A generalised multiple mapping conditioning approach for turbulent combustion

M.J. Cleary, A.Y. Klimenko *

Division of Mechanical Engineering, University of Queensland, Australia

Article

---

* Corresponding author: A.Y. Klimenko, Division of Mechanical Engineering, University of Queensland, AUSTRALIA 4072. fax: +61 (0)7 3365 4799
Email address: a.klimenko@uq.edu.au (A.Y. Klimenko).

Published in Flow, Turb. Comb. 82(4), pp. 477-491, 2009 :: DOI 10.1007/s10494-008-9161-3

Abstract

This paper follows the evolution in understanding of the multiple mapping monditioning (MMC) approach for turbulent combustion and reviews different implementations of MMC models. As the MMC name suggests, the original version represents a consistent combination of CMC-type conditional equations (conditional moment closure) and generalised mapping closure. It seems that the strength of the MMC model, and especially that of its stochastic version, lies in a more general (and much more transparent) interpretation. In this new generalised interpretation, we can replace complicated derivations by physical reasoning and the model appears to be a natural extension of modelling approaches developed in recent decades. MMC can be seen as a methodology for enforcing certain known characteristics of turbulence on a conventional mixing model. This is achieved by localising the mixing operation in a reference space. The reference space variables are selected to emulate the properties of a turbulent flow which have a strong effect on reactive quantities. The best and simplest example is an MMC model which has a single reference variable emulating the mixture fraction. In diffusion flames turbulent fluctuations of reacting quantities are strongly correlated with fluctuations of the mixture fraction. By making mixing local in the reference mixture fraction space a CMC-type mixing closure is enforced. In the original interpretation of MMC the reference variables are modelled as Markov processes. Since the reference variables should emulate properties of turbulent flows as realistically as possible the next step, and the basis of generalised MMC, is to remove the Markovian restriction and set reference variables equal to traced Lagrangian quantities within DNS or LES flow fields. Indeed, no Markov value can emulate the mixture fraction better than the mixture fraction itself. (Using a Markov vector process of dimension higher than the number of conditioning variables represents a more economical alternative for producing reference variables in generalised MMC.) The generalised MMC approach effectively incorporates the mixture fraction-based models, the PDF methods and LES/DNS techniques into a single methodology with possibility of blending useful features developed previously in conventional models. The generalised approach to MMC stimulates a more flexible understanding of simulations using sparsely placed Lagrangian particles as tools that may provide accurate joint distributions of reactive scalars at relatively low computational cost. The physical reasoning behind the new interpretation of MMC is supported by example computations for a partially premixed methane / air diffusion flame (Sandia Flame D). The scheme utilises LES for the dynamic field and a sparse-Lagrangian filtered density function method with MMC mixing for the scalar field. Two different particle mixing schemes are tested. Simulations are performed using only 35,000 Lagrangian particles (of these only 10,000 are chemically active) on a single workstation. The relatively low computational cost allows the use of realistic chemical kinetics containing 34 reactive species and 219 reactions.

Key words: multiple mapping conditioning, Lagrangian simulations, sparse modelling

Published in Flow, Turb. Comb. 82(4), pp. 477-491, 2009 :: DOI 10.1007/s10494-008-9161-3

1 INTRODUCTION

In spite of the rapid progress in computer technology in recent decades, direct numerical simulations (DNS) of fluid flows remains a challenging problem for many practical applications. Simulating reacting flows is even more difficult: in addition to solving for the three velocity components and pressure correction, realistic chemical kinetics requires solutions for tens or, maybe, hundreds of reacting species. Realistic chemical kinetics are always stiff and evaluation of the species is computationally expensive even at a single point. The computer power, which is needed to fully resolve the reactive scalar fields, dramatically exceeds the computer power, which is required for velocity-only simulations. If we can afford to achieve a certain resolution of the reacting species scalar fields, simulating the velocity field with the same resolution is also affordable. Thus, from the perspective of computational affordability, one may have to deal with a well resolved (or even fully resolved) dynamic (velocity) field and a coarser representation of the reacting species on a relatively sparse set of points.

The requirements dictated by the physical characteristics of chemical reactions in turbulent flows do not seem to match these affordability conditions: reactions often occur within a thin reaction zones or are affected by small-scale variations of the flow parameters. A coarse representation of the reacting species seems inadequate. The dynamic field of turbulent flows is characterised by a wide range of scales. Due to a degree of universality of the inertial interval of turbulence, direct simulation of the smallest details of the dynamic field can be avoided and replaced by subgrid modelling as is done in large eddy simulations (LES). In non-premixed combustion considered here, the effect of chemical reactions on dynamics of turbulence is limited: reactions do not radically change the properties of small-scale turbulence and do not invalidate application of conventional LES approaches to the velocity field. Thus there is a fundamental contradiction in conventional LES-DNS simulations of reacting flows: we can afford well-resolved simulations of the velocity field and only coarse simulations of multiple species comprising a realistic chemical kinetic mechanism, while the opposite conditions are generally needed to provide accurate simulations.

The contradiction between availability of computational resources and the need to resolve small-scale properties of reacting scalars is treated differently in different models, which can be broadly divided into three groups: the mixture-fraction-based models (fast chemistry[1], flamelet[2] and conditional moment closure or CMC[3]), the probability density function (PDF) models[4, 5] and

LES/DNS[6]. The mixture-fraction-based models are economical and provide a good resolution in the mixture fraction space but are restrictive (formally or practically) to cases which can be effectively parameterised by the mixture fraction. The PDF and mixture fraction models can also be used for subgrid evaluation of the reactive species as part of LES [6–9]. The subgrid version of PDF models is called filtered density function (FDF) [6, 7]. These three groups of combustion models were recently augmented by the multiple mapping conditioning (MMC) framework [10–18] that, arguably, has links with all three groups discussed above. It seems that the major role of the MMC approach is in bringing a new understanding into non-premixed combustion modelling. This understanding is explored in the present work.

