

THESIS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY IN THERMO AND
FLUID DYNAMICS

Super-grid Linear Eddy Model (SG-LEM)

Efficient mode- and regime-independent combustion closure for Large Eddy Simulation
(LES)

ABHILASH M. MENON

Department of Mechanics and Maritime Sciences

Division of Transport, Energy and Environment

CHALMERS UNIVERSITY OF TECHNOLOGY

Gothenburg, Sweden 2025

Super-grid Linear Eddy Model (SG-LEM)
Efficient mode- and regime-independent combustion closure for Large Eddy Simulation
(LES)
ABHILASH M. MENON
ISBN 978-91-8103-166-9

© ABHILASH M. MENON, 2025

Doktorsavhandlingar vid Chalmers tekniska högskola
Ny serie nr. 5624
ISSN 0346-718X
Department of Mechanics and Maritime Sciences
Division of Transport, Energy and Environment
Chalmers University of Technology
SE-412 96 Gothenburg
Sweden
Telephone: +46 (0)31-772 1419

Cover:
Schematic of the SG-LEM method for LES

Chalmers Reproservice
Gothenburg, Sweden 2025

ABSTRACT

Practical combustion devices, such as ICEs (Internal Combustion Engines), often exhibit characteristics of premixed and non-premixed modes, as well as spatial variation in combustion regimes. Such ‘mixed-mode’/‘multi-regime’ scenarios are a challenge for combustion modelling as models often rely upon assumptions regarding the mode, as well as on the fast vs. slow nature of the underlying chemistry. Newer combustion technologies like ‘lean burn’ for gas turbines introduce further challenges such as ‘blowout’ and complex ignition sequences that require adequate description of turbulence/chemistry interaction (TCI) for accurate simulation. Hence, there is a great need for a mode- and regime-independent combustion model that accounts for TCI for the simulation of next-generation combustion devices. The Linear Eddy Model (LEM) is one that meets these criteria. When used with Large Eddy Simulation (LES), a technique called LES-LEM, it is able to simulate all the processes of reactive flow at their respective scales: large- and small-scale advection, molecular diffusion, and chemical reactions. LES-LEM is computationally intensive as each LES cell is embedded with a highly resolved LEM domain which requires numerical time-advancement, and so the technique finds limited use in industrial simulations.

This dissertation introduces a novel variant of LES-LEM that uses coarse-graining of the LES mesh, and resulting down-scaling of the number of embedded LEM domains, to significantly reduce compute times while still producing LES quality temperature and concentration fields by means of a presumed PDF (probability density function) mapping closure. Large-scale transport between neighbouring LEM domains is simulated through a novel Lagrangian ‘splicing’ scheme. The method, termed super-grid LES-LEM (or SG-LEM), is validated for three distinct flame cases: a premixed backward-facing case, the well-studied Volvo Validation Rig, and a mixed-mode flame produced by the Darmstadt burner. The work also identifies key drawbacks resulting from the effects of coarse-graining on large-scale transport. However, mitigation methods are introduced and tested which also leads to suggestions for future implementations.

Overall, SG-LEM is a significant step in the practical application of LES-LEM, and its many benefits, to real-world technical flames.

Keywords: Large Eddy Simulation; Linear Eddy Model; mapping closure; Lagrangian splicing; mesh coarse-graining; mixed-mode flames

LIST OF PUBLICATIONS

This thesis is based on the following appended papers:

- Paper I** Menon, A. M., Oevermann, M. and Kerstein, A. R. (2023) “A super-grid approach for LES combustion closure using the Linear Eddy Model”, *Combustion Theory and Modelling*, 28(1), pp. 99–126. doi: 10.1080/13647830.2023.2260351.
- Paper II** Menon, A. M., Kerstein, A. and Oevermann, M. (2024) “Investigation of coarse-graining parameters for super-grid LEM closure applied to LES of practical bluff-body flames”, *Combustion Theory and Modelling*, pp. 1–22. doi: 10.1080/13647830.2024.2428156.
- Paper III** Menon, A., Kerstein, A., Oevermann, M., 2024. “Assessing the Multi-Regime Capability of the Super-Grid Linear Eddy Model (SG-LEM) Using the Darmstadt Multi-Regime Burner”, *Flow, Turbulence and Combustion*, doi: 10.1007/s10494-024-00602-x.

ACKNOWLEDGEMENTS

Firstly, I'd like to thank my supervisor Michael Oevermann for giving me this opportunity, his sincere and careful guidance contributed greatly to this work and allowed me to grow both professionally and academically during this time. A special thanks to Alan Kerstein for the spirited discussions we've had, for sharing his seemingly limitless ideas with our research group and for driving us to innovate and seek higher goals.

I couldn't ask for nicer colleagues than the ones the at the division of Transport, Energy and Environment (formerly Energy Conversion and Propulsion Systems) at Chalmers. I shall fondly remember the cooperative spirit, the enthusiastic fika-time discussions and good humour.

Thank you, Nidal Doubiani for being a fantastic office-mate and for the many brainstorming sessions we've had by the whiteboard. Special thanks also to Henrik Ström who brought much needed perspective to our projects during our weekly meetings.

Funding for the project from the Swedish Research Council is gratefully acknowledged as is the compute resources provided by the Swedish National Infrastructure for Computing (SNIC).

I'd like to thank my family for their unwavering support and encouragement. Life in Gothenburg is made immeasurably better by friends like Dani, Nick, Irena, Elin, Malin and Ragnhild – to name a few. Last but not least, thank you, Hanna for being my closest friend and an amazing partner.

Contents

Abstract	i
List of publications	iii
Acknowledgements	v
1 Introduction	1
1.1 Background	1
1.2 Reactive CFD	1
1.3 Challenges and motivation	2
1.4 Challenges and motivation	3
1.5 Objective	5
1.6 Thesis outline	5
2 Theory	7
2.1 The Exact Governing Equations	7
2.1.1 The Enthalpy Equation	8
2.1.2 Specific Heat and Temperature	9
2.1.3 Computational Approaches for Reacting Flows	10
2.1.4 LES vs. RANS for combustion	11
2.2 Governing Equations for LES	11
2.2.1 Momentum Closure	13
2.3 Combustion Modelling	14
2.3.1 Well-Stirred Reactor	15
2.3.2 Partially-Stirred Reactor	16
2.3.3 The Eddy Break-Up model	17
2.3.4 Statistical Approaches	17
2.3.5 Resolving Flame Fronts	18
2.3.6 The Linear Eddy Model	19
2.4 LES-LEM	20
2.5 Governing equations for LES-LEM	22
2.5.1 Small-scale evolution	22
2.5.2 Large-scale transport: splicing	25
2.5.3 Coupling in LES-LEM	26
2.5.4 Features of LES-LEM	27

2.6	Super-grid LES-LEM	28
2.6.1	SG-LEM coupling strategy	30
2.6.2	Secondary Output	33
3	Methodology	35
3.1	Introduction	35
3.2	LEM Domain Implementation	35
3.3	LEM Advancement and Operator Splitting	37
3.3.1	Sampled-sequence operator splitting	37
3.3.2	Blocked-sequence operator splitting	39
3.3.3	Volumetric expansion	39
3.4	Super-Grid Generation	40
3.5	Splicing for SG-LEM	41
3.5.1	Inflow and outflow boundary conditions	43
3.5.2	Fixed wall boundary conditions	44
3.6	Interface between SG-LEM and LES	44
3.7	Conditional Averages	46
3.7.1	Initialisation of Solution Tables	46
3.8	PDF Shapes	47
3.8.1	PDF Moments: numerical integration	48
4	Simulations and Main Results	51
4.1	Introduction	51
4.2	Case I: Backward Facing Step	51
4.3	Case II: Volvo Validation Rig	56
4.4	Case III: Darmstadt Multi-Regime Burner	61
5	Conclusions and Future Work	67
5.1	Future Work	69
	Bibliography	71

1 Introduction

1.1 Background

Humanity has harnessed combustion since time immemorial, it powers the modern world and finds use in sectors ranging from large-scale power generation, aviation and shipping to surface transport and domestic cooking. The use of fossil-fuel-driven combustion technology has, over time, had knock-on detrimental effects to human health and nature. These come by way of pollutants and climate change led by greenhouse gases such as CO_2 and has prompted a global response from governments to try and shift the socio-technical landscape of these activities to cleaner and safer standards. Emission norms for automobiles grow ever stricter on the limits for pollutants like CO and NO_x , and also increasing competition from battery electric cars. Similar trends are seen in aviation and shipping. These factors pressure traditional combustion technologies to continually innovate to remain competitive and in step with global sustainability commitments. In this vein, enormous research effort goes into improving fuel quality, after-treatment systems and design methodology for devices such as internal combustion engines (ICEs) and gas-turbines. Combustion devices are often complex, requiring several years of design and testing by a highly-skilled workforce before entering production. Over the years, simulation methods have emerged as an invaluable tool in the design cycles for these devices. One such is reactive flow simulations using Computational Fluid Dynamics (CFD) and forms the background area of interest for this work. It allows engineers to design and test several aspects of a combustion device, e.g., flow geometry that promotes mixing of air and fuel, the effect of various fuel blends on pollutant formation, flame instabilities and other phenomena.

1.2 Reactive CFD

Simulation of turbulent flows entails the time-advancement of the Navier-Stokes equations on a computational mesh. These equations represent the conservation of mass and momentum in an Eulerian framework. Additionally, transport equations for energy and chemical species conservation are required. Direct Numerical Simulation (DNS) of the above means resolving of all turbulent scales and flame structures and so requires an extremely fine temporal and spatial discretization. This approach is computationally intractable for real-world flames and is instead used for fundamental studies that involve small geometries, low turbulence, and simple fuel chemistries. The solution approaches developed for practical flames involve either time-averaging, termed Reynolds-averaged Navier-Stokes (RANS), or spatial/spectral filtering, termed Large Eddy Simulation (LES), of the exact equations. These approaches generate flow fields at much lower resolution than DNS but are more practical and widely used in engineering. The usefulness of reactive CFD for in designing combustion devices depends upon its ability to reliably simulate the key components of combustion, such as:

1. The breakup behaviour of droplets in fuel sprays.

2. Evaporation (or devolatilization) of fuels to their gaseous forms.
3. Turbulent mixing of air and fuel.
4. Molecular diffusion.
5. Chemical reactions and heat release.
6. Heat exchange between reacting gases and device components.

The importance of each of these aspects can vary greatly depending on the type of fuel, combustion mode, propagation regimes, and lastly, on the desired level of output detail. Combustion CFD involves a large separation of scales as chemistry and diffusion (which determine flame structure) are molecular phenomena while turbulent transport is resolved at much larger *continuum* scales. Turbulence also plays a significant role in the combustion process and, via turbulence/chemistry interaction (TCI), creates transient behaviour like extinction and re-ignition. For ICEs, investigations have shown that the predictive capability of reactive CFD depends strongly on the handling of these nuanced interactions [1]. In a general sense, reactive CFD can be seen as a coupling of flow simulation with computing chemical reactions.

LES has been found to be suitable for ICE simulations due the importance of large-scale flow structures and transient phenomena [2]. LES is also more suited to unsteady phenomena like blow-out or re-ignition in a gas turbine combustor [3] whereas Reynolds-averaged methods are more suited to mean stationary flows. These factors are discussed in more detail in the upcoming chapters.

1.3 Challenges and motivation

The role of combustion modelling in CFD is to simulate the effects of chemical reactions and heat release on the flow-field, i.e., changes in composition, temperature, density and pressure. This is known as chemical closure. Typically, combustion models are designed using certain assumptions regarding the flame being studied, such as:

- The combustion mode, i.e., premixed vs. non-premixed state of the fuel-air mixture.
- The propagation regime, which depends on the relation between chemical timescales (fast vs. slow chemistry) and Kolmogorov turbulent timescales. Regimes are classified as corrugated flamelets, thin reaction zone, or broken reaction zones.
- The treatment of TCI, some models ignore it altogether while other include it in their formulation.

Over the years, several combustion models have been developed using the above assumptions. Simple models like the Eddy Break-up Model (EBU) [4] assume fast chemistry and describes combustion purely as a small-scale mixing process. The mixing term relies on a parameter that is tuned to reproduce experimental or DNS data, it also exhibits case-by-case variation. The fast chemistry assumption cannot be used predict pollutants like CO – an intermediate species that requires a finite-rate treatment of chemistry. In

contrast, the Well-stirred Reactor (WSR) model neglects (sub-grid) turbulent mixing altogether, but allows for such a finite-rate description using arbitrarily complex reaction mechanism. It is often called a “no model” approach and cannot capture TCI unless used with meshes fine enough to resolve all turbulent scales. As for numerics, finite-rate models require temporal integration of highly non-linear reaction rates, which is usually the most time consuming sub-step during simulation, depending on the complexity of the reaction mechanism. Some finite-rate models can account for turbulent mixing, e.g., the Partially-Stirred Reactor (PaSR) [5] and Eddy Dissipation Concept (EDC) [6], they also rely on tuned model coefficients and require input from turbulence closure models. Statistical descriptions of small-scale mixing utilise probability density functions (PDFs). Presumed PDF approaches utilise well-described mathematical functions, they are often used in conjunction with flamelet models [7] and Conditional Moment Closure (CMC) [8]. A more involved alternative is to advance transport equations for the PDFs (termed TPDF) [9, 10] which has the advantage of an exact treatment of the chemical source term. The models listed above can handle both premixed and non-premixed combustion, and are discussed in more detail later in the text. Several models specialise to one or the other, e.g., Bray-Moss-Libby (BML) [11] for premixed flames and Representative Interactive Flamelets (RIF) [12] for non-premixed, in particular for diesel ICEs.

1.4 Challenges and motivation

It has been observed that real-world flames often exhibit characteristics of both premixed and non-premixed flames, and also variation in the propagation regime, i.e., they do not conform to rigid classification that standard combustion models assume. Newer combustion technology like Homogeneous Charge Compression Ignition (HCCI) and ‘lean burn’ for gas turbines bring additional modelling challenges in the form of transient phenomena like ‘blow-out’ which requires adequate modelling of TCI, ‘mixed-mode’ operating conditions, and sometimes features *auto-ignition*. Furthermore, Hydrogen is now being explored as a fossil-fuel alternative for transport and in steel production, which brings its own modelling challenges owing to its high diffusivity and differential diffusion effects which contributes to, e.g., elevated burning temperatures, combustion instabilities and the formation of NO_x – a temperature-sensitive pollutant. These factor motivate the development of mode- and regime-independent combustion models that also account for TCI. Given this background, an ideal combustion model for next-generation devices can be imagined as having the following qualities:

1. Mode-independence, i.e., it does not assume a premixed or non-premixed nature of the flame.
2. Propagation-regime-independence, i.e., it does not require assumptions on the chemical time scales relative to turbulent time scales.
3. Differential diffusion capability for fuels like H_2
4. Computationally efficiency for use in industrial simulations.

These combined goals are not straightforward to achieve. Of the models listed above,

TPDF makes fewer assumptions regarding the combustion process, is quite general in nature, and uses a stochastic processes for small-scale mixing. The standard approach involves recasting the transport equations for high-dimensional PDFs into as stochastic differential equations for notional particles. Micro-mixing, i.e., the combination of small-scale stirring and diffusion, constitutes a significant modelling challenge. It is carried out on these virtual particles using stochastic micro-mixing models like Modified Curl [13] and Euclidean Minimum Spanning Tree [14]. TPDF, on the whole, is computationally expensive.

The Linear Eddy Model (LEM) of Kerstein [15] meets most of the above criteria. Flame structures are simulated on highly resolved one-dimensional (1D) domains and differs from TPDF in that micro-mixing is simulated rather than modelled. This is achieved using physical diffusion coefficients and length-scale breakdowns governed by established scaling laws for isotropic turbulence. Furthermore, it is straightforward to incorporate different diffusivities for each component in the reacting mixture, useful for fuels like H_2 . LEM was originally developed as standalone mixing model [15] and has shown to exhibit scalar power spectra similar to DNS [16]. It was later extended for 1D simulation of diffusion-flames [17] and freely propagating flames [18]. As for CFD simulation, LEM was formulated as a sub-grid combustion closure for LES by Menon et al. [19], a technique known as LES-LEM (or LEM-LES), and also to Reynolds-averaged methods, particularly for ICEs known as Representative Interactive LEM (RILEM) [20, 21].

This work is based on LES-LEM, where each LES cell is embedded with 1D LEM domain that independently evolves reaction-diffusion processes, concurrently with the LES solver. It has been used extensively in literature to for the study of premixed flames [22, 23], spray combustion [24, 14], scalar mixing in supersonic layers [25], and partially-premixed combustion [26] to name a few. Chen et al. [27] had used LEM as a micro-mixing model for TPDF to study the effects of differential diffusion of H_2 in fuel mixtures. More recently, Li et al. [28] used LES-LEM to study the effects of differential diffusion in a swirl stabilised partially premixed syngas flame. Since TCI is simulated based on well-defined physical laws, LEM lends itself to simulating transient phenomena like combustion instabilities, as demonstrated by Srinivasan et al. [29], and also pollutant formation near lean blow-out conditions (broken reaction zone regime) in gas turbines by Eggenpieler et al. [30].

The above examples leverage the combined qualities of LES and LEM: accurate large-scale advection and mixing provided by LES with sub-grid micro-mixing, flame-structures and chemical reactions simulated by LEM. They also demonstrate how LEM closure meets most requirements for an ideal combustion model, the sole exception being that of computational affordability, which is a major drawback of LES-LEM. The computational cost is associated with the chemistry advancement of multiple LEM domains in the LES simulation, which also scales with LEM resolution. Small-scale stirring (especially at high Reynolds numbers), as well as the treatment of large-scale advective fluxes for LEM are additional costs which shall be discussed later in the text. Hence, there is a need for significant computational speed-up which, if the accuracy and generality of LES-LEM can be retained, would make the model far more attractive to industry.

The current work employs coarse-graining of the LES mesh as a speed-up technique for

LES-LEM. Due to the coarse-grained mesh, or ‘super-grid’, the method is called super-grid LES-LEM or SG-LEM.

1.5 Objective

The objective of this work is then to describe, implement and validate SG-LEM as combustion closure for LES. More specifically, the goals are outlined as follows:

1. Primarily, demonstrate computational speed-up relative to standard LES-LEM.
2. Produce LES quality concentration, temperature and density fields, despite the coarse-grained formulation.
3. Demonstrate its predictive capability by comparison with experiments or DNS data.
4. Demonstrate that it can retain the mode- and regime-independent qualities of LES-LEM.

1.6 Thesis outline

The remainder of this thesis is organised as follows: Chapter 2 details the theoretical and mathematical aspects of LES, combustion closure, LEM, LES-LEM and the proposed SG-LEM model. Chapter 4 contains details of the methodology and implementation. Chapter 4 summaries the investigations and findings in the attached papers. Finally, Chapter 5 with concluding remarks and a discussion on future work.

2 Theory

2.1 The Exact Governing Equations

The governing equations for compressible, multi-species, reactive flow in the framework of continuum mechanics are the Navier-Stokes equations for the conservation of mass, momentum, chemical species, and energy. These are written as

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

$$\frac{\partial \rho u_i}{\partial t} + \frac{\partial \rho u_i u_j}{\partial x_j} = -\frac{\partial p}{\partial x_i} + \frac{\partial \tau_{ji}}{\partial x_j}, \quad (2.2)$$

$$\frac{\partial \rho Y_\alpha}{\partial t} + \frac{\partial \rho u_j Y_\alpha}{\partial x_j} = -\frac{\partial J_{\alpha,j}}{\partial x_j} + S_\alpha, \quad \text{where } \alpha = 1, \dots, N-1 \quad (2.3)$$

and

$$\frac{\partial \rho E}{\partial t} + \frac{\partial \rho u_j E}{\partial x_j} = \frac{\partial}{\partial x_j} (-q_j + u_i \tau_{ji} - p u_j). \quad (2.4)$$

Here, ρ is the density, u_i ($i = 1, 2, 3$) is the velocity, p is the pressure, E is the total (kinetic plus internal) energy per unit mass, and τ_{ij} is the viscous stress tensor given by

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

where δ_{ij} is the Kronecker delta function and μ the dynamic viscosity. Body forces and radiative heating are ignored in the above equations. Species transport (mass fraction Y_α) has a diffusion component (Fickian flux) $J_{\alpha,j}$ computed as

$$J_{\alpha,j} = -\rho D_{\text{mix}} \frac{\partial Y_\alpha}{\partial x_j}, \quad (2.6)$$

where D_{mix} is the mixture averaged diffusion coefficient. Species flux due to pressure diffusion (Dufour effect) and temperature diffusion (Soret effect) are ignored here. Similarly, heat flux q_j is computed as

$$q_j = -\kappa \frac{\partial T}{\partial x_j} + \rho \sum_{\alpha=1}^N h_\alpha Y_\alpha J_{\alpha,j}, \quad (2.7)$$

where T is the temperature, κ is the mixture averaged thermal conductivity and h_α is species enthalpy. Temperature dependent viscosity is computed using Sutherland's law [31], which is simplified to

$$\mu = \sum_{\alpha=1}^N Y_\alpha \frac{A_{s,\alpha} T^{3/2}}{T + T_{s,\alpha}}, \quad (2.8)$$

based on a Sutherland coefficients $A_{s,\alpha}$ and reference temperature $T_{s,\alpha}$ for species α . The *exact* chemical source term S_α is given by the ansatz

$$S_\alpha = \rho \sum_{i=1}^M (\nu''_{\alpha,i} - \nu'_{\alpha,i}) \left(k_{f,i} \prod_{j=1}^N Y_j^{\nu'_{j,i}} - k_{b,i} \prod_{j=1}^N Y_j^{\nu''_{j,i}} \right) \quad \alpha = 1, \dots, N \quad (2.9)$$

where the reaction rate k_i is given by the Arrhenius expression

$$k_i = A_i T^{b_i} e^{-\frac{E_i}{R_u T}} \quad (2.10)$$

and subscripts b and f represent ‘backward’ and ‘forward’ reactions. Here, $\nu'_{i,\alpha}$ and $\nu''_{i,\alpha}$ are the stoichiometric coefficients of species α for the i ’th forward and backward reaction (M in total), which has Arrhenius rate, temperature exponent, and activation energy given as A_i , b_i and E_i , respectively. R_u is the universal gas constant. Finally, these equations are complemented by the equation of state for a mixture of ideal gases

$$p = \rho R_u T \sum_{\alpha=1}^N \frac{Y_\alpha}{W_\alpha}. \quad (2.11)$$

2.1.1 The Enthalpy Equation

Energy transport, i.e., Eqn. (2.4), can be recast using total enthalpy h_t and enthalpy h using the procedure below. With

$$E = e + \frac{1}{2} u_i u_i \quad (2.12)$$

where e is the internal energy. Total enthalpy h_t , defined as

$$h_t = E + \frac{p}{\rho}, \quad (2.13)$$

using the continuity equation (2.1), we arrive at the transformation

$$\rho \frac{dh_t}{dt} = \rho \frac{dE}{dt} + \frac{dp}{dt} + p \frac{\partial u_j}{\partial x_j}. \quad (2.14)$$

Using this in the energy equation (2.4) gives the conservation equation for h_t as

$$\frac{\partial \rho h_t}{\partial t} + \frac{\partial \rho u_j h_t}{\partial x_j} = \frac{\partial}{\partial x_j} (-q_j + u_i \tau_{ji} - p u_j) + \frac{dp}{dt} + p \frac{\partial u_j}{\partial x_j}. \quad (2.15)$$

A more convenient form of the energy equation, using enthalpy h , defined as

$$h = h_t - \frac{1}{2} u_i u_i, \quad (2.16)$$

can be obtained as follows. First, obtain the kinetic energy equation, i.e., $u_i u_i/2$, by multiplying the momentum equation (2.2) by u_i as

$$\frac{\partial}{\partial t} \left(\frac{1}{2} \rho u_i u_i \right) + \frac{\partial}{\partial x_j} \left(\frac{1}{2} \rho u_i u_i u_j \right) = -u_i \frac{\partial p}{\partial x_i} + u_i \frac{\partial \tau_{ij}}{\partial x_j}. \quad (2.17)$$

Then, subtract this from (2.15) to arrive at enthalpy transport –

$$\frac{\partial \rho h}{\partial t} + \frac{\partial \rho u_j h}{\partial x_j} = \frac{dp}{dt} - \frac{\partial q_j}{\partial x_j} + \tau_{ji} \frac{\partial u_j}{\partial x_i}, \quad (2.18)$$

where $\Phi = \tau_{ji} \frac{\partial u_j}{\partial x_i}$ is the viscous heating term. Enthalpy h , and its transport equation (2.18) was chosen for this work for the following reasons:

- It does not require a chemical source term, a point of convenience. Tests using *sensible* enthalpy transport, which requires such a term, introduced numerical instabilities when used in conjunction with the closure technique developed in this work.
- The test cases investigated in this work involved low velocities with Mach number $Ma \leq 0.3$. Thus, the effect of kinetic energy on temperature is far lower than that of chemical processes. Total enthalpy, i.e., Eqn. (2.15), ought to be used for high velocity cases.

