



A fast, low-memory, and stable algorithm for implementing multicomponent transport in direct numerical simulations

Aaron J. Fillo^a, Jason Schlup^b, Guillaume Beardsell^c, Guillaume Blanquart^c,
Kyle E. Niemeyer^{a,*}

^a School of Mechanical, Industrial, and Manufacturing Engineering, Oregon State University, Corvallis, OR 97331, USA

^b Graduate Aerospace Laboratories, California Institute of Technology, Pasadena, CA, USA

^c Department of Mechanical Engineering, California Institute of Technology, Pasadena, CA, USA

ARTICLE INFO

Article history:

Received 13 May 2019

Received in revised form 27 October 2019

Accepted 11 December 2019

Available online 17 December 2019

Keywords:

Turbulent flames

Direct numerical simulation

Multicomponent diffusion

Mixture-averaged diffusion

ABSTRACT

Implementing multicomponent diffusion models in reacting-flow simulations is computationally expensive due to the challenges involved in calculating diffusion coefficients. Instead, mixture-averaged diffusion treatments are typically used to avoid these costs. However, to our knowledge, the accuracy and appropriateness of the mixture-averaged diffusion models has not been verified for three-dimensional turbulent premixed flames. In this study we propose a fast, efficient, low-memory algorithm and use that to evaluate the role of multicomponent mass diffusion in reacting-flow simulations. Direct numerical simulation of these flames is performed by implementing the Stefan–Maxwell equations in NGA. A semi-implicit algorithm decreases the computational expense of inverting the full multicomponent ordinary diffusion array while maintaining accuracy and fidelity. We first verify the method by performing one-dimensional simulations of premixed hydrogen flames and compare with matching cases in Cantera. We demonstrate the algorithm to be stable, and its performance scales approximately with the number of species squared. Then, as an initial study of multicomponent diffusion, we simulate premixed, three-dimensional turbulent hydrogen flames, neglecting secondary Soret and Dufour effects. Simulation conditions are carefully selected to match previously published results and ensure valid comparison. Our results show that using the mixture-averaged diffusion assumption leads to a 15% under-prediction of the normalized turbulent flame speed for a premixed hydrogen-air flame. This difference in the turbulent flame speed motivates further study into using the mixture-averaged diffusion assumption for DNS of moderate-to-high Karlovitz number flames.

© 2019 Elsevier Inc. All rights reserved.

1. Introduction

Implementing full multicomponent mass diffusion transport in direct numerical simulation (DNS) can be memory intensive and computationally expensive. This is because calculating diffusion fluxes requires point-wise knowledge of the multicomponent diffusion coefficient matrix, which scales with the number of chemical species squared [1]. The unity Lewis number, non-unity Lewis number, and mixture-averaged diffusion assumptions have been used to reduce the costs associated with mass diffusion by approximating the full diffusion coefficient matrix as a constant scalar value, a constant

* Corresponding author.

E-mail address: kyle.niemeyer@oregonstate.edu (K.E. Niemeyer).

vector, and a matrix diagonal, respectively. In addition, several approaches further reduce the system's complexity by approximating multicomponent diffusion processes in terms of equivalent Fickian processes, such as those used by Warnatz [2] and Coltrin et al. [3]. However, to our knowledge, the accuracy and appropriateness of these assumptions have not been evaluated in turbulent reacting flows against multicomponent diffusion transport due to its high computational expense and a dearth of affordable computing tools.

As further motivation for this study, Lapointe and Blanquart [4] recently investigated the impact of differential diffusion on simulations using unity and nonunity Lewis number approximations. They reported that methane, *n*-heptane, iso-octane, and toluene flames have similar normalized turbulent flame speeds and fuel burning rates when neglecting differential diffusion, but flames using the nonunity Lewis number approximation underpredict the normalized flame speed when including differential diffusion due to reduced burning rates [4]. Building on these results, Burali et al. [5] evaluated the relative accuracy of the nonunity Lewis number assumption relative to mixture-averaged diffusion for lean, unstable hydrogen/air flames; lean, turbulent *n*-heptane/air flames; and ethylene/air coflow diffusion flames. They demonstrated that the relative error associated with the nonunity Lewis number assumption could be minimized with careful selection of the Lewis number vector for a wide range of flames [5]. Similarly, Schlup and Blanquart [6] examined the impact of multicomponent thermal diffusion on DNS of turbulent, premixed, high-Karlovitz hydrogen/air flames. They showed that simulations using the mixture-averaged thermal diffusion assumption underpredict the normalized flame speeds compared with results from simulations using full multicomponent thermal diffusion. In addition, including multicomponent thermal diffusion results in increased production of chemical source terms in regions of high positive curvature [6]. These observed discrepancies in similar flame simulations with different diffusion models warrant a detailed investigation of the fundamental transport phenomena involved.