2 STRONG AND WEAK APPROXIMATIONS

Stochastic simulations may approximate a realistic process in two different ways – strong and weak. In the case of a strong approximation, the simulation approaches the full stochastic realisation of the physical process while a weak approximation matches only stochastic distributions of the physical process but not necessarily its realisation. For example, DNS provides a strong approximation of a turbulent field while PDF methods deal only with weak approximations of the reacting species. The FDF methods weakly approximate the overall distributions of reactive scalars and produce strong approximations only for large-scale structures of turbulent scalar fields. Any strong approximation is, obviously, a weak approximation at the same time but not vice versa. In general, achieving a strong approximation may require much more extensive computational resources compared to a weak approximation.

The terminology of weak and strong convergence is widely used in mathematics but we adopt a less rigorous and more intuitive use of these terms. A numerical scheme may work as a weak approximator and a strong approximator at the same time. In this case, the quality of the weak approximation may be higher than the quality of the strong approximation produced by the same scheme. In this work, we argue that achieving a complete strong approximation of the reactive species fields is, in many cases, practically excessive and computationally prohibitive. Equivalent distributions of reactive species that are targeted in PDF methods are perfectly sufficient to characterise all the complexity of turbulent combustion processes (provided, of course, that these distributions can be accurately simulated).

3 DIRECT LAGRANGIAN SIMULATIONS

In Lagrangian simulations of a turbulent reacting flow, fluid particles move convectively and each particle carries a full set of species. The concentrations of species at each particle change due to chemical reactions within the particle and due to mixing between different particles. If molecular diffusion is to be incorporated into this consideration, the fluid particles are replaced by diffusing particles, which experience small random fluctuations of their position accounting for molecular diffusion. This physical picture of Lagrangian particles carrying concentrations of reacting species and mixing with each other may serve as a foundation for a numerical scheme simulating transport of reacting species. By default, we assume that a fully resolved velocity field is available in simulations. The number of Lagrangian particles must be sufficiently large and mixing between particles must comply with a set of physical constraints to simulate the scalar transport correctly. The physical constraints imposed on mixing, which ensure convergence of the Lagrangian simulations to the realistic scalar fields, are discussed in Ref. [18]. Here, we specifically mention only one of the constraints – localness. In a physical flow and in its numerical model, two particles can be mixed only if they are located close to each other. Enforcing this condition is important to ensure convergence to fully resolved scalar fields. Practically, localness requires a large number of particles in the simulations. In principle, Lagrangian simulations can resolve scales larger than the characteristic distance between particles although it is the characteristic distance between mixed particles that separates the resolved and Lagrangian subgrid scales [18]. Lagrangian simulations which resolve all the smallest scales in a turbulent flow are, in fact, Lagrangian DNS.

Under conditions of increasing numbers of particles and progressively enforced localness of mixing, the specific algorithm used in a mixing model becomes a secondary issue: all reasonable mixing models will behave in a similar manner as the limit of a fully resolved scalar fields is approached (this trend was indeed observed in practical simulations[19].) This limit provides a benchmark for using a "brute force" option in Lagrangian subgrid simulations: if the required number of particles in a method is comparable to (or exceeds) the number of particles required to fully resolve the scalar fields, the method does not offer any savings of computer time and we probably would be better off with complete DNS of all involved scalars.

Even if convergence of Lagrangian simulations to fully resolved scalar fields can not be achieved in practical applications, it has implications for comparing different mixing models. An inferior model can perform better if it uses more particles in the simulations. A good mixing scheme should be able to perform well with relatively fewer particles. Two mixing models can be directly compared only if they use a similar number of particles, otherwise we compare

computational performances rather than the mixing models themselves. The number of particles determines how well the scalar fields are resolved. It could be the case that different models (or different versions of a specific model) may be specifically adapted to deal with a high or low degree of resolution of the scalar fields. It is not inconceivable that even the most sophisticated sparse mixing models may appear to be inadequate to treat most complex combustion cases with a small number of particles. The intensive option of improving simulations by increasing the number of particles is always available for these cases.

It seems, however, that achieving the limit of Lagrangian DNS is more of a hypothetical than practical possibility as this would require enormous computational resources. High quality practical simulations can be performed with a much smaller number of particles. We use the term "sparse-Lagrangian simulation" to refer to simulations with fewer Lagrangian particles than Eulerian grid points while the opposite case of having many Lagrangian particles in a single Eulerian cell can be called "intensive-Lagrangian simulations". It is inevitable, of course, that the sparse simulations can not resolve details of the scalar fields between particles and, thus, can not provide a strong approximation for the small-scale features of these fields. The sparse simulations can characterise the large-scale structures in realisations of the scalar fields but not the small-scale structures. We will argue, however, that sparse simulations may still provide reasonable weak approximations for the reacting flow. These arguments can be deduced from consideration of the MMC approach.

4 THE MMC APPROACH

The MMC approach was originally suggested [10] as an effective combination of CMC, which is used for evaluation of reactive scalars, and generalised mapping closure (MC), which is used for consistent modelling of the conditional dissipation and PDFs of the conditioning variables. The MMC approach allows for a stochastic implementation that converts this model into a PDF model. The MMC links with CMC and MC are obvious but as a PDF method MMC has a notable degree of similarity with the EMST (Euclidean minimum spanning trees) model [20] – both models use the MC concept and localise particle interactions but implement these ideas differently. The MMC uses reference variables simulating properties of turbulent flow that are different from reactive scalars but deemed to be important for combustion processes. The original MMC approach requires that the group of reference variables forms a Markov family (in practice this means that these variables can be simulated by a system of stochastic Ito equations). The MMC-Curl version [14, 15] of the stochastic MMC model, which uses Curl's mixing, has many common properties with the original and modified Curl's models [21]. The MMC model

is a versatile model that can emulate Curl's model, CMC or any effective combination of these two models. In the same way, the MMC-IEM version of model [13, 14], which uses IE(C)M mixing, has many similarities with the IE(C)M model [22–25]. The authors' understanding of MMC has evolved to some extent since the original formulation of the model was suggested [10]. The main difference from old MMC is that the Markov restrictions are no longer applied to the reference variables in the generalised version of the model. This new more general understanding is discussed below.