The *low-Mach-number* assumption also allows us to ignore the viscous heating term. Finally, for open flames dp/dt can also be assumed to be zero; it is, however, important for confined flames such as in an ICE.

2.1.2 Specific Heat and Temperature

Total enthalpy h can be obtained from the caloric equation of state as

$$h = \int_0^T c_p(\theta) d\theta, \quad (2.19)$$

where c_p is the specific heat at constant pressure. Assuming constant c_p allows for a trivial temperature computation from the above. However, c_p is dependent on temperature, which is approximated using the polynomial expression

$$c_{p,\alpha}(\theta) = \sum_{k=0}^4 c_{k,\alpha} \theta^k \quad (2.20)$$

and composition data, which leads to

$$c_p(\theta) = \sum_{\alpha=1}^N Y_\alpha c_{p,\alpha}(\theta), \quad (2.21)$$

with $c_{k,\alpha}$ being the NASA polynomial coefficients [32] which are known and tabulated. This simplifies the integral in (2.19) to

$$h(T) = \sum_{\alpha=1}^N Y_\alpha \sum_{k=0}^4 \frac{c_{k,\alpha}}{k+1} T^{k+1} + c_{5,\alpha}. \quad (2.22)$$

With h from (2.18), temperature T is obtained as a Newton-Raphson iterative solution to the expression

$$T := T - \frac{h(T) - h}{dh/dT} \quad (2.23)$$

with $dh/dT = c_p(T)$ from (2.21), and an initial guess for T . Thermal conductivity κ_α is obtained via the modified Eucken expression (see Eqn. 10-3.5 in Poling et al. [33])

$$\kappa_\alpha = \mu_\alpha c_{v,\alpha} \left(1.32 + 1.77 \frac{R_u}{c_{v,\alpha}} \right), \quad (2.24)$$

where $c_{v,\alpha} = c_{p,\alpha} - R_u$ is the specific heat at constant volume. Thermal conductivity for the mixture is then

$$\kappa = \sum_{\alpha=1}^N Y_\alpha c_{v,\alpha}. \quad (2.25)$$

2.1.3 Computational Approaches for Reacting Flows

The numerical solution to the governing partial differential equations using the finite volume method involves division of the problem domain into discrete volumes, for each of which the discretised integral forms of the governing equations is conservative, requiring initial and boundary conditions. The resulting system of linear equations is then solved numerically to obtain flow variables at the centres of the discrete volumes. Together, these comprise a computational ‘grid’ or ‘mesh’; discrete volumes are often termed ‘cells’ - the terminology will be used throughout this work. Direct numerical simulation (DNS) of the above equations require that the mesh be resolved fine enough to capture the smallest turbulent scales and their effect on combustion, a proposition that is really only practical for small geometries and turbulent flow at low to moderate Reynolds numbers (Re). Simulation of real-world flows requires additional treatment of the governing equations, these typically fall into one of two categories:

1. Reynolds Averaged Navier Stokes (RANS), where the exact equations are Reynolds (ensemble) or Favre (density weighted) averaged to obtain the balance equations for the mean flow variables. Numerical solution to these equations provide mean flow fields, i.e., all turbulent fluctuations are averaged out and hence unresolved.
2. Large Eddy Simulation (LES), where the exact equations are spatially filtered to remove small-scale fluctuations. The solution to the filtered equations provide instantaneous flow field up to the resolved scales whereas turbulent scales smaller than the cut-off frequency are unresolved.

Each approach requires its own, very different, treatment of the unresolved turbulent scales, the so-called “closure problem”. A *turbulence model* is required for the unresolved flow dynamics and a *turbulent combustion model* to provide species conversion rates and heat release which can adequately capture the interplay between the unresolved flow-dynamics and chemistry.

2.1.4 LES vs. RANS for combustion

A majority of reacting flow simulations use RANS due to its affordability. RANS simulates a statistically averaged position for the flame front whereas DNS resolves the instantaneous flame front down to the underlying structure. As a consequence, the mesh requirements vary drastically – from tens of mm for RANS to tens of μm for DNS. LES offers an “in-between” solution as it requires the relevant scales, that contain most of the turbulent kinetic energy, to be resolved by the mesh. The method has been tested extensively for both non-reacting as well as reacting flows. As a result, LES simulation can be tens to hundreds of times more expensive than RANS. Despite this, reactive LES, and development of turbulent combustion models for the same, are active areas of research. Following are some features of reactive LES.

- Turbulence models developed for non-reacting flows, based on a similarity assumption between the resolved and unresolved scales, can often be reliably used for combustion LES.
- The large-scale flow dynamics differs from case to case, but there is a universal nature to the small unresolved scales. This separation is beneficial to turbulent combustion models for LES, whereas such an assumption cannot strictly be made for RANS due to the non-universality of turbulent motion on the integral scale.
- As LES can resolve large coherent structures in the flow field, it is more straightforward to simulate the interaction between the flow field, heat-release, and acoustic perturbations than RANS, which makes it a suitable tool to study combustion instabilities.
- LES is also better suited for turbulence-chemistry interaction (TCI) as it explicitly calculates regions of fresh and burnt gases, and is therefore able to resolve sharp gradients in density, viscosity, and turbulence levels.
- LES has proven to be superior over RANS for non-reacting and mixing problems, under complex flows. Mixing, here of fuel and air, is of utmost importance for combustion. It has also demonstrated more accurate results by resolving fluctuations in scalar dissipation rate [34].
- The treatment of chemical source terms, however, requires modelling just as in RANS.

2.2 Governing Equations for LES

Reactive LES uses a Favre-filtering operation on exact governing equation, decomposing fields into filtered and sub-filter components. This is written for a field ϕ as

$$\phi = \tilde{\phi} + \phi'' \tag{2.26}$$

with Favre-filtered $\tilde{\phi}$ and sub-filter ϕ'' . Favre-filtering is defined as

$$\tilde{\phi} = \overline{\rho\phi}/\bar{\rho} \tag{2.27}$$

where the over-bar represents spatial filtering, i.e.,

$$\overline{\rho\phi} = \int_{\Omega} \rho\phi G(x_i - r_i, \Delta) dr_i, \quad (2.28)$$

where G is the filter kernel. The kernel chosen must satisfy the condition

$$\int_{\Omega} G(x_i - r_i, \Delta) dr_i = 1, \quad (2.29)$$

where Δ is the LES filter width. Three types of filter kernels are generally used in LES i.e. box (top-hat), Gaussian and cut-off filters in spectral- or wave-number-space. This work uses an implicit top-hat filter, written in 1D as

$$G(x - r) = \begin{cases} \frac{1}{\Delta}, & \text{if } |x - r| \leq \frac{\Delta}{2} \\ 0, & \text{otherwise,} \end{cases} \quad (2.30)$$

where Δ equals (locally) the LES cell size, or *implicit* filtering. Filtering effectively removes high frequency (high wave-number) turbulence structures which makes LES far more tractable than DNS of the exact equations which would require all turbulent scales be resolved. Unlike RANS, LES also provides instantaneous flow variables which are important phenomena such as extinction-reignition cycles, or acoustic instabilities. Favre filtering applied to (2.3) and (2.18) results in the LES equations that are formally similar to (2.3),

$$\begin{aligned} \frac{\partial \bar{\rho}}{\partial t} + \frac{\partial \bar{\rho} \tilde{u}_i}{\partial x_i} &= 0, \\ \frac{\partial \bar{\rho} \tilde{u}_i}{\partial t} + \frac{\partial \bar{\rho} \tilde{u}_i \tilde{u}_j}{\partial x_j} &= -\frac{\partial \bar{p}}{\partial x_i} + \frac{\partial}{\partial x_j} (\bar{\tau}_{ji} + \bar{\tau}_{ji}^{\text{sgs}}), \\ \frac{\partial \bar{\rho} \tilde{Y}_\alpha}{\partial t} + \frac{\partial \bar{\rho} \tilde{u}_j \tilde{Y}_\alpha}{\partial x_j} &= \frac{\partial}{\partial x_j} (\bar{J}_{\alpha,j} + \bar{J}_{\alpha,j}^{\text{sgs}}) + \bar{S}_\alpha, \\ \frac{\partial \bar{\rho} \tilde{h}}{\partial t} + \frac{\partial \bar{\rho} \tilde{u}_j \tilde{h}}{\partial x_j} &= \overline{u_j \frac{\partial p}{\partial x_j}} + \frac{\partial}{\partial x_j} (\bar{q}_j + \bar{q}_j^{\text{sgs}}) + \frac{\partial \bar{p}}{\partial t}, \end{aligned} \quad (2.31)$$

where the symbols have the usual meanings but with the filtering operators, and new *sub-grid-scale* terms (superscript ‘sgs’). The filtered equation of state is then

$$\bar{p} = \bar{\rho} R_u \sum_{\alpha=1}^N \frac{\tilde{Y}_\alpha \tilde{T}}{W_\alpha} + \underbrace{\frac{\tilde{Y}_\alpha T - \tilde{Y}_\alpha \tilde{T}}{W_\alpha}}_{T_{\text{sgs}}} \quad (2.32)$$

which is reduced to $\bar{p} = \bar{\rho} R \tilde{T}$, if the species-temperature correlation, T_{sgs} , is ignored. The following terms are unclosed and need to be modelled:

- Unresolved Reynolds stress $\bar{\tau}_{ji}^{\text{sgs}} = \widetilde{u_j u_i} - \tilde{u}_j \tilde{u}_i$, modelling of which is termed *momentum closure*.
- Unresolved *convective* fluxes $\bar{J}_{\alpha,j}^{\text{sgs}} = \widetilde{u_j Y_\alpha} - \tilde{u}_j \tilde{Y}_\alpha$ and $\bar{q}_j^{\text{sgs}} = \widetilde{u_j h} - \tilde{u}_j \tilde{h}$.
- Filtered *diffusive* fluxes for which a gradient assumption can be used as

$$\bar{J}_{\alpha,j} = -\bar{\rho} \bar{D}_{\text{mix}} \frac{\partial \tilde{Y}_\alpha}{\partial x_j}, \quad (2.33)$$

and

$$\bar{q}_j = -\bar{\kappa} \frac{\partial \tilde{T}}{\partial x_j} = -\bar{\alpha} \frac{\partial \tilde{h}}{\partial x_j}, \quad (2.34)$$

which are comparable to (2.6) and (2.7). However, enthalpy transport due to species diffusion is neglected and $\alpha = \kappa/(\rho c_p)$ is the thermal diffusivity.

- The pressure velocity term $\overline{u_j \partial p / \partial x_j}$ is simplified using resolved variables as $\tilde{u}_j (\partial \bar{p} / \partial x_j)$.
- The filtered reaction rate \bar{S}_α , through *chemical closure* - a principal objective of *combustion modelling*.
- Sub-grid diffusive fluxes for temperature and species are usually neglected in high *Re* flows, they also lack well-defined models [35].

2.2.1 Momentum Closure

The unresolved Reynolds' stress, τ_{ij}^{sgs} , is modelled by comparing it with the unfiltered viscous stress, τ_{ij} , in (2.5). It can split into the deviatoric and isotropic parts as $\tau_{ij}^{\text{sgs}} = \tau_{ij,\text{dev}}^{\text{sgs}} + \tau_{kk}^{\text{sgs}}$. The Boussinesq assumption then says that $\tau_{ij,\text{dev}}^{\text{sgs}}$ is proportional to the deviatoric part of the resolved strain rate, \tilde{S}_{ij} , by a constant known as turbulent viscosity, ν_t . For variable density flows, this is written as

$$\tau_{ij,\text{dev}}^{\text{sgs}} = -2\nu_t \left(\tilde{S}_{ij} - \frac{1}{3} \tilde{S}_{kk} \delta_{ij} \right), \quad (2.35)$$

where

$$\tilde{S}_{ij} = \frac{1}{2} \left(\frac{\partial \tilde{u}_j}{\partial x_i} + \frac{\partial \tilde{u}_i}{\partial x_j} \right). \quad (2.36)$$

Using $\tau_{kk}^{\text{sgs}} = 2/3 k^{\text{sgs}}$, we get the expression

$$\tau_{ij}^{\text{sgs}} = -2\nu_t \left(\tilde{S}_{ij} - \frac{1}{3} \tilde{S}_{kk} \delta_{ij} \right) + \frac{2}{3} \bar{\rho} k^{\text{sgs}} \delta_{ij}, \quad (2.37)$$

where

$$\nu_t = C_k \Delta \sqrt{k^{\text{sgs}}} \quad (2.38)$$

using a model constant C_k . Hence, the turbulent kinetic energy, k^{sgs} remains to be modelled. Smagorinsky [36] proposed a simple algebraic model where

$$\nu_t = (C_s \Delta)^2 |\tilde{S}_{ij}|. \quad (2.39)$$

In this work, the Smagorinsky model is realised by substituting the solution to the quadratic equation

$$a(k^{\text{sgs}})^2 + bk^{\text{sgs}} + c = 0 \quad (2.40)$$

into (2.38) using the coefficients

$$a = \frac{C_e}{\Delta} \quad ; \quad b = \frac{2}{3} \tilde{S}_{kk} \quad ; \quad c = 2C_k \Delta (\tilde{S}_{ij} - \frac{1}{3} \tilde{S}_{kk}) : \tilde{S}_{ij}. \quad (2.41)$$

Model coefficients, C_e and C_k , are set to 1.048 and 0.094, respectively. The Smagorinsky model uses an equilibrium assumption for the production and dissipation of sub-grid turbulent kinetic energy, i.e., there is an instantaneous response of k^{sgs} to \tilde{S}_{ij} . This might not hold true for coarser LES meshes used in engineering applications or high Reynolds' number (Re) flows [37]. Non equilibrium models like Yoshizawa et al. [38] use a transport equation of k^{sgs} with production and dissipation terms. Sub-grid fluxes are then closed as

$$\bar{j}_{\alpha,j}^{\text{sgs}} = \frac{-\bar{\rho} \nu_t}{S_{c_t}} \frac{\partial \tilde{Y}_\alpha}{\partial x_j}, \quad (2.42)$$

and

$$\bar{q}_j^{\text{sgs}} = \frac{-\bar{\rho} \nu_t}{Pr_t} \frac{\partial \tilde{h}}{\partial x_j}, \quad (2.43)$$

with S_{c_t} and Pr_t being the turbulent Schmidt and Prandtl number, respectively. The gradient diffusion approach for Y_α has been shown to lead to significant errors when compared to experimental data [37]. The combustion modelling approach presented in this work, using LEM, directly simulates (as opposed to models) sub-grid scalar transport processes, both diffusive and convective, which relaxes many of the assumptions described in this section.

2.3 Combustion Modelling

Closure of the filtered reaction rate \bar{S}_α is not straightforward due to the highly non-linear nature of combustion chemistry. Favre-filtering, when applied to (2.9), leads to unclosed terms (sub-grid fluctuations) that cannot be ignored in turbulent combustion. The closure of these terms is a major research area in combustion modelling, termed turbulence/chemistry interaction TCI. Additionally, combustion models may fall into the reaction-rate-type or mapping-type categories. Reaction-rate type closures provide source terms \bar{S}_α for a direct time-advancement of the \tilde{Y}_α transport equations. On the other hand, mapping closures split the problem into mixing, flame-propagation, and flame structure. Mixing between fuel-oxidiser includes fluid-dynamic as well as molecular mixing. Flame

propagation determines the boundary between unburnt to burnt regions. These two steps require CFD advancement of primitive variable fields, i.e., mixture fraction and progress variable. Lastly, flame structure serves as a one-to-one map between composition and the joint-space of the primitive variables. Thus, the primitive variable fields can be used to map composition and temperature to every point in the CFD domain, informed by the flame structure. Usually, flame structures are computed from 1D canonical set-ups known as ‘flamelets’. There are some advantages to this second approach.

- The flame structure can be pre-computed and tabulated into a *flamelet library*, this decoupling of fluid dynamics and chemistry leads to significant savings in computational requirements.
- The \tilde{Y}_α equations are now replaced with one or two filtered primitive variable transport which help improve numerical stability.

An important assumption here is that the flame-front can be statistically represented as a collection of one-dimensional flame structures that are embedded in (or carried by), and also stretched or distorted by a turbulent flow field [7]. Mapping-type closures using flamelets have been extensively used for both premixed and non-premixed combustion. The accuracy can sometimes be challenged by transient phenomena, such as extinction and reignition, as well as in predicting species that do not adapt instantaneously to manifold solutions [39]. Soot, CO, or NO_x are examples of such species which are important in a real-world context. Mapping closures may be complemented by additional transport equations for such species (e.g. in Popp et al. [40]).

Reaction rate closures do not have these limitations; however, they are more expensive as \tilde{S}_α must be evaluated for every mesh cell. For the time-discretised (integral) form of the species transport equation (2.3), the averaged (over the discrete volumes) source term requires numerical integration of the ODEs (Ordinary differential equations) given by (2.9) and (2.10), for each reaction. The highly non-linear temperature dependence of reaction rate makes for stiff ODEs that need sophisticated numerical integrators. Detailed reaction kinetics, with many species and reactions, also lead to dramatic increases in compute costs.

The approach presented in this work is a hybrid of these two paradigms – for a given three-dimensional (1D) reacting flow problem, flame structure information is generated by representative 1D domains using finite-rate chemistry, which is used to map concentrations to the flow field via primitive scalars. Individual 1D domains correspond to various locations in the problem domain, and are informed by local conditions of composition and turbulence. To form an adequate background, some standard LES combustion closures are described below.

2.3.1 Well-Stirred Reactor

Also known as the “no model” approach, the Well-Stirred Reactor (WSR) model neglects sub-grid mixing altogether and expresses the filtered reaction rate purely as a function of filtered flow variables (shown only for the forward reaction) as

$$\bar{S}_\alpha = \bar{\rho} M_\alpha \sum_{i=1}^M (\nu''_{\alpha,i} - \nu'_{\alpha,i}) k_i \prod_{j=1}^N \tilde{Y}_j^{\nu'_{j,i}} \quad \alpha = 1, \dots, N \quad (2.44)$$

with

$$k_i = A_i \tilde{T}^{b_i} e^{-\frac{E_i}{R\tilde{T}}}. \quad (2.45)$$

The above expression assumes that all chemical timescales, τ_c , are larger than all turbulent mixing timescales τ_m ($\tau_c \gg \tau_m$). This is not valid for most combustion applications and has been shown to produce unacceptable errors [35], but is still widely used in ICE simulations. find limited use in reactive flows for atmospheric boundary layers.

2.3.2 Partially-Stirred Reactor

The Partially-Stirred Reactor (PaSR) is based on the work of Chomiak [5], the leading principle is that the turbulence-reaction problem domain can be divided into a reacting and non-reacting fractions that interact with each other through mass-exchange driven by concentration gradients. The fraction of the reacting part is given by

$$\gamma = \frac{\tau_c}{\tau_c + \tau_m}, \quad (2.46)$$

where the mixing timescales is approximated as

$$\tau_m = C_{\text{mix}} \sqrt{\frac{\mu + \mu_t}{\rho \epsilon}}, \quad (2.47)$$

where C_{mix} is a tunable constant. The turbulent dissipation, ϵ , is approximated for the Smagorinsky LES model as

$$\epsilon = C_e \frac{k_{\text{sgs}}^{3/2}}{\Delta}. \quad (2.48)$$

Chemical timescales (τ_c) present a more challenging computation. Wartha et al. [41] present a systematic comparison of the various methods used to compute τ_c . The method used for PaSR simulations carried out in this work is

$$\tau_c = \sum_{i=1}^M \frac{c_{\text{tot}}}{\sum_{\alpha=1}^N \nu'_{\alpha,i} k_{f,\alpha}}, \quad (2.49)$$

with k_f being the production rate and c_{tot} being the total number of moles. Finally, filtered reaction rates can then be related to (2.10) as

$$\bar{S}_\alpha = \gamma S_\alpha. \quad (2.50)$$

In the limit of perfectly stirred ($\tau_m = 0$) or perfectly segregated ($\tau_m = \infty$) conditions, $\bar{S}_\alpha = S_\alpha$ or 0, respectively. A criticism of the PaSR closure is the use of a single chemical timescale whereas detailed kinetic mechanism comprise both fast and slow chemistries, some works like Ref. [39] try to address this by scaling each species individually determined by the eigenvectors of the chemical Jacobian.

