} \equiv \frac{\overline{\rho Q}}{\overline{\rho}}, \text{ and } \quad \widetilde{Q''} = \frac{\overline{\rho\left(Q - \widetilde{Q}\right)}}{\overline{\rho}} = 0. \tag{2.9b}$$

No simple relation exists between Favre and Reynolds averages – a meaningful relation depends on density fluctuation correlations $\overline{\rho' Q'}$ which are im-

plicit in the Favre-average. More importantly, since experimental techniques determine Reynolds averaged data, comparing experimental results with Favre averaged quantities obtained using numerical simulations might not be obvious [7]. The main reason for Favre-averaging is compactness of resulting equations. The Favre-averaged momentum equations are derived as

$$\overline{\rho}\left(\frac{\partial \widetilde{u}_i}{\partial t}+\widetilde{u}_j\frac{\partial \widetilde{u}_i}{\partial x_j}\right)=-\frac{\partial \overline{p}}{\partial x_i}+\frac{\partial}{\partial x_j}(\widetilde{\tau}_{ij}-\overline{\rho}\widetilde{u_i''u_j''}). \tag{2.10}$$

wherein $\widetilde{u_i''u_j''}$ is not closed within the governing equation system. The RANS or FANS decompositions thus result in the creation of non-linear terms which are unresolved in the equation system. In general, this term describing the average exchange of turbulent momentum across the boundary of each finite volume is described by a simple and inaccurate relation

$$\widetilde{u_i''u_j''}-\frac{2}{3}k\delta_{ij}=\nu_T\left(\frac{\partial u_i}{\partial x_j}+\frac{\partial u_j}{\partial x_i}\right) \tag{2.11}$$

Turbulence modelling then reduces to the art of estimating turbulent viscosity, ν_T. An overwhelming majority of such models calculate this quantity using empirical and scaling laws. Typically, turbulent energy spectrum parameters such as the dissipation rate ε and turbulent kinetic energy k are used. The most popular technique is to use the scaling

$$\nu_t=C_\mu\overline{\rho}\frac{\widetilde{k}^2}{\widetilde{\varepsilon}} \tag{2.12}$$

to obtain a rough estimate. The value of C_μ depends on the flow and hence, this modelling technique is not *complete* in the sense that it requires to be tuned to suit a particular problem. A similar approach solves a separate transport equation for the turbulent frequency $\omega=k/\varepsilon$ and is substantially more accurate in near-wall layers [20]. As a unique strategy, [21] developed a transport equation for turbulent viscosity which is a popular choice in external flow applications such as aircraft simulations. In the recent past, several models have aimed to calculate Reynolds stress using transport equations. Although this method can resolve the stress tensor with greater accuracy it is not robust for industrial application.

In general, the RANS concept must be used with discretion and care. The

approach is constrained by time and spatial resolution requirements dictated by the numerical solution procedure. A RANS simulation is, however, a fast means to obtain basic insight into the flow field. The low computational expense and simplicity makes RANS the most popular approach in industry.

2.1.2 Filtering and Spatial Averaging

The objective of Large Eddy Simulation (LES) is to calculate the largest structures of the flow exactly and model the small-scale behaviour. LES resolves the velocity field obtained by applying a low-pass filter of characteristic width Δ to the underlying turbulent velocity field, $u_i(\vec{x},t)$. Large-scale behaviour of turbulent flows is known to be governed by the domain geometry and motion at small-scales is believed to have universal properties. In turbulent flows and turbulent reactive flows, critical features including instabilities and unsteady mixing are governed by the large structures which can be computed exactly by LES. The general filtering operation is defined by

$$\langle u_i(\vec{x},t)\rangle = \int_V G(\vec{r},\vec{x})\, u_i(\vec{x}-\vec{r},t)\, d\vec{r} \tag{2.13a}$$

where the specified filter function G satisfies the normalization condition

$$\int G(\vec{r},\vec{x}) d\vec{r} = 1 \tag{2.13b}$$

and integration is over the entire flow domain [14]. Standard filters include i) a cut-off filter in spectral space which preserves length scales greater than 2Δ, ii) box filter in physical space which corresponds to a spatial averaging of Q over a box with size Δ, and iii) a Gaussian filter. Unlike Reynolds averaging, the filtering operations do not follow the moment relations, i.e.:

$$\widetilde{\widetilde{Q}} \neq \widetilde{Q}, \quad \widetilde{Q}' \neq 0. \tag{2.14}$$

However, unresolved stresses defined as follows in LES have a similar expression as in RANS and require sub-grid scale modelling:

$$\tau_{ij}^R \equiv \widetilde{u_i u_j} - \widetilde{u}_i \widetilde{u}_j. \tag{2.15}$$

Filtered balance equations coupled with appropriate sub-grid models can be used to numerically solve the behaviour of filtered fields. Although LES lacks the accuracy level of DNS, a drastically more detailed picture of the turbulent flow is obtained in contrast RANS. LES has advanced further as a powerful tool for the simulation of turbulent flows in the past decade [22]. However, application to turbulent reactive flows of even canonical industrial problems are challenging particularly due to the complex geometry that demands finer mesh resolution thereby increasing computational expense [23, 24]. Pope [25] highlights the manifold challenges including completeness and tractability of LES as a precise mathematical modelling tool; most industrial problems remain beyond the scope LES .

2.1.3 Probability Density Function Methods

An averaged quantity, computed using RANS, represents the first moment of the field. However, given the one-point, one-time joint Cumulative Density Function (CDF) of velocity in a turbulent flow-field

$$F(\vec{u};\vec{x},t) \equiv P\{u_i(\vec{x},t) < V_i,\, i = 1,2,3\}, \tag{2.16a}$$

the joint probability density function (PDF) can be computed as

$$f(\vec{u};\vec{x},t) = \frac{\partial^3 F(\vec{u},\vec{x},t)}{\partial u_i \partial u_j \partial u_k}. \tag{2.16b}$$

At each point in in space, the PDF characterizes the random velocity vector without holding any joint information for two or more points in space or time [14]. PDF methods solve a transport equation for these PDFs using stochastic Lagrangian methods. At high Reynolds numbers, molecular diffusion has a negligible contribution to spatial transport and convection dominates the transport of momentum, chemical species and enthalpy. Despite the accurate treatment of convective terms in the Lagrangian PDF method, however, the computational cost involved in Monte Carlo simulations render these methods infeasible for industrial application.

At one end of the spectrum of approaches for the simulation of turbulent flows, DNS provides complete detail at an extremely high computational cost and, hence, a limited range of applicability [19]. RANS-based models such as the $k-\varepsilon$ method

lie at the other end; these offer a diverse range of applicability but with poor accuracy and limited insight [26]. LES and PDF methods form the middle range and are gaining popularity with the rapidly developing high performance computing technology.

2.2 Turbulent Combustion

Turbulent combustion has been a topic of intense scientific inquiry for several decades [27–30]. Several questions remain open due to the formidable complexity of problems involved in a turbulent reactive flow particularly due to heat release and differential diffusion [4, 31].

There are two principal modes of combustion based on whether fuel and oxidiser are mixed prior to combustion (premixed combustion) or not (non-premixed combustion). In reality, most reactive flows exhibit a hybrid behaviour and are categorised under partially-premixed combustion. Premixing provides for effective control of stoichiometry of flames – this allows lean-burning and reduction of unburned hydrocarbons while providing an unsuitable environment for NO_x formation mechanisms as understood today. Due to these reasons premixed combustion is being employed in gas turbines, and is being theoretically investigated in order to develop efficient automobile engines and industrial furnaces. On the other hand, premixed fuel storage is a known safety hazard. Moreover, premixed flames are susceptible to convectively and acoustically coupled instabilities. From a theoretical standpoint, the motion of turbulent premixed flames is a superposition of flame propagation and turbulent fluid convection. This renders their mathematical modelling significantly more challenging than that of non-premixed flames. Greater understanding of premixed flames and development of accurate predictive techniques is, therefore, absolutely necessary. Since the non-reactive aspects in a turbulent reactive flow have been already discussed, the following sub-sections focus on the two main processes that chemical species undergo – transformation and transport.

2.2.1 Combustion Chemistry

In hydrocarbon combustion, 50–7000 reacting species may be actually involved in a complete mechanism describing the transformation of the fuel-air mixture into products. However, known chemical mechanisms for most fuels used with CFD analysis comprise 150–250 [4]. Solving transport equations for even the decreased number of species is a computationally daunting task. Therefore, reduced mechanisms that represent the chemistry in terms of few key species and scalars are sought as a chemistry model for most applications of practical concern.

Species Transport

The generalised conservation equation of chemical species is written as

$$\frac{\partial}{\partial t}\rho Y_k + \frac{\partial}{\partial x_j}(\rho u_j Y_k) = -\frac{\partial}{\partial x_j}\mathscr{J}_{j,k} + \dot{\omega}_k, \tag{2.17}$$

where $\mathscr{J}_{j,k}$ is the molecular diffusion term and $\dot{\omega}_k$ is the production rate of the $k^t h$ species. Species molecular diffusion is generally described using Fick's law:

$$\mathscr{J}_{j,k} = -\rho \mathscr{D}_k \frac{Y_k}{x_j} \tag{2.18}$$

where $\mathscr{D}_k$ is the coefficient of molecular diffusion for the kth chemical species. Premixed flames are adequately and inexpensively defined by a reaction progress variable (c) in low-Mach number flows under adiabatic conditions and unity Lewis number [32]. Typically, this may be defined using a product species as

$$c = \frac{Y_P}{Y_{P,eq}}, \tag{2.19a}$$

or alternatively, using reactant concentrations as

$$c = \frac{Y_R - Y_{R,i}}{Y_{R,eq} - Y_{R,i}}. \tag{2.19b}$$

where $Y_{R,i}$ and $Y_{R/P,eq}$ are the initial and equilibrium concentrations of appropriately chosen reactant, R or product, P. In general, the reaction progress variable

must transform monotonically across a flame, and preferably its gradient is evenly distributed across the reaction. For such a progress variable, the Favre-averaged species conservation equation can be written as

$$\frac{\partial \overline{\rho}\widetilde{c}}{\partial t} + \frac{\partial \overline{\rho}\widetilde{u}_j\widetilde{c}}{\partial x_j} = \frac{\partial}{\partial x_j}\left(\overline{\rho \mathscr{D}_c \frac{\partial c}{\partial x_j}}\right) - \frac{\partial \overline{\rho u_j'' c''}}{\partial x_j} + \overline{\dot{\omega}}_c \tag{2.20}$$

where $\overline{\dot{\omega}}_c$ describes the averaged transformation rate of the progress variable. The density weighted averaged chemical diffusivity remains roughly constant; thus, the following simplification is achieved

$$\overline{\rho \mathscr{D}_c \frac{\partial c}{\partial x_j}} \approx \overline{\rho \mathscr{D}_c} \frac{\partial \widetilde{c}}{\partial x_j}. \tag{2.21}$$

Chemical transport equations in turbulent combustion using RANS decomposition, therefore, results in yet another closure problem – resolving averaged turbulent chemical transport and averaged chemical source-terms. Species turbulent transport is typically modelled using a gradient diffusion assumption

$$\overline{\rho u_j'' c''} = -\overline{\rho} \mathscr{D}_T \frac{\partial \widetilde{c}}{\partial x_j} \tag{2.22}$$

where $\mathscr{D}_T$ is an effective turbulent molecular diffusivity of the species associated with the progress variable, c. Although this model for species turbulent transport is widely debated due to the known existence of counter-gradient diffusion regimes, the aforementioned model provides an inexpensive closure. It is, thus, widely used, particularly in the context of RANS.

Chemical Transformation

The forward reaction rate of a generic chemical transformation of S species via M reactions represented concisely as

$$\sum_{s=1}^{S} \nu_{m,s}^{r} [A]_s \xrightarrow{k_f} \sum_{s=1}^{S} \nu_{m,s}^{p} [A]_s \quad with \quad m = 1, \ldots, M, \tag{2.23}$$

is a non-linear function of the concentration of the participating reacting species ($[A]_s$). The production rate of the k^{th} reacting species is expressed as

$$\dot{\omega}_k = \sum_{m=1}^{M} k_m \left(\nu_{m,k}^{p} - \nu_{m,k}^{r} \right) \prod_{s=1}^{S} [Y_s]^{\nu_{m,s}^{r}} \tag{2.24}$$

wherein k_m is the reaction rate coefficient and Y_s denotes the concentration of the s^{th} species. To aggravate the non-linearity of the calculations, rate coefficients of chemical reactions bear a strong non-linear dependence on temperature, which governs activation energy of the reaction. The relation can be expressed in the Arrhenius form [32] as

$$k_m = A \cdot T^{\alpha} \cdot \exp\left(-\frac{E_a}{RT} \right) \tag{2.25}$$

where the constant A is termed the pre-exponential factor and α is the non-linearity constant of temperature. The activation energy E_a, quantifies an energy barrier that must be overcome for the reaction to proceed in the forward direction.

It is evident that averaged chemical reaction rates are non-linear functions of scalar fields. These cannot be closed using other quantities solved for in the turbulent combustion system, i.e. the averaged source-term for k^{th} species or the corresponding progress variable cannot be described by the averaged flow scalars from the chemistry model.

$$\overline{\dot{\omega}_c(\rho, Y_k, T)} \neq \dot{\omega}_c(\overline{\rho}, \overline{Y_k}, \overline{T}), \tag{2.26}$$

2.2.2 Regimes of Combustion

Given the complex, non-linear law describing the burning rate $\overline{\dot{\omega}}_c$, a physical approach is often sought to derive models for turbulent combustion. Physical analysis of chemical source-terms relies on a knowledge of the chemical and turbulent scales involved in a turbulent reactive flow. Representation of flame behaviour and its descriptive turbulence-chemistry interactions is the subject of the *Borghi diagram* [33, 34]. The objective is to analyse premixed turbulent combustion using characteristic length and time scales corresponding to chemistry and turbulence.

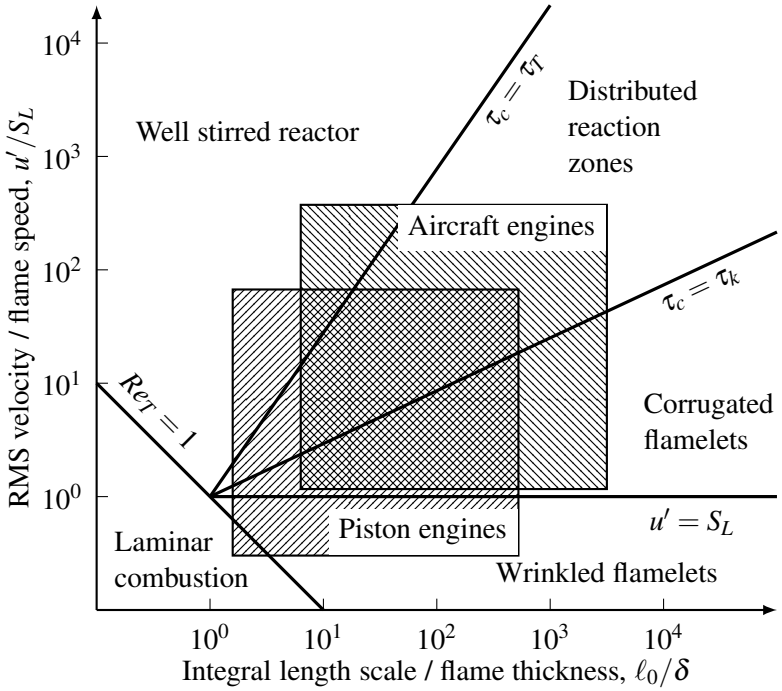

Figure 2.2: Classical turbulent premixed combustion regime diagram.

Fig. 2.2 shows the diagram wherein the oncoming turbulence intensity normalized by the laminar burning velocity u'/S_L is plotted against the largest turbulent eddy length scale normalized by the laminar flame thickness ℓ_0/δ. The turbulent flow is characterised by the turbulent Reynolds number

$$Re_T \equiv \frac{u'\ell_0}{\nu} \tag{2.27}$$

where u' is the turbulent velocity fluctuation which is related to turbulent kinetic energy as $u' \sim \sqrt{k}$, ℓ_0 is the turbulence integral length scale and ν is the kinematic viscosity of the fluid. When the turbulent Reynolds number is smaller than one, laminar combustion is observed. The remaining portion of the diagram represents turbulent combustion which is well scoped by two other dimensionless ratios.

The regime diagram in Fig. 2.2, illustrates the approximate placement of turbulent premixed flames observed in two common combustion systems – aircraft

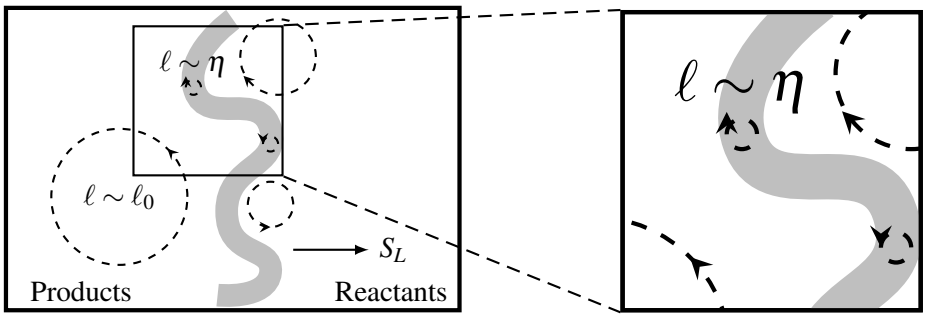

Figure 2.3: Turbulence-chemistry interaction in premixed flames.

and automobile engines. The Damköhler number compares macroscopic turbulent time scales (τ_t) and chemical reaction time scales (τ_c) as

$$Da \equiv \frac{\tau_t}{\tau_c}. \tag{2.28}$$

For large values of the Damköhler number (Da $\gg$1), the flame front is thin and its inner structure is not affected by turbulence motions at small scales. Large scale turbulent structures wrinkle the flame surface imparting strain in different regions of the flame-sheet. At each point on this sheet, the flame may be described as a one-dimensional laminar flame element called *flamelet*. Exceedingly low Damköhler numbers ($Da \ll 1$), on the other hand, correspond to slow chemical reactions wherein mixing always occurs prior to reaction as regime termed *perfectly stirred reactor*. A transition away from the flamelet regime occurs when the smallest turbulence scales have a time scale (τ_k) smaller than (τ_c). The Karlovitz number defines the Klimov-Williams criterion corresponding the value of the ratio

$$\text{Karlovitz number, Ka} \equiv \frac{\tau_c}{\tau_k}. \tag{2.29}$$

being equal to unity. This criterion helps delineate the flamelet regime from the *distributed reaction zones* regime where the inner flame structure is strongly affected by turbulent motions.

Fig. 2.3 illustrates the structure of turbulent premixed flames in a high turbulence-intensity flow wherein Kolmogorov scale turbulent motion may occur within the flame. It is noteworthy that in any given turbulent flow system, a wide range of dissipation occurs – therefore, a turbulent premixed flame is most generally represented by a zone that may comprise smaller turbulent structures within. The true structure of turbulence-chemistry interactions forms a topic of significant debate [4]. Qualitative analyses of turbulent premixed combustion provide a starting point for several quantitative chemical reaction rate modelling techniques. The dimensionless parameters that characterise the location of a premixed flame on the diagram can guide the choice of the most appropriate modelling strategy.

2.2.3 Chemical Closure Methods

Closure of the chemical source-term can be addressed using one of several combustion models; other than basic algebraic correlations, chemical closure techniques comprise three classes: flamelet models, CMC methods and PDF methods.

Algebraic Models

The chemical source-term can be described as a function of easily acquired quantities such as turbulent mixing. However, such models neglect local stochastic behaviour and must be tuned to match the turbulent premixed flame problem considered. Although these descriptions are commonly used in the context of RANS [35], more accurate models have been sought for LES [7]. An important class of models, termed the eddy break-up model (EBU), for example, assumes that chemistry is fast. It is then, the mixing of reactants that is the rate determining process of average reaction rate. If the time scale of mixing is taken to be the of the order of turbulence time scale, the following model is obtained

$$\overline{\dot{\omega}} = C \cdot \frac{\varepsilon}{k} \overline{\rho c''^2} \tag{2.30}$$

wherein the progress variable is used as a measure of scalar fluctuation magnitude [32]. The EBU and other similar models are widely used because of their simplicity. Although such models can be implemented in LES contexts, the incorporation of chemistry effects is difficult besides the fact that these are entirely *ad hoc* and

may result in non-physical solutions [32].

Flamelet Models

The flamelet assumption forms the basis for geometrical analyses of turbulent premixed flames. The flamelet method, conceptualized by [36], resolves turbulence-chemistry interactions by describing the turbulent flame locally with characteristics of a steady, one-dimensional laminar flame. In the context of LES, two widely used techniques comprise flamelet-based chemical closure [23]: the flame surface density approach (FSD) and the G-equation model. In the (FSD) approach, the flame is identified as a surface which can be convected, diffused, curved and strained by the turbulent velocity field [37]. The mean chemical source-term is calculated from the total reaction rate and the flame surface density

$$\overline{\dot{\omega}}_k = \overline{\dot{\Omega}}_k \Sigma \tag{2.31}$$

where $\overline{\dot{\Omega}}_k$ is the integrated reaction rate of the surface of the flame and Σ, the mean flame surface density. A similar approach introduces a flame wrinkling description using the ratio of flame surface and its projection in the direction of projection [38]. However, the accurate closure of the turbulent flux and propagation velocity of the flame surface, and the effects of curvature and strain rate on the flame surface remains a challenge [15, 23, 32].