Simulated mixing between particles (Curl's mixing, for example) may satisfy several important mixing constraints but is different from physical mixing occurring in a turbulent flow. We can, however, enforce additional conditional properties on the mixing model by emulating a selected Lagrangian property (or properties) $\xi(t)$ of turbulence, and requiring that mixing is localised in the ξ-space. In MMC, ξ is called a reference variable and it does not coincide with any of the simulated stochastic reactive scalars. The selection of the reference variable is meaningful only if the property represented by this variable has a strong effect on combustion processes, otherwise the localisation procedure would have trivial results with reactive scalars being independent of ξ. In non-premixed combustion, the most logical choice for ξ is the mixture fraction but this variable must be obtained separately from the mixture fraction Z that is simulated stochastically along with the reactive scalars. (in general, MMC can also use reference variables effectively representing fluctuations of velocity, scalar dissipation, sensible enthalpy or other useful quantities [14, 16, 17]) Thus, MMC has two mixture fractions, one is the mixture-fraction-like reference variable, ξ, and another one is the simulated mixture fraction, Z, which is treated as a real mixture fraction for the purposes of evaluating the chemical source terms. MMC ensures that mixing does not disturb the values of reactive scalars, Y, conditioned on the reference variable. The required expectations of reactive scalars conditioned on the values of reference variables are enforced by accurate emulation of the reference variables [14].

In conventional MMC [10], a Markov process is used to simulate ξ. This corresponds to a diffusion-type approximation of conditional fluxes of reactive scalars in both ξ-space and Z-space [10, 13]. The diffusion-type approximation of the fluxes in the mixture fraction space is effectively CMC. Thus, the MMC model with Markov emulation of the mixture-fraction-type reference variable enforces CMC treatment of conditional expectations on conventional mixing models. It should be noted that in the true MMC regime, the intensity of mixing is not determined by the dissipation rate of the mixture fraction. The intensity of mixing should be determined to match the expected value of conditional fluctuations $Y'' = Y - \langle Y|Z\rangle$. Accurate simulation of conditional fluctuations is important in reacting flows close to extinction [13–15, 18], is not automatically enforced by the MMC model and requires proper selection of the intensity of MMC mixing. This issue is specifically addressed in Section 6.

If we consider real mixture fraction evaluated along trajectories of Lagrangian particles, these stochastic processes do exhibit some features of a Markov process but in a real flow they never become fully Markovian. In general, the reference variable does not have to be represented by a Markov process and removing this restriction corresponds to a generalised understanding of MMC [14, 18]. A better emulation of the mixture-fraction reference variable $\xi(t)$ (i.e. closer to real Lagrangian trajectories in the mixture-fraction space) makes it easier for MMC to simulate the properties of Z correctly since the difference between ξ and Z are kept small. Emulation of the process $\xi(t)$ can be improved by using a Markov vector process of higher dimension while selecting only some of the Markov variables for MMC conditioning. It seems, however, the best logical way to emulate $\xi(t)$ is to trace Lagrangian particles in a realistic mixture fraction field produced by DNS. The best approximation for any process is the process itself! The DNS-simulated mixture fraction is treated in MMC not as a actual mixture fraction but as a reference variable. The DNS (reference) mixture fraction is used only for localisation of the mixing operation: two particles are allowed to mix only if they are close to each other in the DNS mixture fraction space and in physical space. LES can also produce a good reference variable for MMC that is different from DNS-produced variable only in the range of subgrid scales.

5 WHAT DO WE SIMULATE?

The answer for this question may seem trivial – turbulent combustion, of course. It should be noted, however, that various answers for this question are possible and different answers can be correct at the same time. MMC is a weak simulator. From an MMC perspective, we simulate only distributions of reactive scalars but not the scalar fields. The velocity and the reference mixture fraction fields are needed only to produce realistic Lagrangian reference variables which, in principle, could have been produced in some other way without calculating the whole fields. The reference variables are used to enforce desired conditional behavior on the reactive scalars while the conditional variance around these conditional values is adjusted in a more empirical way to match the overall joint PDFs. As a result of MMC simulations we obtain only distributions but not the fields and we wish to see that these distributions are as realistic as possible. The DNS/LES perspective on the same model is different. The strong approximations are provided only for the fields of velocity and reference mixture fraction. The reactive scalars, given only at a sparse set of Lagrangian particles, are under-resolved compared to the dynamic field. Strong approximations for the reactive scalars are obtained only at the largest scales. If we combine the MMC and DNS perspectives then the reactive scalar values given at a sparse set of points can be seen as being an FDF representa-

tion. The method provides a strong approximation for large-scale structures of reactive scalars and a weak approximation for the overall distributions of these scalars. This FDF understanding is not necessarily linked to MMC models but can be seen as a general attribute of all sparse-Lagrangian simulations.

Sparse-Lagrangian simulations are economical when dealing with realistic chemical kinetics. We can not expect that these simulations will produce complete information about the finest details of reacting scalar fields, but we should require that realistic joint distributions of the reactive scalars are produced. The MMC-DNS or MMC-LES models seem to be a natural choice for accurate and economical sparse-Lagrangian simulations.

6 EXAMPLE MMC-LES SIMULATION FOR A JET DIFFUSION FLAME

We demonstrate the philosophy of sparse-Lagrangian simulations with a generalised MMC mixing model by way of example calculations for a piloted methane / air jet diffusion flame (Sandia Flame D) [26, 27]. As the Lagrangian simulation of the scalar distributions requires only 35,000 particles (of these only 10,000 are chemically active) it is possible to use realistic detailed chemistry and perform calculations on a single processor. Flame D is considered to be an appropriate flame for the following reasons. Firstly, a detailed experimental database for this flame is available online [28] and has been widely disseminated. Secondly, Flame D has been the object of numerous modelling attempts including FDF simulations [29] and this allows us to judge the merits of our new method. Thirdly, as the flame exhibits low levels of local extinction then conditioning on a reference variable representing LES-simulated mixture fraction is sufficient.