2.3.3 The Eddy Break-Up model

An assumption opposite to that in 2.3.1 can also be made ($\tau_c \ll \tau_m$) where combustion is said to be *mixing controlled*, such as in the Eddy break-up model (EBU) of Spalding [4]. Considering a simple global mechanism $F + O \rightarrow P$, the production rate can be expressed as

$$\bar{\omega}_F = C_{\text{EBU}} \bar{\rho} \frac{1}{\tau_t^{\text{sgs}}} \tilde{Y}_F (1 - \tilde{Y}_F), \quad (2.51)$$

where the sub-grid mixing timescale is given as

$$\tau_t^{\text{sgs}} \approx \frac{\Delta}{\sqrt{k^{\text{sgs}}}}, \quad (2.52)$$

and C_{EBU} is a heuristic constant. The decoupling of reaction progress from chemical kinetics leads to significant inaccuracies (fuel consumption rates and temperature) and tuning of the model constant not a simple task. More so, the fast chemistry assumption implies a sharp transition from unburnt to burnt states in physical space, which ignores the underlying flame structure and radical/pollutant prediction is not possible.

2.3.4 Statistical Approaches

An exact closure for the filtered reaction rates is obtained by assuming that the local thermo-chemical state can be described using joint-PDFs. Based on the work by Pope [9], these models either assume or compute PDF shapes from which any filtered flow variable, such as \tilde{S}_α , can be computed as a weighted first moment of the joint PDF. This is written as

$$\tilde{S}_\alpha = \int_{Y_1=0}^1 \int_{Y_2=0}^1 \cdots \int_{Y_N=0}^1 \int_{T=T_{\min}}^{T_{\max}} S_\alpha(Y_1, Y_2, \dots, Y_N, T) P(Y_1, Y_2, \dots, Y_N, T) \, dY_1 dY_2 \dots dY_N dT, \quad (2.53)$$

with P representing the joint PDF in composition-temperature space. This can either be evolved through an exact transport equation [9], a method called transported PDF (TPDF), or can be presumed to be an analytical distribution based on experimental observations, known as the presumed-PDF method. TPDF has the advantage of an exact closure for \tilde{S}_α but diffusion and sub-grid mixing are difficult to model, these phenomena play an important role in predicting reaction-rates and heat release. Furthermore, the standard numerical approach to TPDF involves Monte-Carlo simulations of notional particles which can be computationally prohibitive for highly resolved LES and detailed kinetics. More recent methods combine TPDF with flamelet-generated reaction rates (see [42]).

Typically, the β distribution is taken for presumed-PDF methods, although difficulties arise when the state-space contains more than two independent variables. Presumed

PDF methods are also used in mapping closures where filtered composition data (instead of reaction-rates) is arrived at by first moments similar to (2.53), but for a reduced state-space that comprises the primitive variables - mixture fraction and/or progress variable. The approach presented in this work falls under this category and discussed in more detail in the upcoming chapter.

2.3.5 Resolving Flame Fronts

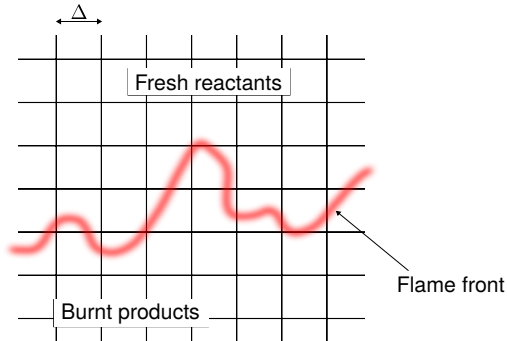


Figure 2.1: LES filter width, Δ , as compared to flame front thickness.

In the context of premixed flames for LES, a major challenge arises from the flame thickness inadequately resolved by the LES mesh, seen here in Fig. 2.1. The flame front is a sharp discontinuity in reactive scalars which can cause numerical issues. The Karlovitz number Ka , defined as the ratio between chemical and turbulent time scales, is commonly used to characterise the turbulence regime for premixed flames. The ‘corrugated flamelet regime’ exhibits $Ka < 1$, here turbulent eddies are too large to distort the laminar flame structure. This could be due to low turbulence, fast chemistry, or a combination of the two. The eddies, however, can wrinkle the overall flame front. Such flames are entirely on the sub-grid scale. The ‘thin reaction zone’ regime exhibits $1 \leq Ka < 100$, here turbulent transport broadens the preheat zone of the laminar flame structure [43] but eddies are too large to disrupt the inner reaction zone. The broadened flame front can be partially resolved by the LES mesh. For $Ka > 100$, i.e., the ‘broken reaction zone’ regime, turbulent eddies are small enough to disrupt the inner reaction zone which can lead to instabilities and local (or even global) extinction. Technical combustion devices try to avoid this regime for these reasons [34].

Since the contribution of reaction rates and heat release is largest at the sub-grid level, techniques have been developed to overcome this difficulty, some of which are given as follows.

- Artificially thickened flames (ATF) of Butler et al. [44] effects an increase in flame thickness by modifying diffusivity and the pre-exponential constant in (2.10) by an arbitrary ‘thickening’ factor so the flame front may be resolved by the LES mesh. Laminar flame speed is preserved; however, the increase in laminar flame thickness (w.r.t turbulent length scales) reduces the interaction of the flame front with small

vortical structures. A ‘wrinkling’ factor known as *efficiency function* was introduced to compensate for this effect, this requires known values for the flame speed and thickness.

- Flame front-tracking, e.g., the G-equation technique of Kerstein et al. [45], associates the flame brush with an iso-level of a filtered field variable– the “G-field”. Here, the flame front is not directly resolved by the LES-mesh but by a smoothed tracking variable. The transport equation for the G-field requires a model for the flame front conditioned burning velocity, which is not well-defined, and limited to fast chemistry.
- The use of Gaussian filters in physical space for reacting scalars with filter widths larger than the LES filter width – a ‘test’ filter. The filtered diffusion and source terms in the reactive scalar transport equations are replaced by a flame front displacement speed. This speed is based on the assumption that volumetric consumption of unburnt gases is proportional to the product of the laminar flame speed and flame surface area. Flame Surface Density (FSD) models provide the later (flame front convolutions of units [area/volume]) [46] using either a transport equation or an algebraic method.

2.3.6 The Linear Eddy Model

The Linear Eddy Model (LEM) approach for LES involves simulating (as opposed to modelling) flame structures, within each LES cell, using embedded highly resolved 1D domains. LEM forms the backbone in this work and addresses many of the issues described above. It involves deterministic simulations of diffusion and heat release combined with stochastic small-scale advection or stirring.

Turbulent fluid flow simulation remains a scientific challenge due to broad range of scales that need be captured - large scale advection results in progressive length-scale breakdown of inhomogeneities in the flow field down to the turbulent scales, followed by smoothing of turbulent fluctuations of composition and temperature by molecular diffusion. Combustion is a molecular phenomenon that creates fluctuations in temperature and density at the small scales which, along with heat release, feed back to the large scale fluid dynamics. This two-way interplay between the large, turbulent, diffusive remains a modelling difficulty. Inaccuracies at the smaller scales could result in significant error in scalar fields.

LEM closure is guided by the principle of scale separation, i.e., the aforementioned scales ought to be distinguished from one another and resolved (spatially and temporally) at their relevant scales. Resolving all scales in 3D (i.e. DNS) is intractable for most applications; instead, a one-dimensional (1D) approach would allow for full resolution at significantly reduced cost. A unique feature is that turbulent mixing in 1D is simulated using established turbulence scaling laws in the inertial range, and so is molecular diffusion using physical diffusivities. As a result, the turbulent flame speed is also simulated rather than requiring a model.

2.4 LES-LEM

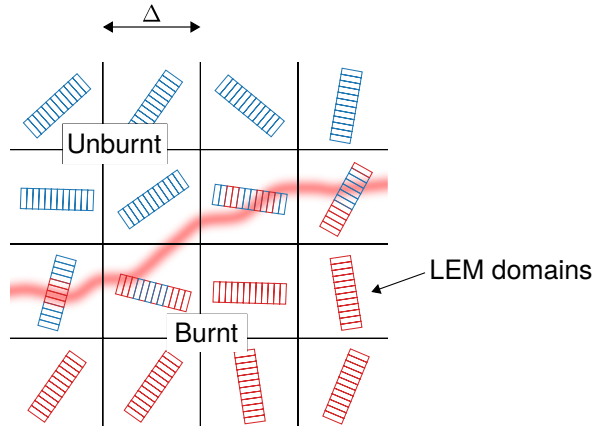


Figure 2.2: Schematic of LES-LEM closure. Δ is the LES cell size. Note that non-premixed flames would also contain interfaces between fuel and oxidiser at the LEM level.

LES-LEM is a two-scale solution methodology for the governing equations of reactive flows. The large scales are simulated using the filtered governing equations (2.31) using an appropriate LES mesh, whereas sub-grid processes are simulated using LEM domains that are embedded in each LES cell, as shown in Fig. 2.2. The large-scale processes consists of LES resolved advection for which SGS momentum closure is provided by the SGS models, such as described in Sec. 2.2.1, or any other SGS model deemed appropriate. The small-scale processes are sub-grid diffusion, sub-grid turbulent mixing (or stirring), and chemical reactions – these are simulated on 1D LEM domains using a Lagrangian fame of reference. In other words, the large scale processes range between the problem geometry and the LES grid resolution (or filter width), and the small scales range between the LES grid and the LEM domain resolution. LEM domains should ideally be resolved down to the flame structure. They are time-advanced independently of one another, but concurrently with the LES solver. The LES and LEM domains are two-way coupled, information exchange occurs at the following levels.

1. **Advection:** LES-resolved advection is coupled to the LEM domains by driving mass exchange. This is carried out using Lagrangian transport of LEM domain ‘fragments’ between neighbouring LEM domains – a process known as ‘splicing’. A fragment here is defined as a contiguous interval of an LEM domain with defined end points. The spliced masses are consistent with Eulerian mass transport between neighbouring LES cells given by the filtered mass and momentum equations.
2. **Turbulence:** While the above method simulates LES-resolved turbulence, sub-grid turbulence levels, as given by the SGS model, is used as an input for small-scale stirring that is simulated by the LEM domains.
3. **Coupling:** Finally, LEM domains provide either filtered chemical source terms or

filtered concentration fields to the LES solver.

Note that the current implementation does not couple the LES pressure field with LEM pressures. Hence, LEM density and temperature fields are computed assuming a constant background pressure. The constant pressure assumption is suitable for open flames. Closed combustion systems or super-sonic flows would require such a coupling.

The two-scale method is described in more detail as follows: consider a reacting scalar ϕ , the exact transport equation would read

$$\rho \frac{\partial \phi}{\partial t} + \underbrace{\rho u_j \frac{\partial \phi}{\partial t}}_{\text{I+II}} = -\rho \underbrace{\frac{\partial}{\partial x_j} \left(D \frac{\partial \phi}{\partial x_i} \right)}_{\text{III}} + \underbrace{\dot{\omega}_\phi}_{\text{source}}, \quad (2.54)$$

consisting of an advective term (I+II), a diffusive term (III) and a source term. Splitting the velocity into large and small scales, i.e., $u_j = \tilde{u}_j + u''$, (2.54) can be split into

$$\rho \frac{\partial \phi}{\partial t} + \underbrace{\rho \tilde{u}_j \frac{\partial \phi}{\partial t}}_{\text{I}} = 0 \quad (2.55)$$

and

$$\rho \frac{\partial \phi^*}{\partial t} + \underbrace{\rho u''_j \frac{\partial \phi^*}{\partial t}}_{\text{II}} = -\rho \underbrace{\frac{\partial}{\partial x_j} \left(D \frac{\partial \phi^*}{\partial x_i} \right)}_{\text{III}} + \underbrace{\dot{\omega}_{\phi^*}}_{\text{source}}, \quad (2.56)$$

which can together describe the complete evolution of ϕ . The notation ‘*’ denotes the small-scale distribution of ϕ and, importantly, the source term which accounts for chemical reactions is relegated to these small scales. As stated earlier, a full 3D resolution down to the small scales is intractable for most combustion applications. Instead, LES-LEM resolves the small scales, i.e., (2.56), in 1D using a ‘‘DNS like’’ resolution, while the aforementioned Lagrangian splicing is used to simulate (2.55). The 1D LEM domains comprise discrete ‘wafers’¹.

LES-LEM uses the following assumptions for low-Mach-number reacting flows:

1. The interior of every LES cell is in a state of Homogeneous Isotropic Turbulence (HIT), which is an idealised state where turbulent fluctuations are statistically uniform in all directions.
2. LEM formalism accurately captures the physics of scalar mixing under HIT.
3. Within each LES cell, statistical distributions of temperature and composition along any *line-of-sight* is independent of orientation.
4. Heat-release and thermal expansion is isotropic at the small scales.

The domain lengths correspond to local filter width, this representative approach lends to spatially accurate simulation of sub-grid diffusion. LEM domains then provide on-the-fly combustion closure *in-situ*, which requires a coupling mechanism between (2.55)

¹The term ‘wafer’ is used to distinguish it from ‘cell’, which is the LES finite volume discretisation.

and (2.56), and the LES-resolved composition (\tilde{Y}_α) and temperature (\tilde{T}). LEM domain orientation is *not* an input parameter, this can be interpreted in different ways:

1. LEM domains are randomly orientated as sub-grid turbulence is isotropic.
2. They are orientated along the strongest scalar gradients, i.e., normal to the flame front.
3. They are orientated along the resolved flux.

The interpretation does not play a role in the coupling algorithm either, but the third view is useful for splicing transport for LES-LEM, as shown in a later section.

2.5 Governing equations for LES-LEM

2.5.1 Small-scale evolution

The small-scale (2.56) is re-written for reaction-diffusion in a Lagrangian frame of reference pertinent to LEM simulation. For a domain length L_{LEM} set to the cells size Δ , this is written as

$$\rho c_p \frac{dT}{dt} + \underbrace{C_T}_{\text{II}} = -\frac{dq}{dx} + S_T, \quad (2.57)$$

$$\rho \frac{dY_\alpha}{dt} + \underbrace{C_{Y_\alpha}}_{\text{II}} = -\frac{dJ_\alpha}{dx} + S_\alpha, \quad (2.58)$$

along the LEM coordinate x , not to be confused with Cartesian x_i used earlier. S_T and S_α are source terms. Terms II (turbulent convection) are symbolically written as C_T and C_{Y_α} . Other symbols have the same meanings as in Sec. 2.1. Diffusive fluxes are modelled using a gradient assumption as

$$Q = -\kappa \frac{dT}{dx} \quad (2.59)$$

and

$$J_\alpha = -\rho D_{\text{mix},\alpha} \frac{dY_\alpha}{dx}. \quad (2.60)$$

Some important assumptions and features:

1. There is no velocity term in (2.57) and (2.58), the effect of small-scale fluctuations is through C_T and C_{Y_α} alone. Oevermann et al. [47] use an Eulerian frame of reference for the LEM equations with a velocity (along the LEM coordinate) that represents flow velocity due to dilational effects.
2. LEM is *not* a turbulence model but a turbulent mixing model, i.e., for diffusion and small-scale stirring. An extension of LEM, known as the ODT (One-dimensional Turbulence) model [48], includes the turbulent diffusion of momentum in its formulation.

3. There are no pressure gradients on the LEM domain, i.e., no pressure-driven flows or resolved shocks (low Ma). Pressure changes, say, to account for piston movement in an ICE using a slider-crank model, must be applied globally. For open flames, reaction rates needed for the source terms are computed assuming a constant pressure (or isenthalpic) process. Confined flames, like in a Diesel-cycle ICE, require a constant volume assumption for the same.
4. Thermally perfect gas law is assumed, supported by the equation of state (2.11), specific heat relations in Sec. 2.1.2 and viscosity as per (2.8). For open flames, a background pressure can be assumed for density and temperature calculations.
5. Radiation and small-scale viscous heating are neglected.

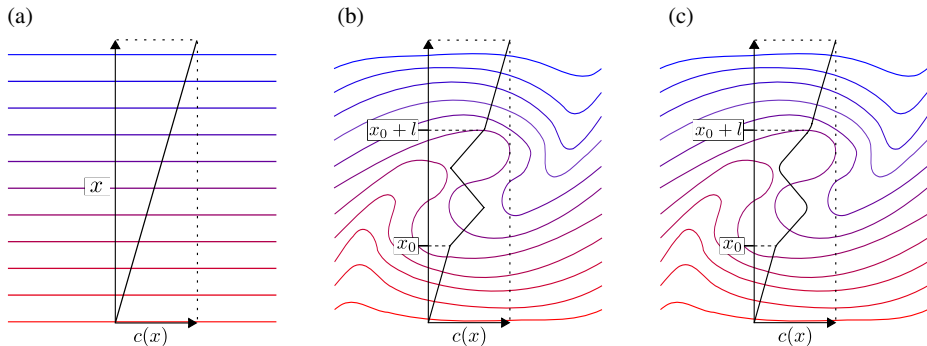


Figure 2.3: Rearrangement events: coloured lines are concentration isopleths; (a) shows an initial linear concentration gradient, (b) a triplet map which approximates the stirring effect of a single eddy turnover, and (c) the sharp gradients created by the map are smoothed by diffusion.

The evolution of the above equations is carried out in a deterministic manner *without* terms II, punctuated by ‘rearrangement events’ that model the effect of turbulent small scale advection (term II). These rearrangements represent the effect of individual eddies on the scalar field and henceforth called ‘eddy-events’. Kerstein [16] defines these as

a measure-preserving map of a spatial segment of the concentration field onto itself.

They are uniquely defined by three random variables: location x_0 , length l and epoch t_0 . The mapping rule adopted is the ‘triplet map’, illustrated for a concentration field $c(x)$ in Fig. 2.3. As the name implies, the map compresses the concentration field in $[x_0, x_0 + l]$ by a factor of three, replaces the interval with three copies of the compressed field, and

then flips the middle copy. Formally, this is written as

$$c(x, t_0) = \begin{cases} c(3x - 2x_0, t_0) & x_0 \leq x \leq x_0 + \frac{l}{3} \\ c(-3x + 4x_0 + 2l, t_0) & x_0 + \frac{l}{3} \leq x \leq x_0 + \frac{2l}{3} \\ c(3x - 2x_0 - 2l, t_0) & x_0 + \frac{2l}{3} \leq x \leq x_0 + l \\ c(x, t_0) & \text{otherwise.} \end{cases} \quad (2.61)$$

Eddy events are statistically independent and instantaneous. Eddy epochs (t_0) are determined by a stochastic Poisson process with a mean rate $\lambda \times \Delta$, where λ is termed the ‘eddy frequency parameter’ (units $[\text{length} \times \text{time}]^{-1}$) and given by

$$\lambda = \frac{54 \nu Re_\Delta [(\Delta/\eta)^{5/3} - 1]}{5 C_\lambda \Delta^3 [1 - (\eta/\Delta)^{4/3}]} \quad (2.62)$$

Here, η is the Kolmogorov scale, C_λ is a rate constant, and the turbulent Reynolds number, Re_Δ , is computed as

$$Re_\Delta = \frac{u' \Delta}{\nu_t} \quad (2.63)$$

Turbulent viscosity, ν_t , and velocity fluctuation, $u' = \sqrt{\frac{2}{3} k_{\text{sgs}}}$, are obtained from LES momentum closure (see Sec. 2.2.1). This can be interpreted as a turbulence level input from the resolved scales to the small-scale LEM evolution. The stirring time, i.e., the mean time *between* eddy events in the Poisson process is

$$\Delta t_{\text{stirr}} = \frac{1}{\lambda L_{\text{LEM}}} \quad (2.64)$$

At epoch t_0 , location x_0 is sampled from a uniform distribution in $[0, \Delta]$ while the length l is based on a size distribution $f(l)$ obtained using inertial range scaling laws of 3D turbulence [15]. Eddy sizes range between the length of the domain to the Kolmogorov scales, the distribution is given as

$$f(l) = \frac{5}{3} \frac{l^{-8/3}}{\eta^{-5/3} - \Delta^{-5/3}} \quad (2.65)$$

which binds l to $[\Delta, \eta]$. N_η is an empirical range scaling constant. In this work, eddy events that exceed the LEM domain, i.e., $x_0 + l \geq L_{\text{LEM}}$, are implemented in a cyclical manner. The Kolmogorov scale is approximated as

$$\eta = N_\eta \frac{\Delta}{Re_\Delta^{3/4}}, \quad (2.66)$$

with N_η being an empirical constant that controls the scaling between Δ and η . The eddy frequency and size parameters relate inertial range scaling to the stochastic processes via

turbulent diffusivity of the fluid element. The full derivation of this can be found in Ref. [16]. The eddy rate parameter in (2.62) is independent of l , which is sampled from (2.65), meaning that stirring times are uniform across eddy sizes. While this is not strictly true, it is a justified assumption for LES-LEM using the arguments given by Sankaran [37] as follows:

1. HIT assumption within the sub-grid implies that eddies are closely spaced in wave-number space.
2. Re_Δ and Δ vary spatially over the resolved grid, this implies λ varies spatially over the resolved grid, but uniform within a sub-grid region.

Open flames, which are assumed to be at constant pressure, necessitates individual LEM wafers be allowed to expand to accommodate the density drop due to combustion. Physically, dilatation arises due to propagating pressure waves that caused local flow acceleration via the momentum equation, however, keeping with the LEM assumptions above, this expansion is simply realised as an expansion to the LEM wafer length assuming constant wafer mass. This, Lagrangian (fixed-mass), treatment of individual LEM wafers leads to non-uniform LEM wafer sizes (lengths). The LEM domain then needs to re-gridded to maintain a specified spatial resolution.

2.5.2 Large-scale transport: splicing




Figure 2.4: A ‘splicing’ operation in 1D.

As mentioned previously, large-scale advection is coupled to LEM domains via a Lagrangian transport model called splicing. Strictly speaking, it does not solve (2.55), but is a means by which neighbouring LEM domains can exchange mass in a manner that reflects the LES-resolved transport of fluid between their respective control volumes. Splicing in its most basic, 1D, form is shown in Fig. 2.4. A splicing event consists of a detachment and an attachment operation. Here, LEM fragment ‘f’ is detached from domain ‘1’ and attached to ‘2’, across LES face ‘F’. The fragment size is determined by the ratio of the resolved flux ϕ over a discretised time step Δt and the volume of the LES cell 1, the ‘donor’ cell. Hence,

$$f = S_{\text{LEM},1} \frac{\phi \Delta t}{S_{\text{LES},1}}, \quad (2.67)$$

where $S_{\text{LEM},1}$ is the size of the donor domain, $S_{\text{LES},1}$ is the size of donor cell, and flux ϕ could be volumetric or mass-flux across F. A fully three-dimensional version of the above is required for LES-LEM, previous studies like Sankaran [37] achieved this for structured grids by employing three 1D splicings operations along the three coordinates,

involving the six faces of a hexahedral LES cell. This was extended by Arshad et al. [49] for unstructured meshes, which can be made of polyhedral or tetrahedral cells. Since each LES cell could have several incoming and outgoing LEM fragments, splicing events require an order.