While data are sparse from three-dimensional reacting-flow simulations with multicomponent transport, several groups have investigated the effects of multicomponent transport in simpler configurations. These studies include one-dimensional [7–12] and two-dimensional flames [13–15] at various unburnt conditions. These works compared the multicomponent model with various levels of diffusion and transport property models, from constant Lewis number to mixture-averaged properties. In general, prior studies found some errors between multicomponent and mixture-averaged formulations for simplified hydrogen/air and methane/air flame configurations, such as unstretched laminar flames. However, these studies did not assess flames where diffusion effects may be more important, such as two- and three-dimensional, unsteady laminar and turbulent flames. Moreover, advancing clean and efficient combustion technology requires incorporating realistic fuel chemistry in large-scale turbulent simulations relevant to practical applications. Thus, there is a clear need for a computationally efficient algorithm capable of modeling full multicomponent diffusion transport [16].

The studies by Lapointe and Blanquart [4], Burali et al. [5], and Schlup and Blanquart [6] each took care to isolate the diffusion assumptions in question by neglecting higher-order terms that may affect diffusion transport. For example, with the exception of Schlup and Blanquart [6], these studies neglected Soret and Dufour diffusion, as it would be difficult to determine the direct cause of an observed effect when including both molecular and thermal diffusion. However, despite this methodical approach, the results of these studies were presented with reference to mixture-averaged diffusion, rather than full multicomponent diffusion. This further highlights the need for a computationally efficient method for implementing full multicomponent transport, and a subsequent examination of the differences between its “true” results and those resulting from the approximations conventionally used.

In this direction, several studies have examined the impact of full multicomponent transport on simplified three-dimensional flame configurations. Giovangigli [14] demonstrated that multicomponent Soret effects significantly impact a wide range of laminar hydrogen/air flames. Specifically, they noted that multicomponent Soret effects influence laminar flame speeds and extinction stretch rates for flat and strained premixed flames, respectively. For high-pressure systems, Borchesi and Bellan [17] developed and analyzed a multi-species turbulent mixing model for large-eddy simulations. They focused on turbulent crossflow mixing of a five-species combustion-relevant mixture of *n*-heptane, O₂, CO₂, N₂, and H₂O. The multi-species transport model significantly improves the accuracy and fidelity of the solution throughout the mixing layer; however, this study only considered non-reacting flows and, as a result, did not assess the impact of multicomponent transport on the chemistry inherent in turbulent combustion. In addition, these simulations implement a simplified diffusion model to approximate multicomponent diffusion but do not directly solve the diffusion terms present in the generalized conservation equations for species and energy [18].

Motivated by the dearth of affordable three-dimensional multicomponent transport models, Ern and Giovangigli [9,19,20] developed the computationally efficient Fortran library EGLIB for accurately determining transport coefficients in gas mixtures. More recently, Ambikasaran and Narayanaswamy [21] proposed an efficient algorithm to compute multicomponent diffusion velocities, which scales linearly with the number of species. This significantly reduces computational cost compared with previous methods that directly invert the Stephan–Maxwell equations and thus scale with the number of species cubed. Although both libraries reduce the computational cost of determining the multicomponent diffusion coefficients, they do not provide a method for reducing the resulting large memory requirements for multidimensional simulations.