In contrast, the G-equation model borrows from the description of a laminar flame surface using a level set [23, 39]. This concept eliminates the need for resolving the flame-front structure, but consistency with LES filtering has only been achieved recently by [40]. Further, the lack of resolution of the flame-front structure challenges the accuracy of chemical closure when the flame structure is adequately affected by turbulence. The thickened flame model was proposed by [41] as a method to resolve the flame structure on the mesh. However, the empirical parameters used to thicken the flame are cannot be applied to arbitrary flame specifications.

As a less restrictive alternative to the geometric analysis of flame-fronts, the statistical behaviour of a turbulent reactive flow is represented using a PDF. In premixed laminar flames, the statistics of scalars are well characterized by the re-

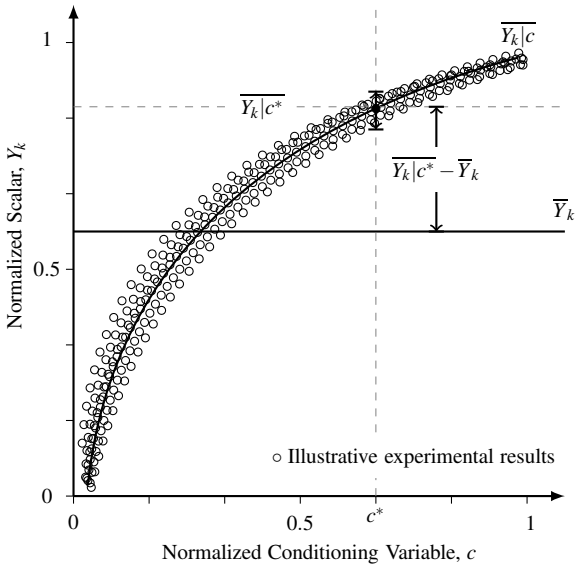

Figure 2.4: Conditional averaging of a scalar in a turbulent reactive flow.

action progress variable defined such that $c = 0$ in fresh gases and $c = 1$ in burnt gas. One line of approach then is to presume the PDF of the progress variable and obtain chemical closure using conditional statistics (as shown in Fig. 2.4) based on the solution of a laminar flame. The Bray-Moss-Libby model combines such an approach, using a bimodal distribution, with a physical analysis which accounts for differences in fresh and burnt gases. Variations of the flamelet/progress variable method have been applied to LES using a presumed PDF [42–44]. As mentioned earlier, and illustrated in Fig. 2.3, these *flamelet* models are only accurate in high Damköhler regimes, but are insufficient for modelling high turbulence-intensity turbulent premixed flames [45].

Conditional Moment Closure

Formulated by two researchers, Klimenko [46] and Bilger [47] independently, Conditional Moment Closure (CMC) is a well established chemical closure tool for the simulation of non-premixed combustion [15, 32]. It has been used successfully for industrial applications including spray combustion in diesel engines [48]. CMC offers a promising closure tool for simulating premixed combustion since it

is devoid of the assumption of a prevailing combustion regime. Recently, Amzin et al. [49] simulated stoichiometric pilot-stabilized Bunsen flames using the RANS-CMC method. Further, an algorithm was recently proposed for the application of CMC for LES of premixed flames [50].

The CMC method is based on the concept that small fluctuations are experienced from the average within an ensemble of realizations which comply with a given condition [47]. Fig. 2.4 illustrates that the deviation about an average across all possible conditions is seldom accurate. Whereas, for at each given condition, fluctuations about the average of the ensemble of realisations is minimal. Conditional averages of scalars are defined as

$$Q(\zeta,\vec{x},t) \equiv \langle Y(\vec{x},t)|\xi(\vec{x},t)=\zeta\rangle \equiv \langle Y|\zeta\rangle \tag{2.32}$$

where the angle brackets denote an ensemble average over an ensemble of realizations of the flow and the vertical bar indicates that this average is conditional on the condition, $\xi(\vec{x},t)=\zeta$ that is, only those members of the whole ensemble that meet this condition are included in the average [47].

Now, instead of presuming the reaction rate that is conditionally averaged on the reaction progress variable, as in conventional flamelet based statistical approaches, CMC calculates this term from the hypothesis

$$\overline{\dot{\omega}_k|\zeta} \approx \dot{\omega}_k(\overline{T|\zeta},\, \overline{Y_k|\zeta},\, \overline{\rho|\zeta}), \tag{2.33}$$

assuming small fluctuations of reaction rate about its conditional average. Transport equations are solved for conditionally averaged scalars required to estimate the conditioned reaction rate. This adds to the dimensionality of the problem and makes CMC computationally expensive. The conditionally averaged reaction rate obtained can be integrated by writing

$$\overline{\dot{\omega}_k(\vec{x},t)} = \int_0^1 \overline{\dot{\omega}_k|c^*}(c^*)\, P(\vec{x},t;c^*)\, dc^* \tag{2.34}$$

where $\overline{\dot{\omega}_k|c^*}(c^*)$ is estimated here through the CMC hypothesis and the presumed PDF, $P(\vec{x},t;c^*)$ for any point in the domain is estimated using the first two mo-

ments of the conditioning variable. As opposed to the flamelet method, CMC has the capability of accurately modelling the chemical source-term for flames beyond the flamelet regime [33]. However, transport equations typically solved for conditionally averaged scalars consist of various unclosed terms which are not well understood and are modelled inaccurately [31].

Transported PDF models

Various other statistical methods have been developed such as the linear-eddy model (LEM) and the one-dimensional turbulence (ODT) model by [51, 52] and the broad class of transported PDF models developed for turbulent reactive flows by Pope [53]. In this formulation, the chemical source-terms are closed exactly, whereas turbulent transport and molecular mixing must be modelled. The relatively recent multiple mapping conditioning (MMC) method combines features of CMC, PDF methods and mapping closure models [54, 55]. Although these methods provide excellent chemical closure, the computational expense of Monte Carlo solution discourages application to industry-scale problems.

Conditional Source-term Estimation

Conditional Source-term Estimation (CSE) was formulated by Bushe and Steiner [6], Steiner and Bushe [56] as a chemical closure model for the LES of turbulent non-premixed combustion. The general concept of CSE follows from the CMC method wherein the conditionally averaged source-term is approximated by conditionally averaged scalars relevant to the chemistry model from the CMC hypothesis (Eq. 2.33). Similar to CMC, the averaged chemical source-term is calculated by integrating the corresponding conditionally averaged source-term with the presumed PDF at the given spatial coordinate with respect to the conditioning variable.

However, while conventional CMC models solve transport equations for the conditionally averaged scalars, CSE employs an integral equation similar to Eq. 2.34. Each Favre-averaged scalar (ϕ) necessary for calculating the conditionally averaged reaction rate can be written in terms of the corresponding conditionally averaged scalar as

$$\widetilde{\phi}_k(\vec{x}) = \int_0^1 \overline{\phi_k|c^*}(\vec{x},c^*)\,\widetilde{P}(\vec{x},c^*)\,dc^*, \tag{2.35}$$

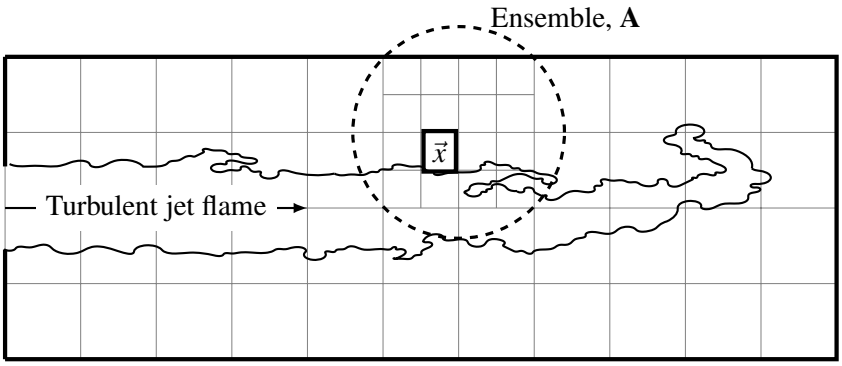

Figure 2.5: Chemical closure in conditional source-term estimation.

for the corresponding spatial coordinate ($\vec{x}$) in the physical domain where the $\widetilde{P}(\vec{x},c^*)$ represents the PDF at that point. This integral equation can be averaged over an ensemble of discrete points (A) in the domain (where the conditional scalar is constant) to yield an inverse problem for the calculation of the conditional scalar. Formally, clusters of points that preserve the statistical homogeneity of conditionally averaged scalars are chosen as ensembles. The ensemble averaged conditional scalar is then independent of the spatial coordinate and the integral equation is written as

$$\widetilde{\phi}_k(\vec{x}_j) = \int_0^1 \overline{\phi_k|c^*}(c^*)\,\widetilde{P}(\vec{x}_j,c^*)\,dc^*, \qquad j \in A \tag{2.36}$$

for each point j within the ensemble. Conditionally averaged scalars have been observed to exhibit statistical homogeneity over spatially localized ensembles of points [56]. Fig. 2.5 highlights a chosen ensemble of localized points in the computational domain of turbulent jet flame. With the conditionally averaged scalar now invariant of the jth spatial coordinate ($\vec{x}_j$) in these ensembles, the previous equation can be treated as the well-known Fredholm integral equation of the first kind [57], written in the discrete form as

$$\vec{b} = \mathbf{A}\vec{\alpha}, \quad A_{ji} = \int_{c_1}^{c_2} \widetilde{P}(c^*;\vec{x}_j)dc^*, \tag{2.37}$$

where $\widetilde{P}(\vec{x}_j,c^*)$ is the kernel, $\vec{b}$ is the known array of scalar averages in the entire ensemble and $\vec{\alpha}$ represents the conditionally averaged scalar. Integrations of the

presumed PDF required for calculating A_{ji} are done *a priori* and tabulated for m bins using lower (c_1) and upper (c_2) limits for the i^{th} bin of conditioning variable. The Tikhonov method is implemented for regularising the solution of the integral equation using an L-curve approach for the choice of an optimal regularisation parameter as discussed by Salehi et al. [58].

The application of CSE to the simulation of turbulent premixed flames was accomplished by Salehi [59]. As such, CSE has only ever been used for simple canonical reference flames, both in the context of non-premixed and premixed flames. The following work focuses on advancing CSE for application to problems of industrial relevance. Such problems generally involve transient phenomena which have not been modelled using CSE before. Moreover, industrial combustion chambers have complex geometries and bounding walls which make the accurate implementation of CSE more challenging. In the next chapter, a robust method is developed to extend the applicability of localized CSE inversions to the complex geometries generally encountered in an industrial problem. The following chapter employs CSE for the simulation of a homogeneous-charge natural gas spark-ignition engine.

3 | Spatial Partitioning Algorithm

3.1 Introduction

Data classification is a frequently encountered solution approach in a wide range of scientific disciplines. Spatial partitioning is a classification technique that is critical for obtaining accurate solutions of physical problems wherein constitutive local phenomena demand segregated treatment. In complicated domain geometries, solutions can be achieved conveniently by partitioning the problem space into localized sub-domains. Moreover, with the advent of parallel processing technology, mesh partitioning has become an essential tool to address large-scale problems by dividing computational load across the available computing resources. In this context, size governs load balance whereas locality designates communication between processors. The choice between load and communication is particularly difficult to make for problems with multiple requirements.

The partitioning technique presented in the current work is motivated by a stringent demand posed by a relatively recent model, termed Conditional Source-term Estimation (CSE), applied for the computational fluid dynamics (CFD) simulation of turbulent reacting flows [6, 58, 60, 61]. The key underlying assumption is that conditionally averaged scalars requisite for the CFD solution are statistically homogeneous in localized ensembles of points. In simple geometries, equi-sized slices of the domain along its major axis have sufficed [56, 62]. However, automated construction of ensembles is absolutely necessary for applying the model

Preprint to be submitted to Journal of Computational Physics.

to the arbitrarily complex geometries that have industrial relevance. Also, spatial proximity *must* be maintained within each cluster to retain the validity of the underlying assumption. In addition, parallel implementation of this large-scale problem demands that a measure of load balance be considered in the partitioning procedure. Finally, frequent mesh changes enforced by industrial demands or by adaptive mesh refinement (AMR) impose a limit on the computational expense of partitioning.

What makes this problem particularly challenging is that locality is not merely desirable from a standpoint of computational efficiency, but is critical for calculating a valid solution. Load balance has relatively less significance, and as such, were there to be a trade-off between load balance and strict locality, strict locality must be preferred.

3.1.1 Relevant Background

The problem of clustering N objects into M groups of objects with similar properties is a persistent challenge in several academic fields. Frequent applications are found in computer vision, pattern recognition, geographical information systems, large-scale database management, cache performance and many other fields [63, 64]. Statistical classification, such as the cognitive mapping developed by Jenks [65], seeks to reduce the variance within classes and maximize the variance between classes. Hartigan [66] discusses the relevant statistical theory for clustering, whereas Davé and Krishnapuram [67] present a unified view of robust clustering methods. These age old methods are efficient for applications with few dimensions, but it is recognized that the *curse of high-dimensionality* pervades the problem of clustering nearest point neighbours [68].

In CFD problems, partitioning corresponds to clustering regions of mesh based elements such as nodes, finite elements or finite control volumes (also called *cells*). For a multi-block finite volume mesh, popularly used in AMR implementations, blocks of cells can be divided into clusters assigned to different processors in the computing architecture. Such computational mesh partitioning has been the focus of extensive research [69–72]. The choice of adopted strategy depends on known geometric information and the trade-off between partition quality and computa-

tional expense that best suits the problem. Dynamic problems are addressed using fast methods, whereas relatively slower but highly customized procedures are formulated to match specific criteria such as those imposed by the turbulent reactive flow model.

Over the last few decades, a variety of mesh partitioning techniques have been developed, majority of which fall into two broad classes. Geometric algorithms including recursive bisection techniques and the SFC approach need only the local geometric information, such as block-centre coordinates. As an example, the recursive coordinate bisection (RCB) developed by Berger and Bokhari [73] is widely accepted due to its simplicity and the wide range of its applicability. The unbalanced recursive bisection strategy provides an advantageous modification of the RCB method [74] and Gilbert et al. [75] investigate yet another geometric method of dividing an irregular mesh into equal-sized pieces with few interconnecting faces. In general, geometric partitioning is conceptually simple, so the corresponding algorithms are fast and easy to implement. Additionally, the space-filling curve (SFC) approach, discussed at greater length in section 3.1.2, provides a unified and scalable data structure relevant for a wide range of CFD operations [76, 77].

The other broad class of partitioning, graph partitioning, has been reviewed by Schloegel et al. [78]. In this method, nodes of a computational graph represent tasks that can be executed concurrently and edges represent the communication required between tasks from one iteration to the next. Domain partitioning using such graphs is challenging and is recognized as an NP-hard problem [79]. The greedy algorithm developed by Farhat and Lesoinne [80], *builds* each partition by starting with a vertex and adding adjacent vertices until the target size or expected computational load has been reached. A relatively advanced technique, termed recursive spectral bisection (RSB), was developed by Simon [69]. It provides excellent quality partitions at exceedingly high computational costs. Many standard graph-based partitioning packages have been developed to employ efficient global and local graph-based methods for domain partitioning. Teresco et al. [81] provide an extensive review of partitioning algorithms.

As mentioned earlier, the choice of adopted strategy depends on known information and the expense-quality trade-off that best suits the problem. The SFC

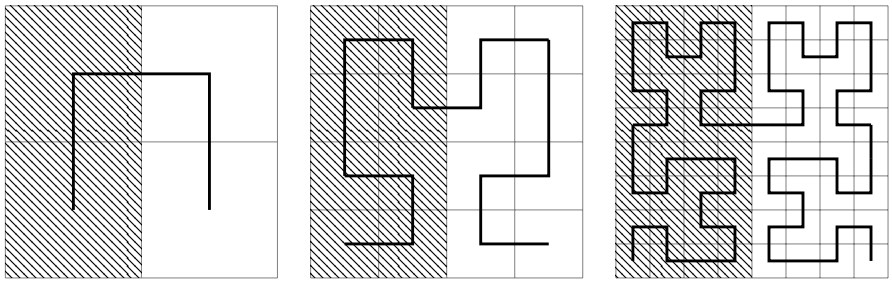

(a) Hilbert-Peano curve

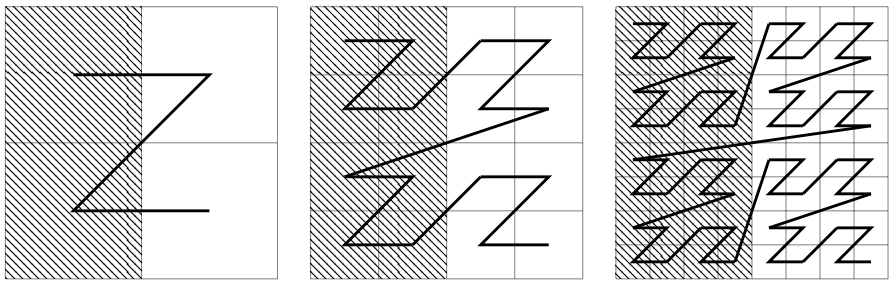

(b) Morton order curve

Figure 3.1: Contiguity of linear mapping within partitions using the (a) Hilbert curve, and the (b) Morton order curve. *Leaps* between consecutive points can be significantly larger in the Morton order curve.

approach not only provides a cheap clustering technique, but also an efficient data structure for performing CFD operations faster. In this work, we employ an SFC for to partition a multi-block finite volume mesh.

3.1.2 Space-filling Curves

An SFC is a continuous mapping of a d-dimensional space onto a one-dimensional discrete parametrized curve. The objective of this mapping is to retain spatial proximity of points in the linear space, i.e. neighbouring points in the physical domain remain close on the SFC. The central operation involved in the SFC partitioning approach is the reordering of mesh elements based on a locality index assigned uniquely in each SFC. The ordered curve can then be partitioned into equal sections wherein spatial locality relies on the accuracy of the original mapping func-

tion. The clustering properties of SFCs inherently ensure a reasonable shape and locality of clusters [82]. Also, in contrast to iterative or recursive partitioning algorithms, an SFC directly obtains a mapping for a given mesh. Additionally, the linear mapping provides a method to access mesh elements for performing SFC operations quickly. Due to these favourable properties, SFC have become an increasingly popular strategy for mesh partitioning.

Among the various, well-documented SFCs, the Hilbert curve and the Morton order curve are standard choices in CFD applications [77, 83]. While the Hilbert curve is known for its excellent proximity preserving behaviour [84], the Morton order curve is widely accepted for its simplicity and adaptability [85, 86]. Fig. 3.1 illustrates the differences between the two mappings shown for grids with progressively increasing refinements. A partition of the entire domain is generally not accessed completely by a contiguous part of the Morton order curve. The *leaps* characteristic to this curve occasionally span the entire physical domain. On the other hand, the Hilbert curve retains measurable and bounded cluster locality [87], subjectively observable in the illustration. However, construction of Hilbert curves is challenging, particularly for higher dimensional problems [88–90]. Moreover, Hilbert curves also contain *leaps* when used in arbitrary geometries. In the present work, we employ the Morton order curve although the partitioning algorithm developed may be used in conjunction with any SFC.

The Morton order curve is widely known for its compact construction rule [91]. A locality index, called Z-value, is generated for every coordinate by interleaving the bits in corresponding binary representations [91]. The Z-value of a coordinate in a 2-dimensional space is, hence, calculated as

$$Z(x_i, y_i) = x_i^1 y_i^1 x_i^2 y_i^2 \ldots x_i^m y_i^m \tag{3.1}$$

where x_i^j and y_i^j represents the j^{th} significant bit of x and y coordinates of the i^{th} point in its m bit representation. The interleaving concept can be extended to d-dimensions in general, at the expense of an increased frequency and size of the characteristic *leaps* between adjacent points [92]. In addition, unstructured grids can be mapped, but bit-interleaving of the representative floating-point coordinates is challenging. Connor and Kumar [93] introduced a Morton order curve construc-

tion mechanism for floating-point coordinates which avoids assigning the Z-value by sorting points based on their relative ordering. The algorithm is extendible to multiple dimensions and easily implemented, but the ordering has been observed to be *sensitive* to coordinate values. We adopt the floating-point Morton ordering algorithm after translating block-centre coordinates to avoid negative values. The spatial partitioning algorithm presented uses properties of the generated curve to obtain highly localized partitions.

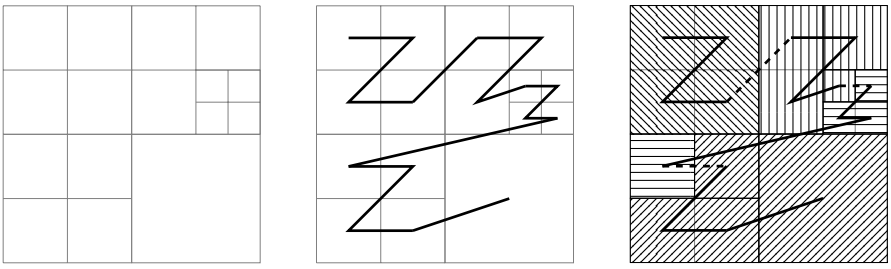

(a) Mesh geometry. **(b)** Morton order curve. **(c)** Obtained clusters.