In general, a splicing ‘scheme’ consists of the following choices:

1. The definition of ‘size’ in (2.67) as it pertains to LEM domains and LES cells, e.g., LEM-fragments could be defined by their masses or lengths.
2. The choice of volumetric or mass fluxes to determine fragment sizes.
3. The *order* of attachments and detachments: this could be random or follow some heuristic reasoning.
4. The *location* of attachments and detachments: this could also be random, fixed (designated inflow and outflow), or follow some other rule.

Evidently, bookkeeping routines are needed to relate these fluxes to splicing between the appropriate LEM domains. The following principles are considered when developing such a scheme.

- **Control volume crossing rate:** LEM fragments should not be able to propagate multiple LES control volumes during a time step.
- **Residence time of LEM fragments:** Fragments should not be stagnant within a control volume if it resolves a non-zero flux value. Conversely, it should not “outpace” the resolved fluxes either.
- **Diffusion between LEM domains:** This could be ignored or modelled in an empirical manner. If ignored, the propagation of a flame front *between* LEM domains is driven solely by the resolved dilatational fluxes, i.e., between a burnt and an unburnt LES cell which should trigger a splicing event.
- **Mass conservation:** Attachment and detachments should not create or destroy mass on the LEM domain. As an extension, it should not lead to nonphysical accumulation or depletion of mass over time.

2.5.3 Coupling in LES-LEM

Finally, chemical closure is achieved by coupling the LEM equations (2.54) and LES covering (2.31). Two straightforward strategies present themselves:

1. Directly map an averaged value of mass fractions, temperature or enthalpy from the LEM domain to the LES cell that it inhabits. This is akin to a sub-grid filtering operation and means that the LES species transport (and optionally the enthalpy transport) are fully replaced by (2.54). A majority of LES-LEM studies employ this approach. E.g., a Favre-filtered coupling for scalar ϕ can be realised as

$$\tilde{\phi} = \frac{\sum_{i=1}^{N_{\text{LEM}}} \phi_i \rho_i}{\sum_{i=1}^{N_{\text{LEM}}} \rho_i}, \quad (2.68)$$

where i denotes an LEM-cell of the embedded domain of N_{LEM} . Other options include using algebraic mean or median values.

2. Extract a filtered source term for the species transport (and for enthalpy if sensible enthalpy is being used) from each LEM domain, this approach was taken by Arshad et al. [23] and, in theory, allows for LES-LEM to be used alongside other reaction-rate closures. This could be written for a time rate ω_ϕ as

$$\bar{\omega}_\phi(\mathbf{x}, t) = \frac{\sum_{i=1}^{N_{\text{LEM}}} \dot{\omega}_{\phi_i} \rho_i}{\sum_{i=1}^{N_{\text{LEM}}} \rho_i}, \quad (2.69)$$

or, as before, using mean or median values.

2.5.4 Features of LES-LEM

The LES-LEM formulation addresses some important assumptions and limitations described in Sec. 2.3.

- Embedded 1D domains allows for fully resolving the flame front (in the flame normal direction), i.e., at resolutions much finer than Δ . This circumvents the aforementioned numerical difficulties (see Sec. 2.3.5), especially for premixed flames.
- Sub-grid diffusion is simulated using molecular diffusion, albeit in 1D, which were ignored in Sec. 2.2 due to the lack of reliable models.
- Sub-grid convective fluxes are simulated using stochastic processes derived from established turbulence scaling laws, which is an improvement over the eddy-diffusivity assumption for sub-grid species diffusive flux. Furthermore, it does not require assumptions for the turbulent Prandtl or Schmidt numbers (see 2.2).
- SGS fluxes *between* LES cells require the addition of a ‘cross-term’ – see Sec. 10.6.2 of Menon et al. [50]. However, Arshad et al. [23] found this to have little to no effect for a stabilised premixed flame.
- The turbulent flame speed is simulated rather than modelled.
- The LEM formulation is unchanged for premixed or non-premixed flames nor for higher or lower Re , i.e., it is mode- and regime-independent.

The multi-scale nature of LES-LEM closure allows it to be much more tractable than full 3D DNS, despite resolving all scales, at least in one spatial dimension. However, when compared to standard LES closures in Sec. 2.3, it is considerably more expensive, and factors in when assessing cost-to-benefit of LES-LEM. The costs originate from the following:

1. Each LEM wafer is, in essences, a well-stirred reactor and thus requires numerical integration of the stiff ODEs, i.e., (2.9). Since each embedded LEM domain is highly resolved, typically tens of wafers per domain, the costs associated with chemistry advancement scales with this factor.

2. Each mapping rearrangement entails an interruption of the small-scale diffusion equation (see Sec. 3.3). Start-restart cycles for numerical integration codes is a computational overhead, this means high Re flows, i.e., more frequent triplet maps, can quickly become cumbersome to simulate with LES-LEM.
3. Lagrangian splicing overheads can be also significant, especially when LEM fragments have to be sent across processor domains for parallel simulations. Splicing also introduces artefacts in the form of sharp discontinuities on the LEM domain, frequent splicing implies that it could become the dominant mode of small-scale stirring over triplet maps. This effect, however, has not been quantified in previous studies.

2.6 Super-grid LES-LEM

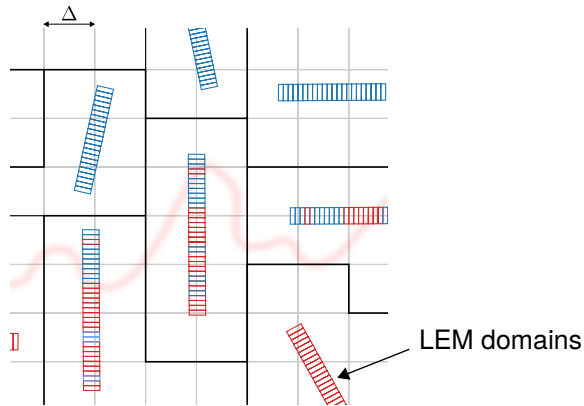


Figure 2.5: Schematic representation of the novel ‘super-grid LES-LEM’ (SG-LEM) closure, grey lines are LES cells, thickened lines show LES cell agglomerations or *clusters*.

The points mentioned above motivate a new version of LES-LEM, which can retain the advantages of an LEM-driven closure but at a reduced computational cost. This work explores such an approach using mesh agglomeration. The rationale is that if clusters of LES cells can be embedded with LEM domains, instead of each LES cell, it would lead to a reduction in the number of LEM domains while retaining representative lengths and resolution. A schematic is shown in Fig. 2.5. The coarse-grained mesh is termed a ‘super-grid,’ while the individual agglomerations are called ‘clusters.’ The use of agglomeration requires some important modifications to standard LES-LEM. These are:

1. LEM domains must be representative of clusters and not of the LES cell size Δ (or filter-width for implicit LES).
2. The use of filtering or averaging in Sec. 2.5.3 would lead to a uniform distribution of scalars for each cluster. To retain the LES-resolved information, a more appropriate coupling routine is required.
3. Top-down methods for cell agglomeration might lead to clusters of non-standard

shapes, and neighbouring clusters having very different volumes. This needs to be accounted for by the splicing scheme.

The above points are addressed in this and upcoming sections. The first of these is achieved by setting the domain length to the cluster size l_t , and replacing Δ in (2.62) and (2.65) with the same. This is written as

$$\Delta := l_t = \left(\sum_{j \in \mathbf{V}} V_j \right)^{\frac{1}{3}}, \quad (2.70)$$

where \mathbf{V} is the set of LES cell indices that comprise the clusters, and the sum is over LES cell volumes V_j . l_t can be considered an integral length scale, LEM resolves scales from here down to η . A cluster averaged value for turbulent Reynolds number Re_t is needed for small-scale evolution, this can be computed using an averaged k_{sgs} as

$$k_c = \frac{\sum_{j \in \mathbf{V}} k_{\text{sgs},j} V_j}{\sum_{j \in \mathbf{V}} V_j}, \quad (2.71)$$

and

$$Re_c = \frac{\sqrt{\frac{2}{3} k_c} l_t}{\nu}, \quad (2.72)$$

where the suffix ‘c’ represents the cluster. Turbulent cluster Reynolds number, Re_c , is used instead of Re_t in (2.62). Alternatively, the root-mean-squared (RMS) value of velocity (within the cluster) can be used as

$$Re_c = \frac{u_{\text{rms}} l_t}{\nu}. \quad (2.73)$$

The kinematic viscosity, ν , here is set to the average value on the corresponding LEM domain.

The larger LEM domains motivates modifications to the LEM constants N_η and C_λ . The effect of transformation (2.70) is illustrated in Fig. 2.6. We assume $l_t = 5\Delta$, $\Delta = 1$ mm, and $l_t = 5$ mm. Fig. 2.6 (a) shows the eddy size distribution (2.65) a function of Re_t – using the black curve (for $l_{\text{LEM}} = \Delta = 1$ mm, and $N_\eta = 4.0$) as a reference, we see that the increased LEM domain size ($l_{\text{LEM}} = l_t = 5$ mm) elicits a much different eddy size sampling, as shown by the dashed-black curve. In order to recover the original distribution, a modified $N_\eta = 0.8$ is required, as shown by the black squares. These curves were obtained for $Re_t = 100$, the behaviour for $Re_t = 200$ is shown in blue where it is evident that the new constant is sufficient to recover the size distribution here as well. Fig. 2.6 (b) shows the eddy sampling frequency as a function of Re_t in a similar manner. This time, using the blue curve (for $l_{\text{LEM}} = \Delta = 1$ mm, and $C_\lambda = 15.0$) as a reference, the increased domain size gives a much lower sampling rate as shown by the black curve. Applying $N_\eta = 0.8$ shows a slight increase (dashed-black curve), however, this must be combined with a modified $C_\lambda = 8.5$ to recover the original sampling rate behaviour. Such a comparative procedure can be used to modify literature values of N_η and C_λ from previous LES-LEM studies to accommodate the larger domain sizes used in SG-LEM.

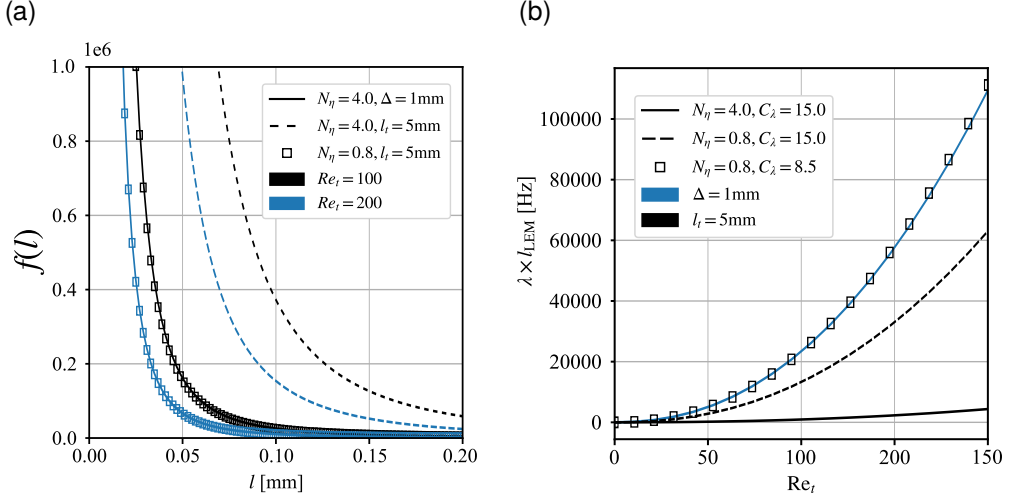


Figure 2.6: The effect of LEM constants C_χ and N_η on (a) eddy size distribution, and on (b) eddy sampling rate (b), using turbulent Reynolds number Re_Δ , and LEM domain length $l_{LEM} = \Delta$ or l_t , as input parameters.

2.6.1 SG-LEM coupling strategy

A new coupling strategy is needed as mentioned above, this work adopts the strategy developed by Lackmann et. al. [51] known as RILEM (Representative Interactive Linear Eddy Model). RILEM is a mapping closure similar to RIF (Representative Interactive flamelets) [12], that uses statistical moments using presumed PDFs to map LEM derived scalars to the flow field, and specifically formulated for ICEs. The computational domain for RILEM typically represents an ICE cylinder, or sector, with piston movement modelled by mesh deformations. Chemical closure for the domain is supported by a single LEM domain, this was improved upon by Doubiani et al. [21] by using multiple LEM domains (MRILEM) to increase statistical fidelity of LEM-derived scalars. While the configuration bears similarities with SG-LEM, an important difference is noted as follows: all LEM domains in MRILEM represent the same (large) domain, they differ from one another in that they are independent instances of the stochastic processes that govern small-scale stirring. This implies that large-scale advection, i.e., (2.55), cannot be simulated by splicing but requires an additional model for the tumbling fluid motion seen during an engine cycle. RILEM is based on RANS and tailored to non-premixed, auto-ignited flames. Despite these, the RILEM strategy is used here with minor modifications for premixed flames. It is a mapping closure using presumed PDFs (see Sec. 2.3.4 for comparison), and can be summarised, for *each* cluster, as:

1. Conditionally average (bin) thermo-chemical scalars generated by LEM advancement using primitive variables, here: mixture fraction and reaction progress variable.
2. Generate unique (here: presumed) PDFs for said primitives for each LES cell within the cluster.

3. Perform a mapping closure using presumed PDFs and conditionally averaged chemistry data for each LES cell. Mathematically, this is the first moment of the conditionally weighted PDFs.

The details of the coupling strategy are as follows: LEM simulates combustion processes in physical space, a thermo-chemical scalar, ϕ , that results from LEM simulation is conditionally binned on primitive variables, i.e., mixture fraction, Z , and reaction progress variable, c , which can be computed from the thermo-chemical state of every LEM wafer. Conditionally binned data is represented as $\langle \phi | Z, c \rangle$.

Mixture fraction, as defined by Bilger [52], is computed as follows: elemental mass fraction of element j is computed as

$$Y_j = \sum_i^{N_s} \frac{a_{j,i} M_j}{M_i} Y_i, \quad (2.74)$$

where $a_{j,i}$ is the number of j atoms in species i , M_i the *molecular* weight of species i , and M_j is the *atomic* weight of element j – for elements C, H and O. A coupling function can then be defined, such as

$$\beta = \frac{Y_C}{mM_C} + \frac{Y_H}{nM_H} - 2 \frac{Y_O}{\nu'_{O_2} M_O}, \quad (2.75)$$

using the stoichiometric coefficients for the global reaction of the general hydrocarbon $C_m H_n$ in air. Similarly, elemental compositions of the fuel- and oxidiser-stream is used to compute β_1 and β_2 , respectively. Using these, Z for an LEM wafer is computed as

$$Z = \frac{\beta - \beta_2}{\beta_1 - \beta_2}. \quad (2.76)$$

Secondly, the reaction progress variable c is defined as

$$c = \frac{\psi - \psi_u}{\psi_b - \psi_u}, \quad (2.77)$$

where suffixes ‘b’ and ‘u’ denote the fully burnt and unburnt states, and ψ is a choice of scalar, usually temperature or combination of mass fractions. For premixed combustion, ψ_u is known from the initial temperature or equivalence ratio of the investigated problem while, for non-premixed combustion, it depends on the local Z value as

$$\psi_u = Z\psi_1 + (1 - Z)\psi_2, \quad (2.78)$$

where suffixes 1 and 2 are the fuel and oxidiser stream values. From this, ψ_b is set to the equilibrium value computed at constant pressure and enthalpy.

Once Z and c are computed for each LEM, and thermo-chemical scalars are conditionally-binned, the LES filtered scalar values are determined using the first moment of the Favre joint-PDF of Z and c as

$$\tilde{\phi} = \int_0^1 \int_0^1 \langle \phi | Z, c \rangle \tilde{P}(Z, c) dZ dc. \quad (2.79)$$

Assuming statistical independence of Z and c gives

$$\tilde{\phi} = \int_0^1 \int_0^1 \langle \phi | Z, c \rangle \tilde{P}(Z) \tilde{P}(c) dZ dc. \quad (2.80)$$

RILEM is a presumed PDF method, in the context of LES this is interpreted as a presumed sub-grid distributions for Z and/or c . PDF shapes typically used in combustion modelling, such as the β -PDF, are uniquely defined by their first and second moments. Here, Favre-mean mixture fraction \tilde{Z} ; Favre-mean progress variable \tilde{c} (first moments); and respective variances (second moments) $\widetilde{Z'^2}$ and $\widetilde{c'^2}$. This requires LES advancement of the transport equations for \tilde{Z} and \tilde{c} . These are given by

$$\frac{\partial(\bar{\rho}\tilde{Z})}{\partial t} + \frac{\partial}{\partial x_j} (\bar{\rho}\tilde{u}_j\tilde{Z}) = \frac{\partial}{\partial x_j} \left(-\frac{\mu_t}{Sc_t} \frac{\partial\tilde{Z}}{\partial x_j} \right) \quad (2.81)$$

and

$$\frac{\partial(\bar{\rho}\tilde{c})}{\partial t} + \frac{\partial}{\partial x_j} (\bar{\rho}\tilde{u}_j\tilde{c}) = \frac{\partial}{\partial x_j} \left(-\frac{\mu_t}{Sc_t} \frac{\partial\tilde{c}}{\partial x_j} \right) + \bar{\rho}\tilde{c}, \quad (2.82)$$

where \tilde{c} is the equivalent of a filtered chemical source term, and other terms retain their meanings from in Sec. 2.2. Spray models are straightforward to include using an evaporative source term (not shown) for (2.81), this would also require a sink term for latent heat of evaporation in the filtered energy equation. The first moment also yields the source term for (2.82):

$$\tilde{c} = \int_0^1 \int_0^1 \langle \dot{c} | Z, c \rangle \tilde{P}(Z) \tilde{P}(c) dZ dc, \quad (2.83)$$

where \dot{c} at the LEM level is given by the time derivative of (2.77):

$$\dot{c} = \frac{1}{\psi_b - \psi_u} \frac{d\psi}{dt}. \quad (2.84)$$

For premixed flames, (2.80) is modified to

$$\tilde{\phi} = \int_0^1 \langle \phi | Z, c \rangle \tilde{P}(c) dc, \quad (2.85)$$

where $\langle \phi | c \rangle$ represents scalars that are conditionally averaged (solely) in c -space. Finally, LES-resolved temperatures are iterated from enthalpy \tilde{h} , the transport equation for which is given an eddy-diffusivity treatment. The presumed shapes of the Favre-PDFs used in this work are:

1. For premixed cases: β -PDFs for $\tilde{P}(c)$.
2. For Mixed-mode or partially-premixed cases: the ‘extended top-hat’ PDF of Floyd et al. [53] for $\tilde{P}(Z)$, and Dirac- δ function for $\tilde{P}(c)$.

2.6.2 Secondary Output

The coupling procedure described above is the primary output of the model, however, a secondary output can be reported, for diagnosis, at super-grid resolution using (2.68), i.e., all cells in a cluster are simply set to the corresponding Favre-mean value of scalars over the embedded LEM line. This is denoted as $\tilde{\phi}_{\text{LEM}}$.

Chapter 3 describes the numerical methods behind SG-LEM, including LEM advancement, treatment of presumed PDFs, and the modified splicing scheme.

3 Methodology

3.1 Introduction




Figure 3.1: A 1D LEM domain comprising discrete elements known as wafers. Symbols $\Delta x_i, \rho_i, T_i$ and $Y_{\alpha,i}$ are the length, density, temperature and composition of wafer index i .

This chapter details the methodology behind LEM simulation, super-grid generation and the SG-LEM coupling procedure. The programming language C++ was used extensively in this work, and its object-oriented features proved useful for constructing and simulating LEM domains. The coupled LES-SG-LEM solver was implemented using the `OpenFOAM` (v. 9) library [54].

3.2 LEM Domain Implementation

Each LEM domain, as shown in Fig. 3.1, is represented in memory as a double-linked list. Each ‘node’ in the list is a data-structure that represents LEM wafers. Wafers (nodes) contain all data necessary for a full description of chemical-states and sizes. A distinct feature of list containers is that all nodes contain pointers to the previous and next nodes in the sequence, as shown in Fig. 3.2.

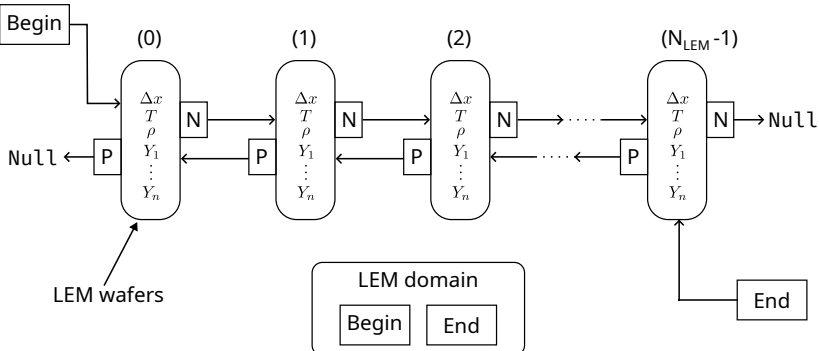


Figure 3.2: LEM domain represented in memory as a double-linked list. ‘P’ (meaning previous), ‘N’ (meaning next), ‘Begin’ and ‘End’ are pointers. N_{LEM} is the number of wafers in the domain, i.e., nodes in the list container.

Other possibilities for representing LEM domains exist, such as an array of LEM wafers or multiple arrays for each thermo-chemical scalar. The list container method, using

non-contiguous memory, offers some advantages and disadvantages:

- Insert and erase operations are inexpensive, with overhead time scaling with $\mathcal{O}(1)$, as they only require the reassignment of previous and next pointers on the appropriate nodes in the container. Therefore, the insertion and erasure of LEM wafers are inexpensive. This, however, does not consider the memory allocation overhead for the actual wafer data structure, such as the creation of LEM wafers from the inlets. When contiguous memory containers (arrays or vectors) are used, an insertion or erasure at any index would involve shifting subsequent nodes to maintain contiguity. The overhead scales with $\mathcal{O}(n)$, where n is the number of subsequent nodes from the point of insertion or erasure in the container. The scaling is similar for the other two representations mentioned above, i.e., an array of LEM wafers (objects) or an array for each scalar. The former requires shifting large data structures, while the latter would require shifting all scalar arrays. This cost would also scale with the number of chemical species.
- The point above regarding list containers extends itself to triplet mapping and splicing, i.e., these operations only require pointer reassignments. It is assumed that triplet maps align with the grid spacing, i.e., there is no compression of wafers necessary. An exception would be splicing across processor boundaries when using distributed computation, where LEM fragments are moved between neighbouring processor domains. This is explained further in the upcoming Sec. 3.5.
- A disadvantage of using non-contiguous list containers is the time of traversal, which scales with $\mathcal{O}(n)$, where n is the number of nodes to traverse (in either direction) from any given position, as opposed to $\mathcal{O}(1)$ for arrays. They also lack random-access capability, e.g., accessing the beginning of a triplet map. There is also the additional overhead that stems from the machine architecture, i.e., reduced cache locality. While list traversals are quite frequent during LEM simulation, in particular for diffusion advancement, the slow-down incurred from the use of lists is more than compensated for by the speed at which rearrangements and splicing can be performed. Furthermore, chemistry advancement is the most expensive sub-step of LEM simulation,




Figure 3.3: Mapping rearrangements over the discretised domain: (a) shows the initial field and (b), after a triplet map. Blue lines indicate scalar levels, and numbers indicate wafer indices.