Overall, these prior studies provide compelling evidence that multicomponent transport is important and can affect the accuracy of combustion models. However, none assessed how multicomponent transport impacts three-dimensional turbulent systems with detailed chemistry. In this article, we demonstrate and analyze an efficient, dynamic algorithm that reduces the computational expense of calculating the multicomponent diffusion fluxes. We demonstrate the approach is accurate and stable for a wide range of time-step sizes. In addition, we present a comprehensive assessment of the numerical

costs associated with this method. To verify the proposed algorithm we present one-dimensional freely propagating, laminar hydrogen/air flames and compare with results from Cantera. Finally, we simulate three-dimensional, turbulent, premixed, hydrogen/air flames. As a preliminary comparison of the mixture-averaged and multicomponent diffusion models, we perform an a posteriori assessment of how the choice of diffusion model impacts the turbulent statistics of the three-dimensional hydrogen simulation.

2. Governing equations

This section presents the low-Mach number reacting Navier–Stokes equations used in this study. In addition, this section outlines the method used to determine the mass diffusion fluxes for both the mixture-averaged and multicomponent approaches, abbreviated here as MA and MC, respectively.

2.1. Low Mach-number equations

In this work we solve the variable-density, low-Mach number, reacting-flow equations [22,23]. The conservation equations are

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0, \quad (1)$$

$$\frac{\partial \rho \mathbf{u}}{\partial t} + \nabla \cdot (\rho \mathbf{u} \otimes \mathbf{u}) = -\nabla p + \nabla \cdot \boldsymbol{\tau}, \quad (2)$$

$$\frac{\partial \rho T}{\partial t} + \nabla \cdot (\rho \mathbf{u} T) = \nabla \cdot (\rho \alpha \nabla T) + \rho \dot{\omega}_T - \frac{1}{c_p} \sum_i c_{p,i} \mathbf{j}_i \cdot \nabla T + \frac{\rho \alpha}{c_p} \nabla c_p \cdot \nabla T, \quad (3)$$

$$\frac{\partial \rho Y_i}{\partial t} + \nabla \cdot (\rho \mathbf{u} Y_i) = -\nabla \cdot \mathbf{j}_i + \dot{\omega}_i, \quad (4)$$

where ρ is the mixture density, $\mathbf{u}$ is the velocity vector, p is the hydrodynamic pressure, $\boldsymbol{\tau}$ is the viscous stress tensor, T is the temperature, α is the mixture thermal diffusivity, $c_{p,i}$ is the constant-pressure specific heat of species i , c_p is the constant-pressure specific heat of the mixture, $\mathbf{j}_i$ is the diffusion flux of species i , Y_i is the mass fraction of species i , and $\dot{\omega}_i$ is the production rate of species i . In Equation (3), the temperature source term $\dot{\omega}_T$ is given by

$$\dot{\omega}_T = -c_p^{-1} \sum_i h_i(T) \dot{\omega}_i, \quad (5)$$

where $h_i(T)$ is the specific enthalpy of species i as a function of temperature. The density is determined from the ideal gas equation of state

$$\rho = \frac{P_o W}{RT}, \quad (6)$$

where P_o is the thermodynamic pressure, R is the universal gas constant, and W is the mixture molecular weight determined via $W = \left(\sum_i^N Y_i / W_i \right)^{-1}$, where W_i is the molar mass of the i th species and N is the number of species.

The diffusion fluxes are calculated with either the mixture-averaged [1] or multicomponent [24] models, which are both based on Boltzmann's equation for the kinetic theory of gases [24,25]. The baro-diffusion term is commonly neglected in reacting-flow simulations under the low Mach-number approximation [26]. We have also neglected thermal diffusion because our objective in this work is to investigate the impact of mass diffusion models; Schlup and Blanquart [6,27] previously explored the effects of thermal diffusion modeling.