Figure 3.2: Equi-sized partitioning of complex or refined mesh structures results in poor spatial locality within clusters. Regions of similar shading are assigned the same cluster.

3.2 Theory

SFC-based partitioning algorithms generally divide an SFC into equi-sized sections which represent point-clusters in physical space. In Fig. 3.2, an irregularly refined grid (Fig. 3.2(a)) is used to illustrate partitions obtained for complex geometry or a refined mesh by dividing the Morton order curve (Fig. 3.2(b)) into sections of equal size (Fig. 3.2(c)). Since two adjacent points on the Morton order curve may not be spatially local, equi-sized partitioning may allocate these points to a single partition regardless of the size of the *leap* (distance in the spatial domain) between the points. Such losses in partition quality may be acceptable for problems wherein spatial locality is not essential to the solution. However, in the context of our reactive flow problem, this would imply that statistical homogeneity of conditionally averaged scalars is not retained within a cluster containing a *leap*. Hence, the focus

Algorithm 3.1: Spatial partitioning algorithm for multi-block meshes.

Require: Morton ordered d-dimensional point set P of size n_b.
Ensure: C contains k sequential cluster sizes c_j in point set P.

1: **function** CLUSTER(**set** P, **int** k)
2: $R \leftarrow \{r_i\} \,|\, r_i \equiv |\vec{p}_i - \vec{p}_{i+1}| \forall\, i \in \{0,1,\ldots,n_b-2\}$
3: $j \leftarrow 1; \quad c_0 \leftarrow 1; \quad c_j \leftarrow 0 \forall\, j \in \{1,2,\ldots,k-1\}$
4: $r_{th} \leftarrow$ THRESHOLD(R,k) ▷ using Algorithm 3.2
5: **for all** $i = 0 \rightarrow n_b - 1$ **do**
6: **if** $r_i / r_{max} < r_{th}$ **then**
7: $c_j \leftarrow c_j + 1$
8: **else**
9: $j \leftarrow j+1; \quad n_c \leftarrow j$
10: **end if**
11: **end for**
12: $C \leftarrow$ BALANCE(C,R)
13: $C \leftarrow$ PARTITION(C,k)
14: **return** C
15: **end function**

16: **function** PARTITION(**set** C, **int** k)
17: **while** $n_c < k$ **do**
18: $n_c \leftarrow n_c + 1; \quad j_0 \leftarrow j \,|\, A_j = \text{MAX(A)}$
19: **for all** $j = n_c \rightarrow j_0 + 1$ **do**
20: $c_{j+1} \leftarrow c_j; \quad j \leftarrow j-1$
21: **end for**
22: $c_j, c_{j+1} \leftarrow c_j / 2$ ▷ halve largest cluster
23: **end while**
24: **return** C
25: **end function**

of our spatial partitioning algorithm is to avoid such locality losses by intelligently partitioning an SFC .

The guiding principle of our partitioning heuristic is to partition the SFC between points located across a *leap*. As presented in Algorithm 3.1, a normalization scheme is employed to identify *leaps* between consecutive points on a Morton order curve. The distance ($r_{i,i+1}$) between the i^{th} point and its next adjacent point is normalized by the maximum distance (r_{max}) between consecutive points on the

Algorithm 3.2: Threshold consecutive point distance for partitioning.

Require: Array R of consecutive point distances in Morton ordered set.
Ensure: r_{th} is maximum threshold distance that generates $n_c < k$.

```
1:  function THRESHOLD(set R, int k)
2:      r_min ← MIN(R),   r_max ← MAX(R);   r_step ← r_min/(2 × r_max)
3:      n_c ← 1;   r_th ← 1
4:      while n_c < k do
5:          for all i = 0 → n − 2 do
6:              if r_i/r_max > r_th then
7:                  n_c ← n_c + 1
8:              end if
9:          end for
10:         r_th ← r_th − r_step
11:     end while
12:     return r_th + r_step                          ▷ generate n_c < k
13: end function
```

curve as

$$\widehat{r_i} = \frac{r_{i,i+1}}{r_{max}}. \tag{3.2}$$

The largest *leaps* are identified first and the curve is partitioned at those locations. Using the step size defined in Eq. 3.3, progressively smaller *leaps* are identified, creating more partitions. The Morton order curve has several identically spaced consecutive points throughout the mapping due to its repetitive Z-like pattern. Therefore, partitioning the curve at all *leaps* which are larger than a certain threshold distance (r_{th}) results in more clusters than the required number (k). Since this would add an additional challenging task of merging neighbouring clusters, Algorithm 3.2 provides a method to obtain the maximum possible threshold distance which results in fewer clusters than required. The exact number of divisions can then easily be produced by halving the largest clusters.

$$r_{step} = \frac{r_{min}}{2 \cdot r_{max}} \tag{3.3}$$

The partitioning algorithm has been illustrated in Fig. 3.3 for a Morton order curve that has been constructed for a geometry identical to Fig. 3.2. The difference in

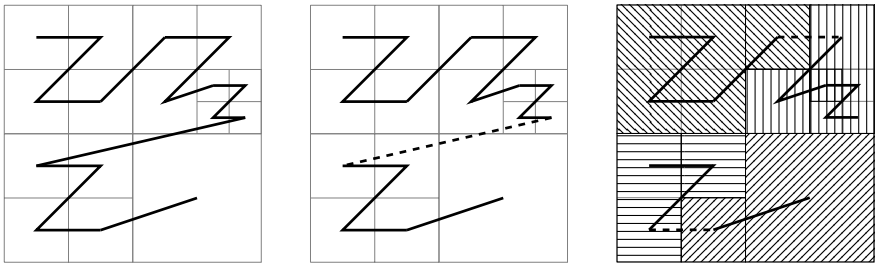

(a) Morton order curve. **(b)** Partition at *leaps*. **(c)** Obtained clusters.

Figure 3.3: Partitioning of the Morton order curve between points located across *leaps* increases cluster locality at the expense of losing cluster size balance.

spatial locality of obtained clusters is obvious. A quantitative analysis of a Morton order curve constructed for a simple geometry is performed in Section 3.4. A range of point-clouds representing relevant geometries (described in the following section) were partitioned using the algorithm.

Normalization of distances between adjacent points serves to maintain generality of the algorithm. However, a strict adherence to this metric for the identification of *leaps* ignores the size of clusters obtained. Often such a partitioning algorithm produces exceedingly wide range of cluster sizes, particularly for irregularly spaced points, resulting in a severe load imbalance. As a method to limit cluster size imbalance, Algorithm 3.3 merges clusters smaller than a threshold size (c_{th}) with their neighbouring clusters. A cluster is merged with its closest adjacent cluster only if the inter-cluster distance is smaller than a pre-defined value.

Discrete SFCs are not uniquely defined. In general, additional criteria are required to specify the curve that best preserves locality. Appropriately defined measures of spatial locality allow the construction of curves that minimize the metric in each cluster. Additionally, this provides a tool for meaningful comparison between two SFCs. The measures defined by Perez et al. [94] and by Gotsman and Lindenbaum [87] calculate the weighted sum of distances within a cluster, where the weights are inversely proportional to the spatial distance between the points. Another formal analytical measure for parallel domain decomposition using SFCs has been established for Hilbert curves by Tirthapura et al. [95]. A universal metric

Algorithm 3.3: Cluster coarsening for reviving computational load balance.

Require: Cluster size set C with n_c elements. Point-to-point distance set.
Ensure: C is devoid of small clusters that skew load balance.

```
1:  function BALANCE(set C, set R, int k)
2:      i ← c_0;   c_th ← n_p/(2 × k)
3:      r_max ← MAX(R)   r_co ← r_max/2
4:      for all j = 1 → n_c − 1 do
5:          i ← i + C_j
6:          if c_j < c_th then
7:              if r_i/r_max < r_co or r_{i−C_j}/r_max < r_co then
8:                  j_co ← {r_i < r_{i−C_j} ? j : j − 1}
9:                  C_{j_co} ← C_{j_co} + C_{j_co+1}
10:                 for all k = j_co + 1 → n_c do
11:                     C_k ← C_{k+1}
12:                 end for
13:                 n_c ← n_c − 1;   j ← j + 1
14:             end if
15:         end if
16:     end for
17:     return C
18: end function
```

for the Morton order curve is more challenging because of the irregular frequency and length of *leaps*, particularly in unstructured meshes.

We define a simple metric to quantify the locality preserved within clusters obtained using our spatial partitioning algorithm. Within a cluster of points, the root-mean-square distance (r_{rms}) from the centroid ($\vec{r_c}$) is a common method of represent the locality preserved for a given number of points. The maximum distance (r_{max}) between any points within the cluster is an estimate of the skewness. We define an index for each cluster as the ratio

$$L_R = \frac{r_{max}}{r_{rms}}, \tag{3.4a}$$

with the RMS distance written as

$$r_{rms} = \sqrt{\frac{\sum_{i=1}^{n_b} |\vec{r}_i - \vec{r}_c|}{n_b}}, \tag{3.4b}$$

where vectors $\vec{r}_i$ represent block-centre coordinates and n_b is the number of blocks in the cluster. The ratio L_R quantifies how skewed a given cluster is from the ideal spherical distribution of points considering the number of points in the cluster. In our locality analysis, we use L_R to compare the performance of partitioning methods. A similar ratio ($L_{R,c}$) is applied to quantify inter-cluster locality using the coordinates of cluster centroids instead of block-centres.

3.3 Test Geometries

The partitioning algorithm has been tested on three different geometries pertinent to turbulent reactive flow simulations. The experimentally studied cases comprise a Bunsen burner [96] and a bluff-body burner [97]. The third case is the combustion chamber in a Ricardo Hydra internal combustion engine, including a relatively simple piston [98]. Details of each geometry including prominent geometrical features have been listed in Table 3.1.

Specification		Dimensions (m)			
Case	Type	ℓ_x	ℓ_y	ℓ_z	Geometry Feature
Gulder	Burner	$5 \cdot 10^{-2}$	$5 \cdot 10^{-2}$	0.1	-
Cambridge	Burner	$8 \cdot 10^{-2}$	0.15	0.25	bluff-body
Ricardo Hydra	SI-Engine	$8 \cdot 10^{-2}$	$4 \cdot 10^{-3}$	0.21	piston-bowl

Table 3.1: Specification and characteristic dimensions of test cases.

Computational meshes for the two burner cases were constructed using an in-house multi-block meshing code. An axi-symmetric mesh representing the Ricardo Hydra engine geometry (at 330° crank angle) was constructed using *blockMesh* utility of OpenFOAM [99]. Details of all three meshes have been summarised in Table 3.2. These meshes form a broad range of point distributions suitable for a rigorous testing of the partitioning algorithm. All tests were performed on a

desktop computer with 64-bit Intel®Xeon®CPU E31225 with 3.10GHz $\times$ 4 cores and 8GB memory.

3.4 Results

We analyse the performance of our spatial partitioning algorithm primarily in the context of *load* and *locality*. Qualitative illustrations have been used to highlight the inaccuracy involved in equal sized partitions along with an assessment of computational expense. Finally, validation using an *a priori* case has been reported.

As an illustration of the partitioning mechanism employed, properties of the Morton order mapping have been plotted in Fig. 3.4 and Fig. 3.5 for various stages of Algorithm 3.1. A simple cylindrical mesh geometry with 3200 blocks representing Bunsen burner is being partitioned here to obtain 128 localized clusters. Fig. 3.4(a) graphs the normalized distance between consecutive points (which represent blocks in this geometry) on the Morton order curve. The conspicuous peaks denote *leaps* between adjacent points – the largest *leaps* have unity normalized length by definition. 125 point-clusters are obtained on partitioning the curve at all locations where the consecutive block distance is greater than the threshold (r_{th}). In Fig. 3.4(b) the inter-cluster separations (which are the lengths of corresponding *leaps*) have been plotted or each cluster. Whereas equi-sized partitions would contain 25 blocks each, partitioning at *leap* positions produces some clusters (groups of closely placed points in Fig. 3.4(b)) with fewer than 10 blocks. Fig. 3.5(a) shows the result of the coarsening procedure implemented using Algorithm 3.3. Closely placed clusters below a threshold distance are merged together, decreasing the number of clusters to 101. Finally, $k = 128$ clusters are obtained by halving the largest clusters. In Fig. 3.5(b), the troughs represent regions where halving was

Mesh Type		Mesh Specifications		
Case	Geometry	Case	Blocks	Cells
Bunsen	cylindrical	block-based	3,200	$1.63 \cdot 10^6$
Bluff-body	cylindrical	block-based	588	$4.7 \cdot 10^5$
Ricardo Hydra	axi-symmetric	cell-based	1	$2.1 \cdot 10^3$

Table 3.2: Mesh type and features for different test geometries.

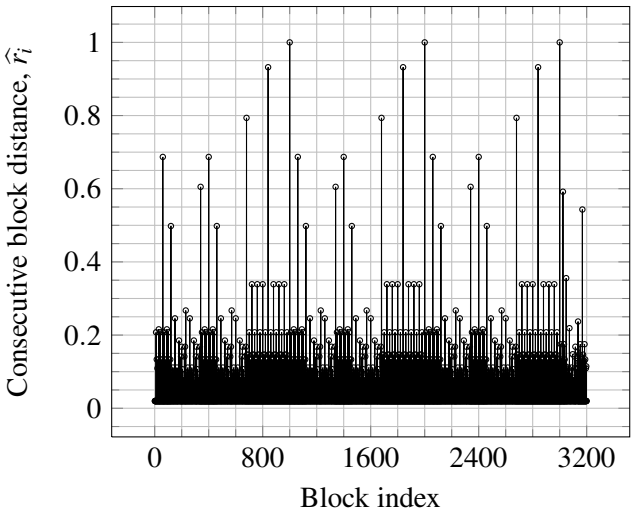

(a) Separations between adjacent blocks on the Morton order curve.

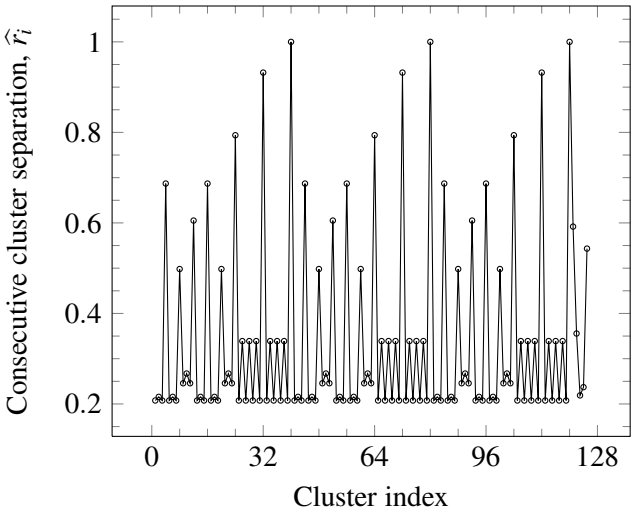

(b) Separations between adjacent clusters of blocks on the Morton order curve.

Figure 3.4: Distances between consecutive points (block-centres) on the Morton order curve displayed for the first two steps of the partitioning algorithm. The normalized distance between points ($\widehat{r}_i \equiv r_i/r_{max}$) has been employed for the case with $n_b = 3200$, and $k = 128$.

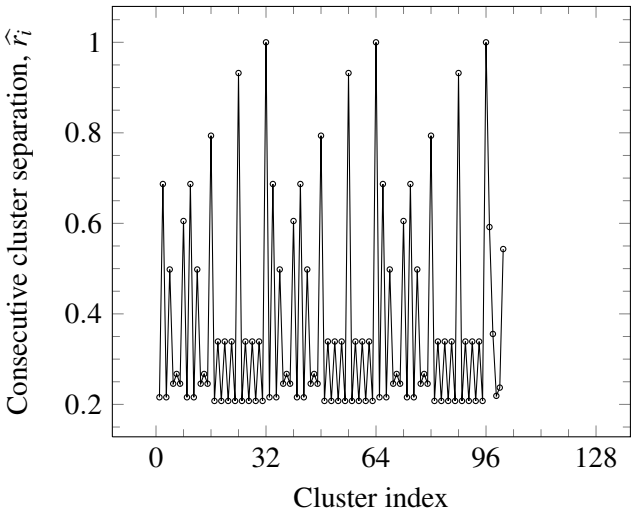

(a) Separations between adjacent clusters after the merging procedure.

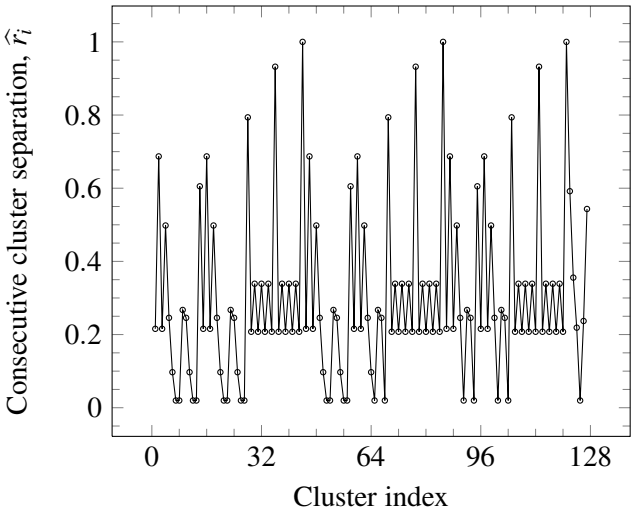

(b) Resultant adjacent cluster separations after halving the largest clusters.

Figure 3.5: Distances between consecutive points (block-centres) on the Morton order curve displayed for final steps of the partitioning algorithm. The normalized distance between points ($\widehat{r_i} \equiv r_i/r_{max}$) has been employed for the case with $n_b = 3200$, and $k = 128$.

done – clusters that are halved generally have high spatial locality within them resulting in very small inter-cluster distance between the halved partitions. Based on the strictness of spatial-locality desired, the threshold size of cluster can be changed which is expected to affect coarsening results and hence, the cluster-size balance.

Load balance is assessed by comparing resulting cluster sizes (c_i) with sizes obtained from equal partitions ($c_{eq,i}$). The ratio

$$\widehat{c_i} \equiv \frac{c_i}{c_{eq,i}} \tag{3.5}$$

then provides a relative workload comparison for problems where each point contributes equally to the computational load. The locality metric (L_R) defined in Section 3.2 has been used to quantify preserved spatial locality within each cluster. A similar locality metric ($L_{R,c}$) estimates the packing of clusters within a geometry.

Through Fig. 3.6 and Fig. 3.7, a comparison of the load balance and spatial locality within clusters has been made between equi-sized partitioning and our partitioning algorithm for the simple Bunsen burner geometry. In addition, the comparison is extended to manually partitioned axial slices obtain for the cylindrical mesh in Fig. 3.8(a). Relative cluster sizes obtained by the partitioning algorithm are roughly bounded between 0.5 and 2 with the exception of a few partitions constructed for the sparse case, $k = 32$. Alongside, the packing of clusters has been graphed in Fig. 3.8(b) using the inter-cluster locality metric. Equi-sized partitioning results in a number of poor quality partitions particularly for $k = 128$.

To visualize the poor quality of partitions often obtained from equi-sized SFC partitioning, a section of a such partitioned Bunsen burner mesh has been shown in Fig. 3.9. Although several partitions may be mapped by a single contiguous part of the curve, Fig. 3.9(c) and Fig. 3.9(d) clearly illustrate that several partitions contain *leaps*. Such clusters are unfavourable in terms of minimizing communication since they comprise of two (or sometimes more) sub-clusters of blocks. A parallel computation (with each processor assigned a cluster) would, therefore, have higher communication overhead for these clusters when compared to the cluster such as Fig. 3.9(b). More importantly, for the model implemented in our turbulent reactive flow simulations, these clusters would be entirely unacceptable since conditional averages of reactive scalars (such as temperature, for example) are expected to be

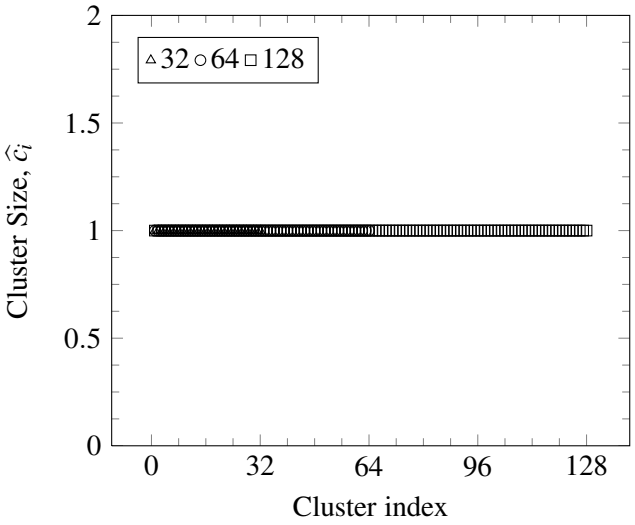

(a) Cluster sizes from strictly equal partitioning of the Morton order curve.