Triplet mapping using the list container is done through pointer reassignments. A minimum of six wafers are needed for a triplet map. An example of a triplet map using nine wafers is shown in Fig. 3.3.

3.3 LEM Advancement and Operator Splitting

LEM domains are time-advanced concurrently to LES. Rewriting (2.56) as

$$\rho \frac{\partial \phi}{\partial t} = F_{\phi, \text{stirr}} + D_{\phi} + \Omega_{\phi}, \quad (3.1)$$

where $F_{\phi, \text{stirr}}$, D_{ϕ} , and Ω_{ϕ} represent the stirring (i.e. triplet map), diffusion and chemical reaction operators. Due to the various sub-steps involved in this operation, i.e., reaction, diffusion and triplet mapping, some sort of operator splitting is required. In particular, the time-accurate implementation of triplet maps introduces some complexity as each entails an interruption to the combined reaction-diffusion advancement. Also, the chemical reaction operator requires its own, separate, treatment due to chemical stiffness. Two splitting methods are used in this work. The first of which is a time-accurate eddy implementation, and described as follows.

3.3.1 Sampled-sequence operator splitting

1. Let the LES time-step be Δt_{LES} .
2. Begin time-marching of LEM domains, let the elapsed time for a given domain start at $t = 0$.
3. Sample the time interval to the first eddy event as a Poisson process. This is written as

$$\Delta t_e = -\log(X)\Delta t_{\text{stirr}}, \quad (3.2)$$

where X is a random integer between $(0, 1]$ and Δt_{stirr} is the stirring interval (cf. (2.64)).

4. Let the diffusion-reaction time step be Δt_{diff} , perform the following operations:

```

while  $t < \Delta t_{\text{LES}}$  do
  if  $t + \Delta t_e < \Delta t_{\text{LES}}$  then
    |  $\Delta t_{\text{diff}} = \Delta t_e$ 
  else
    |  $\Delta t_{\text{diff}} = \Delta t_{\text{LES}} - t$ 
  end
   $\phi := \phi + \frac{\Delta t_{\text{diff}}}{2} D_\phi$ 
   $\phi := \phi + \int_{\Delta t_{\text{diff}}} \Omega_\phi dt$ 
   $\phi := \phi + \frac{\Delta t_{\text{diff}}}{2} D_\phi$ 
  Compute conditional averages  $\langle \phi | Z, c \rangle$  or  $\langle \phi | c \rangle$ 
   $t = t + \Delta t_e$ 

  if  $\Delta t_e < \Delta t_{\text{LES}}$  then
    |  $\phi := \phi + F_{\phi, \text{stirr}}$ 
    | Generate random  $X \in (0, 1]$ 
    |  $\Delta t_e = -\log(X) \Delta t_{\text{stirr}}$ 
  end
end

```

As mentioned before, a Lagrangian (fixed-mass) discretisation is applied to the LEM domain. Diffusion advancement is carried out using a finite difference approach. Temporal discretisation uses a second-order fully implicit scheme, while second-order central differencing is used to compute gradients. The Thomas algorithm for tri-diagonal matrices is used to solve the resulting system of linear equations. Chemistry advancement, i.e. $\int_{\Delta t_{\text{diff}}} \Omega_\phi dt$, is carried out using the numerical integrator `CVODE` from the `SUNDIALS` [55] package. It employs adaptive fractional time-marching to integrate the chemical Jacobian, and uses implicit Backward Differentiation Formula (BDF). Reaction rates, transport coefficients and thermodynamic data are computed using the `Cantera` [56] software package. An implicit scheme is necessary for detailed chemical mechanisms consisting of fast and slow reactions that lead to a stiffness of differential equations. The alternative, i.e., explicit schemes, would require extremely small time-steps to ensure numerical stability which would make computation very expensive.

A sophisticated package like `CVODE` brings memory and processing overheads. In particular, there is a small penalty for start-restart cycles, i.e., for computing a fresh integral in the interval $[t, t + \Delta t_{\text{diff}}]$ using high order BDF schemes. The operator splitting described above would result in frequent interruptions (or very small Δt_{diff}) for high Re_t , leading to a significant slowdown of the overall simulation. To counteract this, an alternative procedure was developed to minimise the restart overhead. This is termed ‘blocked-sequence’ operator splitting and differs from the previous method in that all eddies are implemented at the start of LEM domain advancement. The procedure is described as

follows.

3.3.2 Blocked-sequence operator splitting

1. Sample the time interval to the first eddy event Δt_e
2. Perform the following operations:

```

while  $t < \Delta t_{\text{LES}}$  do
   $t = t + \Delta t_e$ 
  if  $\Delta t_e < \Delta t_{\text{LES}}$  then
     $\phi := \phi + F_{\phi, \text{stirr}}$ 
    Generate random  $X \in (0, 1]$ 
     $\Delta t_e = -\log(X)\Delta t_{\text{stirr}}$ 
  end
end

 $\phi := \phi + \frac{\Delta t_{\text{LES}}}{2} D_\phi$ 
 $\phi := \phi + \int_{\Delta t_{\text{LES}}} \Omega_\phi dt$ 
 $\phi := \phi + \frac{\Delta t_{\text{LES}}}{2} D_\phi$ 
Compute conditional averages  $\langle \phi | Z, c \rangle$  or  $\langle \phi | c \rangle$ 

```

The above method introduces a splitting error (or bias) which reduces with smaller Δt_{LES} , it must be determined if this produces a large error in the conditionally binned data; standalone LEM simulations can be used for this purpose where a first guess for domain lengths (in a full SG-LEM run) and expected turbulence levels are input parameters.

3.3.3 Volumetric expansion

Combustion results in local density variations and local flow acceleration. Physically, this is due to pressure waves that propagate at the speed of sound. Since LEM simulations assume constant pressure, this dilatational advection is explicitly modelled as an expansion of LEM wafers to account for changes in the LEM density field. Conserving wafer mass $m_i = \rho_i \Delta x_i$, this is expressed as

$$\Delta x_i^{\text{new}} = \Delta x_i \left(\frac{\rho_i}{\rho_{i, \text{new}}} \right), \quad (3.3)$$

where $\rho_{i, \text{new}}$ is the density of wafer i after a reaction-diffusion advancement and Δx_i^{new} is the updated wafer length. This procedure results in LEM domain lengths that are neither constant nor tied to the control volume they embed (LES cells in the case of LES-LEM, and clusters in the case of SG-LEM). For confined flames, e.g., in an ICE, the expanded domain is then compressed to a length that is representative of the current state of the CFD domain, adjusted for piston movement, using isentropic compression. Alternatively, constant volume assumption can be used to compute the LEM source terms, following

which, pressure in each wafer can be relaxed to the target pressure. This increases pressure uniformly for all wafers, accounting for both piston movement and the pressure rise due to combustion. Such a procedure is not used in this work, as investigations were conducted exclusively for open flames. For more details on LEM closure for ICEs, as it pertains to pressure and volumetric coupling of the CFD and LEM domains, the reader is directed to Doubiani et al. [21], and Lackmann et al. [20] respectively.

3.4 Super-Grid Generation




Figure 3.4: An LES cell cluster showing agglomerated faces including internal faces (in grey), and a boundary face that is split into two inlet faces for fuel (red) and oxidiser (blue).

The `MGRIDGEN` library [57] was used to perform mesh agglomeration. `MGRIDGEN` is capable of handling unstructured meshes, which are used in `OpenFOAM`. This enables easy discretisation and coarse-graining of complex geometries. Cluster sizes (minimum and maximum) as well as type (in this case, ‘globular’) are given as input parameters. The algorithm tries to minimise the summation of (volume-weighted) aspect ratios of all clusters in the domain. Bookkeeping is necessary to record the relationships between the super-grid and LES mesh in the form of lists, e.g., LES cells that constitute a cluster, LES faces that form each cluster face, and clusters indices on either side of a cluster face, i.e., ‘owner-neighbour’ lists.

In addition to this, cell-clusters near the LES domain boundary have to be identified, and agglomerated faces need to be split in a way that reflects the CFD domain boundary types or ‘patches’¹. This step is necessary as the agglomeration routine is independent of the boundary types – inlets, outlets, cyclic boundaries, walls, etc. An example is shown in Fig. 3.4, where an agglomerated face is split to accommodate fuel and oxidiser patches – part of the larger CFD domain which is not shown. Also shown are inlet LEM domains that act as a reservoirs for the splicing procedure discussed in the upcoming section.

¹In `OpenFOAM` terminology, a ‘patch’ is defined as one or more enclosed areas of the boundary surface which may or may not be proximal.

Cluster faces at processor boundaries require a similar treatment, however, they bring an added complexity: `MGRIDGEN` operates independently in each processor domain and thus, cluster faces may not be contiguous across processor boundaries. This must also be addressed for correct splicing. Hence, cluster faces are split in such a way that each resulting face shall interface with exactly two clusters, one on either side of the boundary. An example is shown in Fig. 3.5, albeit in two dimensions. `Message Passing Interface` (MPI) for distributed memory architectures is used to communicate face-owner lists between neighbouring processor domains so that this operation may be performed.

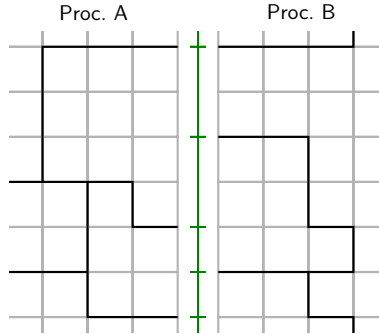


Figure 3.5: Cluster faces at processor boundaries, green line segments show the face splits.

3.5 Splicing for SG-LEM

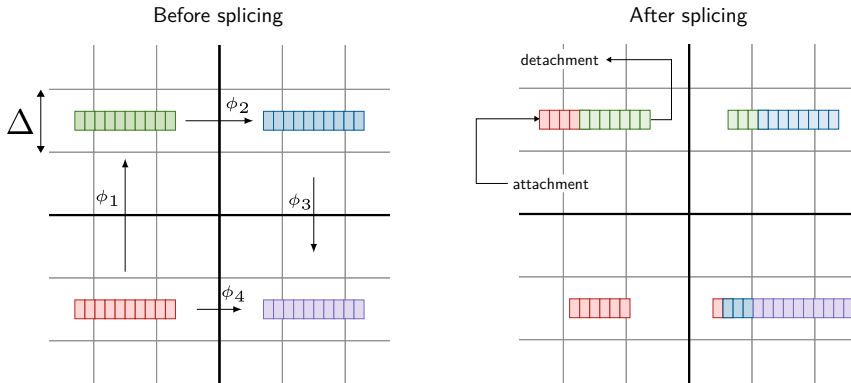


Figure 3.6: Schematic showing flux ordered splicing. Colours indicate individual LEM domains before splicing, arrows of different length indicate unequal fluxes between clusters.

A novel splicing method was developed for this work, it bears many similarities with previous studies [58, 37, 49, 22]; however, modifications were required to work with the super-grid. Salient features are discussed below (cf. Sec. 2.5.2).

- LEM fragments are determined by their *length*, cluster sizes are determined by their

aggregate *volumes*. This differs from the *mass*-based splicing schemes of previous LES-LEM studies, and allows for splicing between clusters of starkly differing volumes.

- Volumetric fluxes are used to determine splicing fractions, units $[\text{m}^3\text{s}^{-1}]$.
- Attachments are ordered by the descending order of associated fluxes, conversely detachments are ordered by the ascending order of associated fluxes. For a discrete time-step, this should lead to larger fluxes resulting in larger displacements (in three-dimensional space) of LEM fragments.
- Each LEM domain has a designated ‘in-splice’ (attachment) and ‘out-splice’ (detachment) end. This preserves the concept of control volume crossing-rate that was discussed previously.
- LEM domain are given a nominal cross-section given by

$$A = \frac{\left(\sum_{j \in \mathbf{V}} V_j\right)}{l_t}, \quad (3.4)$$

where l_t is the LEM domain length given by (2.70).

The procedure for SG-LEM splicing is as follows:

For a given cell cluster k , let \mathbf{O}_k be the set of super-grid faces belonging to the cluster which have a net outward flux and \mathbf{I}_k be the set of faces with a net inward flux. Note that *flux* here refers to volumetric flux with \cdot . Perform the following steps:

1. Cycle through LEM domains $k \in [1, K]$, where K is the total number of cell clusters (and hence LEM domains) in a given processor domain.
 - (a) Sort \mathbf{O}_k in ascending order of flux $\phi(N)$ for $N \in \mathbf{O}_k$ where N refers to the face numbers in the (per-processor) cell-cluster numbering scheme; $\phi(N)$ is computed as the algebraic sum of CFD-prescribed volume fluxes for the CFD faces that make up cluster face N .
 - (b) For each N in the sorted list, splice *out* an LEM fragment of length

$$L_{N,k}^{frag} = L_k^{\text{LEM}} \frac{\phi(N) \Delta t_{\text{LES}}}{V_k} \quad (3.5)$$

from the designated *out-splice* end, where L_k^{LEM} is the length of the LEM domains and V_k is the volume of the corresponding cell cluster.

- (c) For each N , send the fragment to location N of a global ‘fragment list’ which temporarily holds fragments to be spliced for all faces in a processor domain. Also send the LEM cross-section A_k (c.f. (3.4)) to location N of a global ‘cross-section’ list \mathbf{A} .
2. Cycle through LEM domains $k \in [1, K]$
 - (a) Sort \mathbf{I}_k in descending order of flux $\phi(N)$ for $N \in \mathbf{I}_k$.

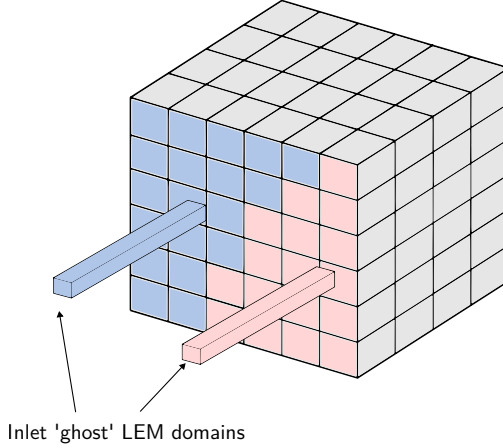


Figure 3.7: A cell-cluster near the inlet boundary with attached ‘ghost’ LEM domains.

(b) For each N in the sorted list:

- i. The LEM fragment stored in the fragment list at location N is modified to account for differences in cross-section and to maintain a consistent 1D representation. The length of each cell i of the fragment is modified as

$$\Delta x_i^{new} = \Delta x_i \frac{A[N]}{A_k}, \quad (3.6)$$

note that this operation does not affect mass consistency.

- ii. Splice *in* the modified fragment to the designated *in-splice* end of LEM domain k .

Note that step 2(b)i above models compression and super-grid approach presented in this study.

3.5.1 Inflow and outflow boundary conditions

The general scheme outlined above is modified slightly for inflow boundary faces, depending upon whether the case is premixed or non-premixed. For premixed combustion a ‘ghost’ LEM domain is initialised with the necessary equivalence ratio for each inflow patch (per-processor). For the latter, fuel and oxidiser inlet patches are each assigned a ghost LEM domain, as shown in Fig. 3.7. These are initialised, in turn, with fuel and oxidiser stream conditions. From these ghost domains, LEM fragments are spliced out (step 1b) proportional to fuel and oxidiser fluxes into relevant locations of the fragment list (step 1c). Ghost domains are constantly replenished such that their lengths and resolution remain roughly constant, i.e., they remain unaffected by splicing.

For outflow boundary faces, the fragment list is simply cleared at the required locations. Special care must be taken for splicing across processor boundaries. Here, step 1c is

modified where, in addition to sending LEM fragments to the fragment list at processor out-flux locations, LEM fragments from neighbouring processors populate the list at processor in-flux locations, a similar approach is taken for processor cyclic boundaries. The parcel exchanges are handled using MPI. The splicing scheme described above allows for the splicing of fractional LEM wafers to ensure mass conservation during the process. Fig. 3.6 shows the re-configuration of LEM domains after a flux ordered splicing operation. LEM fragments corresponding to higher fluxes are displaced by larger amounts for a given time step. The reader is directed to Ref. [59] for further details on splicing, particularly on the analogy between Lagrangian splicing and Eulerian transport algorithms.

3.5.2 Fixed wall boundary conditions




Figure 3.8: Wall treatment, dotted lines indicate splicings.

Modelling of near wall effects can be necessary for accurate flame structure and radical predictions. LEM, by its very formulation, can only handle gaseous reacting mixtures. This necessitates transfer functions which convey heat and mass information from sources such sprays or walls to the LEM level, if these features are present in the CFD domain [20]. Consider a near-wall scenario as in Fig. 3.8, a straightforward wall treatment is obtained by identifying clusters that are adjacent to the walls and attaching one end of the corresponding LEM domains to the wall patch, i.e., we effectively impose wall temperature as a Dirichlet boundary condition to one end of the LEM domain for temperature diffusion. Splicing (both *in*- and *out*-splicing) are now performed exclusively at the remaining free end.

3.6 Interface between SG-LEM and LES

The solver-level implementation for the coupling procedure in Sec. 2.5.3 is illustrated in Fig. 3.9. The interface is shown between three levels, i.e., LES, super-grid and LEM domains, indicated by the colours grey, blue and green, respectively. Note that figure shows both operator-splitting options.

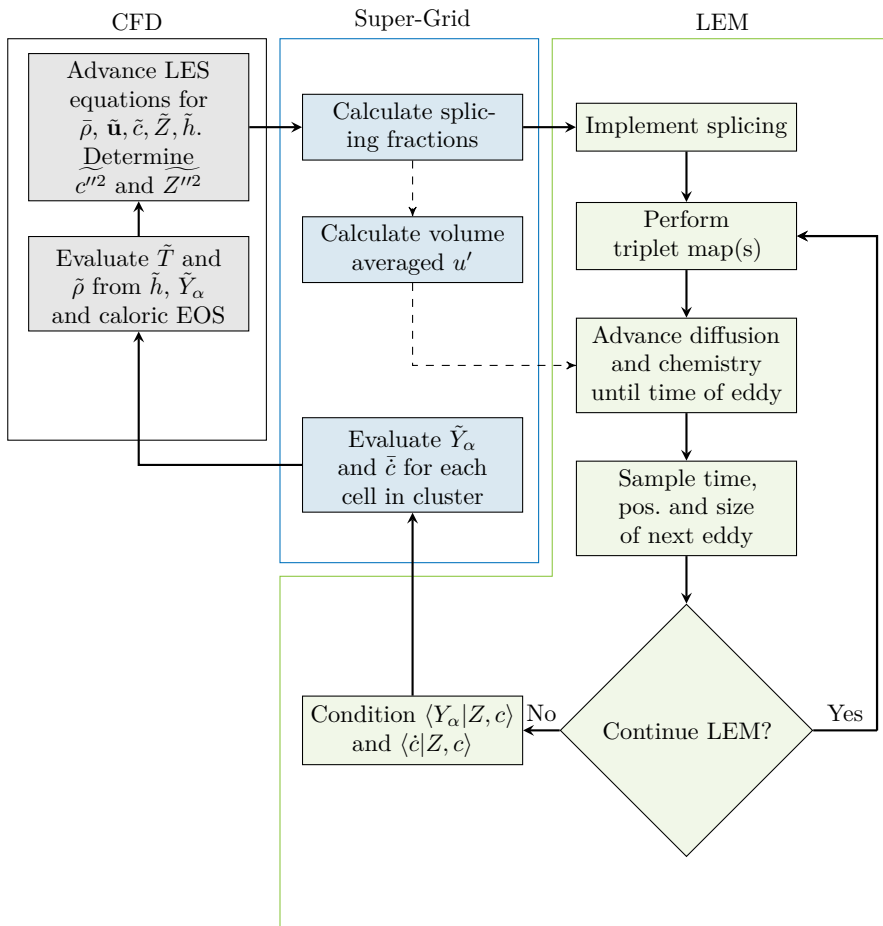


Figure 3.9: Interface between LES, super-grid and LEM.

3.7 Conditional Averages

The presented mapping closure, i.e., (2.80), requires that LEM fields be conditionally averaged into discrete, equally spaced, bins for Z and c . This conditioning procedure for a given LEM domain (or cluster) is described as follows.

1. Conditionally averaged scalar ϕ , i.e., $\langle \psi | Z, c \rangle$, is stored in the solution table $S[j, k]$, where j and k are indices of the discrete Z - and c -bins.
2. The total number of bins are J and K , and bin-centred values are $Z[j]$ and $c[k]$.
3. For the bin j, k ,

$$S[j, k] = \frac{\sum_{i \in \mathcal{B}} \phi_i}{|\mathcal{B}|}, \quad (3.7)$$

where \mathcal{B} is the set of LEM wafers that satisfies the condition: $Z[j] \leq Z < Z[j + 1]$ and $c[k] \leq c < c[k + 1]$, and $|\mathcal{B}|$ is the number of elements in \mathcal{B} .

It is unlikely that a single execution of the above will update every bin in the complete solution table. This is due to the small number of wafers per LEM domain, and the fact that it can only see the current LEM state. Hence, each value in the solution table $S[i, j]$ is made to persist until a suitable (combined value of) Z and c state are encountered on the domain during a future conditioning event. Persistence allows for the continual update of the solution table that reflects local conditions of reaction, large-scale transport, and turbulence. For a statistically stationary flame, each solution table could, in theory, attain a stable set of values. For premixed flames, Z is ignored and the solution table has one fewer dimension, i.e., $S[k]$.

Solution tables can be large, and an SG-LEM simulation might require several thousand solution tables (depending on the level of coarse-graining) in memory at any given time. Hence, they are created using the `multi_array` data structure from the `boost` library for optimal memory management.