6.1 The FDF model

The Lagrangian particle scheme is governed by the following system of stochastic differential equations

$$d\boldsymbol{x}^* = \left[\widetilde{\boldsymbol{u}} + \frac{1}{\overline{\rho}}\nabla\overline{\rho}\,(D + D_t)\right] dt + \sqrt{2\,(D + D_t)}\, d\boldsymbol{\omega}^* \tag{1}$$

$$d\boldsymbol{\phi}^* = \left[\boldsymbol{W}(\boldsymbol{\phi}^*) + \boldsymbol{S}(\boldsymbol{x}^*, \boldsymbol{\phi}^*)\right] dt \tag{2}$$

The asterisk superscript denotes stochastic values, $\widetilde{\boldsymbol{u}}$ are the filtered velocity vector, $\boldsymbol{W}$ is the reaction rate for the multi-scalar composition space $\boldsymbol{\phi}$, D is the molecular diffusivity, D_t is the Smagorinski turbulent diffusivity and $\boldsymbol{\omega}$

is a Gaussian Weiner process. Conditional scalar dissipation is simulated by the particle mixing operator denote by S. In addition to the evolving scalar composition, each particle carries a stochastic reference variable and mixing is enforced between Lagrangian particles with similar $\widetilde{\xi^*}$. Here we use a one-dimensional reference space set equal to the filtered mixture fraction from the LES (i.e. $\xi^* = \widetilde{f}^* = \widetilde{f}(\boldsymbol{x}^*, t)$). At each time step all particles in the domain are grouped in pairs without replacement. Particle pairs (p and q) are selected so that the normalised square distance

$$\hat{d}^2_{(p,q)} = \frac{1}{1+\lambda^2} \times \left[\sum_{j=1}^{3} \left(\frac{x_j^{*(p)} - x_j^{*(q)}}{L_x} \right)^2 + \lambda^2 \left(\frac{\widetilde{f}^{*(p)} - \widetilde{f}^{*(q)}}{L_f} \right)^2 \right] \tag{3}$$

is minimised. In Eq.(3) L_x and L_f are characteristic physical and reference scales, and λ is a parameter which determines the relative localisation in physical and reference spaces respectively. Pair selection is illustrated schematically in Fig.**??**. The contour plot shows a 2D region of the instantaneous $\widetilde{f}$-field and A, B, C and D are the locations of four Lagrangian particles. Each particle has a physical location, $\boldsymbol{x}^*$, and a reference location, $\widetilde{f}^*$. The latter is interpolated between LES cell centres. All particles in the field must be grouped in pairs but to illustrate the process we concentrate on the selection of a partner for particle A. Note that particle density is much higher in the simulations. If $\lambda = 0$ then localisation is in physical space only and particle B, being the closest in $\boldsymbol{x}$-space is selected. For sparse modelling localisation in physical space does not enforce compositional locality (note the wide $\widetilde{f}$-space separation of particles A and B). Thus reactive species predictions are not expected to be good. Our sparse simulations for Flame D with $\lambda = 0$ produce global extinction. For $\lambda = \infty$ localisation is in $\widetilde{f}$-space only and particle C is chosen to mix with particle A. This imposes a flamelet-type closure on the mixing model. Real mixing obviously occurs in physical space. For particle mixing to mimic real fluid mixing, pairs should be as physically local as possible without compromising compositional locality. This is achieved in Eq.(3) by choosing finite λ resulting in mixing between particles A and D. In the present simulations we set $\lambda = 1$. Further work is required to determine if modelling is sensitive to alternative (finite) λ.

Selected particle pairs mix linearly according to

$$\begin{aligned} \boldsymbol{\phi}^{*(p)\text{ new}} &= \boldsymbol{\phi}^{*(p)} + \alpha \left(\hat{\boldsymbol{\phi}}^{*(p,q)} - \boldsymbol{\phi}^{*(p)} \right) \\ \boldsymbol{\phi}^{*(q)\text{ new}} &= \boldsymbol{\phi}^{*(q)} + \alpha \left(\hat{\boldsymbol{\phi}}^{*(p,q)} - \boldsymbol{\phi}^{*(q)} \right) \end{aligned} \tag{4}$$

where $\hat{\boldsymbol{\phi}}^{*(p,q)}$ is the two particle average and α is the mixing extent ($\alpha = 1$ gives complete mixing, and $\alpha = 0$ gives no mixing). We test two different mixing schemes for α:

Scheme 1 – at each time-step all particles mix with extent $\alpha = \alpha_0$;
Scheme 2 – at each time-step all particles mix such that $\alpha = 1$ with probability α_0 and $\alpha = 0$ with probability $1 - \alpha_0$.

Here we denote $\alpha_0 \equiv \Delta t/\tau_L$. Scheme 2 is more conventional for Curl-type mixing while Scheme 1 was used in the first sparse-Lagrangian simulations [30]. According to ref.[14] the generation of conditional fluctuations Y"is proportional to $\langle\gamma^2\rangle / \langle\gamma\rangle^2$ where $\gamma = 2\alpha - \alpha^2$. Since $\langle\gamma^2\rangle / \langle\gamma\rangle^2 = 1$ for Scheme 1 and $\langle\gamma^2\rangle / \langle\gamma\rangle^2 = 1/\alpha_0 \geq 1$ for Scheme 2, the generation of conditional fluctuations is expectably greater for the second scheme. The mixing time scale, τ_L, controls the dissipation of minor fluctuations (i.e. fluctuations with respect to averages conditioned on $\widetilde{f}$). Minor fluctuations are mathematical artifacts only. It is the conditional fluctuations (i.e. conditioning with respect Z) which are the physical quantities we wish to predict. Although there is not a fixed ratio between minor and conditional fluctuations the trends for each are the same [13–15]. Lower values of τ_L reduce the extent of mixing resulting in increased levels of both minor and conditional fluctuations. Conversely greater τ_L reduces those fluctuations. Following convention τ_L is made proportional to the turbulence time scale, $\tau = k/\varepsilon$, where k is the turbulent kinetic energy and ε is its dissipation. For a certain physical distance between mixing particles, Δ_L, it follows that $\tau_L \sim (\Delta_L^2/\varepsilon)^{1/3}$. Similarly we can define the Eulerian mixing time scale, τ_E, which scales as $(\Delta_E^2/\varepsilon)^{1/3}$ where Δ_E is the LES filter width. τ_E is the mixing time scale associated with the reference mixture fraction and can be defined as the ratio of the filtered mixture fraction variance and its dissipation so that we additionally have the scaling $\tau_E \sim \Delta_E^2/2\,(D + D_t)$. With these scaling relations and assuming Δ_L and Δ_E are within the inertial interval where ε is constant, the model for τ_L is

$$\tau_L = C_L^{-1} \left(\frac{\Delta_L}{\Delta_E}\right)^{2/3} \frac{\Delta_E^2}{2\,(D + D_t)} \tag{5}$$