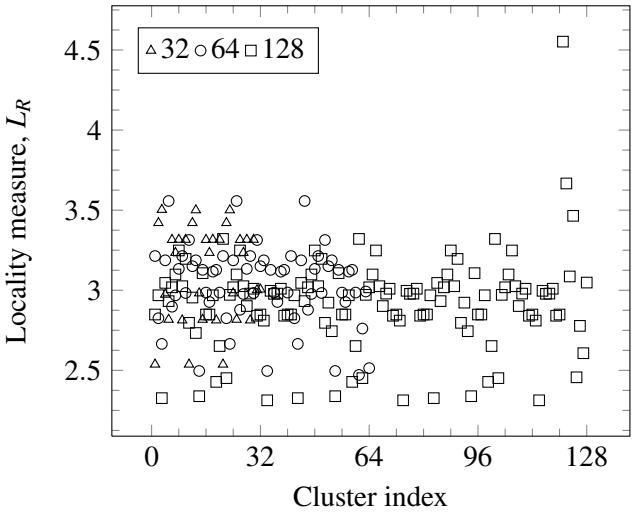

(b) Spatial locality of equi-sized clusters on the Morton order curve.

Figure 3.6: Degree of computational load balance for the Bunsen burner geometry gauged using relative cluster ($\widehat{c}_i$) for strictly equal partitioning. Interprocessor communication is indicated by the locality metric ($L_{R,c}$).

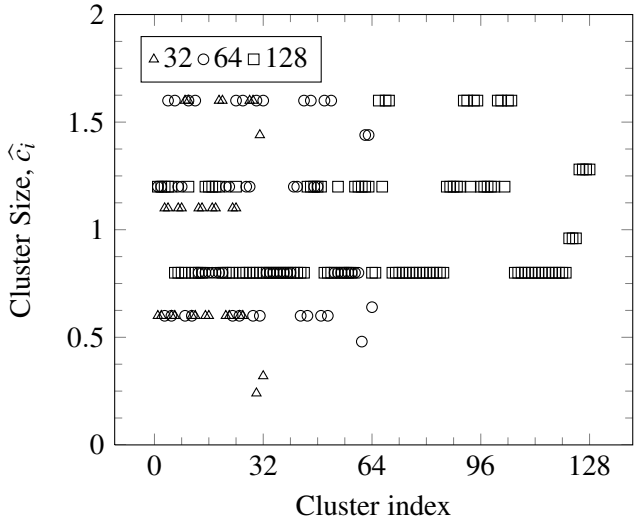

(a) Cluster sizes obtained from the spatial partitioning algorithm.

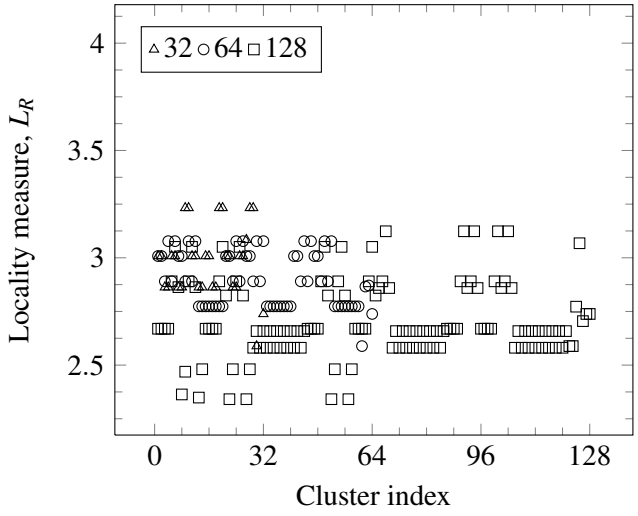

(b) Spatial locality of corresponding clusters on the Morton order curve.

Figure 3.7: Degree of computational load balance for the Bunsen burner geometry gauged using relative cluster ($\widehat{c}_i$) for *leap*-based partitioning. Interprocessor communication is indicated by the locality metric ($L_{R,c}$).

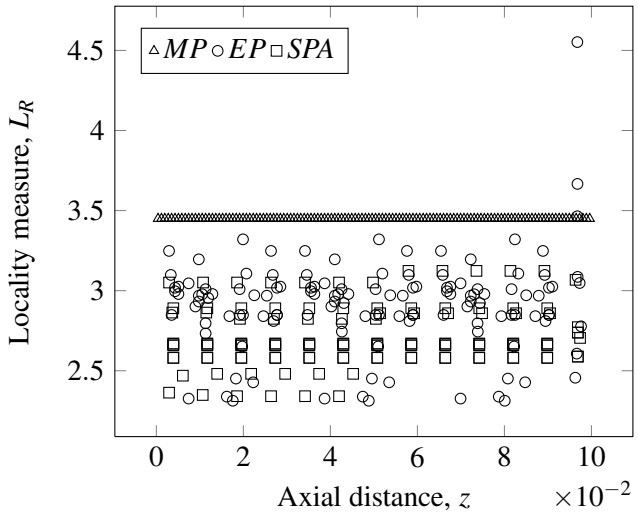

(a) Spatial locality within each cluster.

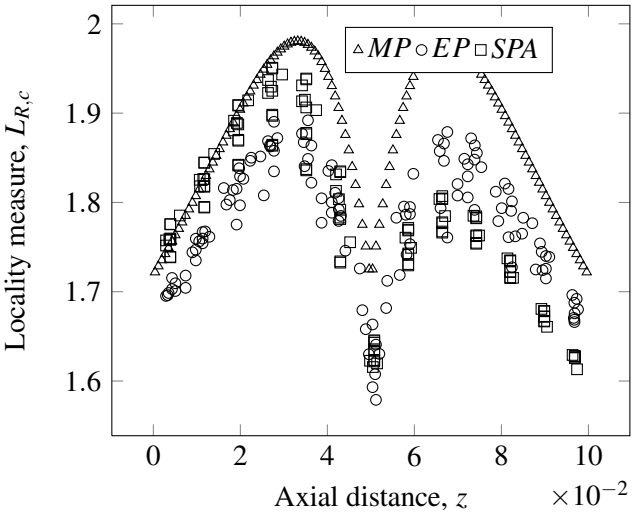

(b) Inter-cluster spatial distribution.

Figure 3.8: Comparison of intra-cluster (L_R) and inter-cluster ($L_{R,c}$) locality measures for manual partitioning (*MP*), equi-sized partitioning (*EP*) and using the spatial partitioning algorithm (*SPA*) on the Morton order curve in the Bunsen burner geometry for $k = 128$ divisions.

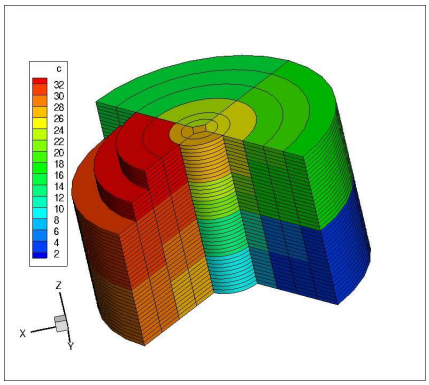

(a) Clusters obtained from equal partitions of Morton order curve.

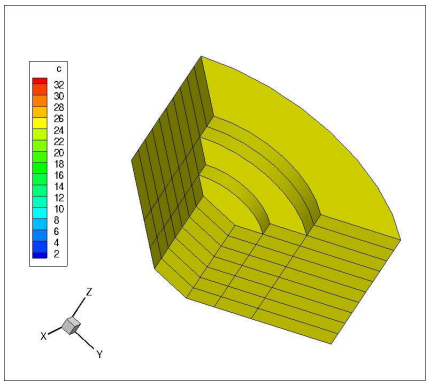

(b) Cluster with no *leaps*.

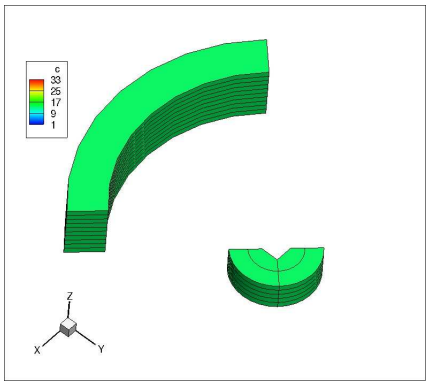

(c) Pathological cluster with *leap*.

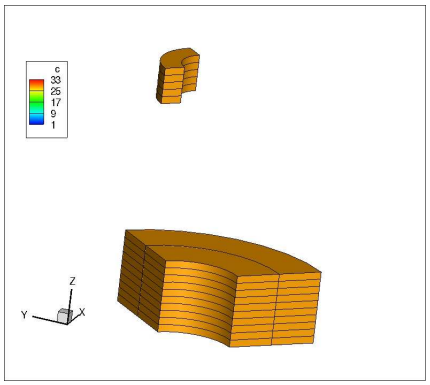

(d) Pathological cluster with *leap*.

Figure 3.9: Clusters produced from the equi-sized partitioning of Morton order curve often contain *leaps* that reduce spatial locality within. Even for a simple cylindrical Bunsen burner geometry, *leaps* occur at frequent and irregular intervals.

completely different in the individually localized, but separated, sub-clusters. Essentially, for these two pathological cases, the underlying assumption made in the formulation of the combustion model would be violated.

Results of *leap*-based partitioning have been illustrated using sections of partitioned geometries in Fig. 3.10 for the Bunsen burner, and in Fig. 3.13 for the relatively complex geometries accompanied with respective sectional views. Spa-

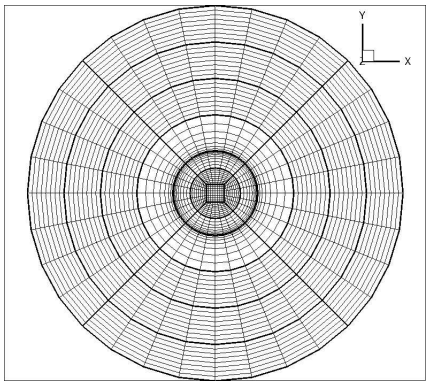

(a) Cross-section view.

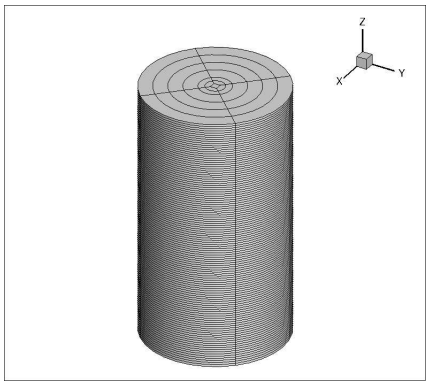

(b) Perspective view.

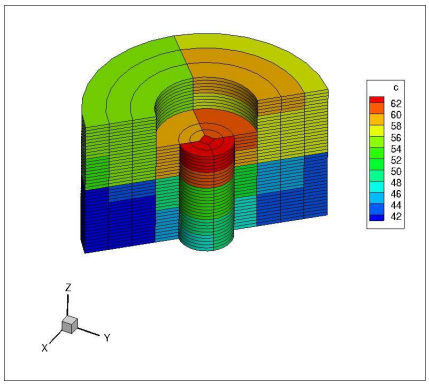

(c) Section of clusters obtain from equi-sized partitioning.

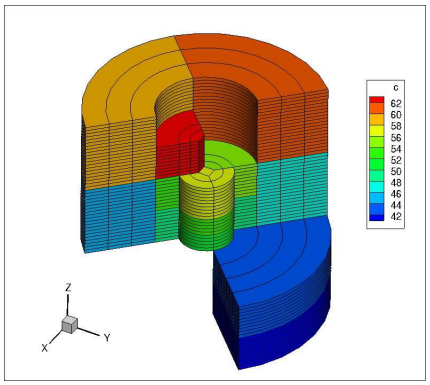

(d) Clusters as obtained from the spatial partitioning algorithm.

Figure 3.10: A qualitative picture of the difference between regular equi-sized partitioning and our spatial partitioning algorithm applied to Bunsen burner geometry shown above.

tial locality is well preserved in the axi-symmetric mesh despite a low number of cells. This emphasises the weakness of space-filling curves, especially the Morton order curve, in accurately preserving locality in 3-dimensional domains. In the geometries used in our study, points on the Morton order curve designate centres of blocks. Due to the varying shape and size of each block, an SFC mapping is expected to have lower quality for these cases than cell-based construction.

Fig. 3.11 and Fig. 3.12 assess the performance of the partitioning algorithm for

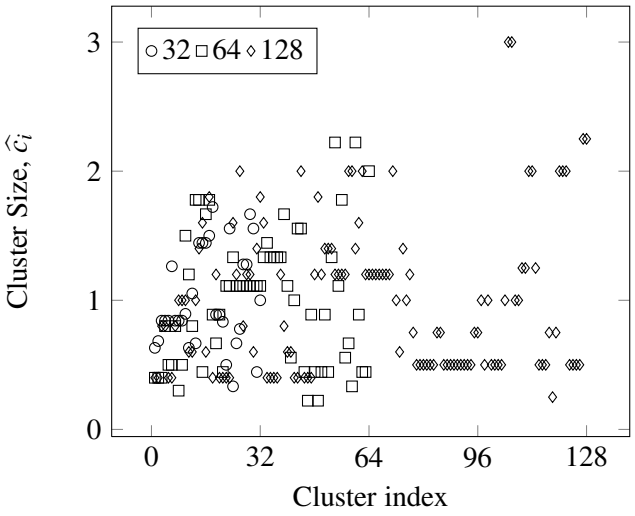

(a) Cluster sizes obtained for the bluff-body burner geometry.

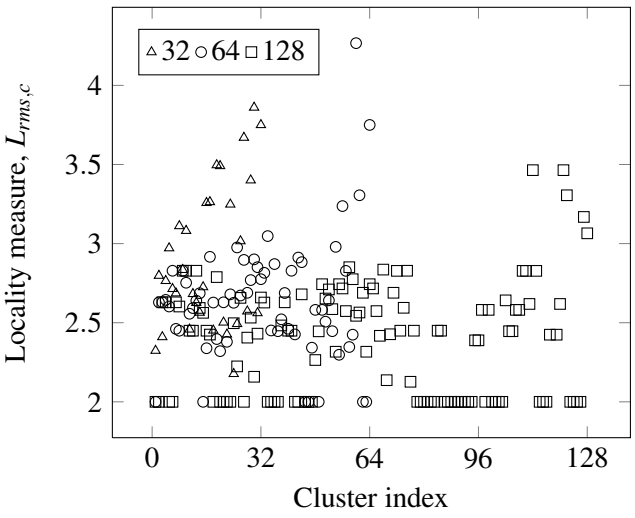

(b) Spatial locality of corresponding clusters on the Morton order curve.

Figure 3.11: Computational load balance measures ($\widehat{c}_i$ and $L_{rms,c}$) estimated for the bluff-body burner for different cluster divisions, $k = 32, 64$ and 128.

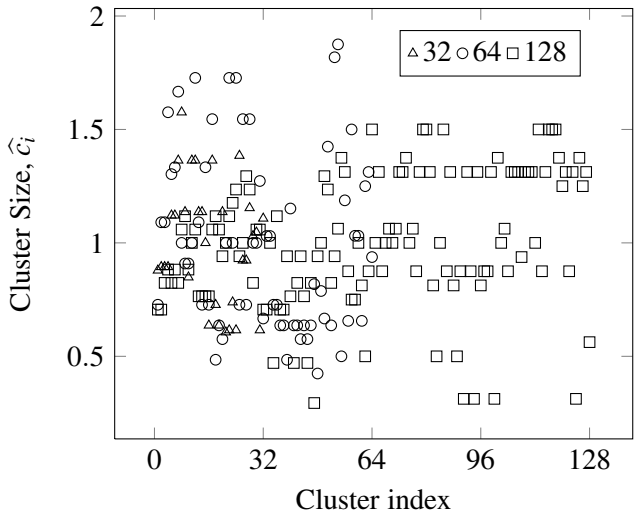

(a) Cluster sizes obtained for the Ricard Hydra engine geometry.

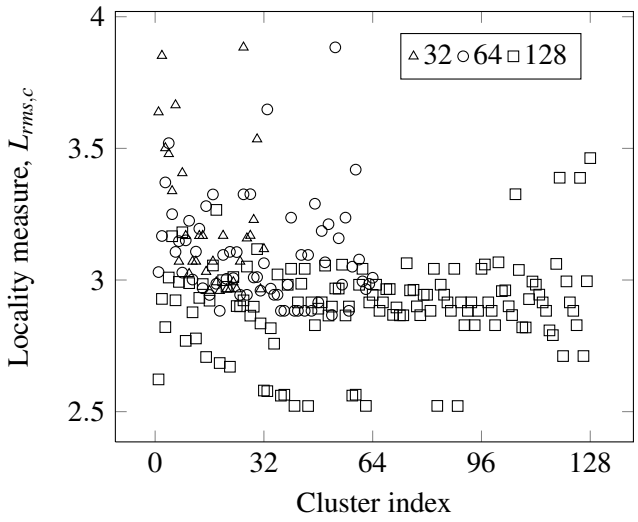

(b) Spatial locality of corresponding clusters on the Morton order curve.

Figure 3.12: Computational load balance measures ($\widehat{c}_i$ and $L_{rms,c}$) estimated for the Ricardo Hydra engine for cluster divisions, $k = 32, 64$ and 128.

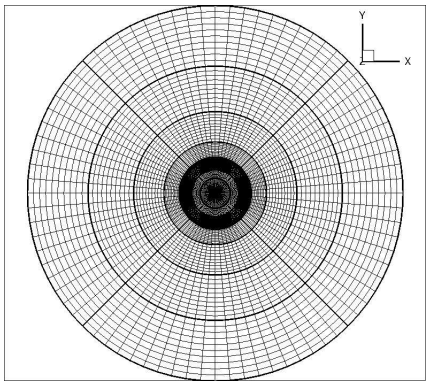

(a) Cambridge bluff-body burner.

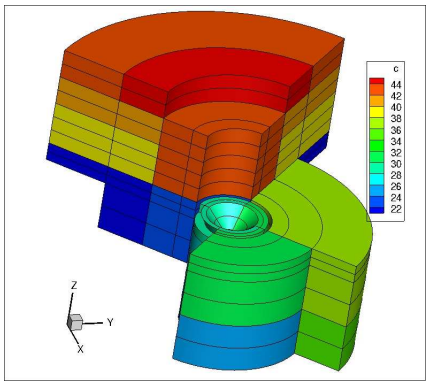

(b) Section of burner showing clusters obtained from algorithm.

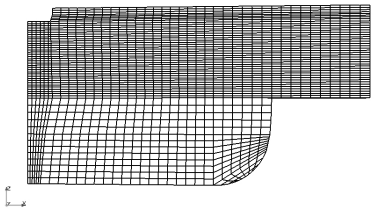

(c) Ricardo Hydra engine.

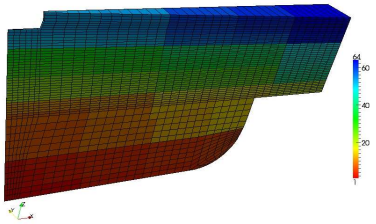

(d) Perspective of engine showing clusters obtained from algorithm.

Figure 3.13: Sections of complex geometries clustered using the spatial partitioning algorithm along with appropriate sectional views.

the two relatively complex geometries, the bluff-body burner and the Ricardo Hydra engine combustion chamber. It is observed that on average, clustering quality decreases with progressively more challenging geometrical features in the domain. Computational load is, however, maintained within the $[0.5, 2]$ bounds barring a few exceptions. This is expected since, the floating-point Morton order algorithm is *sensitive* to domain irregularities. Moreover, it is observed that, an optimum number of partitions, $k = k_{opt}$ provide the best load balance for a given geometry. Often in large-scale parallel computing, one is constrained to use a number of processors that is a power of 2 to take maximal advantage of the available re-

Case	Sensitivity Parameters (m)		Computational Expense (*ms*)		
	r_{max}	r_{min}	$k = 32$	$k = 64$	$k = 128$
Bunsen	0.03	0.0008	72.54	41.6	35.8
Bluff-body	0.25	0.0019	7.1	7.2	7.61
Ricardo Hydra	0.04	0.0002	33.9	28.9	27.9

Table 3.3: Performance of spatial partitioning algorithm in different geometries and inter-point distance parameters defining the corresponding Morton order curve.

sources. As such, the distribution of cluster-sizes for each case considered here has been graphed for 32, 64 and 128 clusters. For serial processing architectures, any desirable number of partitions (preferably, $k \approx k_{opt}$) may be constructed for serial calculations.

The space-filling curve partitioning approach is sensitive to the coordinate values provided since inter-point distances can significantly affect the algorithm. Several parameters such as the maximum and minimum inter-point distances have been tabulated in Table 3.3. In addition the run-times have been tabulated for different cluster numbers. A consistent correlation of execution time with mesh parameters such as cell numbers and inter-point distances is not be observed. However, the overall execution time is negligible, particularly for an operation that is expected to be done only once on a simulation with no moving mesh and without any mesh adaptation. On simulations with moving mesh or mesh adaptation, while this computational cost would also be negligible, it is likely that there would be considerable redistribution of data among nodes on a parallel machine; nevertheless, the computational overhead of recalculating the distribution would be very small and the redistribution is almost certain to be needed regardless of locality considerations, simply to maintain load balance. In comparison to graph-partitioning methods, the computational expense is low and encourages the use of SFCs for low to medium end parallel processing architectures.