3.7.1 Initialisation of Solution Tables

The evaluation of the PDF integrals using (2.80) requires a completely filled solution table for numerical stability. Hence, it must be initialised with a reasonable guess for mass fractions, temperature, and production rates (to compute \bar{c}). In this work, two methods are used for the initial guess prior to running the SG-LEM solver: a zero-dimensional reactor solution, and a premixed laminar flamelet solution. For the former approach, the `constantPressureReactor` class from `cantera` was used to create standalone zero-dimensional reactors. These were initialised with the appropriate fuel-air mixture at unburnt conditions. H radicals were periodically injected in small quantities for mixtures that do not auto-ignite. During runtime, scalars were conditioned between reactor time-steps. Once equilibrium conditions are reached, the solution table is almost fully populated in c -space, linear interpolation can be used for unfilled bins.

The latter approach, i.e., premixed flamelets, were simulated using Lagrangian 1D reaction-diffusion solvers, similar to those used for LEM advancement. The 1D domain is initially

filled with the appropriate fuel-air mixture. A portion of the domain is then brought to equilibrium conditions, creating a flame-front and establishing the initial conditions for a freely propagating flame. After every time step, a domain fragment of length

$$l_{\text{out}} = l_1 - l_0 \quad (3.8)$$

is removed from the burnt end, where l_0 is an initial domain length and l_1 is current length after advancement. l_{out} then represents dilation. This mass is replaced by introducing (splicing) a domain fragment (of fresh charge) of length

$$l_{\text{in}} = \frac{\rho_b}{\rho_u} l_{\text{out}} \quad (3.9)$$

which then holds the flame-front roughly stationary in the 1D domain. Hence, a statistically stationary state can be reached which can be used to initialise the solution tables. If solution tables contain Z as an independent variable, the above procedure is repeated for each Z -bin by initialising the fuel-air mixture using the central bin value. This stabilising method is also used to run standalone LEM simulations of freely propagating flames that are subject to triplet-map stirring.

3.8 PDF Shapes

The mapping closure process requires calculation of PDF values at each $Z[j]$ and $c[k]$. The β PDF is given by

$$P(x) = P_\beta(x; \tilde{x}, \widetilde{x''^2}) = \underbrace{\frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)}}_{\beta(a,b)} x^{(a-1)}(1-x)^{(b-1)}, \quad (3.10)$$

for the random variable x , using a prescribed mean \tilde{x} and variance $\widetilde{x''^2}$. It uses the well-defined Γ -function, and coefficients a and b computed as

$$\gamma = \tilde{x}(1-\tilde{x})/\widetilde{x''^2} - 1 \quad ; \quad a = \gamma\tilde{x} \quad ; \quad b = \gamma(1-\tilde{x}). \quad (3.11)$$

β -PDFs are valid for variances $\widetilde{x''^2} < \tilde{x}(1-\tilde{x})$.

The top-hat PDF defines a uniform frequency between two limits f_a and f_b ,

$$P(x) = T(x) = \begin{cases} t_0 & f_a \leq x \leq f_b \\ 0 & \text{otherwise} \end{cases}, \quad (3.12)$$

where

$$t_0 = \frac{1}{f_b - f_a}, \quad f_a = \tilde{x} - \frac{l}{2}, \quad f_b = \tilde{x} + \frac{l}{2}, \quad l = \sqrt{12\widetilde{x''^2}}. \quad (3.13)$$

This is valid for small variances ($< \frac{1}{3} \min[\tilde{x}^2, (1-\tilde{x})^2]$) and needs modifications for larger variances and when \tilde{x} is near 0 or 1. The modifications include the addition of *weighted* δ

peaks at the extrema, i.e., $x = 0$ and $x = 1$. The modified top-hat-PDF is summarised as

$$P(x) = \begin{cases} \delta(x - \tilde{x}) & \widetilde{x''^2} = 0 & \mathbf{a)} \\ T(x) & 0 < \widetilde{x''^2} \leq \frac{1}{3}\min(\tilde{x}^2, (1 - \tilde{x})^2) & \mathbf{b)} \\ w_0\delta(x) + T(x) & \frac{\tilde{x}^2}{3} < \widetilde{x''^2} \leq \tilde{x}(\frac{2}{3} - \tilde{x}) & \mathbf{c)} \\ w_1\delta(x - 1) + T(x) & \frac{1}{3}(1 - \tilde{x})^2 < \widetilde{x''^2} \leq (1 - \tilde{x})(\frac{2}{3} - (1 - \tilde{x})) & \mathbf{d)} \\ w_0\delta(x) + T(x) + w_1\delta(x - 1) & \frac{2}{3}\tilde{x} - \tilde{x}^2 < \widetilde{x''^2} \leq \tilde{x}(1 - \tilde{x}) & \mathbf{e)} \end{cases} \quad (3.14)$$

For the cases marked **a)** to **e)** above: limits (f_a and f_b), δ -weights (w_0 and w_1), and top-hat-weight, t_0 , are given in Table 3.1. Fig. 3.10 compares the two PDF assumptions, a weighted δ -peak can be seen added to the left side of the top-hat function at $\tilde{x} = 0.1$. The variable x can be mixture fraction Z or progress variable c . A notable advantage of the extended-top-hat function over β -functions is the speed of computation, in particular as it does not require the evaluation of the Γ -function. As the PDF profiles in LES are narrow, the exact shape is not important [34], and so the top-hat profile is a viable alternative to the β -function.

3.8.1 PDF Moments: numerical integration

Case	t_0	w_0	w_1	f_a	f_b	l
a)	-	-	-	-	-	-
b)	$\frac{1}{l}$	-	-	$x - \frac{l}{2}$	$x + \frac{l}{2}$	$\sqrt{12x''^2}$
c)	$\frac{1 - w_0}{l}$	$1 - \frac{4\tilde{x}^2}{3(\widetilde{x''^2} + \tilde{x}^2)}$	-	0	l	$\frac{2\tilde{x}}{1 - w_0}$
d)	$\frac{1 - w_1}{l}$	-	$1 - \frac{4(1 - \tilde{x})^2}{3(\widetilde{x''^2} + (1 - \tilde{x})^2)}$	$1 - l$	0	$\frac{2(1 - \tilde{x})}{1 - w_1}$
e)	$6(\tilde{x} - \tilde{x}^2 - \widetilde{x''^2})$	$1 - t_0 - w_0$	$\widetilde{x''^2} - \frac{t_0}{3} + \tilde{x}$	0	1	-

Table 3.1: Weights and limits for the extended-top-hat PDF of Floyd et al. [53], cf. (3.12) and (3.14).

Small variances lead to narrow PDF shapes which means that the peak may not be




Figure 3.10: β -PDFs and extended-top-hat PDFs, calculated using the mean and variance values written in the boxes.

correctly resolved in discretised space. Hence, Cumulative Distribution Functions (CDFs²) are preferred for the integration (2.80). For the β -PDF, the equivalent CDF is the ‘incomplete’- β function, which is written as

$$C(x) = C_{\beta}(x; \tilde{x}, \tilde{x}''^2) = \beta(a, b) \int_0^x t^{(a-1)}(1-t)^{(b-1)} dt. \quad (3.15)$$

The equivalent CFD for the top-hat PDF would be

$$C(x) = \begin{cases} w_0 & x < f_a \\ 1 - w_1 & x > f_b \\ xT(x) + w_0 & \text{otherwise.} \end{cases} \quad (3.16)$$

The first moment in single variable-space using discrete CDFs can be performed as

$$\tilde{\phi} = \int \langle \phi | c \rangle dc \approx \sum_k \frac{1}{2} (S[k+1] + S[k]) (C[k+1] - C[k]), \quad (3.17)$$

where $C[k]$ is the bin-centred CDF value. CDFs are also convenient for the Dirac- δ function, however, in this work, they are realised through linear interpolation. The joint PDF integration, when Dirac- δ functions are assumed for c , is performed as follows.

²A CDF $C(x) = \int_0^x P(t) dt$, where $P(t)$ is the PDF.

1. Given \tilde{Z} and \tilde{c} , linear interpolate the solution table $S[j, k]$ as

$$I[j] = S[j, k] + \frac{S[j, k+1] - S[j, k]}{c[k+1] - c[k]} (\tilde{c} - c[k]) \quad (3.18)$$

by setting $k = \tilde{c}/(K - 1)$.

2. Compute bin-centred $C[j]$ using \tilde{Z} and $\widetilde{Z''^2}$ using the procedure above.
3. Compute the first moment of the weighted joint PDF as

$$\tilde{\phi} = \int \int \langle \phi | Z, c \rangle P(Z) P(c) dZ dc \approx \sum_j \frac{1}{2} (I[j+1] + I[j]) (C[j+1] - C[j]). \quad (3.19)$$

Chapter 4 describes the validation cases and main simulation results for SG-LEM.

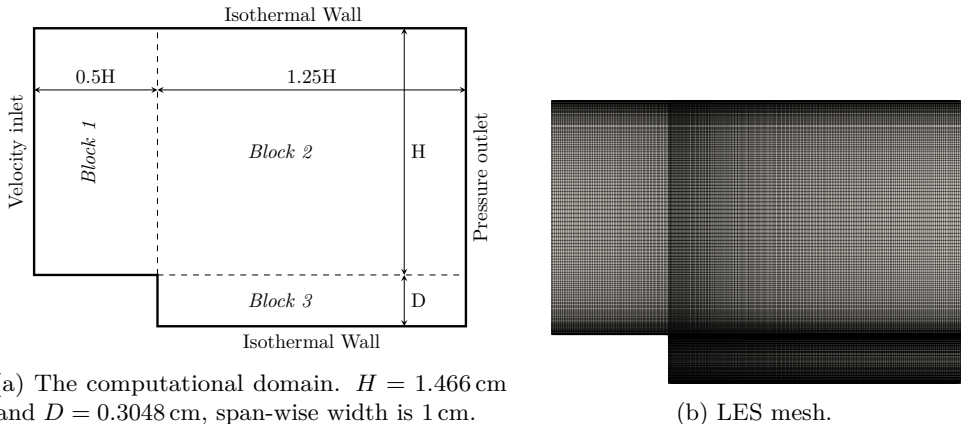
4 Simulations and Main Results

4.1 Introduction

The results from the main investigations carried out during this work are described in detail in the appended papers. This chapter presents a general overview and highlights the evolution of the SG-LEM methodology using a summary of the papers. Three different flame cases are used for validation: a backward-facing step, the Volvo Validation Rig, and the Darmstadt Multi-Regime Burner. The three cases correspond to Papers I, II and III in the manuscript. The investigations evaluate SG-LEM's ability to produce concentration, density, and temperature fields at LES resolution, as well as validation using available data. These cases entail a progressive increase in complexity with the first being a canonical flame for which only DNS data was available, the second having dimensions and parameters more akin to practical flames for which experimental data was available, and the third fully leveraging the mode-independent nature of LEM simulation for which, again, experimental data was used for validation.

4.2 Case I: Backward Facing Step

The first of the validation cases for of SG-LEM was a canonical stabilised flame setup involving a backward facing step, as shown in Fig. 4.1a. The premixed formulation of the mapping closure assuming β -PDFs for progress variable was used here.



(a) The computational domain. $H = 1.466$ cm and $D = 0.3048$ cm, span-wise width is 1 cm.

(b) LES mesh.

Figure 4.1: Numerical setup for the test case.

Validation data was obtained from a reactive DNS for premixed ethylene-air flame for this setup, performed by Aditya et al. [60]. A 32 species skeletal mechanism was used for the SG-LEM simulations, while the DNS solution used a reduced mechanism of 22 species. The computational domain, created with a multi-block approach and hexagonal cells,

is shown in Fig. 4.1b. An ethylene-air mixture with an equivalence ratio of 0.42 and a temperature of 1125 K was introduced at an inlet velocity of 200 m s^{-1} (Mach number 0.3). A pressure-based solver is justified here due to the moderate Mach number. Solution tables were initialised using a zero-dimensional reactor and LEM domains resolved at $26 \mu\text{m}$, which is comparable to the resolution used by the DNS study. Two cluster resolutions were tested – 125 and 1000 (LES cells per cluster), and LEM constants were set to $N_\eta = 4.0$ and $C_\lambda = 15.0$ based on literature values for LES-LEM. Eddies were implemented using the time-accurate operator splitting (cf. Sec. 3.3). The results of the investigation is summarised in the following text.




Figure 4.2: Instantaneous progress variable: LEM-level \tilde{c}_{LEM} (secondary output), and LES resolved \tilde{c} .

SG-LEM approach relies on an accurate LES advancement of mean reaction progress variable, and of course solution tables for each cluster. The source term for progress variable, also derived from LEM and the mapping closure, is thus vital to obtain the correct flame front. Snapshots, such as Fig. 4.2, illustrates that the method is able to capture overall flame behaviour for this case, i.e., maintain a nearly horizontal flame due to the high inlet velocity. Differences in resolved structures, i.e., between the primary and secondary outputs, are attributed to inaccuracies in splicing transport.



Figure 4.3: Instantaneous temperature: LEM-level \tilde{T}_{LEM} (secondary output), LES resolved \tilde{T} (iterated from \tilde{h}), and LES resolved $\tilde{T}_{\text{mapped}}$ (from mapping closure).

Temperature snapshots, Fig. 4.3, show a time-accurate consistency between the three methods, i.e., the secondary output (\tilde{T}_{LEM}), mapping closure ($\tilde{T}_{\text{mapped}}$), and temperatures iterated from the enthalpy field (\tilde{T}), which is used to compute the density field. The mapping output temperatures also shows slight blocky artefacts, revealing the underlying

coarse mesh, and also the effect of wall-treatment on wall-adjacent clusters. Mapping closure using β -PDFs works well for major species like CO_2 , but blocky artefacts for radical OH can be observed, as shown in Fig. 4.4, which also shows a deviation from time-accurate consistency between the two here shows differences in mapping closure (c -conditioned values together with the β -PDFs) and LEM Favre-mean values for a sensitive radical specie.




Figure 4.4: Instantaneous OH results LEM-level \tilde{Y}_{OH}^{LEM} (secondary output), and LES resolved \hat{Y}_{OH} .




Figure 4.5: Mean streamwise velocity.

Next, time-averaged data is compared to DNS results that were provided by the authors of Ref. [60], as well as a PaSR simulation using the same mesh and chemical mechanism for a fair comparison. Velocity profiles in Fig. 4.5 shows that the LES setup accurately captures mean the flow features, including the recirculation zone at $x = 0.6H$, there is little difference between the PaSR and SG-LEM outputs here.

Mean temperature plots in Fig. 4.6 show that SG-LEM reports lower temperatures near the recirculation region ($x = 0.6H$) compared to DNS; however, it is significantly more accurate than PaSR. Both methods show much better agreement downstream.

Mean mass fractions also demonstrate SG-LEM's superior predictive capability for this setup, in particular for radicals CO and OH shown in Fig. 4.7. The shape of the peak for CO a predicted downstream and so is the overall profile for radical OH. This could be due to the high (LEM) resolution in the near-wall regions, combined with the isothermal




Figure 4.6: Mean temperature profiles.




Figure 4.7: Mean mass fractions.

wall-treatment described previously. The PaSR simulation can only resolve down to the LES mesh in the near-wall regions. The differences between LEM and LES values for OH is consistent with Fig. 4.4, also the presence of blocky artefacts does not adversely affect the mean field. The two cluster sizes tested showed some differences in the LEM-level




Figure 4.8: Comparison of cluster sizes 125 and 1000 w.r.t. mean progress variable.

flame distribution, as seen in Fig. 4.8 in the downstream region. This is due to the increased inaccuracies in splicing transport with the larger cluster size, transporting burnt LEM fragments upwards, into the free stream. However, the overall LES flame (\tilde{c}) remains qualitatively unchanged for this setup. This suggests a degree in robustness in the approach to mean reaction progress advancement, even though the source term is conditionally binned at the super-grid level.

Finally, computational performance was evaluated for the two clusters sizes and compared to the PaSR simulation. SG-LEM using cluster size 125 was able to outperform the standard `reactingFoam` solver, part of `OpenFOAM`, using PaSR closure, requiring around 50% of the compute time (27.6K core-hours vs 12.8K core-hours). The savings are due to scaling-down of the number of elements that need chemistry integration, i.e., LES cells vs. LEM wafers, and also by replacing transport equations for 32 individual species with the single mean progress variable transport. The 1000-sized cluster required even lesser compute-time, primarily due to fewer LEM domains – around 40%. While it is desirable to obtain a one-to-one comparison with standard LES-LEM, performing a full simulation using the same mesh and mechanism was found infeasible, given resource constraints. However, a partial run enabled gathering of time statistics from which it was extrapolated that a full LES-LEM it would require approximately 248K core-hours to run a full simulation, i.e., SG-LEM delivers a speedup of around 20x for this case and setup – a clear demonstration of the method’s performance potential.

This first investigation and application of SG-LEM showed encouraging results for the method. It is able to produce reasonably accurate concentration and temperature fields, at i) LES resolution and ii) at compute costs that are competitive with standard models, with chemical-state validation using DNS data – a substantial improvement over traditional LES-LEM. Even though some sensitivity to cluster size was observed, there was little to no change in the overall flame behaviour in terms of stability.

4.3 Case II: Volvo Validation Rig




Figure 4.9: The Volvo validation rig setup, spanwise width is 0.24 m. Not shown are the pre-inlet section and round exhaust pipe. Dashed lines show measurement planes named P1, P2, and P3 at $x = 0.95h$, $3.75h$, and $9.4h$ measured from the base of the bluff-body.

The second investigation aimed to broaden the validation of SG-LEM for premixed flames and stress test the sensitivity to clustering parameters. The Volvo Validation Rig presents a larger domain, higher turbulence and a lean propane-air mixture (equivalence ratio 0.62) at atmospheric pressure and 288 K. As mentioned before, this case has dimensions and parameters more akin to practical flames. Sjunnesson et al. [61, 62, 63] provide experimental data for validation in the form of mean and RMS scalars. This case is a step forward to ascertain if SG-LEM is indeed capable of handling practical flames, larger geometries, and a fuel-air mixture that is more difficult to ignite than the (high temperature) ethylene-air mixture tested previously. The test domain, shown in Fig. 4.9, consists of a rectangular flow channel and a triangular bluff body that acts as a flame holder. A multi-block LES mesh of uniform cell size 1 mm was used. The case parameters are given in Table 4.1, these are known by ‘Case 1’ in the referenced literature.

Table 4.1: Volvo validation rig, Case 1.

Parameter		Value
Inflow pressure	p_0 [atm]	1.0
Inflow temperature	T_0 [K]	288
Inflow velocity	v_0 [m s ⁻¹]	17.6
Reynolds number	$Re = v_0 h / \nu$	46.59×10^3
Equivalence ratio	ϕ	0.62
Laminar flame speed	s_u [m s ⁻¹]	0.15
Integral length scale	l_t [m]	4×10^{-3}
Velocity fluctuations	v_{RMS} [m s ⁻¹]	7.8
Karlovitz Number	Ka	62.0
Damköhler Number	Da	0.7
Temperature ratio	T_b/T_u	5.90

A more detailed analysis on the effect of clustering parameters was carried out in this

investigation than for Case I. Cluster sizes 64, 125 and 250 were tested. Note that the wall-treatment was not applied here as all walls were assumed to be adiabatic, this is in line with existing studies. The finest cluster resolution, denoted ‘Sim. 1’, was computed using a two-step mechanism from Ghani et al. [64]. Solution tables were initialised for these conditions using a premixed laminar flamelet solution produced by an in-house code. The initialisation prevents early numerical problems and also tabulated values for the progress variable source term, naturally these values are updated during runtime by LEM domains. The SG-LEM simulation shows reasonably good agreement with experimental data for mean velocity and temperature, as shown in Fig. 4.10. It even compares with PaSR simulation data from Zettervall et al. [65] shown in the same figure.

Instantaneous snapshots, such as Fig. 4.11, show a spreading artefact in the secondary output which is not consistent with the LES solution – the LEM level flame shows greater spread in the cross-stream direction. This is due to inaccuracies in large-scale splicing transport, which shall be discussed in more detail shortly. The PaSR simulation data taken from literature was based on a more detailed ‘Z66’ skeletal mechanism, which is able to correctly predict ‘CO’ levels near the bluff body as well as downstream in the referenced simulation. Using this mechanism with the finest cluster sizes turned out to be impractical with the available compute resources and so it was first tested with the cluster size 125, denoted ‘Sim 2’. Initial tests did not produce a stable flame. It was observed that splicing inaccuracies due to the irregular cluster shapes led to excessive entrainment of the freestream fluid behind the recirculation zone, resulting in cooling and eventual quenching of the flame. This is the same mechanism that produced the spreading artefact observed earlier. Two methods were devised to mitigate this problem:

1. An *ad-hoc* subroutine called Intra-Cluster Mixing (ICM) that introduces large (of cluster dimensions) breakdowns to the LEM domains that modelled advective structures that are smaller than the cluster size, but larger than sub-grid turbulence. One eddy is implemented per time step (size and location are sampled as usual) assuming $\eta = l_{\text{LEM}}/2$.
2. A method to exercise some control over the cluster shapes using processor decomposition to produce flow-aligned processor domains, and hence, flow-aligned clusters. This is based on `MGRIDGEN` working independently in each process/thread (cf. Sec. 3.4). A ‘projection’ decomposition method was used where the outlet face was first decomposed using the ‘scotch’ method [66], followed by projection in the streamwise coordinate. This was performed using the `blockMesh` utility that is distributed with `OpenFOAM`.