The value of the constant C_L which gives the best results will vary for different values of the localisation parameter λ. In the case of Scheme 1 and $\lambda = 0$, a reasonable agreement with experimental data occurs for $C_L = 5$. For $\lambda = 1$ the passive scalar predictions are less sensitive to changes in C_L and values between 25 and 500 are possible. The reactive scalars, in particular their conditional fluctuations with respect to Z, are sensitive to the choice of C_L [30]. Conditional fluctuations in Flame D are significantly under-predicted with $C_L = 500$. For $C_L = 25$ conditional fluctuations are under-predicted by about half in the near nozzle region and only slightly under-predicted further downstream. In the present paper we use $\lambda = 1$ and $C_L = 25$ and compare results for mixing schemes 1 and 2.

6.2 Numerical details

The Eulerian filtered flow field equations are solved using the Flowsi LES code [31, 32]. A cylindrical domain 0.5m (70 jet diameters) in length and with radius 0.25m has 1024 axial, 55 radial and 32 azimuthal finite volume cells. The smallest cells are 0.5mm x 0.5mm x $\pi/32$ radians. Details of the discretisation and integration schemes, and the accuracy of velocity and filtered mixture fraction predictions for Flame D can be found in ref. [32].

For the composition field there are nominally 35,000 Lagrangian particles distributed within a domain extending 0.1m radially. To reduce computational effort, reaction rates are calculated for $0.05 < Z^* < 0.95$ leaving about 10,000 reacting particles. To ensure sufficient Lagrangian resolution in the high-temperature, low-density flame regions the numerical scheme uses weighted particles and a particle number control algorithm [29]. Particle pairs for the mixing model are selected to minimise $\hat{d}$ in Eq.(3) by way of a low-cost 'divide and conquer' method similar to the k-d tree [33]. The characteristic physical length scale, L_x is set to the experimental steady-state stoichiometric flame length and the corresponding value for L_f is $1-Z_{st}$. Chemical source terms are determined from a detailed kinetics scheme (GRI-3.0) [34] containing 34 reactive species and 219 reactions (NOx excluded) and radiation is modelled using the optically-thin assumption and absorption coefficients given in ref.[35]. As particle densities are stochastic values, direct density coupling back to the LES solver can be destabilising. Here density is obtained from a flamelet table and full coupling has not been achieved. This can be justified for the present flame case where local extinctions are low and flamelet approximations are reasonable. Simulations of flames with higher levels of local extinction will require density coupling which should be obtained through local conditional averaging or methods such as solving an Eulerian equivalent enthalpy equation [29].

Computations are made on an AMD64 Opteron workstation with two dual-core processors operating in parallel. Both the Eulerian and Lagrangian domains are divided into four regions and boundary information is passed between processors. After the fields reach steady-state samples are taken over approximately one domain flow-through requiring approximately 50 hours of computing time.

6.3 Results

Figures **??** through **??** show experimental and simulated scatter plots of temperature, and mass fractions of O_2 and H_2 versus mixture fraction, Z, at 7.5, 15 and 30 jet diameters downstream of the nozzle. MMC predictions are shown

for both mixing schemes 1 and 2. The first observation we make is that the overall trends are correctly predicted for the quantities shown and also for other major and minor reactive scalars. A closer analysis reveals that peak rich-side H_2 at $x/d = 7.5$ and 15 is over-predicted by about 40% which is consistent with earlier FDF [29] and filtered CMC [8] predictions for this flame. At $x/d = 30$ the peak H_2 is only slightly over-predicted. Although not shown, CO predictions are generally in very good agreement with experimental data.

At the three axial locations the level of conditional fluctuations measured in experiments are of similar magnitude. In contrast the two different mixing schemes yield increasing levels of conditional fluctuations with downstream distance. At $x/d = 7.5$ Scheme 1 under-predicts conditional fluctuations considerably while for Scheme 2 conditional fluctuations are of similar order to (or slightly greater than) experimental fluctuations. The rms of conditional fluctuations (not shown) by Scheme 1 is approximately half that of experiments at $x/d = 15$ while Scheme 2 over-predicts them by about 10 or 20%. Further downstream Scheme 1 performs better predicting the level of conditional fluctuations quite well at $x/d = 30$ while Scheme 2 produces approximately twice that level of local extinction.

A number of observations are made based on the scatter plots discussed above. The first observation is that MMC is a quality mixing model which is able to produce physically realistic results with a relatively small number of particles and thus at lower computational cost than conventional mixing models. Secondly, MMC is able to realistically model a range of turbulent flame conditions between those that are fully burning and flames with significant local extinction. The third observation is that for mixing governed by the time scale in Eq.(5), C_L does not appear to be universal; not even within the one flame. In addition the correct choice of C_L will vary for different mixing schemes. Given that τ_L is a (minor) dissipation time scale, a better approach [14] may also consider the generation of conditional fluctuations in its formulation. The current analysis suggests that Scheme 1, whereby all particles mix to a certain extent, generates approximately half the level of conditional rms than does Scheme 2, where a certain number of particles mix to a full extent and the remainder are unmixed.

As many other models, the presented version does not incorporate the influence of Reynolds number on combustion processes. The turbulent fields in these simulations are very far from being resolved to the level that can allow for direct incorporation of changes in Reynolds number. Nevertheless, it is possible to account for some of the effects that these changes may cause in combustion processes by adjusting the model. For example, higher Reynolds numbers correspond to larger fluctuations of the scalar dissipation and this is likely to generate larger conditional fluctuations of reactive species. In MMC simulations, this can be accounted for by relaxing τ_L. Effectively, we assume

that the minor dissipation time τ_L depends on the Reynold number.