2.2. Mixture-averaged (MA) species diffusion flux

The i th species diffusion flux for the mixture-averaged diffusion model is related to the species gradients by a Fickian formulation and is expressed as

$$\mathbf{j}_i = -\rho D_{i,m} \frac{Y_i}{X_i} \nabla X_i + \rho Y_i \mathbf{u}'_c, \quad (7)$$

where X_i is the i th species mole fraction, $D_{i,m}$ is the i th species mixture-averaged diffusion coefficient as expressed by Bird et al. [1]:

$$D_{i,m} = \frac{1 - Y_i}{\sum_{j \neq i} X_j / \mathcal{D}_{ji}}, \quad (8)$$

where $\mathcal{D}_{ji}$ is the binary diffusion coefficient between the i th and j th species. Finally, $\mathbf{u}'_c$ is the correction velocity used to ensure mass continuity:

$$\mathbf{u}'_c = \sum_{i=1}^N D_{i,m} \frac{Y_i}{X_i} \nabla X_i. \quad (9)$$

The expression for species diffusion flux can be re-stated in terms of mass fraction Y_i as

$$\mathbf{j}_i = -\rho D_{i,m} \left(\nabla Y_i - Y_i \sum_{k=1}^N \nabla Y_k \frac{W}{W_k} \right) + \rho Y_i \mathbf{u}'_c, \quad (10)$$

where $D_{i,m}$ corresponds to the i th element of the diagonal mixture-averaged diffusion coefficient matrix, defined herein as $\mathbf{D}^{\text{MA}}$.

2.3. Multicomponent (MC) species diffusion flux

The multicomponent diffusion model for the i th species diffusion flux is

$$\mathbf{j}_i = \frac{\rho Y_i}{X_i W} \sum_{k=1}^N W_k D_{ik} \nabla X_k, \quad (11)$$

where D_{ik} is the ordinary multicomponent diffusion coefficient (computed using the `MCMDIF` subroutine of CHEMKIN II [28] with the method outlined by Dixon-Lewis [29]). Equation (11) can be restated in terms of mass fraction as

$$\mathbf{j}_i = \rho \sum_k -D_{ik}^{\text{MC}} \nabla Y_k, \quad (12)$$

where

$$D_{ik}^{\text{MC}} = -\frac{W_i}{W} \left[D_{ik} - \frac{W}{W_k} \left(\sum_j D_{ij} Y_j \right) \right]. \quad (13)$$

The diagonal of the ordinary multicomponent diffusion matrix, D_{ii} , is zero. As will be shown later, the $\mathbf{D}^{\text{MC}}$ matrix is singular with a kernel of dimension one. Interestingly, the vector of species mass fractions is in the kernel:

$$\sum_{k=1}^N D_{ik}^{\text{MC}} Y_k = -\frac{W_i}{W} \left[\left(\sum_k D_{ik} Y_k \right) - \left(W \sum_k \frac{Y_k}{W_k} \right) \left(\sum_j D_{ij} Y_j \right) \right] = 0. \quad (14)$$

This property will be important later (in Section 3.4) for the stability analysis.

The multicomponent diffusion coefficients, thermal conductivities, and thermal diffusion coefficients are computed by solving a system of equations defined by the $\mathbf{L}$ matrix, composed of nine sub-matrices:

$$\begin{bmatrix} \mathbf{L}^{00,00} & \mathbf{L}^{00,10} & \mathbf{0} \\ \mathbf{L}^{10,00} & \mathbf{L}^{10,10} & \mathbf{L}^{10,01} \\ \mathbf{0} & \mathbf{L}^{01,10} & \mathbf{L}^{01,01} \end{bmatrix} \begin{bmatrix} \mathbf{a}_1^{00} \\ \mathbf{a}_1^{10} \\ \mathbf{a}_1^{01} \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{X} \\ \mathbf{X} \end{bmatrix}, \quad (15)$$

where the right-hand side is composed of the one-dimensional mole fraction arrays $\mathbf{X}$. Based on this system of equations, the inverse of the $\mathbf{L}^{00,00}$ block provides the multicomponent diffusion coefficients:

$$D_{ij} = X_i \frac{16T}{25P} \frac{W}{W_j} (q_{ij} - q_{ii}), \quad (16)$$

where

$$\mathbf{q} = \left(\mathbf{L}^{00,00} \right)^{-1}. \quad (17)$$

The $\mathbf{L}^{00,00}$ sub-matrix block is given by

$$L_{ij}^{00,00} = \frac{16T}{25P} \sum_{k=1}^N \frac{X_k}{W_i \mathcal{D}_{ik}} \{ W_j X_j (1 - \delta_{i,k}) - W_i X_i (\delta_{i,j} - \delta_{j,k}) \}, \quad (18)$$

where $\delta_{i,j}$ is the reduced dipole moment corresponding to the i th component of the vector of dipole moments.