The ICM routine was used with Sim. 2 and produced a stable flame, but still showed the spreading artefact. It is reasoned that the additional breakdowns produced by ICM creates more flame fronts which counteracts the excessive fresh charge that is spliced in, this applies particularly to LEM domains downstream of the recirculation zone. The shape-control method was used with the coarsest cluster resolution (250) and did not require ICM to produce a stable flame, this is denoted as ‘Sim. 3’. Time averaged results are shown in Fig. 4.12, where the spreading artefact is clearly seen in the (LEM) CO levels for Sim. 2, but not for Sim. 3. Both setups still produce elevated CO levels near

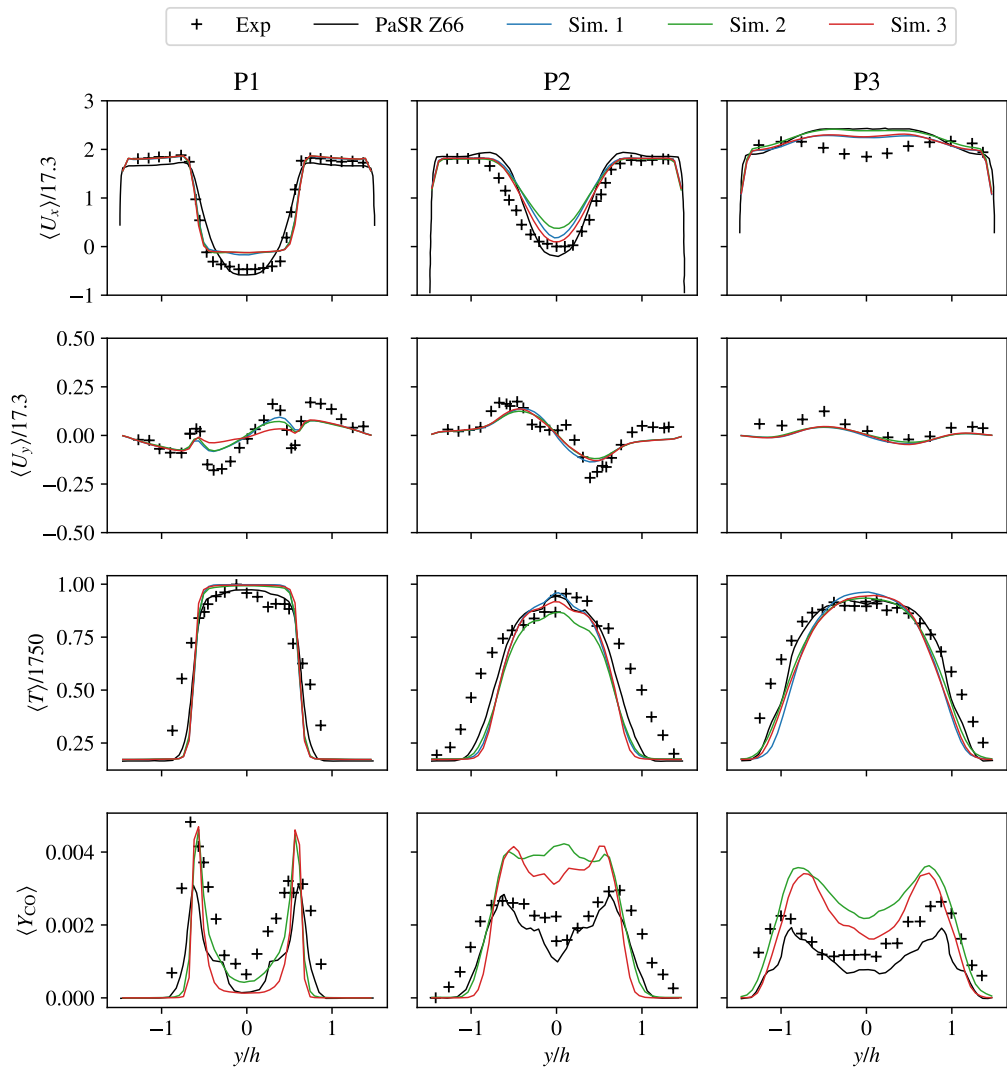


Figure 4.10: Mean scalar profiles, experimental data from Ref. [61], and PaSR data from Ref. [65].

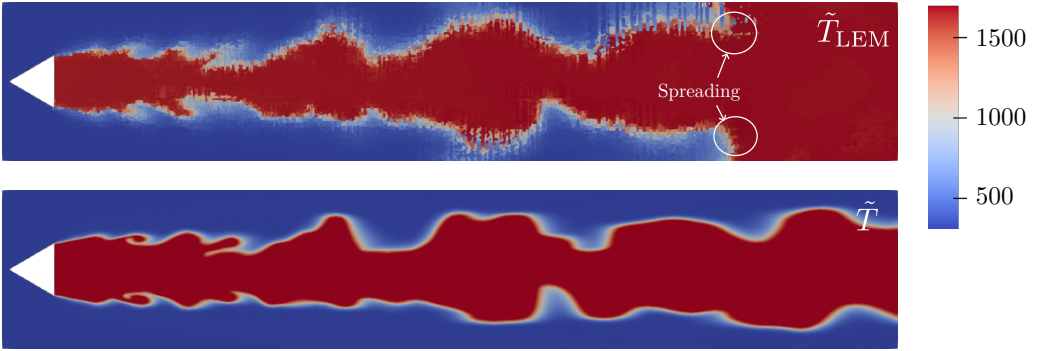


Figure 4.11: Instantaneous temperature, secondary output at coarse resolution \tilde{T}_{LEM} , and LES resolved \tilde{T} .

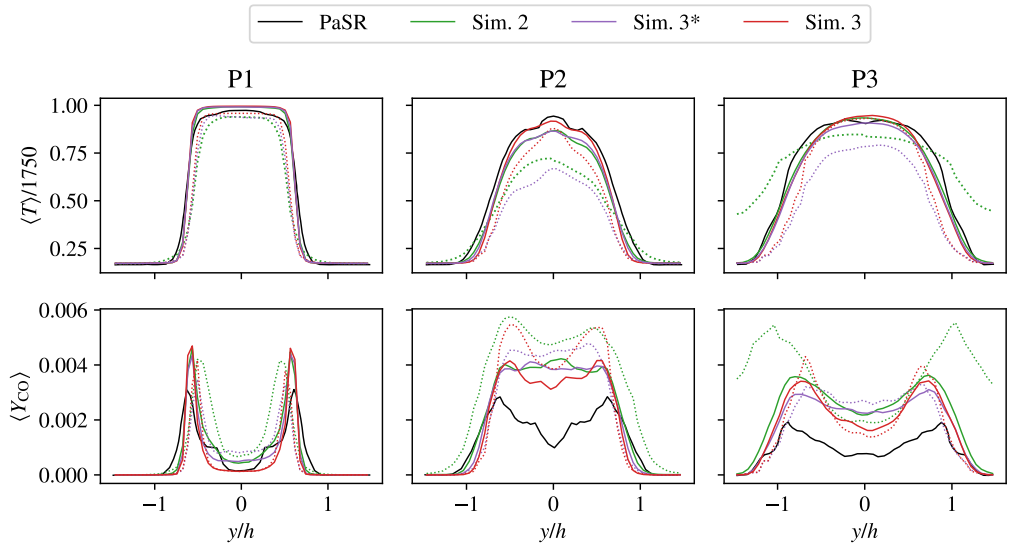


Figure 4.12: Comparison of mean LES results and LEM states. Dotted lines are obtained from Eqn. (2.68), reported at SG resolution and time-averaged.

the axis, behind the recirculation zone, this is thought to be due to excessive entrainment of freestream fluid, i.e., inaccuracies in the transport of LEM fragments.

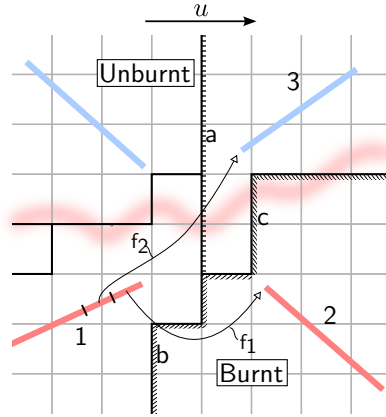


Figure 4.13: Large scale transport between LEM domains for agglomerated meshes.

The mechanism by which irregular cluster shapes lead to such inaccuracies is illustrated in Fig. 4.13. The cluster numbered 3 has a protrusion into the burnt region that is primarily covered by clusters 1 and 2. Hence, a spliced fragment ‘ f_2 ’ will reside in the LEM domain in 3 before it is spliced back to the burn region, unlike ‘ f_1 ’ which remains in the . The splicing scheme (cf. Sec. 3.5), with designated attachment and detachment points, results in a “bucket-brigade” movement of the fragment f_2 where it will travel the length of LEM domain 3, creating additional flame fronts, before exiting. This is the mechanism behind the spreading artefact and a similar explanation can be given for the spurious transport of cold fluid which causes extinction.

All simulations used the adjusted LEM constants as described in Sec. 2.6, and the blocked-sequence operator splitting for eddies as described in Sec. 3.3. Larger LEM domains result in less frequent triplet mapping, to such an extent that no triplet maps were observed in the flow field when the LEM parameters from Case I were used. The adjusted parameters retain the same frequency and size distribution under the assumption that LEM domain sizes have been scaled up by a constant (cf. Fig. 2.6). This second investigation showed a clear dependence on cluster generation for the overall simulation accuracy, which was not observed to the same extent for the Case I. Cluster shapes and sizes have a large impact on (coarse-grained) large-scale transport. In extreme cases, this can lead to nonphysical flame quenching when large cluster sizes are setup improperly. Additional mixing can, to a limited extent, compensate for the unresolved advective structures that promote flame stability.

Overall, this investigation provides a realistically constrained assessment of the model’s capabilities and limitations by using the more challenging Volvo rig as a stress-test, in particular regarding clustering parameters and how they affect the overall solution. The current SG-LEM framework is aimed at general unstructured meshes and complex geometries in engineering applications, which motivates the use of mesh agglomeration to

create the super-grid. It is reasonable to expect improved results with regular super-grid clusters using a pure Cartesian grid, or at least a structured mesh, which are often used in academic test cases such as this.

4.4 Case III: Darmstadt Multi-Regime Burner

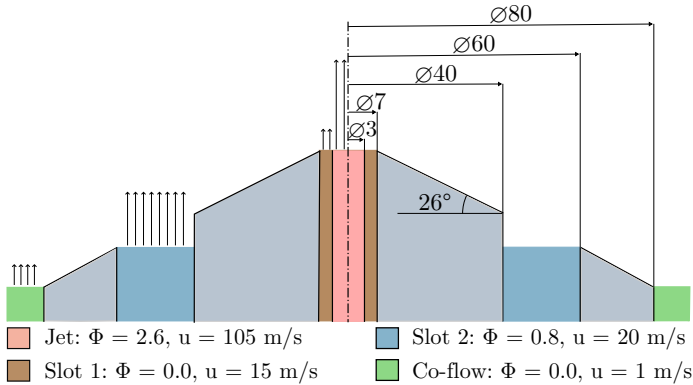


Figure 4.14: Schematic of the Darmstadt Multi-Regime Burner (MRB).

The third validation case in this work involves utilises the ‘full’ mapping closure, i.e., using both mixture fraction and progress variable, to simulate a multi-regime scenario that has characteristics of both a premixed and partially-premixed (diffusion) flame. Note that the term ‘multi-regime’ is used here instead of ‘multi-mode’ to be consistent with the referenced literature. The setup investigated is the Darmstadt Multi-Regime Burner (MRB) as presented by Butz et al. [67], the schematic for which is shown in Fig. 4.14, where inlet velocities (u) and equivalence ratios (ϕ) for premixed methane-air streams are specified for the burner streams. These correspond to the operating point ‘MRB26b’ in Ref. [67]. The three streams are marked as ‘Jet’, ‘Slot 1’ and ‘Slot 2’, along with a co-flow. This condition for the burner head is characterised by a premixed flame from the annular jet emerging from Slot 2 which is stabilised behind the conical bluff body, while the high equivalence ratio for the central jet results in a lifted diffusion flame that is supported by burnt products from the premixed flame.

The LES domain is cylindrical with a radius of 200 mm and an axial length of 183 mm. The mesh comprises hexahedral cells with sizes ranging from 0.1 to 0.4 mm. Initial tests revealed that the quenching problem, as an artefact of cluster shapes as described above, was present in this setup, just as in the Case II. Hence, the projection-decomposition method was used to control cluster shapes. Flow-aligned clusters, using size of 64, were generated with the `blockMesh` utility as before. The agglomerated mesh is shown in Fig. 4.15. Time-resolved inlet velocities for the jet and Slot 2 are obtained from precursor LES simulations for a pipe-cyclic and an annular-cyclic setup, respectively. From these, inlet velocities for fully developed turbulent conditions at the appropriate bulk velocities were used for the simulation. A laminar velocity profile was used for Slot 1 due to its low

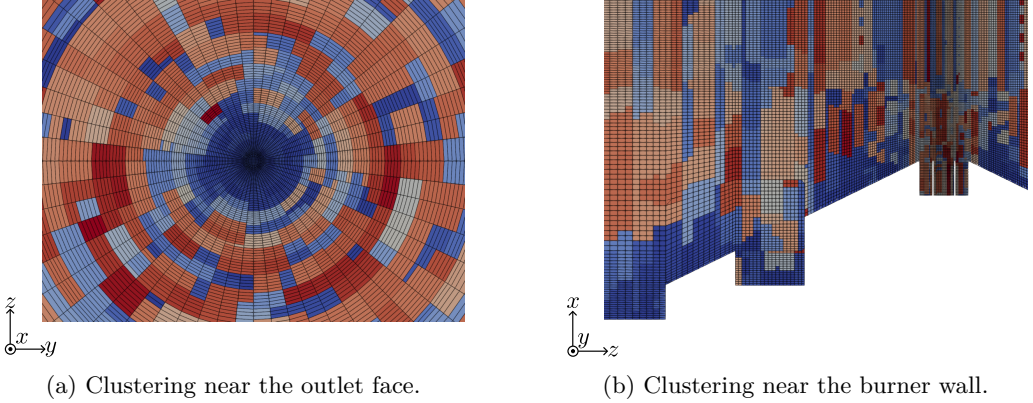


Figure 4.15: LES mesh (black lines) and clusters (colours).

velocity and narrow diameter. The boundary conditions for \tilde{Z} for the jet and Slot 2 are 0.131 and 0.044, respectively, corresponding to the ϕ values in Fig. 4.14. For Slot 1 and the co-flow, $\tilde{Z} = 0$. As with Case II, i.e., the Volvo Rig, solution tables were initialised using premixed unsteady laminar flamelets. The results for the simulation are summarised as follows.

As before, instantaneous scalar fields like Fig. 4.16 showed a high similarity in the structures seen in the primary and secondary outputs, i.e., a degree of adjacency in the LEM and LES flame fronts which supports the notion that SG-LEM produces locally relevant solution tables.

Radial profiles for time-averaged fields are shown in Figs. 4.17 along with comparisons to experimental data, at four axial positions. The top row is mean axial velocity, showing good agreement with experiments near the burner. It shows the jet breakup region as well as the recirculation zone ($u_x < 0$). The accuracy diminishes downstream with lower centreline velocity and greater radial dispersion. Mixture fraction fields (third row) show slight excessive mixing between Slot 1 and hot gases from the outer flame near the burner, which correlates with slight excessive radial velocity (second row, $x = 15$ mm). This is shown by leaner mixture fraction values, which is also observed between the jet and Slot 1 at the downstream locations. The consistency in mixing behaviour between the LEM (secondary output) and LES levels is encouraging. The differences in mixture fraction between experimental data and SG-LEM explains the observed differences in reactive scalars, i.e., slightly lower temperatures (fourth row) as well as lower CO levels (last row $x = 15$ mm) for the inner flame. The excessive radial transport is likely to have originated from fluctuations in the time-resolved inflow conditions used for the central jet.

The outer flame exhibits accurate mixture fraction and temperature profiles; however, there is a noticeable under-production of CO. This contrasts with the results from Case I, where SG-LEM was able to reproduce accurate levels. An analysis using c -conditioned data, shown in Fig. 4.18, narrowed the source of this error down to the mass-exchange between LEM domains. The lack of inter-cluster (inter-LEM) diffusion domains in the

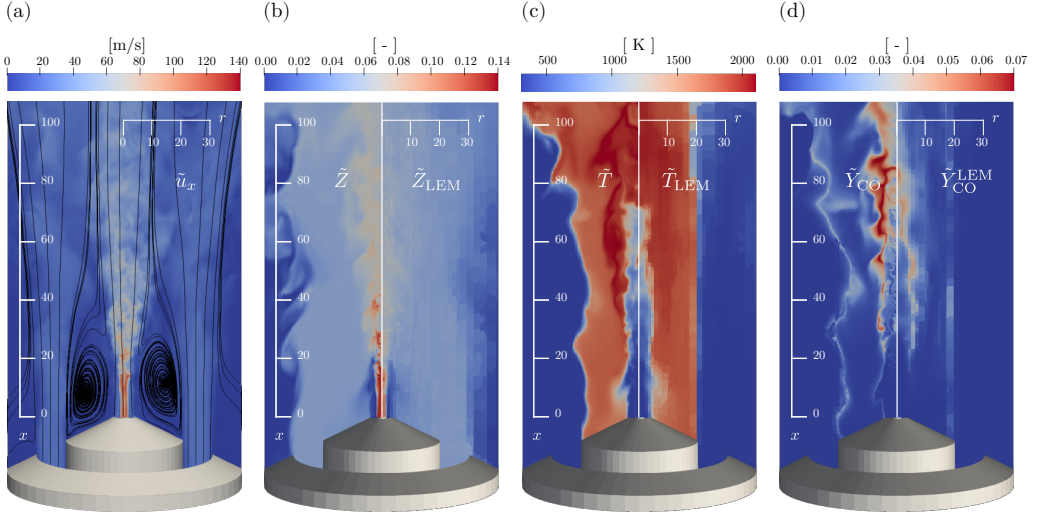


Figure 4.16: Instantaneous (a) velocity, (b) mixture fraction, (c) temperature, and (b) CO mass fraction. Streamlines in panel (a) computed with time averaged velocity. Panels (b), (c) and (d) at LES resolution (left half), and secondary output (right half) showing LEM-level fields.

current framework is thought to play an important role here. Figs. 4.16 (c) and (d) reveal that clusters in the outer region align with the (LEM) flame front, hence the only interaction between the fresh charge from Slot 2 and the outer flame is via large-scale transport. This is unphysical as diffusion plays an important role in determining flame structure, which is aligned normal to the flame-front. Also shown in Fig. 4.18 are investigations using standalone LEM simulations (see Sec. 3.7.1) marked ‘SA1-SA3’, which replicate conditions near Slot 2. The standalone LEM data further supports the observed under-production of CO as originating from mass exchange between LEM domains and not from LEM advancement or the operator splitting used.

Nonetheless, the LEM-level outer flame is stable and provides hot gases (LEM fragments) that are transported into the recirculation zone, via splicing, and stabilises the inner flame which is reproduced with reasonable accuracy by SG-LEM.

For more insight into the mixing and flame behaviour, conditional CO data for the inner flame ($r \leq 10$ mm) was also analysed. Fig. 4.19 shows instantaneous scatter plots of CO mass fractions in the Z - Y_C -plane, with

$$Y_C = Y_{\text{CO}_2} + Y_{\text{H}_2\text{O}}. \quad (4.1)$$

Mixing trajectories are marked as follows: a represents pure mixing between the jet and Slot 1, c between the burnt products from Slot 2 and air from Slot 1, and b_r and b_l are the rich and lean combustion limits for mixing between the jet and Slot 1, respectively. The primary and secondary outputs capture these trajectories well, with LEM showing a decreased range owing to the coarser resolution.




Figure 4.17: Time averaged velocity, mixture fraction, temperature and CO mass fraction.




Figure 4.18: Conditional CO statistics for the outer flame. Laminar profile and standalone LEM simulations (SA1-3) using $Z = 0.044$ shown for comparison. SA1: $Re_t = 100$; SA2: $Re_t = 50$; SA3: $Re_t = 50$, time-accurate eddy implementation.

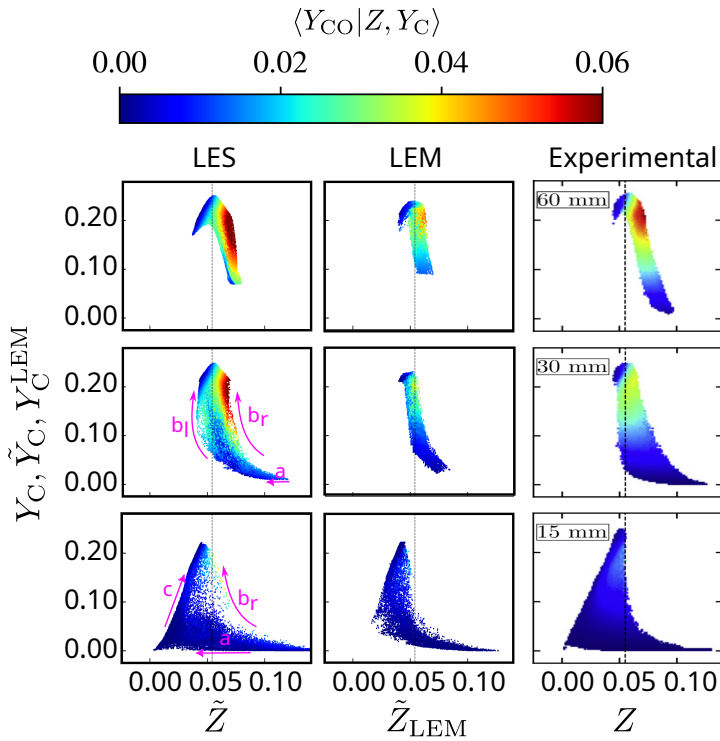


Figure 4.19: Conditional CO statistics of inner flame using LES data (left column), LEM-level data (middle column) and experimental images (right column) from [40] with permission from Elsevier. Rows correspond to streamwise locations indicated in the right column. Dotted lines shows stoichiometric mixture fraction of 0.055.

Overall, Case III leverages the mode-independent nature of LEM, and the ‘full’ SG-LEM mapping closure using mixture fraction and progress variable. LEM domains are able to reproduce the mixing trajectories (between the streams) for the inner flame, as well as flame liftoff. Inaccuracies in the outer flame region is argued to stem from the lack of inter-LEM diffusion that is accentuated by the radial mesh that leads to larger LES clusters. Importantly, it represents progress towards simulating practical flames, which often exhibit regions of a partially-premixed nature.

5 Conclusions and Future Work

The primary focus of this thesis was to develop an efficient combustion model for Large Eddy Simulation (LES) that can be used, with minimal modification, for premixed, non-premixed or mixed-mode flames. The Linear Eddy Model (LEM) has been used as sub-grid closure for LES, a technique known as LES-LEM, and proven reliably accurate for a wide range of conditions including the above modes, as well as complex flow fields, e.g., swirl (for stabilisation), high Re number flow, compressible flows, etc. LES-LEM utilises several features of LEM that are beneficial to reactive flow simulation: highly resolved flame structures, physically accurate sub-grid diffusion, and simulated (as opposed to modelled) sub-grid stirring and turbulent flame speed. LES-LEM incurs significant computational costs, primarily due to chemistry integration needed for the highly resolved LEM domains that are embedded in each LES cell, as well as overheads related to splicing – the routine used in LES-LEM to model large-scale transport between LEM domains. This has posed practical limitations to the use of LES-LEM in real-world applications.

This work describes and validates a novel variant of the LES-LEM technique that uses mesh agglomeration to drastically cut simulation times. Coarse-graining of the LES mesh results in a super-grid comprising clusters of LES cells which are embedded with LEM domains. This results in a drastic down-scaling in the number of LEM domains while retaining resolutions and representative domain lengths, and should lead to proportionate savings in computational costs. Chemistry tabulation and presumed PDF mapping closure is used to produce LES resolved concentration fields. Each cluster reads an individual flame-structure solution table that is continuously updated during runtime by its embedded LEM domain. A modified LES-LEM splicing scheme was developed for this coarse-grained variant. The method, named super-grid LES-LEM, or simply SG-LEM, was tested for three reacting flow cases for which either experimental or DNS data was available for validation.