The model presented here is dependent on the number of particles in its numerical implementation. In this implementation, the characteristic distance between mixed particles d_m is dependent on the characteristic distance between two closest particles as the used algorithm is aimed at selecting the closest particles for mixing (as determined by the distance (3)). The parameters of the model will change as number of particles increase and the model will approach DNS as the number of particles becomes very large (provided the velocity field is fully resolved in simulations) [18]. This scheme provides for most efficient use of the limited number of particles available in simulations. The version of MMC model, which is independent of the number of particles, can be constructed by fixing d_m. As the number of particles increases, a mixing partner for a particle is selected randomly among many particles located at a distance less than d_m from the particle. The limit achieved by this version of the model is independent of the number of particles but it does not coincide with the limit of DNS as the fields are smoothed by mixing over distances less than d_m and d_m does not tend to 0. This situation is very similar to the well-known dilemma of connecting or disconnecting the filter scale and the computational grid size in LES. In practice, the former choice is often preferred to provide the best possible resolution within the limits of given computational resources while the latter choice is consistent with the theoretical principle of independence of a model and its finite-difference implementation.

7 CONCLUSIONS

This paper follows the evolution in understanding of the MMC approach for turbulent combustion and reviews different implementations of MMC models. In stochastic form, MMC is a joint PDF method which enforces mixing between Lagrangian particles closely located in a reference space. The reference space variables are selected to emulate the properties of a turbulent flow which have a strong effect on reactive quantities, such as the mixture fraction. In the original interpretation of MMC the reference variables are modelled as Markov processes. Here a new generalised interpretation of MMC has been implemented where the reference variables are traced Lagrangian properties from a well resolved (LES) dynamic field. The most wonderful feature of the generalised understanding of MMC is that this model naturally blends and unifies nearly all major existing approaches to modelling of turbulent reacting flows (CMC, PDF, LES).

The quality of a mixing model determines the computational effort required. Localness of mixing is of particular importance. Lagrangian simulations resolve turbulence scales down to the level of separation between particles.

At the Lagrangian DNS limit, where all scales of turbulence are resolved, then all conventional mixing models become local. As a high quality mixing model, MMC can perform well with fewer particles making practical combustion applications with realistic chemistry computationaly affordable. Here, we have introduced sparse-Lagrangian simulations where the characteristic particle spacing is larger than the grid size of the DNS or LES dynamic field. With such a coarse representation of the reactive scalars we do not intend to model the finest details of their turbulent fields, but rather to accurately model the joint distributions of scalars. With this perspective the sparse-Lagrangian simulations represent an FDF model.

The generalised MMC concept is tested for a partially premixed jet diffusion flame (Sandia Flame D). The scheme utilises LES for the dynamic field and a sparse-Lagrangian MMC method for the scalars. The reference space consists of a single variable set equal to the LES filtered mixture fraction. Simulations are performed using only 35,000 Lagrangian particles (of these only 10,000 are chemically active) on a single workstation. The relatively low computational cost allows the use of realistic chemical kinetics containing 34 reactive species and 219 reactions. Two different particle-pair mixing schemes are tested: the first mixes all particle-pairs to a certain extent; and the second (which is more similar to conventional versions of Curl's model) mixes a certain number of particle-pairs fully and leaves the remaining pairs unmixed. The results show that sparse-Lagrangian MMC is able to produce realistic conditional distributions of reactive scalars at a relatively low computational cost. The mixing scheme can be tuned to produce varying levels of local extinction; here we showed that the second mixing scheme produces about twice the level of conditional rms fluctuations despite having the same minor dissipation time scale. Although nominating the optimal set of parameters that enforces the "correct" level of conditional fluctuations can be achieved only on the basis of extensive MMC/LES simulations of different test cases and is evidently premature at this stage, the presented work demonstrates the MMC ability of controlling these fluctuations in a very realistic manner and using a small number of particles.

8 Acknowledgements

This work is supported by the Australian Research Council.

References

[1] Bilger, R.W.: Turbulent flows with nonpremixed reactants. In: Libby, P.A., Williams, F.A. (eds.) Turbulent Reacting Flows, pp. 65–113. Springer Verlag, Berlin, New York (1980)

[2] Peters, N.: Turbulent Combustion. Cambridge University Press, Cambridge (2000)

[3] Klimenko, A.Y., Bilger, R.W.: Conditional moment closure for turbulent combustion. Prog. Energy Combust. Sci. **25**, 595–687 (1999)

[4] Pope, S.B.: PDF methods for turbulent reactive flows. Prog. Energy Combust. Sci. **11**, 119–192 (1985).

[5] Dopazo, C.: Recent developments in PDF methods. In: Libby, P.A., Williams, F.A. (eds.) Turbulent Reacting Flows, pp. 375–474. Academic Press, London, (1994)

[6] Pitsch, H.: Large-eddy simulations of turbulent combustion. Annu. Rev. Fluid Mech. **38**, 453–482 (2006)

[7] Colucci, P.J., Jaberi, F.A., Givi, P., Pope, S.B.: Filtered density function for large eddy simulation of turbulent reacting flows. Phys. Fluids. 10, 499–515 (1998)

[8] Navarro-Martinez, S., Kronenburg, A., Di Mare, F.: Conditional moment closure for large eddy simulations. Flow, Turbul. Combust. **75**, 245–274 (2005)

[9] Kim, S.H., Pitsch, H.: Conditional filtering method for large-eddy simulation of turbulent nonpremixed combustion. Phys. Fluids. **17**, 105103 (2005)

[10] Klimenko, A.Y., Pope, S.B.: A model for turbulent reactive flows based on multiple mapping conditioning. Phys. Fluids. **15**, 1907–1925 (2003)

[11] Klimenko, A.Y.: Modern modelling of turbulent non-premixed combustion and reduction of pollution emission. In: Clean Air VII. Lisbon, 2003