3.5 Conclusions

A new algorithm for partitioning mesh for arbitrarily complex geometries has been proposed that favours maximizing locality at the possible expense of increased load imbalance. The method is based on the floating-point Morton order algorithm, which provides a computationally inexpensive alternative for the dynamic mesh partitioning in arbitrary geometries. Careful translation of negative coordinates is required prior to the mapping in order to avoid inconsistencies. Also, the lack of a unique Morton order curve for a given domain lends unwanted geometric sensitivity to the partitioning algorithm. Our spatial partitioning algorithm can be used in conjunction with the floating-point Morton ordering algorithm or with any existing SFC-based domain decomposition. Cluster locality is improved against strictly equal partitioning based on the Morton order curve. The technique can be applied to geometries relevant to industrial problems. An acceptable degree of computational load balance is achieved for large-scale parallel processing environments. The L_R locality metric quantifies spatial locality within clusters and roughly indicates inter-processor communication time.

4 | Engine Simulations

4.1 Introduction

Lean premixed combustion is increasingly employed by modern engines in order to reduce fuel consumption and NO_x emissions. Design and development of robust premixed combustion systems hinge on predictive computational tools that not only reduce dependence on experiments but may offer additional information. The choice of computational approach is a product of the trade-off between solution accuracy and the associated computational time. In this regard, LES is gradually becoming a computationally feasible approach for understanding cyclic variations in spark-ignition engines. In the meantime, RANS continues to offer an inexpensive alternative and provide acceptable design solutions in a time-frame suitable to the industry.

Closure issues arise inevitably from both these solution strategies; in particular, the closure of species turbulent fluxes ($\overline{\rho u_i'' u_j''}$) and chemical source-terms ($\overline{\dot{\omega}}_k$) is the objective of combustion modelling. Few models account for the complexities involved in turbulent chemical transport such as counter-gradient transport in reactive flows [100]. Although models based on the gradient transport hypothesis are relevant only for specific premixed combustion regimes, these are widely used due to their simplicity. Estimation of chemical source-terms has received greater interest due to its contribution to the overall accuracy of the solution [15, 32].

Preprint to be submitted to Combustion Science and Technology.

4.1.1 Relevant Background

Chemical closure for RANS can be broadly divided into three different approaches that respectively address the non-linear chemical source-term using an algebraic, a geometrical or a statistical analysis. Algebraic models represent the source-term as a function of easily acquired quantities such as turbulent mixing. Such models involve inaccuracies and must be tuned for the behaviour of turbulent premixed flame in context. These models have been hitherto commonly preferred in the context of RANS [35]. However, more accurate models have been developed and gaining popularity in recent years.

Geometrical analyses of turbulent premixed flames are based on the flamelet assumption conceptualized by [36]. Turbulence-chemistry interactions are resolved by describing the turbulent flame as an ensemble of locally steady, one-dimensional laminar flames. Two widely used techniques comprise flamelet-based chemical closure: the Flame Surface Density (FSD) approach and the G-equation model. In the FSD approach, the flame is identified as a surface which can be convected, diffused, curved and strained by the turbulent velocity field [37]. A similar approach introduces a flame wrinkling description using the ratio of flame surface and its projection in the direction of projection [38]. In an experimental investigation by Veynante et al. [101], the closures generally used for the flame surface density equation have been debated due to their inaccurate representation of propagation and curvature terms. The accurate closure of the turbulent flux and propagation velocity of the flame surface and the effects of curvature and strain rate on the flame surface continues to be a persistent challenge [15, 32].

In contrast, the G-equation model borrows from the description of a laminar flame surface using a level set [39]. This concept eliminates the need for resolving the flame-front structure, but the lack of resolution of the flame-front structure disrupts the accuracy of chemical closure when the flame structure is adequately affected by turbulence. The thickened flame model was proposed by [41] as a method to resolve the flame structure on the mesh. However, the empirical parameters used to thicken the flame are cannot be applied to arbitrary flame specifications. As an advancement, a three-zone coherent flame model based on a flame surface density equation and a conditional averaging technique that allows precise reconstruction

of local properties has been used with success in gasoline engines [102].

As a less restrictive alternative to the geometric analysis of flame-fronts, the statistical behaviour of a turbulent reactive flow is represented using a PDF of a single variable that characterizes the one-point, one-time state of a turbulent reactive flow. In premixed flames, the reaction progress variable defined as $c = 0$ in fresh gases and $c = 1$ in burnt gas is employed. One line of approach then is to presume the PDF of the progress variable and obtain the averaged chemical source-term by integrating the source-term conditionally averaged on the progress variable with its local PDF as

$$\overline{\dot{\omega}_k(\vec{x},t)} = \int_0^1 \overline{\dot{\omega}_{k,lam}(c^*)}\, P(\vec{x},t;c^*)\, dc^* \tag{4.1}$$

where $P(\vec{x},t;c^*)$ is the PDF and $\overline{\dot{\omega}_{k,lam}}(c^*)$ is the presumed conditional moment (PCM) written as a function of c^*, the statistical random variable associated with the reaction progress variable. The Bray-Moss-Libby (BML) model combines such an approach using a bimodal distribution with a physical analysis that accounts for differences in fresh and burnt gases. Veynante et al. [101] observe that the BML model has good overall trends but lacks precise closure of the flame wrinkling length scale.

Although variations of the progress variable approach exist, most such models describe constant pressure premixed combustion; advancements to accurately represent varying pressure conditions have been relatively recent [103]. Recently, a model has been proposed for simulating turbulent premixed flames in the corrugated flamelet regime by combining the BML approach in a PDF framework with elements of the FSD equation using a reaction progress variable [104]. Also, the Eulerian particle flamelet model has been applied to the simulation of non-premixed combustion in diesel engines [105] and bluff-body flames [106]. This approach includes unsteady effects in stretched laminar flamelet models by introducing marker particles that correspond to flamelet histories associated with the path of the particle in the turbulent flow field. The range of flamelet PDFs required for simulating a particular combustion process may differ from the ensemble chosen manually, thereby demanding an *a priori* knowledge of the solution. Experimental observations of [45] highlight that the flamelet assumption is unwarranted

for high turbulence intensity turbulent premixed flames.

Various other statistical methods have been developed such as the broad class of transported PDF models developed for turbulent reactive flows by Pope [53]. In these models, the complete statistical description of the state of the flow is provided through a velocity-composition joint-PDF. A transport equation is written for the joint-PDF and solved using Monte Carlo methods. The relatively recent, Multiple-mapping Conditioning (MMC) method combines features of CMC methods, PDF methods and mapping closure models [54, 55]. These methods provide excellent chemical closure; but, the associated computational expense discourages their application to industry-scale problems.

Originally formulated by Klimenko [46] and Bilger [47], CMC has been well established as a chemical closure tool for the simulation of non-premixed combustion [15, 32]. It has been used successfully for industrial applications including spray combustion in diesel engines [48]. CMC offers a promising closure tool for simulating premixed combustion since it is devoid of the assumption of a prevailing combustion regime. Instead of presuming the reaction rate that is conditionally averaged on the reaction progress variable, as in conventional flamelet based statistical approaches, CMC calculates this term for the kth filtered species transport equation from the hypothesis

$$\overline{\dot{\omega}_k|\zeta} \approx \dot{\omega}_k(\overline{T|\zeta}, \overline{Y_k|\zeta}, \overline{\rho|\zeta}), \tag{4.2}$$

assuming small fluctuations of scalars about their conditional averages. The conditionally averaged reaction rate obtained is integrated in a similar fashion to the PCM approach by writing

$$\overline{\dot{\omega}_k(\vec{x},t)} = \int_0^1 \overline{\dot{\omega}_k|c^*}(c^*)\, P(\vec{x},t;c^*)\, dc^* \tag{4.3}$$

where $\overline{\dot{\omega}_k|c^*}(c^*)$ is estimated here through the CMC hypothesis. As opposed to the flamelet method, CMC has the capability of accurately modelling the chemical source-term for flames beyond the flamelet regime [33]. However, transport equations typically solved for conditionally averaged scalars consist of various unclosed terms which are not well understood and are modelled inaccurately [31].

CSE is a novel CMC-based chemical closure model. In CSE, an inverse problem is solved over spatially localized ensembles of finite volumes to obtain conditionally averaged scalars. Therefore, CSE is devoid of the closure issues associated with conventional CMC methods. Recent advances have allowed CSE to be applied to RANS simulations of premixed combustion in a Bunsen burner [58]. However, CSE has never been applied to practical problems which are governed by transient phenomena and characterised by complex geometries. The following sections present the theory and results for application of CSE to natural gas fuelled spark-ignition engines.

4.2 Theory

4.2.1 Conditional Source-term Estimation

The general concept of CSE follows from the CMC method wherein the conditionally averaged source-term is approximated by conditionally averaged scalars relevant to the chemistry model from the CMC hypothesis (Eq. 4.2). The averaged chemical source-term is calculated by integrating the corresponding conditionally averaged source-term with the presumed PDF at the given spatial coordinate with respect to the conditioning variable as follows

$$\overline{\dot{\omega}}_k = \int_0^1 \overline{\dot{\omega}_k|c^*} P(c^*)\, dc^*, \tag{4.4}$$

where c^* is the statistical random variable associated with the conditioning variable. The presumed PDF for any point in the domain can be estimated with reasonable accuracy using the first two moments of the conditioning variable. As described in the following section, a normalized oxygen based reaction progress variable has been used in this work with the presumed PDF shape chosen as the β function. While conventional CMC models solve transport equations for the conditionally averaged scalars that are necessary for obtaining the conditionally averaged reaction rates using the CMC hypothesis [31], CSE employs an integral equation similar to Eq. 4.3. Each Favre-averaged scalar (ϕ) necessary for calculating the conditionally averaged reaction rate can be written in terms of the corresponding

conditionally averaged scalar as

$$\widetilde{\phi}_k(\vec{x}) = \int_0^1 \overline{\phi_k|c^*}(\vec{x},c^*)\,\widetilde{P}(\vec{x},c^*)\,dc^*, \tag{4.5}$$

for the corresponding spatial coordinate ($\vec{x}$) in the physical domain. The conditionally averaged scalar can then be averaged over an ensemble of discrete points (A) in the domain. If the ensemble of discrete points is chosen in a way to preserve statistical homogeneity of conditionally averaged scalars, then the ensemble average is independent of the spatial coordinate and the integral equation is written as

$$\widetilde{\phi}_k(\vec{x}_j) = \int_0^1 \overline{\phi_k|c^*}(c^*)\,\widetilde{P}(\vec{x}_j,c^*)\,dc^*, \qquad j \in A \tag{4.6}$$

for each point j within the ensemble. Typically, conditionally averaged scalars are statistically homogeneous over spatially localized ensembles of points [56]. In this work, however, the entire computational domain is considered as a single ensemble of points where the above equation is assumed to be accurate. With the conditionally averaged scalar now invariant of the jth spatial coordinate ($\vec{x}_j$) in these ensembles, the previous equation can be treated as a Fredholm integral equation of the first kind with $\widetilde{P}(\vec{x}_j,c^*)$ as the kernel. The conditionally averaged scalar can be obtained for the ensemble from the deconvolution of this integral equation. Using m bins for the conditioning variable, the integral equation is written in the discrete form as

$$\vec{b} = \mathbf{A}\vec{\alpha}, \quad A_{ji} = \int_{c_1}^{c_2} \widetilde{P}(c^*;\vec{x}_j)dc^*, \tag{4.7}$$

where $b_j = \widetilde{\phi}_k(\vec{x}_j)$ and $\alpha_i = \overline{\phi_k|c_i^*}$ is the conditional average in the i^{th} bin. Integrations of the presumed PDF required for calculating A_{ji} are done *a priori* and tabulated for m bins using lower (c_1) and upper (c_2) limits for the i^{th} bin of conditioning variable. As in the work of Salehi et al. [58], the Tikhonov method is implemented for regularising the solution of the integral equation using an unstrained 1-dimensional laminar flame solution appropriate for the mixture properties and initial conditions. Fig. 4.1 shows different conditionally averaged temperatures obtained from laminar solutions of mixtures with varying relative air-fuel ratio (λ) used as regularisation solutions. An L-curve approach is employed for the choice

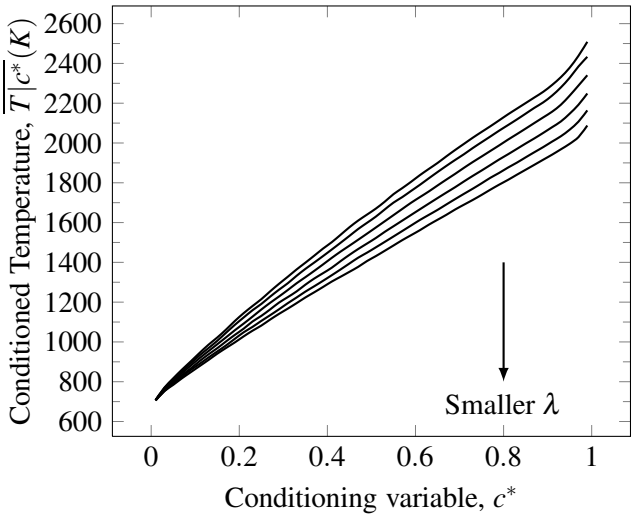

Figure 4.1: Laminar flame solutions of mixtures with varying relative air-fuel ratio ($\lambda = 1, 1.1, \ldots, 1.5$) used for the regularisation of CSE solutions.

of an optimal regularisation parameter as discussed by Salehi et al. [58]. In this work, the entire physical domain is treated as a single ensemble of points for the inversion procedure.

4.2.2 Presumed Probability Density Function

The PDF is central to the accuracy of CSE related integral inversions and the choice of conditioning variable is important for the overall description of the solution. Swaminathan and Bilger [107] have discussed various choices appropriate for turbulent premixed flames. It is critical for the progress variable to vary linearly across the turbulent flame so that the computational description is accurate. In this work, a reaction progress variable defined using O_2 mass fraction is used due to the roughly linear transformation of oxygen across an ignition process. The reaction progress variable is defined as

$$c = \frac{Y_R^0 - Y_R}{Y_R^0 - Y_R^\infty} \tag{4.8}$$

where O_2 is the chosen reactant (R), Y_R^∞ is its equilibrium mass fraction, and Y_R^0 is the initial mass fraction of oxygen in the premixed mixture. The β function

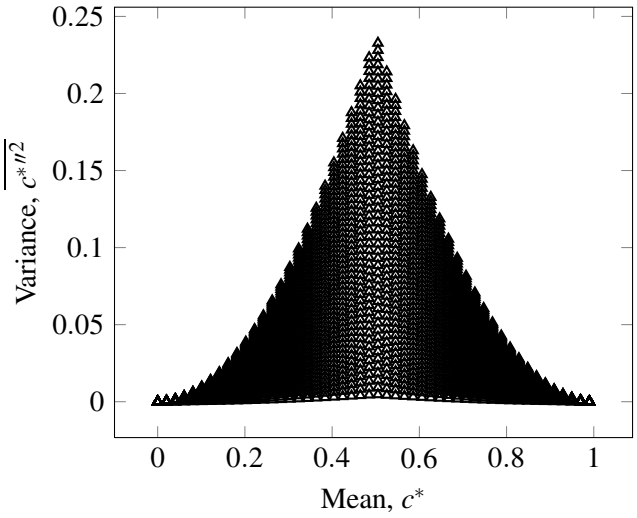

(a) Map of means and variances of stored PDFs.

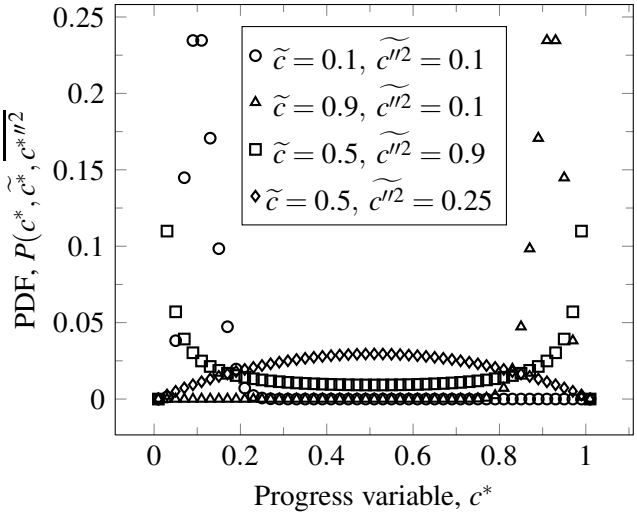

(b) Sample distributions of the β-PDF.

Figure 4.2: The β-PDF has been tabulated *a priori* for different means and variances. A sample distribution of the two presumed PDFs employed has also been illustrated.

has been used as the presumed shape of the PDF despite its failure in reproducing the true shape of the PDF for premixed flames as discussed by Jin et al. [61]. As a first implementation of the CSE framework in an industrial context, the β-PDF offers an acceptable *ad hoc* solution and recovers the extreme properties expected of the true PDF such as having δ functions appear at 0 and 1 for maximal variance or a single δ appear at the mean for zero variance. In practice, the cumulative distribution function is used instead of a PDF because of its monotonic properties which allows better numerical accuracy.

Given the reaction progress variable, its mean and variance characterise the presumed PDF in an approximate sense. Therefore, the one-time, one-point PDF is calculated as

$$P(c^*;\vec{x},t) \approx P(c^*;\widetilde{c},\widetilde{c''^2}) \tag{4.9}$$

where c^* is the statistical random variable associated with the progress variable $\widetilde{c}$ and $\widetilde{c''^2}$ are the mean and variance of the progress variable. In the *a priori* tabulation process, the binning is performed as

$$(Pdc)_1 = \int_0^{\Delta c/2} Pdc; \quad (Pdc)_m = \int_{1-\Delta c/2}^{1} Pdc \tag{4.10a}$$

for the first and last bins, and as

$$(Pdc)_i = \int_{c_i-\Delta c/2}^{c_i+\Delta c/2} Pdc \tag{4.10b}$$

for the general i^{th} bin. Careful tabulation of the first and last bins is done using half-bins in order to preserve the precise placement of the δ functions at 0 and/or 1. Fig. 4.2 shows the mean-variance map for which the β-PDF is tabulated prior to simulations; several sample distributions of the β-PDF have been shown alongside to demonstrate the range of behaviour described by β-PDF.

4.2.3 Chemistry Reduction

The conditionally averaged reaction rate is calculated using Eq. 4.2 from the knowledge of conditionally averaged scalars obtained via the inversions discussed previously. In general, it is possible to consider the entire state space of the turbulent

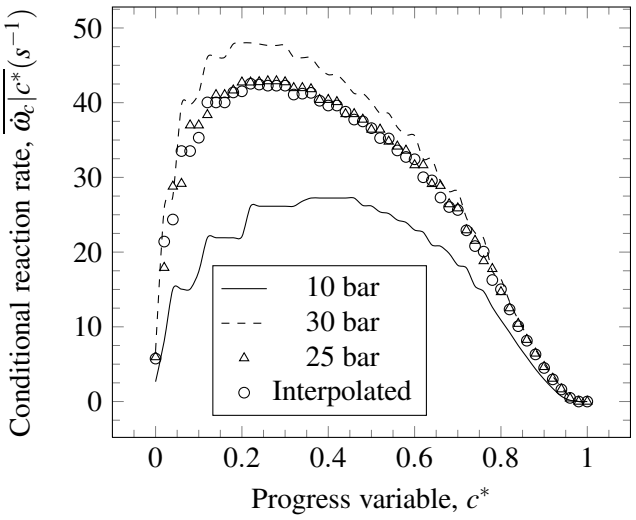

Figure 4.3: Conditional reaction rates at 25 bar obtained directly compared with an interpolation between tabulated values at 10 bar and 30 bar.

reactive flow involving detailed chemistry into the chemistry model necessary for the CMC hypothesis. In such a case, CSE can be used to calculate all required conditionally averaged scalars which can then be used to obtain the conditionally averaged reaction rate. However, in practice the chemistry model is simplified by reducing the chemistry space in order to decrease computational expense. Further, based on the gradients involved in the reaction rate dependence with a scalar, the exact number of bins may be varied. In this work, the Trajectory Generated Low-dimensional Manifold (TGLDM) proposed by Wang et al. [62] is implemented to tabulate a modified GRI-Mech chemistry developed by Huang et al. [108] using temperature, reaction progress variable and pressure. Due to the near-linear dependence of reaction rate on average pressure at any given point, only 3 manifolds of pressure (10, 30 and 60 bar) were tabulated; any conditional reaction rate is well approximated from an interpolation between the available manifolds as illustrated in Fig. 4.3. Therefore, conditionally averaged reaction rate is expressed as

$$\overline{\dot{\omega}_k|c^*}(c^*) \approx \dot{\omega}_k(P,\, \overline{T|c^*}(c^*)) \tag{4.11}$$

where P is the average pressure and $\overline{T|c^*}$ is the conditionally averaged temperature obtained via the CSE model. While 51 bins of the reaction progress variable are used (including the first and last half-bins), 100 bins of temperature have been used due to the sensitivity of reaction rate to temperature. Since the tabulation is done using the reaction progress variable, the conditioning variable is known for each bin. As a result, the only unknown parameter required for closure of chemical reaction source is the conditional average of temperature. This is in contrast to the work of Salehi et al. [58] wherein a Bunsen burner flame was simulated using CO_2 and H_2O mass fractions to tabulate the chemistry with normalized CO_2 fraction used as the conditioning variable.