3. Methods

As discussed previously, multicomponent mass diffusion has not yet been incorporated into three-dimensional turbulent flame simulations due to its high computational expense. This section presents the discretized equations, numerical algorithm, and preconditioner proposed. The method is based on the semi-implicit time-marching scheme for species mass-fraction fields proposed by Savard et al. [23].

3.1. Multicomponent model implementation

This work was completed using the structured, multi-physics, and multi-scale finite-difference code NGA [22,23]. NGA can solve a wide range of problems, including laminar and turbulent flows [30–32], constant- and variable-density flows [22, 33,34], large-eddy simulation [31,35], and DNS [33,34,36]. NGA discretely conserves mass, momentum, and kinetic energy with an arbitrarily high-order spatial accuracy [22].

NGA's variable-density flow solver uses both spatially and temporally staggered variables, storing all scalar quantities (ρ , P , T , Y_i) at the volume centers and velocity components at their respective volume faces [22,37]. The convective term in the species transport equation is discretized using the bounded, quadratic, upwind-biased, interpolative convective scheme (BQUICK) [38]. The diffusion source term is discretized using a second-order centered scheme and the variables are advanced in time using a second-order semi-implicit Crank–Nicolson scheme [39].

An iterative procedure is applied to fully cover the nonlinearities in the Navier–Stokes equations and the species diffusion terms. Prior studies demonstrated this iterative process to be critically important for stability and accuracy [22,23,39,40]. Savard et al. [23] fully detailed the numerical algorithm sequence; we summarize this method here. This summary is independent of the preconditioning strategy employed in NGA, to which propose modifications in Section 3.2.

The algorithm for advancing one time step follows, using a uniform time-step size Δt . The density, pressure, and scalar fields are advanced from time level $t^{n+1/2}$ to $t^{n+3/2}$, and the velocity fields are advanced from time t^n to t^{n+1} , where t^n is the current time. Each iteration (i.e., time step) consists of Q sub-iterations and follows this procedure:

0. Upon convergence of the previous time step, the algorithm stores the density ($\rho^{n+1/2}$), pressure ($P^{n+1/2}$), velocity fields ($\mathbf{u}^n$), and scalar fields ($\mathbf{Y}^{n+1/2}$), where $\mathbf{Y}$ represents the vector of species mass fractions ($Y_1, \dots, Y_N$). The solutions for pressure, species mass fraction, and momentum from the previous time step are used as an initial guess for the iterative procedure:

$$\rho_0^{n+3/2} = \rho^{n+1/2}, \quad \mathbf{Y}_0^{n+3/2} = \mathbf{Y}^{n+1/2}, \quad \text{and} \quad (\rho \mathbf{u})_0^{n+1} = (\rho \mathbf{u})^n. \quad (19)$$

An Adams–Bashforth prediction evaluates the initial density:

$$\rho_0^{n+3/2} = 2\rho^{n+1/2} - \rho^{n-1/2}, \quad (20)$$

which ensures that the continuity equation is discretely satisfied at the beginning of the iterative procedure.

1. For the sub-iterations $k = 1, \dots, Q$, the semi-implicit Crank–Nicolson method advances the scalar fields in time [39,41]:

$$\begin{aligned} \rho_k^{n+3/2} \mathbf{Y}_{k+1}^{n+3/2} &= \rho^{n+1/2} \mathbf{Y}^{n+1/2} + \Delta t (\mathbf{C}_k^* + \mathbf{Diff}_k^* + \mathbf{\Omega}_k^*) \\ &+ \frac{\Delta t}{2} \left(\frac{\partial \mathbf{C}}{\partial \mathbf{Y}} + \frac{\partial \mathbf{Diff}}{\partial \mathbf{Y}} + \frac{\partial \mathbf{\Omega}}{\partial \mathbf{Y}} \right)_k^{n+1} \cdot (\mathbf{Y}_{k+1}^{n+3/2} - \mathbf{Y}_k^{n+3/2}), \end{aligned} \quad (21)$$