[12] Klimenko, A.Y.: MMC modelling and fluctuations of the scalar dissipation. In: Australian Combustion Symposium. Monash University, Melbourne, December 2003

[13] Klimenko, A.Y.: Matching the conditional variance as a criterion for selecting parameters in the simplest MMC models. Phys. Fluids. **16**, 4754–4757 (2004)

[14] Klimenko, A.Y.: Matching conditional moments in PDF modelling of non-premixed combustion. Combust. Flame. **143**, 369–385 (2005)

[15] Wandel, A.P., Klimenko, A.Y.: Testing MMC mixing for monte-carlo PDF simulations. Phys. Fluids. **17**, 128105 (2005)

[16] Cleary, M.J., Kronenburg, A.: Multiple mapping conditioning for extinction and reignition in turbulent diffusion flames. Proc. Combust. Inst. **31**, 1497 –1505 (2007)

[17] Cleary, M.J., Kronenburg, A.: Hybrid multiple mapping conditioning on passive and reactive scalars. Combust. Flame. **151**, 623–638 (2007)

[18] Klimenko, A.Y.: On simulating scalar transport by mixing between La-

grangian particles. Phys. Fluids. **19**, 31702 (2007)
[19] Bisetti, F., Chen, J.Y., Kempf, A., Janicka, J.: LES of Sandia/TUD Flame D with Eulerian PDF and finite-rate chemistry. In: Eighth International Workshop on Measurement and Computation of Turbulent Nonpremixed Flames. Heidelberg, 3–5 August 2006
[20] Subramaniam, S., Pope, S.B.: A mixing model for turbulent reactive flows based on Euclidean minimum spanning trees. Combust. Flame. **115**, 487–514 (1998)
[21] Curl, R.L., Miller, R.S., Ralph, J.L., Towell, G.D.: Dispersed phase mixing. AIChE J. **9**, 175–181 (1963)
[22] Dopazo, C.: Probability density function approach for a turbulent axisymmetric heated jet centreline evolution. Phys. Fluids. **18**, 397–410 (1975)
[23] Pope, S.B.: On the relationship between stochastic Lagrangian models of turbulence and second-moment closures. Phys. Fluids. **6**, 973–985 (1994)
[24] Fox, R.O.: On velocity–conditioned scalar mixing in homogeneous turbulence. Phys. Fluids. **8**, 2678–2691 (1996)
[25] Sawford, B.: Micro-mixing modelling of scalar fluctuations for plumes in homogeneous turbulence. Flow, Turbul. Combust. **72**, 133–160 (2004)
[26] Barlow, R.S., Frank, J.: Effects of turbulence on species mass fractions in methane/air jet flames Proc. Combust. Inst. **27**, 1087–1095 (1998)
[27] Schneider, Ch., Dreizler, A., Janicka, J., Hassel, E.P.: Flow field measurements of stable and locally extinguishing hydrocarbon-fuelled jet flames. Combust. Flame. **135**, 185–190 (2003)
[28] International Workshop on Measurement and Computation of Turbulent Nonpremixed Flames. http://public.ca.sandia.gov/TNF//abstract.html
[29] Raman, V., Pitsch, H.: A consistent LES/filtered-density function formulation for the simulation of turbulent flames with detailed chemistry. Proc. Combust. Inst. **31**, 1711–1719 (2007)
[30] Cleary, M.J., Klimenko, A.Y., Janicka, J., Pfitzner, M.: LES / sparse-Lagrangian MMC in a piloted jet flame. In: Australian Combustion Symposium. University of Sydney, December 2007
[31] Kempf, A., Sadiki, A., Janicka, J.: Prediction of finite chemistry effects using large eddy simulation. Proc. Combust. Inst. **29**, 1979–1985 (2002)
[32] Kempf, A., Flemming, F., Janicka, J.: Investigation of lengthscales, scalar dissipation, and flame orientation in a piloted diffusion flame by LES. Proc. Combust. Inst. **30**, 557–565 (2005)
[33] Friedman, J.H., Bentley, J.L., Finkel, R.A.: An algorithm for finding best matches in logarithmic expected time. ACM Transactions of Mathematical Software. **3**, 209–226 (1977)
[34] Smith, G.P., Golden, D.M., Frenklach, M., Moriarty, N.W., Eiteneer, B., Goldenberg, M., C. Bowman, T., Hanson, R.K. Song, S. Gardiner, W.C.Jr., Lissianski, V.V., Qin, Z. GRI-mech 3.0. http://www.me.berkeley.edu/gri_mech/
[35] Barlow, R.S., Karpetis, A.N., Frank, J.H., Chen, J.Y.: Scalar profiles and

NO formation in laminar opposed-flow partially premixed methane/air flames. Combust. Flame. **127**, 2101–2118 (2001)

9 List of Figure Captions

(1) Schematic of generalised MMC localisation. The contour plot is a 2D region of the MMC reference field given by the filtered mixture fraction. Points A through D denote the location of four Lagrangian particles.
(2) Scatter plots of temperature, O_2 and H_2 versus Z at $x/d = 7.5$. Experimental data – left; MMC with mixing scheme 1 – middle; MMC with mixing scheme 2 – right.
(3) Scatter plots of temperature, O_2 and H_2 versus Z at $x/d = 15$. Experimental data – left; MMC with mixing scheme 1 – middle; MMC with mixing scheme 2 – right.
(4) Scatter plots of temperature, O_2 and H_2 versus Z at $x/d = 30$. Experimental data – left; MMC with mixing scheme 1 – middle; MMC with mixing scheme 2 – right.