Case I involved a backward-facing step where a premixed ethylene-air flame is stabilised by a recirculating flow field. Here, SG-LEM demonstrated its capacity to produce accurate LES-resolved concentration fields, albeit with slight blocky artefacts for the intermediate species OH, revealing the super-grid structure in some areas. A comparison of mean values with DNS data revealed that SG-LEM can produce more accurate temperature and intermediate species fields than a tested PaSR simulation. The estimated speedup, compared to standard LES-LEM for the tested setup (mesh and skeletal chemistry), was shown to be around around 20. This initial test is characterised as a strong success for the method. Although two different cluster sizes (125 and 1000) were tested, LES concentration fields did not show significant differences.

Case II was the Volvo Validation Rig, which also features a recirculation-stabilised premixed flame, but with a propane-air mixture and features a larger geometry with characteristics more akin to practical flames. The investigation revealed a strong dependence on clustering parameters owing to the large geometry when tested with a skeletal mechanism. The finest cluster resolution (64) was tested with a global chemical mechanism and demonstrated accurate temperature and velocity fields. To more accurately capture

CO levels, a more detailed ‘Z66’ skeletal mechanism was used with two coarser cluster resolutions: 125 and 250. Initial tests did not produce a stable flame. An investigation into this phenomena showed excessive large-scale transport, i.e., splicing, of fresh charge into the flame regions led to cooling and eventual quenching of the flame. Splicing accuracy has been identified to be strongly dependent on cluster sizes and shapes. These have a direct impact on residence times of the fragments that are spliced into an LEM domain, given that the splicing scheme is derived from LES-LEM.

Two methods were developed to counteract the issue of splicing inaccuracies, and promote stability of the simulated flame: an *ad-hoc* mixing enhancement that tried to model advective structures of cluster proportions, and a method to control cluster shapes using a projection-type processor decomposition. While both methods resulted in stable flames, the LEM-level flame-spread was noticeably larger downstream when shape control was not in use. However, the LES-resolved flame did not show this spreading artefact. Both coarser cluster resolutions produced slightly elevated CO near the flow-axis downstream, which could be due to excess entrainment of bulk fluid, i.e., large-scale transport inaccuracies.

Finally, Case III was the Darmstadt Multi-Regime Burner – a complex setup that comprises an inner lifted diffusion flame stabilised by an outer premixed flame that is attached to the burner body. The setup includes two methane-air streams, inner and outer, as well as two co-flows. Solution tables were initialised using laminar flamelets and flow-aligned cluster shapes were produced using a similar routine as in Case II. This was needed to stabilise the flame. Simulation data showed that LEM domains can capture both the inner and outer flame behaviour and produce LES-resolved fields via mapping closure. Comparison of mean and RMS data with experimental values demonstrated that SG-LEM is capable of capturing mixing behaviour between the streams at both levels and producing accurate temperature fields. While CO for the inner flame showed good agreement, the outer flame exhibited under-production. This discrepancy, despite similarities between the outer flame and Case I where SG-LEM produced very accurate results, is attributed to under-resolved transport combined with the lack of diffusion between LEM domains on either side of the outer flame-front. Here, the radial LES mesh combined with the cluster shapes, despite providing stability, might have led to a decreased interaction between the flame and the freestream fluid. Nevertheless, the case clearly demonstrates SG-LEM’s ability to handle multi-mode flames.

With this, we can conclude that the major goals outlined previously have been achieved: speedup of LES-LEM, production of LES-resolved fields, and predictive thermo-chemical statistics. These achievements were demonstrated across all three test cases. Another highlight is that SG-LEM works on unstructured meshes, which makes it immediately applicable to practical cases. However, Cases II and III revealed a potentially significant drawback of the model – inaccuracies in large-scale transport that can alter flame behaviour. Another drawback was found to be memory usage. Each cluster requires an individual solution table which imposes a minimum cluster size limit for a given mesh and chemical mechanism.

5.1 Future Work

Based on the discussion above, it is recommended that future implementations of SG-LEM address the deficiencies in large-scale transport. The following suggestions are made for potential avenues of improvement.

1. **Residence-time-informed splicing:** The splicing routine developed for this work was guided by the concept of control volume crossing-rate, consistent with standard LES-LEM. This was achieved by assigning each LEM domain designated attachment and detachment ends. As a result, large-scale transport is uni-directional along the LEM coordinate. This was shown to work poorly for large, irregularly-shaped clusters that result from mesh agglomeration. Residence times for spliced fragments may not be physical, which is important across flame-fronts. This is neither an issue for LES-LEM nor for SG-LEM using small cluster sizes (high cluster resolution), although both methods impose limits on the use of large chemical mechanisms. A more sophisticated splicing scheme is required, should mesh agglomeration be the chosen method of generating super-grids. Such a scheme could, say, use geometric information regarding the cluster shape, as well the instantaneous flux directions to determine attachment and detachment points such that residence times are more accurately reflected.
2. **Cluster shape control:** The current work used an indirect method to control cluster shapes, i.e., using the shapes processor domain boundaries, which worked well for the tested cases but still produced some inaccuracies. Instead, adaptive cluster refinement routines could be used where a refinement criteria, such as velocity gradient, can be used to break up clusters in regions that would benefit from more accurate splicing transport. This way, the current SG-LEM splicing routine could be used without modification.
3. **Overset super-grid:** The basis of the present SG-LEM methodology is mesh agglomeration, automated coarse-graining techniques are not guaranteed to produce flow-aligned, or even uniform, cluster shapes, even when used on structured/Cartesian meshes. An alternative would be to directly design a super-grid based on physical criteria, and superimpose this grid on to the LES grid much like an overset mesh. While it is reasonable to expect such a technique to promote flame stability, it is not clear if the unresolved intra-cluster advection will still need to be modelled as it may affect sensitive scalars like intermediate species.
4. **Inter-cluster diffusion:** Practical simulation times motivates the use of large cluster sizes. In the case of flow-aligned clusters, and in particular for regions where cluster boundaries are aligned with the mean flame brush, the interaction between burnt and unburnt gases is purely through large-scale transport. This is not representative, as diffusion of heat and chemical species determines the flame structure normal to the flame-front. Future implementations of SG-LEM could include a diffusion model for such regions, this requires one of the key model assumptions – of LEM domains having indeterminate orientation – to be re-evaluated. Another, more practical, approach is the incorporation of a ‘cross-term’ as discussed

in Sec. 10.6.2 of Menon et al. [50] to the splicing scheme using resolved LES fluxes so as to increase the interaction of burnt and unburnt LEM domains in such regions.

Bibliography

- [1] J. Janicka et al. “Large Eddy Simulation of Turbulent Combustion Systems”. In: *Proceedings of the Combustion Institute* 30.1 (Jan. 2005), pp. 537–547. ISSN: 15407489.
- [2] C J Rutland. “Large-Eddy Simulations for Internal Combustion Engines – a Review”. In: *International Journal of Engine Research* 12.5 (Oct. 2011), pp. 421–451. ISSN: 1468-0874.
- [3] M. Boileau et al. “LES of an Ignition Sequence in a Gas Turbine Engine”. In: *Combustion and Flame* 154.1 (July 2008), pp. 2–22. ISSN: 0010-2180.
- [4] D. B. Spalding. “Mixing and Chemical Reaction in Steady Confined Turbulent Flames”. In: *Symposium (International) on Combustion*. Thirteenth Symposium (International) on Combustion 13.1 (Jan. 1971), pp. 649–657. ISSN: 0082-0784.
- [5] J. Chomiak. *Combustion A Study in Theory, Fact and Application*. Philadelphia, USA: Abacus Press, Jan. 1990.
- [6] Ivar S. Ertesvåg et al. “The Eddy Dissipation Turbulence Energy Cascade Model”. In: *Combustion Science and Technology* 159.1 (Oct. 2000), pp. 213–235. ISSN: 0010-2202.
- [7] N. Peters. “Laminar Diffusion Flamelet Models in Non-Premixed Turbulent Combustion”. In: *Progress in Energy and Combustion Science* 10.3 (Jan. 1984), pp. 319–339. ISSN: 03601285.
- [8] A. Y. Klimenko et al. “Conditional Moment Closure for Turbulent Combustion”. In: *Progress in Energy and Combustion Science* 25.6 (Dec. 1999), pp. 595–687. ISSN: 0360-1285.
- [9] S.B. Pope. “PDF Methods for Turbulent Reactive Flows”. In: *Progress in Energy and Combustion Science* 11.2 (1985), pp. 119–192. ISSN: 03601285.
- [10] D. C. Haworth. “Progress in Probability Density Function Methods for Turbulent Reacting Flows”. In: *Progress in Energy and Combustion Science* 36.2 (Apr. 2010), pp. 168–259. ISSN: 0360-1285.
- [11] K. N. C. Bray et al. “Unified Modeling Approach for Premixed Turbulent Combustion—Part I: General Formulation”. In: *Combustion and Flame* 61.1 (July 1985), pp. 87–102. ISSN: 0010-2180.
- [12] C. Hasse et al. *Modelling the Effect of Split Injections in Diesel Engines Using Representative Interactive Flamelets*. SAE Technical Paper 1999-01-3547. Warrendale, PA: SAE International, Oct. 1999.
- [13] Rane L Curl. “Dispersed Phase Mixing: I. Theory and Effects in Simple Reactors”. In: *AIChE journal* 9.2 (1963), pp. 175–181.
- [14] S. Subramaniam et al. “A Mixing Model for Turbulent Reactive Flows Based on Euclidean Minimum Spanning Trees”. In: *Combustion and Flame* 115.4 (Dec. 1998), pp. 487–514. ISSN: 0010-2180.

- [15] A. R. Kerstein. “A Linear- Eddy Model of Turbulent Scalar Transport and Mixing”. In: *Combustion Science and Technology* 60.4-6 (Aug. 1988), pp. 391–421. ISSN: 0010-2202, 1563-521X.
- [16] Alan R. Kerstein. “Linear-Eddy Modelling of Turbulent Transport. Part 6. Microstructure of Diffusive Scalar Mixing Fields”. In: *Journal of Fluid Mechanics* 231 (Oct. 1991), pp. 361–394. ISSN: 0022-1120, 1469-7645.
- [17] A. R. Kerstein. “Linear-Eddy Modeling of Turbulent Transport. Part 4. Structure of Diffusion Flames”. In: *Combustion Science and Technology* 81.1-3 (Jan. 1992), pp. 75–96. ISSN: 0010-2202, 1563-521X.
- [18] T. M. Smith et al. “One-Dimensional Simulations of Freely Propagating Turbulent Premixed Flames”. In: *Combustion Science and Technology* 128.1-6 (Oct. 1997), pp. 99–130. ISSN: 0010-2202.
- [19] Suresh Menon et al. “A Linear Eddy Subgrid Model for Turbulent Combustion - Application to Premixed Combustion”. In: *31st Aerospace Sciences Meeting*. American Institute of Aeronautics and Astronautics.
- [20] Tim Lackmann et al. “A Representative Linear Eddy Model for Simulating Spray Combustion in Engines (RILEM)”. In: *Combustion and Flame* 193 (July 2018), pp. 1–15. ISSN: 00102180.
- [21] Nidal Doubiani et al. “A Pressure-Coupled Representative Interactive Linear Eddy Model (RILEM) for Engine Simulations”. In: *Fuel* 355 (Jan. 2024), p. 129423. ISSN: 0016-2361.
- [22] I. Porumbel et al. “Large Eddy Simulation of Bluff Body Stabilized Premixed Flame”. In: *44th AIAA Aerospace Sciences Meeting and Exhibit*. Reno, Nevada: American Institute of Aeronautics and Astronautics, Jan. 2006. ISBN: 978-1-62410-039-0.
- [23] S. Arshad et al. “Subgrid Reaction-Diffusion Closure for Large Eddy Simulations Using the Linear-Eddy Model”. In: *Flow, Turbulence and Combustion* 103.2 (Aug. 2019), pp. 389–416. ISSN: 1573-1987.
- [24] Haiqiao Wei et al. “Large Eddy Simulation of the Low Temperature Ignition and Combustion Processes on Spray Flame with the Linear Eddy Model”. In: *Combustion Theory and Modelling* 22.2 (Mar. 2018), pp. 237–263. ISSN: 1364-7830.
- [25] V. Sankaran et al. “LES of Scalar Mixing in Supersonic Mixing Layers”. In: *Proceedings of the Combustion Institute* 30.2 (Jan. 2005), pp. 2835–2842. ISSN: 1540-7489.
- [26] Shaoshuai Li et al. “Investigation of Dilution Effects on Partially Premixed Swirling Syngas Flames Using a LES-LEM Approach”. In: *Journal of the Energy Institute* 91.6 (Dec. 2018), pp. 902–915. ISSN: 1743-9671.
- [27] J.-Y. Chen et al. “Modeling Differential Diffusion Effects in Turbulent Nonreacting/Reacting Jets with Stochastic Mixing Models”. In: *Combustion Science and Technology* 133.4-6 (Apr. 1998), pp. 343–375. ISSN: 0010-2202, 1563-521X.
- [28] Shaoshuai Li et al. “A LES-LEM Study of Preferential Diffusion Processes in a Partially Premixed Swirling Combustor With Synthesis Gases”. In: *Volume 4B: Combustion, Fuels and Emissions*. Seoul, South Korea: American Society of Mechanical Engineers, June 2016, V04BT04A001. ISBN: 978-0-7918-4976-7.
- [29] S. Srinivasan et al. “Flame Dynamics During Combustion Instability in a High-Pressure, Shear-Coaxial Injector Combustor”. In: *Flow, Turbulence and Combustion* 94.1 (Jan. 2015), pp. 237–262. ISSN: 1573-1987.

- [30] G. Eggenpieler et al. “Modeling of Pollutant Formation near Lean Blow-Out in Gas Turbine Engines”. In: *Direct and Large-Eddy Simulation V*. Ed. by Rainer Friedrich et al. Dordrecht: Springer Netherlands, 2004, pp. 351–358. ISBN: 978-1-4020-2313-2.
- [31] William Sutherland. “LII. The Viscosity of Gases and Molecular Force”. In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 36.223 (Dec. 1893), pp. 507–531. ISSN: 1941-5982, 1941-5990.
- [32] Bonnie J McBride et al. *NASA Glenn Coefficients for Calculating Thermodynamic Properties of Individual Species*. Tech. rep. 2002.
- [33] Bruce E. Poling et al. *Properties of Gases and Liquids, Fifth Edition*. Fifth edition. McGraw-Hill’s AccessEngineering. New York, N.Y: McGraw-Hill Education, 2020. ISBN: 978-0-07-011682-5 978-0-07-149999-6.
- [34] Heinz Pitsch. “Large-Eddy Simulation of Turbulent Combustion”. In: *Annual Review of Fluid Mechanics* 38.1 (Jan. 2006), pp. 453–482. ISSN: 0066-4189, 1545-4479.
- [35] T. Poinsot et al. *Theoretical and Numerical Combustion*. Edwards, 2001. ISBN: 978-1-930217-05-8.
- [36] J. Smagorinsky. “General Circulation Experiments with the Primitive Equations: I. The Basic Experiment”. In: *Monthly Weather Review* 91.3 (Mar. 1963), pp. 99–164. ISSN: 1520-0493, 0027-0644.
- [37] Vaidyanathan Sankaran. “Sub-Grid Combustion Modeling for Compressible Two-Phase Flows”. PhD thesis. Georgia Institute of Technology, Nov. 2003.
- [38] Akira Yoshizawa et al. “A Statistically-Derived Subgrid-Scale Kinetic Energy Model for the Large-Eddy Simulation of Turbulent Flows”. In: *Journal of the Physical Society of Japan* 54.8 (Aug. 1985), pp. 2834–2839. ISSN: 0031-9015.
- [39] Erica Quadarella et al. “A Generalized Partially Stirred Reactor Model for Turbulent Closure”. In: *Proceedings of the Combustion Institute* 39.4 (Jan. 2023), pp. 5329–5338. ISSN: 1540-7489.
- [40] Sebastian Popp et al. “Assessing Multi-Regime Combustion in a Novel Burner Configuration with Large Eddy Simulations Using Tabulated Chemistry”. In: *Proceedings of the Combustion Institute* 38.2 (2021), pp. 2551–2558. ISSN: 15407489.
- [41] Eva-Maria Wartha et al. “Characteristic Chemical Time Scales for Reactive Flow Modeling”. In: *Combustion Science and Technology* 193.16 (Dec. 2021), pp. 2807–2832. ISSN: 0010-2202, 1563-521X.
- [42] J. Kuehne et al. “Analysis of Sub-Grid PDF of a Progress Variable Approach Using a Hybrid LES/TPDF Method”. In: *Proceedings of the Combustion Institute* 33.1 (Jan. 2011), pp. 1411–1418. ISSN: 1540-7489.
- [43] N. Peters. “The Turbulent Burning Velocity for Large-Scale and Small-Scale Turbulence”. In: *Journal of Fluid Mechanics* 384 (Apr. 1999), pp. 107–132. ISSN: 0022-1120, 1469-7645.
- [44] T. D. Butler et al. “A Numerical Method for Two Dimensional Unsteady Reacting Flows”. In: *Symposium (International) on Combustion* 16.1 (Jan. 1977), pp. 1503–1515. ISSN: 0082-0784.
- [45] Alan R. Kerstein et al. “Field Equation for Interface Propagation in an Unsteady Homogeneous Flow Field”. In: *Physical Review A* 37.7 (Apr. 1988), pp. 2728–2731.
- [46] M. Boger et al. “Direct Numerical Simulation Analysis of Flame Surface Density Concept for Large Eddy Simulation of Turbulent Premixed Combustion”. In: *Sym-*

- posium (International) on Combustion*. Twenty-Seventh Symposium (International) on Combustion Volume One 27.1 (Jan. 1998), pp. 917–925. ISSN: 0082-0784.
- [47] M. Oevermann et al. “Investigation of Autoignition under Thermal Stratification Using Linear Eddy Modeling”. In: *Combustion and Flame* 155.3 (Nov. 2008), pp. 370–379. ISSN: 0010-2180.
- [48] Alan R. Kerstein. “One-Dimensional Turbulence: Model Formulation and Application to Homogeneous Turbulence, Shear Flows, and Buoyant Stratified Flows”. In: *Journal of Fluid Mechanics* 392 (Aug. 1999), pp. 277–334. ISSN: 1469-7645, 0022-1120.
- [49] S. Arshad et al. “A Strategy for Large-Scale Scalar Advection in Large Eddy Simulations That Use the Linear Eddy Sub-grid Mixing Model”. In: *International Journal of Numerical Methods for Heat & Fluid Flow* 28.10 (Oct. 2018), pp. 2463–2479. ISSN: 0961-5539.
- [50] S. Menon et al. “The Linear-Eddy Model”. In: *Turbulent Combustion Modeling: Advances, New Trends and Perspectives*. Springer, 2011, pp. 221–247.
- [51] Tim Lackmann et al. “Investigation of Turbulence–Chemistry Interactions in a Heavy-Duty Diesel Engine with a Representative Interactive Linear Eddy Model”. In: *International Journal of Engine Research* 21.8 (Oct. 2020), pp. 1469–1479. ISSN: 1468-0874, 2041-3149.
- [52] R. W. Bilger. “Reaction Rates in Diffusion Flames”. In: *Combustion and Flame* 30 (Jan. 1977), pp. 277–284. ISSN: 0010-2180.
- [53] J. Floyd et al. “A Simple Model for the Filtered Density Function for Passive Scalar Combustion LES”. In: *Combustion Theory and Modelling* 13.4 (Sept. 2009), pp. 559–588. ISSN: 1364-7830.
- [54] Hrvoje Jasak. “Error Analysis and Estimation for the Finite Volume Method with Applications to Fluid Flows”. PhD thesis. Imperial College of Science, Technology and Medicine, June 1996.
- [55] Alan C. Hindmarsh et al. “SUNDIALS: Suite of Nonlinear and Differential/Algebraic Equation Solvers”. In: *ACM Transactions on Mathematical Software* 31.3 (Sept. 2005), pp. 363–396. ISSN: 0098-3500.
- [56] David G. Goodwin et al. *Cantera: An Object-Oriented Software Toolkit for Chemical Kinetics, Thermodynamics, and Transport Processes*. 2023.
- [57] I. Moulitsas et al. *Serial/Parallel Library for Generating Coarse Grids for Multigrid Methods*. Tech. rep. May 2001.
- [58] W. H. Calhoun et al. “Subgrid Modeling for Reacting Large Eddy Simulations”. In: *34th Aerospace Sciences Meeting and Exhibit*. Reno, NV, U.S.A.: American Institute of Aeronautics and Astronautics, Jan. 1996.
- [59] Alan R Kerstein. “Reduced Numerical Modeling of Turbulent Flow with Fully Resolved Time Advancement”. In: ().
- [60] Konduri Aditya et al. “DNS of a Turbulent Premixed Flame Stabilized over a Backward Facing Step”. In: ().
- [61] A Sjunnesson et al. “Validation Rig- A Tool for Flame Studies”. In: *10th International Symposium on Air Breathing Engines*. Nottingham, England, 1991, pp. 385–393.
- [62] A. Sjunnesson et al. *LDA Measurements of Velocities and Turbulence in a Bluff Body Stabilized Flame*. Tech. rep. Dec. 1991.

- [63] A. Sjunnesson et al. “CARS Measurements and Visualization of Reacting Flows in a Bluff Body Stabilized Flame”. In: *28th Joint Propulsion Conference and Exhibit*. American Institute of Aeronautics and Astronautics.
- [64] A. Ghani et al. “LES of Longitudinal and Transverse Self-Excited Combustion Instabilities in a Bluff-Body Stabilized Turbulent Premixed Flame”. In: *Combustion and Flame* 162.11 (Nov. 2015), pp. 4075–4083. issn: 00102180.
- [65] N. Zettervall et al. “Large Eddy Simulation of a Premixed Bluff Body Stabilized Flame Using Global and Skeletal Reaction Mechanisms”. In: *Combustion and Flame* 179 (May 2017), pp. 1–22. issn: 0010-2180.
- [66] François Pellegrini et al. “Scotch: A Software Package for Static Mapping by Dual Recursive Bipartitioning of Process and Architecture Graphs”. In: *High-Performance Computing and Networking*. Ed. by Heather Liddell et al. Berlin, Heidelberg: Springer, 1996, pp. 493–498. isbn: 978-3-540-49955-8.
- [67] David Butz et al. “Local Flame Structure Analysis in Turbulent CH₄/Air Flames with Multi-Regime Characteristics”. In: *Combustion and Flame* 210 (Dec. 2019), pp. 426–438. issn: 00102180.