4.2.4 Governing Equations

As discussed earlier, the CSE method employs a presumed shape for the PDF of the conditioning variable. Construction of the simple β-PDF relies on the knowledge of mean and variance of the conditioning variable at the given point and time. These quantities can be obtained using transport equations: since the conditioning variable is normalized oxygen mass fraction, the equation is identical to that written for any species mass fraction:

$$\frac{\partial \bar{\rho}\widetilde{c}}{\partial t}+\frac{\partial \bar{\rho}\tilde{u}_i\widetilde{c}}{\partial x_i}=-\frac{\partial \bar{\rho}\widetilde{u_i''c''}}{\partial x_i}+\frac{\partial}{\partial x_i}\left(\bar{\rho}\mathscr{D}_c\frac{\partial \widetilde{c}}{\partial x_i}\right)+\overline{\dot{\omega}}_c, \tag{4.12}$$

where $\mathscr{D}_c$ is the Fick's coefficient of molecular diffusivity for the species related to the progress variable. Two unclosed terms exist in the equation mentioned above: i) the turbulent scalar flux, $\widetilde{u_i''c''}$ and ii) the chemical source-term $\overline{\dot{\omega}}_c$ which is estimated using CSE. The widely applied method of modelling the turbulent scalar flux using gradient transport hypothesis has been employed in this work. This is applicable for flows where counter-gradient diffusion is not dominant [100]. A transport equation for the variance of progress variable is also solved:

$$\frac{\partial \bar{\rho}\widetilde{c''^2}}{\partial t}+\frac{\partial \bar{\rho}\tilde{u}_i\widetilde{c''^2}}{\partial x_i}=\frac{\partial}{\partial x_i}\left(\bar{\rho}\frac{\nu_T}{\mathrm{Sc}_1}\frac{\partial \widetilde{c''^2}}{\partial x_i}\right)+2\bar{\rho}\frac{\nu_T}{\mathrm{Sc}_2}\frac{\partial \widetilde{c}}{\partial x_i}\frac{\partial \widetilde{c}}{\partial x_i}-2\overline{\rho\mathscr{D}_c\frac{\partial c''}{\partial x_i}\frac{\partial c''}{\partial x_i}}+2\overline{c''\dot{\omega}_c}$$

where ν_T is turbulent viscosity. Similar to the work of Salehi et al. [58], all Schmidt numbers are set to 0.7 based on the work of Yimer et al. [109]. The two unclosed terms in the equation namely, the correlation between source term and fluctuations of the conditioning variable ($\overline{c''\dot{\omega}_c}$) and the Favre-averaged scalar dissipation rate $\left(\overline{\rho D \frac{\partial c''}{\partial x_i} \frac{\partial c''}{\partial x_i}}\right)$ have been modelled in an identical fashion as used for the RANS simulation of Bunsen burners by Salehi et al. [58]. Vervisch et al. [110] explain these models at great length; the discussion is beyond the scope of this work.

In comparison to a steady state Bunsen burner flame, an SI engine consists of walls on the domain boundaries. Heat transfer from these walls contributes significantly to the governing enthalpy equation and dictates the decay of flame speed near the wall. In this work, walls have been maintained at constant temperature ($T = 500K$). The overall wall heat transfer is calculated using an effective turbulent thermal diffusivity and the surface normal gradient of enthalpy. However, the isothermal assumption disrupts the validity of conditional averages calculated over a single domain-wide ensemble of points. Even though the chemical composition may change and result in temperature fluctuations, the assumption suppresses any such deviations close to the wall. A more accurate method of calculation of the scalar conditional averages would be to construct ensembles of points that are equi-distant from the wall. This method is avoided here since in an axi-symmetric case, with few finite volumes, this method inevitably results in a rank-deficient kernel for the inverse problem; a fine mesh simulation in a 3D domain is appropriate for such ensemble construction methods.

A summary of the structure of CSE-TGLDM model has been illustrated in Figure 4.4. The structure shows the flow of control in the CFD solver which originally provides the mean and variance of conditioning variable through the solution of transport equations. Once the β-PDF is constructed for all points in the mesh from the corresponding mean and variance values, the kernel containing all PDFs is passed along with necessary scalars (Favre-averaged temperature and pressure) calculated in the previous time step to the CSE routine. The desired regularisation procedure is employed to perform inversions in the CSE routine to calculate the conditional averages of scalars. Since the conditioning variable is an oxygen based reaction progress variable, also used in the TGLDM tabulation, the conditional averages of this species are already *known*. All requisite conditional averages are then

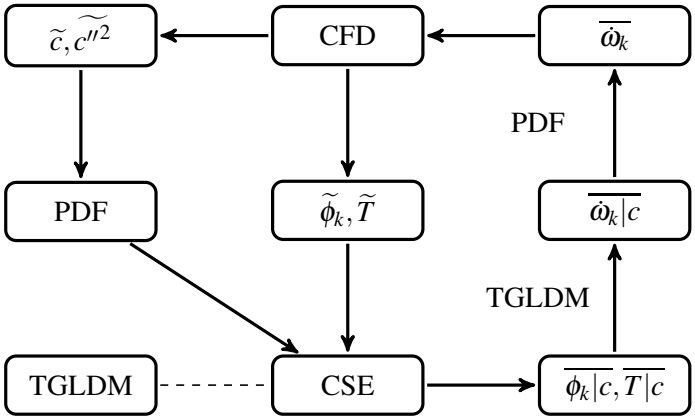

Figure 4.4: Iterative algorithm of the CSE-TGLDM chemical closure model.

used to retrieve conditionally averaged reaction rates from the TGLDM chemistry tabulation; this vector of values is integrated with the PDF at each point to obtain the average reaction rate in the corresponding cell.

4.2.5 Spark-Ignition Model

For a mixture with given relative air-fuel ratio (λ), the experimentally observed initial conditions are matched by using the reported pressure at intake valve closing IVC. Another variable important for matching combustion relevant conditions is the timing of spark-ignition. A realistic high-temperature charge plasma is challenging to implement in the context of numerical simulations because placing a tiny zone of extremely high temperature is constrained by mesh resolution and tabulation of chemical kinetic constants and species chemistry at the corresponding scale of temperature. Thiele et al. [111] have noted that due to the fast expansion of the plasma channel, a complicated flow-field is developed after the emission of a shock wave. Development of a propagating flame after the initiation of the flame kernel is dominated by reactive and diffusive processes [111, 112].

Pischinger and Heywood [113] derive a first-law-based model for the growth of spark-generated flame kernel. They focus on the effect of heat losses on the rate of flame kernel development and observe significant cycle-by-cycle variations. Tan and Reitz [114] explain that the commonly used method of increasing internal en-

ergy in specified ignition cells during ignition by a factor each time step is very sensitive to the computational mesh size. Instead, they derive an equation to calculate spark-ignited kernel growth rate by considering the effects of the spark ignition discharge energy and flow turbulence on the ignition kernel growth. In a more recent study, Enaux et al. [115] advance an Arc and Kernel Tracking Ignition Model (AKTIM) using a three phase progress variable tracking procedure for simulating the kernel development. Another detailed physical model developed by Yaşar [116] accurately describes the momentum and energy exchange between the gas and plasma.

A simpler approach referred by Mastorakos [117] and explained in detail in many texts including that by Law [118] is adopted for this first implementation of CSE in an industrial context. Typically, ignition is forced in the uniform mixture by a sudden insertion of a layer of burnt gas. In this work, instead of initiating a spark at the ignition timing, we place a developed progress variable profile ($\widetilde{\Sigma}_c$) in a small volume (typically, $\sim 2.5\%$ of the chamber volume) around the region of spark plug. This developing kernel is made highly reactive by setting the mean ($\widetilde{c} = 0.5$), variance ($\widetilde{c''^2} = 0.25$) and temperature ($T \approx 0.5 T_{ad}$) to appropriate conditions. The radius of the flame kernel is then calculated from

$$\frac{2\pi R_{kernel}^3}{3} = \Xi \times V_{\Xi} \tag{4.13}$$

where Ξ is the fraction of chamber volume occupied. Special care must be taken in choosing Ξ since a small kernel may not be able to withstand the turbulent strain induced in the engine at the corresponding timing and may, therefore, lead to extinction of the kernel. The timing of kernel placement is decided by the displacement required to achieve 5% fuel mass fraction burned (MFB) at the same timing as reported in the experimental study of Reynolds [98].

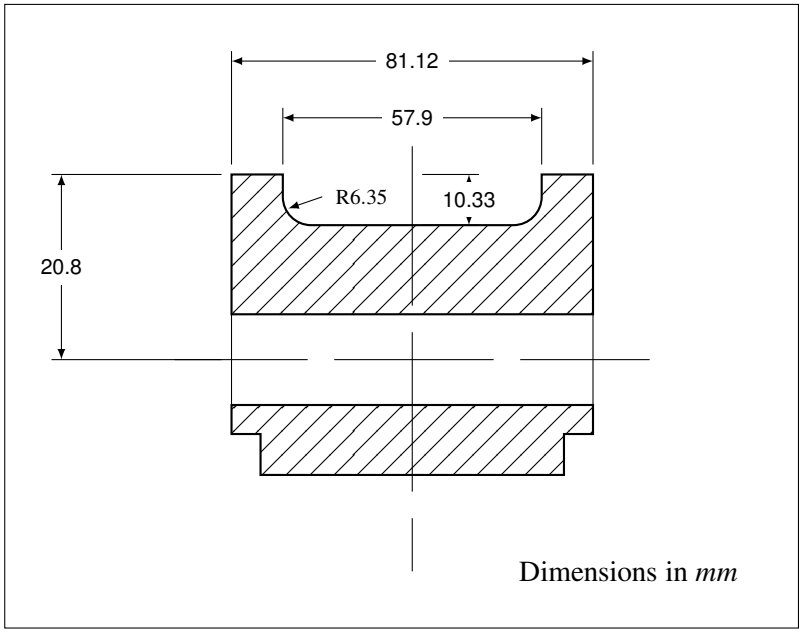

Figure 4.5: Bowl-in piston geometry of the Ricardo Hydra engine.

4.3 Problem Specifications

4.3.1 Experimental Validation Case

The engine used for experimental validation is a Ricardo Hydra single-cylinder research engine rebuilt for natural gas fuelling. The piston geometry is a modified form of the standard Ford Festiva bowl-in piston with dimensions as specified in Fig. 4.5. Other important specifications of the engine geometry have been listed in Table 4.1. Measurements have been provided with varying relative air-fuel ratios ($\lambda = 1, 1.1, \ldots, 1.5$) of the premixed homogeneous natural gas-air mixture. The tests were reported for different engine speeds, but we have made comparisons at 2500 rpm, in particular for a lean mixture with $\lambda = 1.5$. The natural gas composition reported in the experimental study was scaled to provide a fuel composition that matched the chemistry tabulation employed.

In the experimental work, heat release measurements calculated from the in-cylinder pressure trace measured using a pizeo-electric pressure transducer were

Serial	Geometrical Feature	Value	Unit
1	Number of cylinders	1	–
2	Bore	81.12	mm
3	Stroke	88.90	mm
4	Connecting rod length	158.01	mm
5	Swept volume	459.46	cm^3
6	Clearance volume	41.988	cm^3
7	Compression ratio	11.94:1	–
12	Inlet closes	56	°ABDC
13	Exhaust opens	56	°BBDC

Table 4.1: Specifications of the Ricardo Hydra research engine.

used to estimate combustion duration and the fuel Mass Fraction Burned (MFB) as a function of crank angle. No turbo-charger is employed in the experiments; however, reported pressures have large values in the low-pressure zones away from the Top Dead Center (TDC). These have been attributed to the inaccuracies involved in measuring low pressures using a piezoelectric sensor. A corrected intake manifold pressure ($P_{IVC} = 2.32\ bar$) has been used in our calculations in order to match the early pre-combustion pressure trace.

The primary emissions examined in the experimental study were the brake specific exhaust emissions of NO_x, tHC and CO; in our work, NO_x emissions have been reported. Experimentally observed emissions were reported as the mass flow rate of exhaust emissions normalized with the brake power produced by the engine:

$$NO_x(g/kW \cdot hr) = \frac{NO_x(g/h)}{P_b(kW)} \tag{4.14}$$

where P_b is the brake specific power generated. In our study, the emission mass output has been normalized with the Gross Indicated Work (GIW) generated.

4.3.2 Numerical Solver

Simulations have been performed on a 2-dimensional axi-symmetric mesh shown in Fig. 4.6 generated using the *blockMesh* utility of the OpenFOAM CFD software [99]. The cell sizes have a gradient to provide a fine mesh structure near the spark

ignition region and the entire mesh was refined at various pre-specified intervals (crank angle, $CA = 330, 390, 420$) to avoid cell skewness as much as possible. A maximum Courant number, $C_{max} = 0.25$ was used with adjustable time step starting with an initial value of $0.05\ s$. Solution was written at unity intervals of crank angle.

Limited Gauss linear schemes were used for the solution of Laplacian and divergence operators in transport equations of species, enthalpy and momentum. The Gauss upwind scheme was used for the solution of transport equations relevant to the $k - \varepsilon$ turbulence model used from the OpenFOAM library.

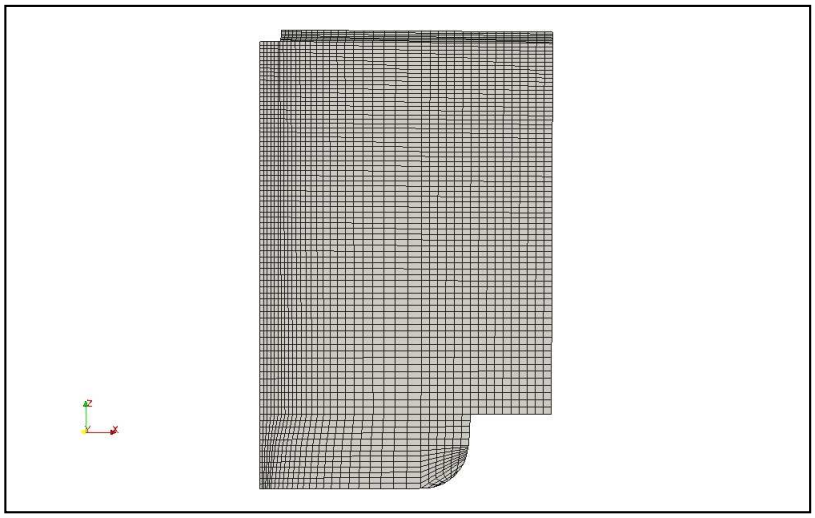

Figure 4.6: Axi-symmetric mesh constructed at 270° crank angle.

4.4 Results

The following section describes our results; flame kernel development has been highlighted and the pressure map obtained has been used to calculate heat release rates and fuel mass fraction burned. Further, trends of major pollutants, NO_x and CO have been discussed for mixtures of varying λ.

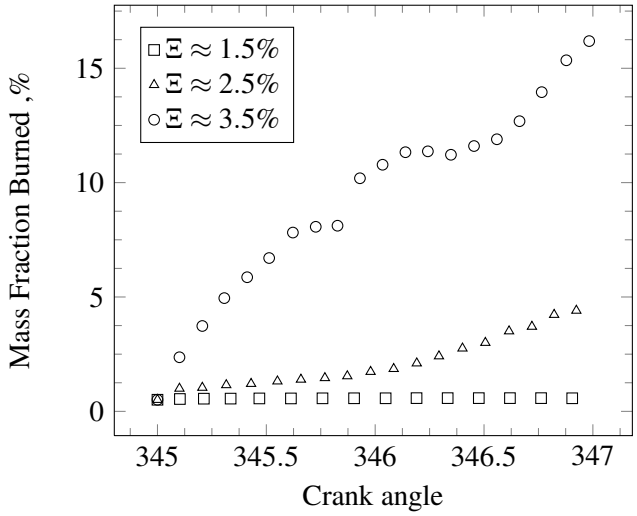

(a) Kernel growth from initiation to $t_{5\%\ mfb}$.

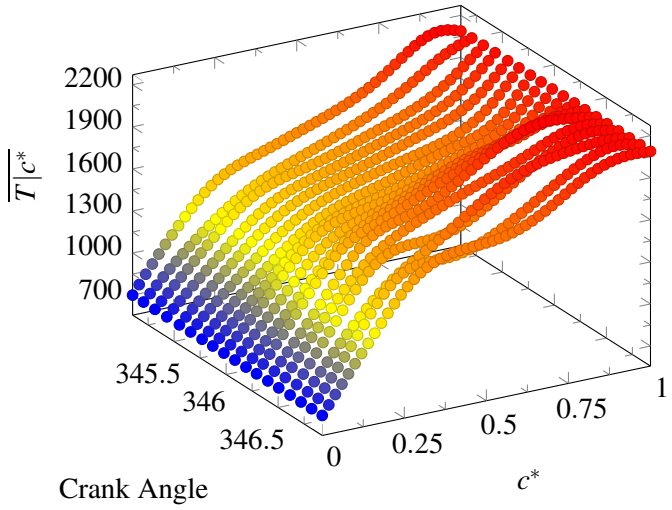

(b) Conditional averages post kernel initiation.

Figure 4.7: Development of flame kernel shown as the degree of MFB until 5% MFB and corresponding conditionally averaged temperatures obtained for a lean mixture with $\lambda = 1.5$ at 2500 rpm.

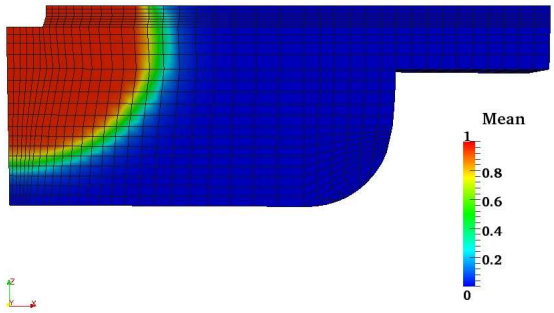

(a) Mean of progress variable, $\widetilde{c}$

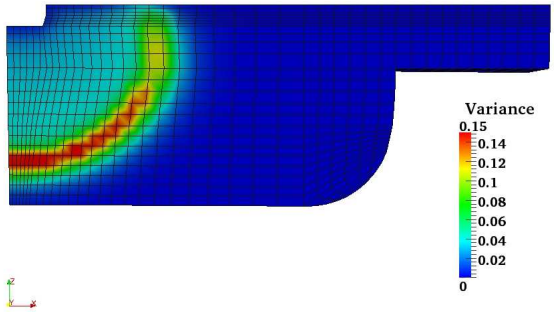

(b) Variance of progress variable, $\widetilde{c''^2}$

Figure 4.8: Mean and variance contours of progress variable at 5% MFB for a mixture with $\lambda = 1.5$ at 2500 rpm.

4.4.1 Flame Kernel Development

Upon the imposition of the impulsive high reactivity field on the turbulent reactive flow, quantities averaged over the geometry, such as the MFB, exhibit sensitivity towards the size of the initial kernel. Fig. 4.7(a) shows the increase in MFB from initiation of the flame kernel until 5% fuel MFB ($t_{5\%MFB}$) is attained in the combustion chamber for different initial volumes (as a fraction of chamber volume) of the kernel. This characterises the growth of the flame kernel and has different pro-

files for varying mixture fractions. The profile of MFB growth has been used as a marker for judging whether the kernel will propagate as a flame. A tiny kernel does not provide enough energy sufficient for successful propagation [117] and is often extinguished by turbulent strain induced by the flow as outlined by Eichenberger and Roberts [119]. In contrast, a large kernel may be characterised by rapid kernel growth and may require that the kernel be placed too close to $t_{5\%MFB}$ thereby allowing little time for the evolution of kernel properties and conditional averages. For the case studied in Fig. 4.7(a), the experimentally observed displacement between the spark and 5% MFB was approximately 31°. The effects of the spark on the turbulent flow field from the initial discharge to the timing of kernel placement cannot be adequately modelled within the 2° displacement permitted by our empirical progress variable approach. Therefore, several measurable parameters in the study such as the timings of peak pressure and the associated crank angle cannot be accurately estimated.

Fig. 4.7(b) shows the variation of conditionally averaged temperature obtained in the flame kernel development phase between initiation and $t_{5\%MFB}$. Significant variations are noted from the steady laminar flamelet behaviour highlighting the effects of turbulence on the flame. It is evident that the CSE model captures such deviations from flamelet behaviour of the flame which might be characterised by the distributed reaction zone regime. In Fig. 4.8, contours of reaction progress variable variance highlight the regions that roughly characterise the flame surface. Flame progress is observed to be slower close the engine walls as expected due to heat transfer away from the flame and possibly due to flame quenching.