Figures

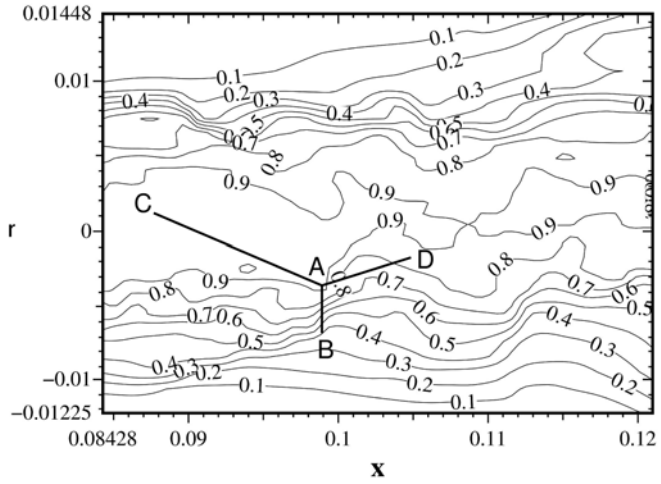

Fig. 1. Schematic of generalised MMC localisation. The contour plot is a 2D region of the MMC reference field given by the filtered mixture fraction. Points A through D denote the location of four Lagrangian particles.

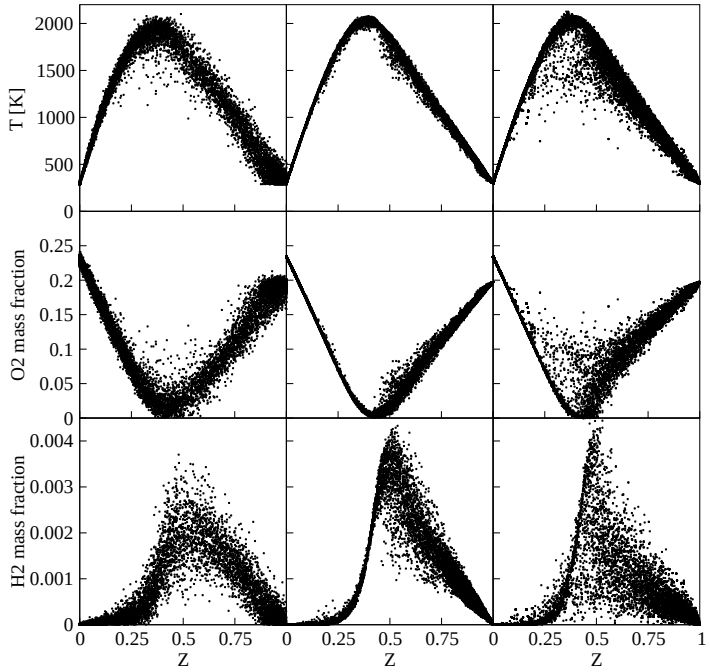

Fig. 2. Scatter plots of temperature, O_2 and H_2 versus Z at $x/d = 7.5$. Experimental data – left; MMC with mixing scheme 1 – middle; MMC with mixing scheme 2 – right.

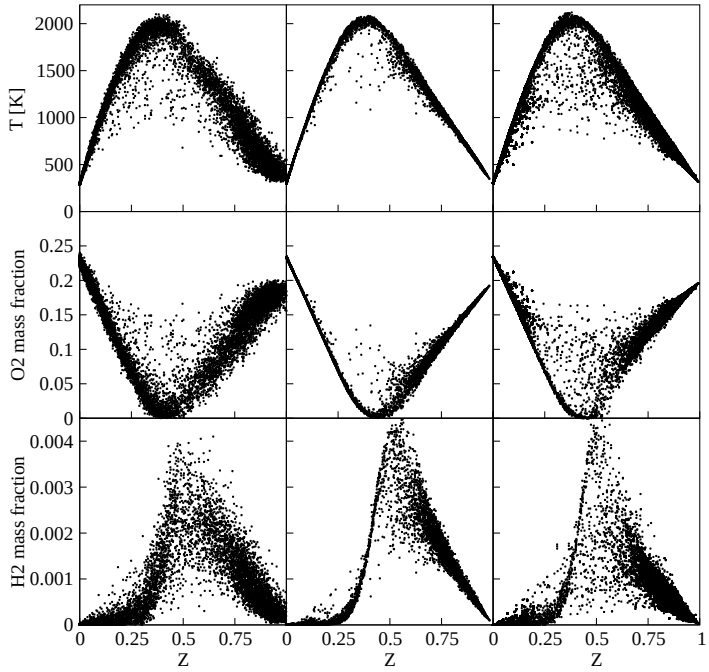

Fig. 3. Scatter plots of temperature, O_2 and H_2 versus Z at $x/d = 15$. Experimental data – left; MMC with mixing scheme 1 – middle; MMC with mixing scheme 2 – right.

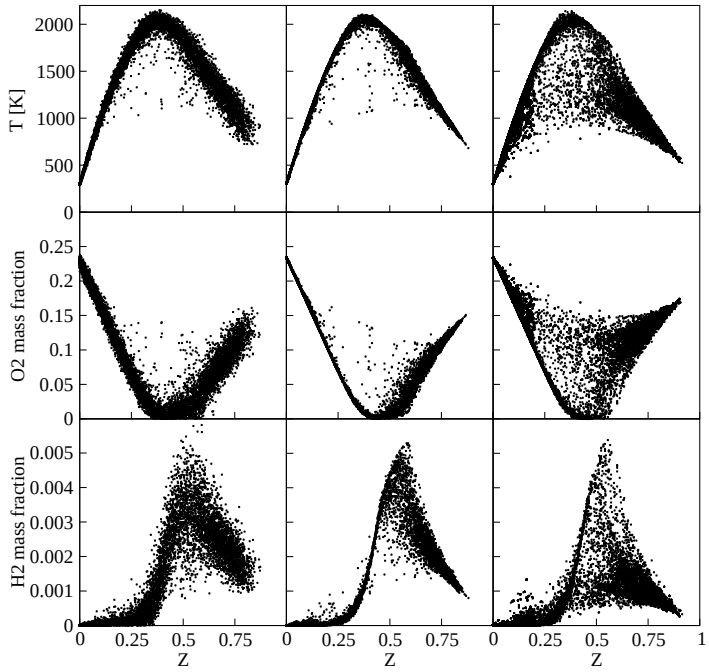

Fig. 4. Scatter plots of temperature, O_2 and H_2 versus Z at $x/d = 30$. Experimental data – left; MMC with mixing scheme 1 – middle; MMC with mixing scheme 2 – right.