where $\mathbf{Diff} = -\nabla \cdot \mathbf{j}_i$ and $\mathbf{Y}_k^*$, $\mathbf{C}_k^*$, $\mathbf{Diff}_k^*$, and $\mathbf{\Omega}_k^*$ are the mass fraction, convection, diffusion, and chemical terms evaluated on the mid-point (or half time-step) scalar field $\mathbf{Y}_k^*$:

$$\mathbf{Y}_k^* = \frac{\mathbf{Y}^{n+1/2} + \mathbf{Y}_k^{n+3/2}}{2}. \quad (22)$$

To simplify the discrete notations for spatial differentiation, the operators corresponding to the convective and diffusive terms in Equation (4) are written as $\mathbf{C}$ and $\mathbf{Diff}$, respectively [23]. $\frac{\partial \mathbf{C}}{\partial \mathbf{Y}}$ and $\frac{\partial \mathbf{Diff}}{\partial \mathbf{Y}}$ are the Jacobian matrices corresponding to the convective and diffusive terms with respect to the species mass fractions, respectively. $\mathbf{C}$ and $\frac{\partial \mathbf{C}}{\partial \mathbf{Y}}$ are functions of the density and velocity, while $\mathbf{Diff}$ and $\frac{\partial \mathbf{Diff}}{\partial \mathbf{Y}}$ are functions of the density, diffusivity, and molar weight. They are consistently updated at each sub-iteration [23].

2. The density field, $\rho_{k+1}^{n+3/2}$, is evaluated from the new scalar fields using Equation (6). We do not rescale the scalar fields as proposed by Shunn et al. [40]. However, upon convergence of the sub-iterations, this method is equivalent to the density treatment they proposed [23].
3. The momentum equation is advanced in time using a similar semi-implicit Crank–Nicolson method for the scalar fields as described by Savard et al. [23].
4. A Poisson equation is then solved for the fluctuating hydrodynamic pressure using a combination of HYPRE [22,42], BICGSTAB [43], and/or FFTW [44]. The predicted velocity field is then updated.

5. Upon convergence of the sub-iterations, the solutions are updated.

The procedure summarized above becomes equivalent to the fully implicit Crank–Nicolson time-integration scheme upon convergence of the sub-iterations [39].

3.2. Preconditioning

We expand the above numerical procedure to incorporate multicomponent diffusion by modifying the time-marching step for species mass fraction fields. Specifically, this method modifies the treatment of the mass-diffusion source term in the species mass fraction fields. All other intermediate steps are unchanged.

3.2.1. Preconditioning iterative method

For simpler implementation, Equation (21) is solved in its residual form:

$$\begin{aligned} & \left[\rho_k^{n+3/2} \mathbf{I} - \frac{\Delta t}{2} \left(\frac{\partial \mathbf{C}}{\partial \mathbf{Y}} + \frac{\partial \mathbf{Diff}}{\partial \mathbf{Y}} + \frac{\partial \mathbf{\Omega}}{\partial \mathbf{Y}} \right)_k^{n+1} \right] \cdot (\mathbf{Y}_{k+1}^{n+3/2} - \mathbf{Y}_k^{n+3/2}) \\ &= \rho_k^{n+1/2} \mathbf{Y}_k^{n+1/2} - \rho_k^{n+3/2} \mathbf{Y}_k^{n+3/2} + \Delta t (\mathbf{C}_k^{n+1} + \mathbf{Diff}_k^{n+1} + \mathbf{\Omega}_k^*) . \end{aligned} \quad (23)$$