Kernel placement timings have been compared for a range of natural gas-air mixtures (varying λ). These have been graphed along with the respective timings of 5% MFB and of spark ignition. Fig. 4.9(b) compares the combustion duration (defined as the displacement between 5% and 95% fuel (MFB) trends observed experimentally and using simulations. In addition to decreasing fuel content, the size of reactive flame kernel zones is also expected to affect the duration of combustion. Although a slight increase in the combustion duration is noted for leaner mixtures, the expected values are generally overestimated. To some extent, the errors may be attributed to the initial uniform progress variable profile of the flame kernel and its timing which collectively result in an incorrect estimate of the initial kernel growth

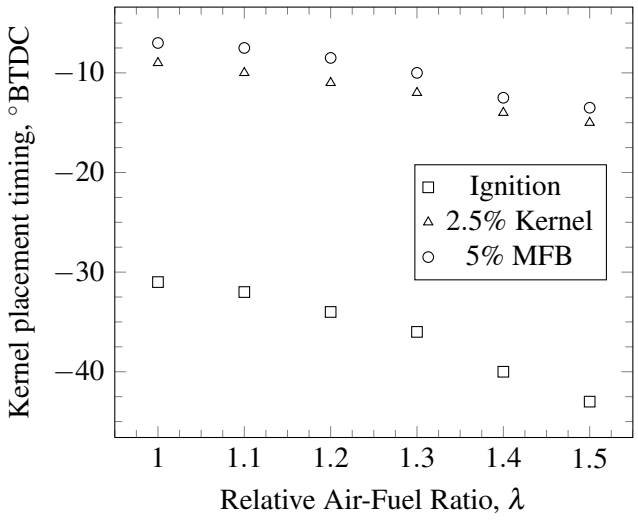

(a) 2.5% volume kernel initiation timing.

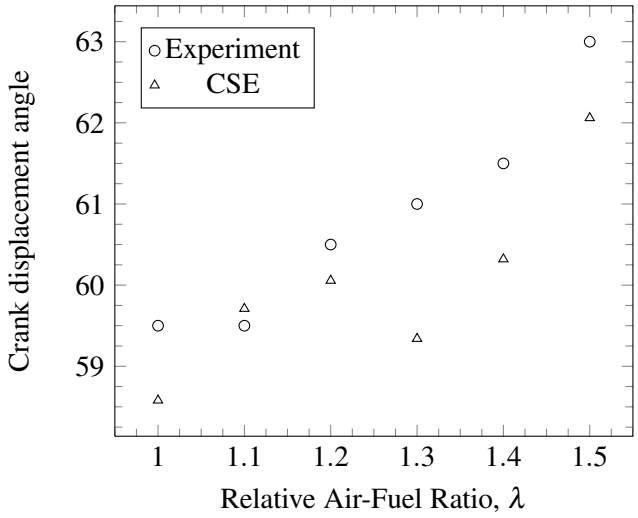

(b) Combustion duration ($5\% - 95\%$ MFB).

Figure 4.9: Initiation timing of high-reactivity flame kernel with $\sim 2.5\%$ chamber volume for varying λ and their combustion durations.

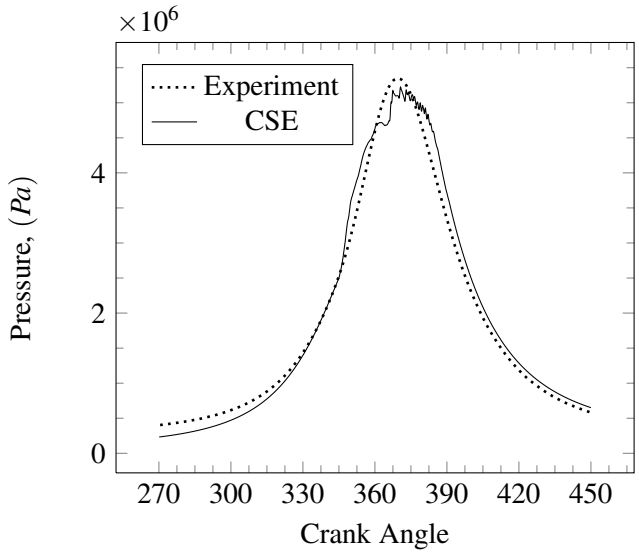

Figure 4.10: Pressure map for a lean mixture ($\lambda = 1.5$) at 2500 rpm.

rate. Moreover, the profile of progress variable and temperature in the initiated kernel does not model the flame-generated turbulence within the kernel leading to an inaccurate growth rate estimate.

4.4.2 Pressure Trace

The pressure trace has been obtained from intake valve opening (IVC) at 270° to exhaust valve opening (EVO) at 484° crank angle (CA). Fig. 4.10 shows the comparison of pressure profile with values obtained experimentally. The piezoelectric pressure transducer used in the experiment by Reynolds [98] is susceptible to error at low pressures. Therefore, the intake manifold pressure observed by [98] has been used as the initial pressure at 270° CA instead in order to achieve accurate matching of the pressure profile prior to combustion. Despite kernel placement at the appropriate timing as discussed previously, calculated pressures are slightly higher until close to TDC. Severe fluctuations in the pressure map are observed post TDC; these fluctuations are attributed to the skewness of moving mesh cells at TDC.

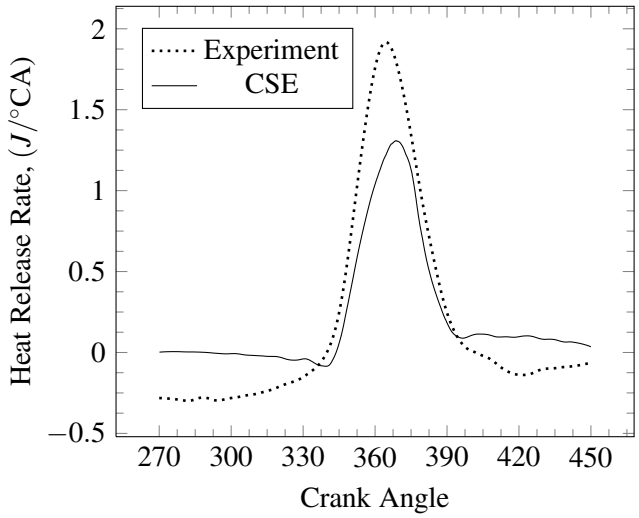

Figure 4.11: Heat release rate for a lean mixture ($\lambda = 1.5$) at 2500 rpm.

Heat release rates (HRR) have been calculated using the pressure profile by employing the following relation obtained using the first law of thermodynamics and assuming uniform and identical properties of reactants and products:

$$\frac{dQ}{d\theta} = \frac{1}{\gamma - 1}\left(V\frac{dP}{d\theta} + \gamma P\frac{dV}{d\theta}\right) \tag{4.15}$$

where γ is the ratio of specific heats for natural gas. The HRR profile obtained from both the experimental and simulation pressure trace were filtered using a low-pass filter in MATLAB followed by a smoothing function. Significant negative heat release is observed for the experimental profile in Fig. 4.11; this may be attributed to incorrect pressure values read by the transducer at earlier crank angles. Due to the fluctuations of pressure near the TDC, calculation of heat release is prone to extreme error in this region.

The technique used here for computing the mass fraction burned (MFB) is based on the method developed by Rassweiler and Withrow [120]: the pressure rise due to combustion is estimated from the total pressure rise for small crank angle intervals as

$$\Delta P = \Delta P_v + \Delta P_c \tag{4.16}$$

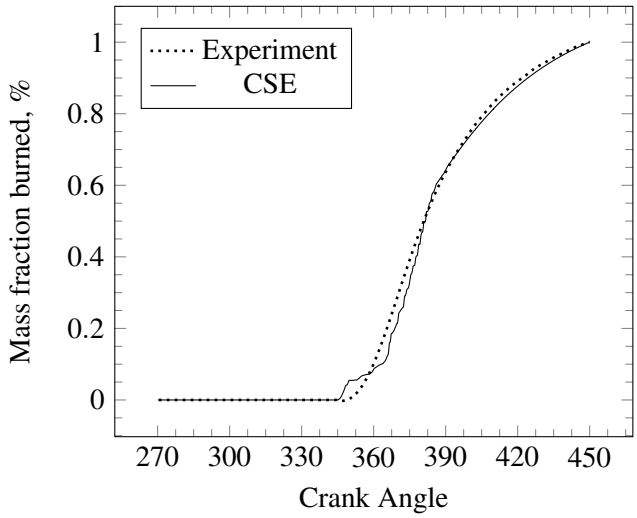

Figure 4.12: Variation of MFB for a lean mixture ($\lambda = 1.5$) at 2500 rpm.

where ΔP_v and ΔP_c are the change in pressure due to volume change and combustion respectively. MFB can then be calculated as a fraction of the pressure change due to combustion up to a given crank angle with respect to the entire range from the Intake Valve Closes (IVC) to Exhaust Valve Opens (EVO):

$$MFB = \frac{\sum_{\theta_i}^{\theta} \Delta P_c}{\sum_{\theta_i}^{\theta_{EOC}} \Delta P_c} \tag{4.17}$$

where θ_i is the initial crank angle (270°) chosen to be close to IVC. Assuming polytropic compression in each small interval of crank angle change, the pressure rise attributed to combustion can be calculated for the i^{th} interval as

$$\Delta P_{c,i} = P_i - P_{i-1}\left(\frac{V_{i-1}}{V_i}\right)^n \tag{4.18}$$

where n is the polytropic index. The mass fraction burned for a lean mixture with $\lambda = 1.5$ is shown in Fig. 4.12. It is to be expected that the profile is non-smooth because the pressure oscillates severely in the region of TDC.

In addition to the pressure profile, the peak pressure and corresponding crank

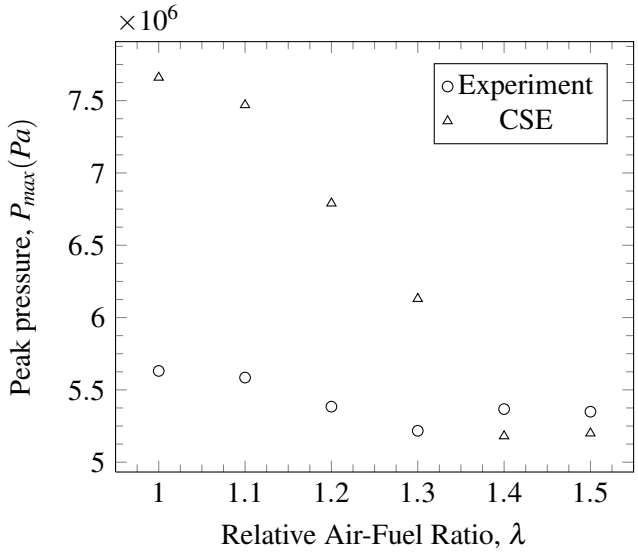

(a) Maximum average cylinder pressure.

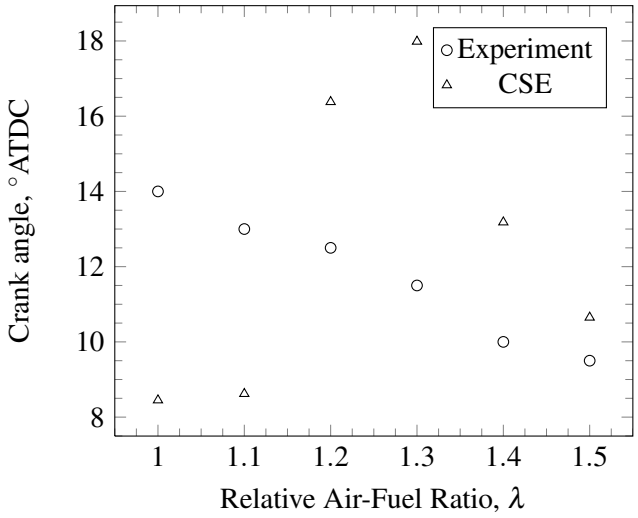

(b) Crank angle of maximum average pressure.

Figure 4.13: Peak pressure and associated crank angle for varying λ.

angles are calculated and have been compared with the experimental estimates in Fig. 4.13. Although the pressure maxima show a decreasing trend similar to experiments, the exact values are overestimated by the CFD calculation. On the other hand, the crank angle of peak pressure does not show a consistent trend as in experiments perhaps due to small mismatches in the 5% MFB timing.

A sensitivity test was performed to gauge the effect of intake manifold pressure on the simulated peak pressures and corresponding crank angles. Fig. 4.14 shows the different maximum pressure values obtained for slightly different intake pressures. Matching the exact values from experimentally obtained pressure profiles means using erroneous results as opposed to noting the corrected intake manifold pressure that is measured with higher precision in experiments. Similarly, crank angle of maximum pressure shows reasonably large variations with the choice of intake manifold pressure. The mismatch between experimentally observed pressures at IVC and the pressure used for simulations ($P = 2.32\ bar$) causes significant deviations in estimated combustion duration and MFB among other parameters.

4.4.3 Emissions

NO_x pollutant emissions have been calculated using the chemistry mechanism proposed by Hanson and Salimian [121]. A transport equation for NO is solved and the chemical source-term is closed using the CSE method, an approach similar to that used by Wang et al. [62]. Fig. 4.15 shows that CSE predicts the general decreasing trend of NO_x calculated at EVO with decreasing fuel content in the homogeneous charge. For an appropriate comparison with the experimental measurement of NO_2, a scaling factor ($44/30$, based on the molecular masses of the two compounds) is applied to the NO estimates calculated by the numerical solver. These calculations have been normalized with the gross indicated work (GIW) estimated from the pressure obtained via simulations, whereas the experimental data consists of calculations made in the exhaust flow normalized with the brake power. The values of normalized emissions are only expected to increase (thereby, matching with experimental results better) if we were to account for friction in the GIW. Experimental error in NO_x mass measurement and exhaust flow rate are considered – error in brake power (P_b) measurement is not considered in the graphed error

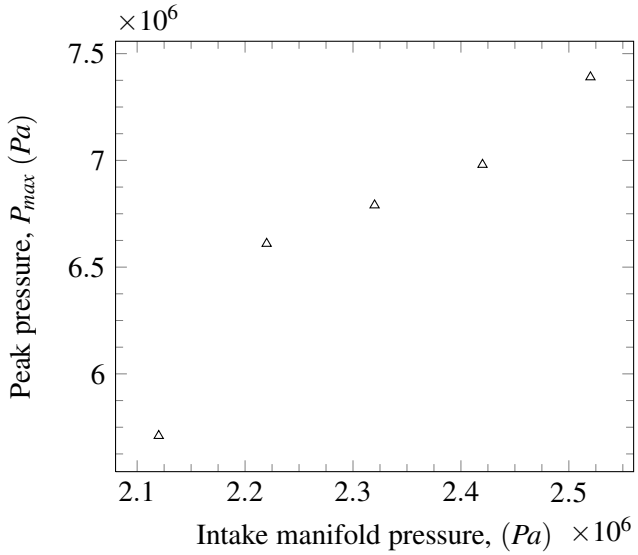

(a) Maximum average cylinder pressure.

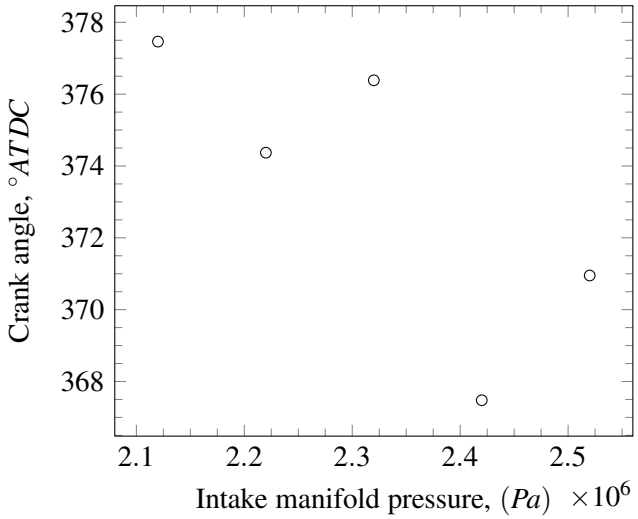

(b) Crank angle of maximum pressure.

Figure 4.14: Sensitivity of maximum cylinder pressure and the corresponding crank angle to the pressure at IVC in the simulations.

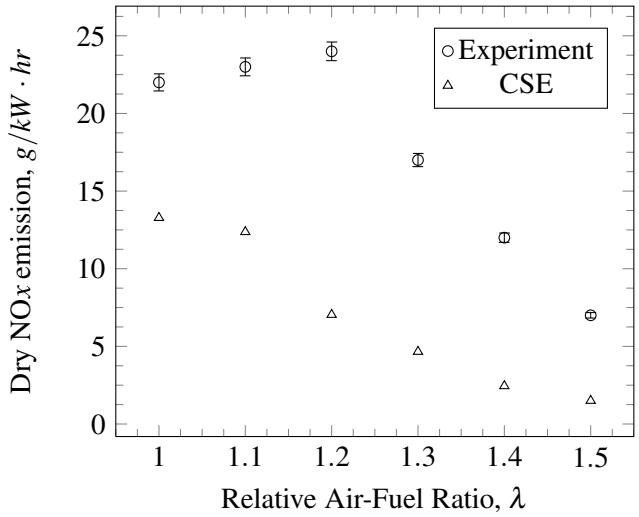

Figure 4.15: Pollutant emission trends with varying λ. Numerical estimates have been scaled appropriately.

bars. Given that a simplified chemistry mechanism is used in a RANS cotnext, the CSE trends show promising agreement with experiments despite considerable underestimation.

4.4.4 Discussion

Accurate calculation of conditionally averaged temperature plays a significant role in chemical closure using CSE. In a moving-mesh scenario in complex geometries, the statistical homogeneity of conditional averages might be sacrificed near TDC due to the skewness of finite volume cells. In addition, the application of β-PDF to describe the stochastic behaviour of reaction progress variable in premixed flames is known to be erroneous and is expected to affect the CSE solution.

4.5 Conclusions

CSE has been established as a combustion model for simulating industry relevant turbulent reactive flow problems. In view of its first application to simulating combustion in a spark-ignition engine, the obtained results are promising and accept-

able in the context of RANS. The Pressure profile, through which heat release and combustion duration is calculated, is an important quantity of interest along with emissions. The pressure profile is matched to a fair extent, particularly for leaner mixtures where more accurate matching is possible for the initial flame kernel evolution. CSE has predicted identical trends for NO_x emissions despite the use of a simplified chemistry. In conclusion, this work validates the consistency of CSE for transient turbulent combustion phenomena in an industrial setting. Further work on precise ignition modelling and the implementation of a robust regularisation parameter choice method will be needed to improve the predictive accuracy.

5 | Conclusions

5.1 Inverse Problem Implementation

A novel methodology has been developed to partition finite volume meshes of arbitrarily complex geometries into spatially localized clusters of points.

- In the context of parallel processing architectures, spatial locality of partitions is chosen over computational load balance in order to maintain the validity of CSE inversion within the partitions treated as ensembles. Although the distribution of mesh elements is not equal, appreciable gains are made in inter-processor communication by preserving locality.
- The spatial partitioning method relies on the Morton order space-filling curve generated based on a floating-point algorithm which is applicable to unstructured finite volume meshes. In conjunction with the mapping algorithm, the developed algorithm automates the process of localized ensemble construction by finding locality destructive leaps.
- While cluster locality is improved against equi-sized partitioning of the Morton order curve, computational load balance is sacrificed to some degree despite algorithms in place to maintain equal distribution of finite volumes across partitions. The degree of computational load balance achieved is acceptable for large-scale parallel processing environments and is expected to improve with mesh refinement.

Therefore, a robust technique has been developed that can be applied for inexpensive and unsupervised partitioning of industrial geometries.

5.2 Industrial Modelling Application

CSE has been established as a combustion model for simulating industry relevant turbulent premixed combustion problems.

- As a first application of CSE to combustion simulations in spark-ignition engines, the pressure trace has been reasonably well predicted. Other quantities of interest such as heat release rate are derived from the pressure trace and their inaccuracies are attributed to the deviations from experimentally measured pressure traces. In particular, the severe oscillations of the pressure profile close to TDC are expected to result in extreme errors in the calculation of heat release rates.

- The NO_x pollutant emissions calculated have shown a good agreement with the experimentally observed trends for varying mixture fractions. Considering the simplified chemical mechanism employed, emission predictions are of an acceptable standard.

- Accurate calculation of conditionally averaged temperature plays a significant role in chemical closure using CSE. The inversions are affected by the wall; in the case of using a single domain-wide ensemble, the underlying assumptions are, in fact, invalidated. Moreover, the application of β PDF to describe the stochastic behaviour of reaction progress variable in turbulent premixed flames is known to be erroneous and is expected to affect the CSE solution.

- Ignition modelling is critical to the simulation of combustion in a homogeneous-charge SI engine. While matching the initial kernel growth rate is challenging with the simplified empirical model in place, it nevertheless gives reasonable estimates of key variables.

This investigation validates the applicability of CSE for transient turbulent combustion phenomena in an industrial setting.