This equation can be restated as

$$\mathbf{Y}_{k+1}^{n+3/2} = \mathbf{Y}_k^{n+3/2} - \Delta t \mathbf{J}^{-1} \cdot \Theta_k , \quad (24)$$

where the matrix $\mathbf{J}$ is

$$\mathbf{J} = \rho_k^{n+3/2} \mathbf{I} - \frac{\Delta t}{2} \left(\frac{\partial \mathbf{C}}{\partial \mathbf{Y}} + \frac{\partial \mathbf{Diff}}{\partial \mathbf{Y}} + \frac{\partial \mathbf{\Omega}}{\partial \mathbf{Y}} \right)_k^{n+1} \quad (25)$$

and the vector

$$\Theta_k = \frac{\rho_k^{n+3/2} \mathbf{Y}_k^{n+3/2} - \rho_k^{n+1/2} \mathbf{Y}_k^{n+1/2}}{\Delta t} - [\mathbf{C}_k^{n+1} + \mathbf{Diff}_k^{n+1} + \mathbf{\Omega}_k^*] \quad (26)$$

is the residual of the species transport equation at the previous sub-iteration, which asymptotes to zero as the sub-iterations fully converge.

Written in its residual form, the time advancement of the species transport equations described here resembles the standard preconditioned Richardson-type iterative method [23,45], where the matrix $\mathbf{J}$ acts as a preconditioner. The choice of $\mathbf{J}$ as a preconditioner is arbitrary and only affects the convergence characteristics of the iterative method [23]. For example,

$$\mathbf{J} = \rho_k^{n+3/2} \mathbf{I} \quad (27)$$

is equivalent to the fully explicit integration of the convective, diffusive, and chemical source terms in the species transport equations. Alternatively,

$$\mathbf{J} = \rho_k^{n+3/2} \mathbf{I} - \frac{\Delta t}{2} \left(\frac{\partial \mathbf{C}}{\partial \mathbf{Y}} + \frac{\partial \mathbf{Diff}}{\partial \mathbf{Y}} + \frac{\partial \mathbf{\Omega}}{\partial \mathbf{Y}} \right)_k^{n+1} \quad (28)$$

is equivalent to fully implicit integration of the convective, diffusive, and chemical source terms [23].

There is a clear tradeoff in selecting the preconditioner. Since preconditioning is applied to each step of the iterative methods, the form of matrix $\mathbf{J}$ should be optimized for low computational and inversion cost while maintaining strong convergence. The fully explicit preconditioner provides the cheapest option but in our experience results in poor convergence performance, requiring extremely small time steps. Alternatively, the fully implicit preconditioner would provide excellent convergence criteria and unconditional stability; however, the Jacobian matrices for the chemical and diffusion source terms are typically dense [1,28,46]. Thus, constructing a fully implicit preconditioner is prohibitively expensive for large kinetic models.

To achieve strong convergence while maintaining a low-cost form for the preconditioner, we propose an approximation of the diffusion Jacobian that lies between the fully implicit and fully explicit extremes: a semi-implicit preconditioner. Savard et al. [23] previously implemented a similar approach for preconditioning the chemical Jacobian.

3.2.2. Semi-implicit preconditioner

In Equation (28), the Jacobian of the diffusion source term depends on the multicomponent diffusion flux, which is proportional to the multicomponent diffusion coefficient matrix, $\mathbf{D}^{\text{MC}}$. However, $\mathbf{D}^{\text{MC}}$ is a dense matrix and would be a computationally expensive approximation for the Jacobian. Alternatively, the mixture-averaged diffusion coefficient matrix, $\mathbf{D}^{\text{MA}}$, is a simplified approximation of $\mathbf{D}^{\text{MC}}$ and thus may provide a reasonable, low-cost approximation of the fully implicit Jacobian.

The mixture-averaged diffusion coefficient matrix, $\mathbf{D}^{\text{MA}}$, and the multicomponent diffusion coefficient matrix, $\mathbf{D}^{\text{MC}}$, are of a similar order and depend on the underlying species diffusivities. In addition, since $\mathbf{D}^{\text{MA}}$ is computed from the local species and temperature values rather than global changes, it is inexpensive to compute. Finally, since $\mathbf{D}^{\text{MA}}$ is strictly diagonal and thus inexpensive to invert, it provides a low-cost approximation to the diffusion Jacobian. In practice the approximate diffusion Jacobian is a tri-diagonal block matrix, where each block is the diagonal $\mathbf{D}^{\text{MA}}$ matrix. In other words, for each species the part of the Jacobian corresponding to that species is tri-diagonal and described by $\mathbf{D}^{\text{MA}}$.