5.3 Future Directions

- Simulation of turbulent combustion in complex industrial geometries with swirl motion may be investigated as the next step. With the development of the spatial partitioning algorithm, an appropriate choice of ensembles can be made to retain the validity of CSE in industrial burner domains. In particular, specific attention could be given to wall effects such as heat transfer that otherwise sacrifice the accuracy of CSE with a single domain-wide ensemble.

- The stochastic behaviour of the conditioning variable significantly affects the accuracy of CSE and, hence, the precision with which transient phenomena such as spark-ignition can be modelled. For premixed flames, better descriptions of the presumed PDF shape are available. While the β-PDF captures the placement of delta functions, the probability distribution is matched better by models that account for the flame reactivity and turbulence. The implementation of such improved PDFs could lead to sizeable improvements in the predictive accuracy in SI engines.

- The regularisation of CSE solutions done using the Tikhonov method is prone to error due to the L-curve method of optimum regularisation parameter choice. A variety of parameter choice methods are available and can be explored to lend robustness to the computational implementation of CSE.

- In the context of SI engines, a more detailed physical representation of the plasma charge could replace the employed simplistic model. If the evolution of the flame kernel is matched better, critical parameters such as maximum pressure and the corresponding crank angle timing can be predicted with greater accuracy. Moreover, to improve predictions of emissions, a transport equation is necessary for CO, whereas an improved chemical mechanism must be implemented for NO_x chemistry.

- The CSE approach could be implemented for non-premixed combustion simulations in diesel engines. Modelling of sprays and their ignition would be the next challenging leap of this promising chemical closure technique.

References

[1] N. Stern. Review on the economics of climate change. *London HM Treasury*, 2006.

[2] J. Tollefson and R. Monastersky. The global energy challenge: awash with carbon. *Nature*, 491:654–655, 2012.

[3] P. Moin and J. Kim. Tackling turbulence with supercomputers. *Scientific American*, 1997.

[4] S. B. Pope. Small scales, many species and the manifold challenges of turbulent combustion. *Proceedings of the Combustion Institute*, 2012.

[5] C. L. Fefferman. Existence and smoothness of the navier-stokes equation. *The millennium prize problems*, pages 57–67, 2000.

[6] W. Kendal Bushe and Helfried Steiner. Conditional moment closure for large eddy simulation of nonpremixed turbulent reacting flows. *Physics of Fluids*, 11:1896, 1999.

[7] D. Veynante. Turbulent combustion modeling. *Progress in Energy and Combustion Science*, 28, 2002.

[8] J. A. Miller and C. T. Bowman. Mechanism and modeling of nitrogen chemistry in combustion. *Progress in Energy and Combustion Science*, 15 (4):287–338, 1989.

[9] D. B. Olson and W. C. Gardiner Jr. An evaluation of methane combustion mechanisms. *The Journal of Physical Chemistry*, 81(25):2514–2519, 1977.

[10] U. Warnatz, J.and Maas and Robert W. Dibble. *Combustion: physical and chemical fundamentals, modeling and simulation, experiments, pollutant formation*. Springer, 2006.

[11] C. K. Westbrook and F. L. Dryer. Chemical kinetic modeling of hydrocarbon combustion. *Progress in Energy and Combustion Science*, 10 (1):1–57, 1984.

[12] W. C. Gardiner. *Gas-phase Combustion Chemistry*. Springer, 2000.

[13] F. M. White. Fluid mechanics. 5th. *Boston: McGraw-Hill Book Company*, 2003.

[14] S. B. Pope. *Turbulent flows*. Cambridge University Press, 2000.

[15] T. Poinsot and D. Veynante. *Theoretical and Numerical Combustion*. RT Edwards Inc., 2011.

[16] G. K. Batchelor. *The theory of homogeneous turbulence*. Cambridge university press, 1982.

[17] 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.

[18] A. A. Townsend. On the fine-scale structure of turbulence. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 208(1095):534–542, 1951.

[19] P. Moin and K. Mahesh. Direct numerical simulation: A tool in turbulence research. *Annual Review of Fluid Mechanics*, 30(1):539–578, 1998.

[20] F. R. Menter, M. Kuntz, and R. Langtry. Ten years of industrial experience with the sst turbulence model. *Turbulence, heat and mass transfer*, 4: 625–632, 2003.

[21] P. R. Spalart. Strategies for turbulence modelling and simulations. *International Journal of Heat and Fluid Flow*, 21(3):252 – 263, 2000.

[22] R. Bouffanais. Advances and challenges of applied large-eddy simulation. *Computers & Fluids*, 39(5):735–738, 2010.

[23] H. Pitsch. Large-Eddy Simulation of Turbulent Combustion. *Annual Review of Fluid Mechanics*, 38(1):453–482, January 2006.

[24] T. Poinsot and D. Veynante. *Theoretical and numerical combustion*. RT Edwards Incorporated, 2005.

[25] S. B. Pope. Ten questions concerning the large-eddy simulation of turbulent flows. *New Journal of Physics*, 6(1):35, 2004.

[26] B. E. Launder and D. B. Spalding. The numerical computation of turbulent flows. *Computer Methods in Applied Mechanics and Engineering*, 3(2):269 – 289, 1974.

[27] F. J. Weinberg. The first half-million years of combustion research and today's burning problems. *Symposium (International) on Combustion*, 15 (1):1 – 17, 1975.

[28] F. A. Williams. The role of theory in combustion science. *Symposium (International) on Combustion*, 24(1):1 – 17, 1992.

[29] K. N. C. Bray. The challenge of turbulent combustion. *Symposium (International) on Combustion*, 26(1):1 – 26, 1996.

[30] N. Peters. Multiscale combustion and turbulence. *Proceedings of the Combustion Institute*, 32(1):1 – 25, 2009.

[31] 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 (1):21–42, January 2005.

[32] T. Echekki and E. Mastorakos. *Turbulent combustion modeling: Advances, new trends and perspectives*, volume 95. Springer Science, 2011.

[33] R. Borghi. On the structure and morphology of turbulent premixed flames. In C. Casci and C. Bruno, editors, *Recent Advances in the Aerospace Sciences*, pages 117–138. Springer US, 1985.

[34] T. Poinsot and D. Veynante. *Theoretical and Numerical Combustion*. RT Edwards Inc., 2001.

[35] D. B. Spalding. Mixing and chemical reaction in steady confined turbulent flames. *Symposium (International) on Combustion*, 13(1):649 – 657, 1971.

[36] N. Peters. Laminar diffusion flamelet models in non-premixed turbulent combustion. *Progress in Energy and Combustion Science*, 10(3):319–339, 1984.

[37] E. R. Hawkes and R. S. Cant. A flame surface density approach to large-eddy simulation of premixed turbulent combustion. *Proceedings of the Combustion Institute*, 28(1):51–58, 2000.

[38] H. G. Weller, G. Tabor, A. D. Gosman, and C. Fureby. Application of a flame-wrinkling les combustion model to a turbulent mixing layer. In *Symposium (International) on Combustion*, volume 27, pages 899–907. Elsevier, 1998.

[39] 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.

[40] H. Pitsch. A consistent level set formulation for large-eddy simulation of premixed turbulent combustion. *Combustion and Flame*, 143(4):587–598, 2005.

[41] O. Colin, F. Ducros, D. Veynante, and T. Poinsot. A thickened flame model for large eddy simulations of turbulent premixed combustion. *Physics of Fluids*, 12:1843, 2000.

[42] J. A. van Oijen and L. P. H. de Goey. Modelling of premixed laminar flames using flamelet-generated manifolds. *Combustion science and technology*, 161(1):113–137, 2000.

[43] M. Ihme and H. Pitsch. Prediction of extinction and reignition in nonpremixed turbulent flames using a flamelet/progress variable model: Part 2. application in LES of sandia flames d and e. *Combustion and Flame*, 155(1–2):90 – 107, 2008.

[44] M. Ihme and H. Pitsch. Prediction of extinction and reignition in non-premixed turbulent flames using a flamelet/progress variable model. part 1: A priori study and presumed PDF closure. *Combustion and Flame*, 155(1–2):70 – 89, 2008.

[45] F. T. C. Yuen and Ö. L. Gülder. Turbulent premixed flame front dynamics and implications for limits of flamelet hypothesis. *Proceedings of the Combustion Institute*, 34(1):1393 – 1400, 2013.

[46] A. Y. Klimenko. Multicomponent diffusion of various admixtures in turbulent flow. *Fluid Dynamics*, 25(3):327–334, 1990.

[47] R. W. Bilger. Conditional moment closure for turbulent reacting flow. *Physics of Fluids A: Fluid Dynamics*, 1993.

[48] G. De 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.

[49] 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.

[50] B. Thornber, R. W. Bilger, A. R. Masri, and E. R. Hawkes. An algorithm for LES of premixed compressible flows using the conditional moment closure model. *Journal of Computational Physics*, 230(20):7687 – 7705, 2011.

[51] A. R. Kerstein. One-dimensional turbulence: model formulation and application to homogeneous turbulence, shear flows, and buoyant stratified flows. *Journal of Fluid Mechanics*, 392(1):277–334, 1999.

[52] A. R. Kerstein. A linear- eddy model of turbulent scalar transport and mixing. *Combustion Science and Technology*, 60(4–6):391–421, 1988.

[53] S. B. Pope. PDF methods for turbulent reactive flows. *Progress in Energy and Combustion Science*, 11(2):119 – 192, 1985.

[54] A. Y. Klimenko and S. B. Pope. The modeling of turbulent reactive flows based on multiple mapping conditioning. *Physics of Fluids*, 15:1907, 2003.

[55] M. J. Cleary and A. Y. Klimenko. A generalised multiple mapping conditioning approach for turbulent combustion. *Flow, turbulence and combustion*, 82(4):477–491, 2009.

[56] 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, 2001.

[57] P. C. Hansen. Numerical tools for analysis and solution of Fredholm integral equations of the first kind. *Inverse problems*, 849, 1992.

[58] 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, April 2012.

[59] M. M. Salehi. *Numerical Simulation of Turbulent Premixed Flames with Conditional Source-Term Estimation*. PhD thesis, 2012.

[60] W. K. Bushe and H. Steiner. Laminar flamelet decomposition for conditional source-term estimation. *Physics of Fluids*, 15(6):1564, 2003.

[61] B. Jin, R. Grout, and W. K. Bushe. Conditional source-term estimation as a method for chemical closure in premixed turbulent reacting flow. *Flow, turbulence and combustion*, 81(4):563–582, 2008.

[62] M. Wang, J. Huang, and W. K. 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.

[63] A. K. Jain, M. N. Murty, and P. J. Flynn. Data clustering: a review. *ACM computing surveys (CSUR)*, 31(3):264–323, 1999.

[64] A. K. Jain. Data clustering: 50 years beyond k-means. *Pattern Recognition Letters*, 31(8):651–666, 2010.

[65] G. F. Jenks. The data model concept in statistical mapping. In Konrad Frenzel, editor, *International Yearbook of Cartography*, volume 7, pages 186+, USA, 1967. ICA, Rand McNally & Co.

[66] J. A Hartigan. Statistical theory in clustering. *Journal of Classification*, 2 (1):63–76, 1985.

[67] R. N. Davé and R. Krishnapuram. Robust clustering methods: a unified view. *Fuzzy Systems, IEEE Transactions on*, 5(2):270–293, 1997.

[68] P. Indyk and R. Motwani. Approximate nearest neighbors: towards removing the curse of dimensionality. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, STOC '98, pages 604–613, New York, NY, USA, 1998. ACM.

[69] H. D. Simon. Partitioning of unstructured problems for parallel processing. *Computing Systems in Engineering*, 2(2–3):135 – 148, 1991.

[70] B. Hendrickson and K. Devine. Dynamic load balancing in computational mechanics. *Computer Methods in Applied Mechanics and Engineering*, 184(2):485–500, 2000.

[71] B. Smith, P. Bjorstad, and W. Gropp. *Domain decomposition*. Cambridge University Press, 2004.

[72] M. Wittmann, T. Zeiser, G. Hager, and G. Wellein. Domain decomposition and locality optimization for large-scale lattice boltzmann simulations. *Computers and Fluids*, 80(0):283 – 289, 2013.

[73] M. J. Berger and S. H. Bokhari. A partitioning strategy for nonuniform problems on multiprocessors. *Computers, IEEE Transactions on*, C–36(5): 570–580, 1987.

[74] M. T. Jones and P. E. Plassmann. Computational results for parallel unstructured mesh computations. *Computing Systems in Engineering*, 5 (4–6):297 – 309, 1994.

[75] J. R Gilbert, G. L Miller, and S. Teng. Geometric mesh partitioning: Implementation and experiments. *SIAM Journal on Scientific Computing*, 19(6):2091–2110, 1998.

[76] T. Asano, D. Ranjan, T. Roos, E. Welzl, and P. Widmayer. Space filling curves and their use in the design of geometric data structures. In *LATIN'95: Theoretical Informatics*, pages 36–48. Springer, 1995.

[77] M. J. Aftosmis and M. J. Berger. Applications of space-filling curves to Cartesian methods for CFD, 2004.

[78] K. Schloegel, G. Karypis, and V. Kumar. Sourcebook of parallel computing. chapter Graph partitioning for high-performance scientific simulations, pages 491–541. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2003.

[79] T. N. Bui and C. Jones. Finding good approximate vertex and edge partitions is np-hard. *Information Processing Letters*, 42(3):153 – 159, 1992.

[80] C. Farhat and M. Lesoinne. Automatic partitioning of unstructured meshes for the parallel solution of problems in computational mechanics. *International Journal for Numerical Methods in Engineering*, 36(5): 745–764, 1993.

[81] J. D. Teresco, K. D. Devine, and J. E. Flaherty. Partitioning and dynamic load balancing for the numerical solution of partial differential equations. In A. Bruaset and A. Tveito, editors, *Numerical Solution of Partial Differential Equations on Parallel Computers*, volume 51 of *Lecture Notes in Computational Science and Engineering*, pages 55–88. Springer Berlin Heidelberg, 2006.

[82] G. Zumbusch. On the quality of space-filling curve induced partitions. *Z. Angew. Math. Mech*, 81:25–28, 2000.

[83] M. Bader. *Space-Filling Curves: An Introduction with Applications in Scientific Computing*, volume 9. Springer, 2012.

[84] B. Moon, H. V. Jagadish, C. Faloutsos, and J. H. Saltz. Analysis of the clustering properties of the hilbert space-filling curve. *Knowledge and Data Engineering, IEEE Transactions on*, 13(1):124–141, 2001.

[85] M. Parashar and J. C. Browne. On partitioning dynamic adaptive grid hierarchies. In *System Sciences, 1996., Proceedings of the Twenty-Ninth Hawaii International Conference on,*, volume 1, pages 604–613, 1996.

[86] J. R. Pilkington and S. B. Baden. Dynamic partitioning of non-uniform structured workloads with spacefilling curves. *Parallel and Distributed Systems, IEEE Transactions on*, 7(3):288–300, 1996.

[87] C. Gotsman and M. Lindenbaum. On the metric properties of discrete space-filling curves. *Image Processing, IEEE Transactions on*, 5(5): 794–797, 1996.

[88] X. Liu and G. F. Schrack. An algorithm for encoding and decoding the 3-D Hilbert order. *IEEE transactions on image processing: a publication of the IEEE Signal Processing Society*, 6(9):1333–7, 1997.

[89] X. Liu. Four alternative patterns of the hilbert curve. *Applied Mathematics and Computation*, 147(3):741–752, 2004.

[90] J. K. Lawder and P. J. H. King. Querying multi-dimensional data indexed using the hilbert space-filling curve. *ACM Sigmod Record*, 30(1):19–24, 2001.

[91] H. Sagan. *Space-Filling Curves*. Springer-Verlag, New York, 1994.

[92] P. M. Campbell, K. D. Devine, J. E. Flaherty, L. G. Gervasio, and J. D. Teresco. Dynamic octree load balancing using space-filling curves. *Williams College Department of Computer Science, Tech. Rep. CS-03-01*, 2003.

[93] M. Connor and P. Kumar. Fast construction of k-nearest neighbor graphs for point clouds. *Visualization and Computer Graphics, IEEE Transactions on*, 16(4):599–608, 2010.

[94] A. Perez, S. Kamata, and E. Kawaguchi. Peano scanning of arbitrary size images. In *Pattern Recognition, 1992. Vol.III. Conference C: Image, Speech and Signal Analysis, Proceedings., 11th IAPR International Conference on*, pages 565–568, 1992.

[95] S. Tirthapura, S. Seal, and S. Aluru. A formal analysis of space filling curves for parallel domain decomposition. *International Conference on Parallel Processing*, pages 505–512, 2006.

[96] F. T. C. Yuen and Ö. L. Gülder. Dynamics of Lean-Premixed Turbulent Combustion at High Turbulence Intensities. *Combustion Science and Technology*, 182(4–6):544–558, 2010.

[97] J. Kariuki, J. R. Dawson, and E. Mastorakos. Measurements in turbulent premixed bluff body flames close to blow-off. *Combustion and Flame*, 159 (8):2589 – 2607, 2012.

[98] C. Reynolds. *Performance of a partially stratified-charge natural gas engine*. PhD thesis, University of British Columbia, 2001.

[99] OpenCFD. Openfoam software, 2013. URL http://www.openfoam.org/.

[100] N. Chakraborty and R. S. Cant. Effects of Lewis number on scalar transport in turbulent premixed flames. *Physics of Fluids*, 21(3):035110, 2009.

[101] D. Veynante, J. M. Duclos, and J. Piana. Experimental analysis of flamelet models for premixed turbulent combustion. In *Symposium (International) on Combustion*, volume 25, pages 1249–1256. Elsevier, 1994.

[102] O. Colin and A. Benkenida. The 3-zones extended coherent flame model (ECFM3Z) for computing premixed/diffusion combustion. *Oil & gas science and technology*, 59(6):593–609, 2004.

[103] V. Mittal and H. Pitsch. A flamelet model for premixed combustion under variable pressure conditions. *Proceedings of the Combustion Institute*, 34 (2):2995–3003, 2013.

[104] Benjamin T. Zoller, M. L. Hack, and P. Jenny. A PDF combustion model for turbulent premixed flames. *Proceedings of the Combustion Institute*, 34 (1):1421–1428, 2013.

[105] H. Barths, C. Hasse, G. Bikas, and N. Peters. Simulation of combustion in direct injection diesel engines using a eulerian particle flamelet model. *Proceedings of the Combustion Institute*, 28(1):1161 – 1168, 2000.

[106] A. Odedra and W. Malalasekera. Eulerian particle flamelet modeling of a bluff-body CH_4/H_2 flame. *Combustion and Flame*, 151(3):512–531, 2007.

[107] N. Swaminathan and R. W. Bilger. Analyses of conditional moment closure for turbulent premixed flames. *Combustion Theory and Modelling*, 5(2):241–260, 2001.

[108] J. Huang, P.G. Hill, W.K. Bushe, and S.R. Munshi. Shock-tube study of methane ignition under engine-relevant conditions: experiments and modeling. *Combustion and Flame*, 136(1–2):25 – 42, 2004.

[109] I. Yimer, I. Campbell, and L.-Y. Jiang. Estimation of the turbulent schmidt number from experimental profiles of axial velocity and concentration for high-reynolds-number jet flows. *Canadian Aeronautics and Space Journal*, 48(3):195–200, 2002.

[110] 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: 1–36, 2004.

[111] M. Thiele, S. Selle, U. Riedel, J. Warnatz, and U. Maas. Numerical simulation of spark ignition including ionization. *Proceedings of the combustion institute*, 28(1):1177–1185, 2000.

[112] T. Kravchik and E. Sher. Numerical modeling of spark ignition and flame initiation in a quiescent methane-air mixture. *Combustion and flame*, 99(3): 635–643, 1994.

[113] S. Pischinger and J. B. Heywood. A model for flame kernel development in a spark-ignition engine. In *Symposium (International) on Combustion*, volume 23, pages 1033–1040. Elsevier, 1991.

[114] Z. Tan and R. D. Reitz. An ignition and combustion model based on the level-set method for spark ignition engine multidimensional modeling. *Combustion and flame*, 145(1):1–15, 2006.

[115] B. Enaux, V. Granet, O. Vermorel, C. Lacour, C. Pera, C. Angelberger, and T. Poinsot. LES study of cycle-to-cycle variations in a spark ignition engine. *Proceedings of the Combustion Institute*, 33(2):3115–3122, 2011.

[116] O. Yaşar. A new ignition model for spark-ignited engine simulations. *Parallel computing*, 27(1):179–200, 2001.

[117] E. Mastorakos. Ignition of turbulent non-premixed flames. *Progress in Energy and Combustion Science*, 35(1):57 – 97, 2009.

[118] C. K. Law. Heat and mass transfer in combustion: fundamental concepts and analytical techniques. *Progress in energy and combustion science*, 10 (3):295–318, 1984.

[119] D. A. Eichenberger and W. L. Roberts. Effect of unsteady stretch on spark-ignited flame kernel survival. *Combustion and flame*, 118(3): 469–478, 1999.

[120] G. Rassweiler and L. Withrow. *Motion pictures of engine flames correlated with pressure cards*. Warrendale, PA : Society of Automotive Engineers, 1980.

[121] R. K. Hanson and S. Salimian. Survey of rate constants in the N/H/O system. In *Combustion Chemistry*, pages 361–421. Springer US, 1984.