3.3. Dynamic memory algorithm

As mentioned previously, high-fidelity simulations with full multicomponent mass diffusion will have a high computational expense. Thus, to facilitate a cost-effective implementation of full multicomponent diffusion we propose a simple dynamic memory algorithm that significantly reduces the computational resources needed for such simulations.

The cost of simulating full multicomponent diffusion comes from evaluating the $\mathbf{D}^{\text{MC}}$ matrix. Thus, we can reduce computational cost significantly by limiting the evaluation of $\mathbf{D}^{\text{MC}}$ to strictly once per grid-point. (In contrast, a naive implementation would involve repeated and redundant evaluations when calculating the species diffusion flux vector and its gradient.) This is possible because the central-difference scheme used is linear and thus additive and commutative by nature. In other words, the terms in the discretized equation are simply added together, and thus are strictly independent of each other and require no information from the surrounding grid points.

Recognizing this, it follows that the order of addition does not matter so long as all of the appropriate terms are included in the discretization. Thus, we can calculate the $\mathbf{D}^{\text{MC}}$ matrix once per grid point, and calculate and store for each species the discrete terms of the discretized scalar field corresponding only to the information available at that grid point. The process then repeats at the next grid point and fills in the remaining information. This approach is simply a memory-efficient rearrangement of the floating-point operations and does not alter the final result. Moreover, this dynamic memory scheme avoids the need to calculate local gradients at each grid point.

In practice, we calculate and store the portions of the enthalpy and species-diffusion source terms (in Equations (3) and (4), respectively) that can be computed from the information available at the i th grid-point for the $(i - 1/2)$ and $(i + 1/2)$ flux vectors. For example, the discretized form of the diffusion source term is

$$\begin{aligned} \text{Diff}_i &= -\nabla \cdot \mathbf{j}_i = \frac{-\dot{j}_{i+1/2} + \dot{j}_{i-1/2}}{\Delta x} \\ &= \left[(\rho_i D_i + \rho_{i+1} D_{i+1}) \frac{Y_{i+1} - Y_i}{\Delta x} - (\rho_{i-1} D_{i-1} + \rho_i D_i) \frac{Y_i - Y_{i-1}}{\Delta x} \right] \frac{1}{2\Delta x^2}, \end{aligned} \quad (29)$$

where the diffusion source term contributions from the $i - 1$, i , and $i + 1$ grid points are

$$\text{Source}_{i-1} = \frac{\rho_{i-1} D_{i-1}}{2\Delta x^2} (Y_i - Y_{i-1}), \quad (30)$$

$$\text{Source}_i = \frac{\rho_i D_i}{2\Delta x^2} (Y_{i+1} + Y_{i-1} - 2Y_i), \quad \text{and} \quad (31)$$

$$\text{Source}_{i+1} = \frac{\rho_{i+1} D_{i+1}}{2\Delta x^2} (Y_{i+1} - Y_i), \quad (32)$$

respectively.

At the i th grid point, information on the diffusion coefficients at the $i - 1$ and $i + 1$ grid points is not available; thus, only the diffusion coefficients for the i th grid point can be stored. However, by recognizing that D_i at the i th grid point is equal to D_{i+1} and D_{i-1} at the $i - 1$ and $i + 1$ grid points, respectively, it is possible to solve Equation (30), Equation (31), and Equation (32) for the $i + 1$, i , and $i - 1$ grid points, and store them in their respective memory locations. At the next grid point ($i + 1$) the process repeats and the remaining information for the i th grid point is calculated and added to the previously stored partial solution, thus completing the information needed at the i th grid point. Fig. 1 summarizes this process; fluxes are located at cell faces while source terms are at cell centers.

This approach reduces the number of $\mathbf{D}^{\text{MC}}$ evaluations from one per species per grid point to strictly one per grid point. Finally, it reduces temporary memory requirements from an array sized $n_x \times n_y \times n_z \times N^2$ to a 1×7 array corresponding to only the information needed at the current grid point (i, j, k) and its six surrounding points, where n_x , n_y , and n_z are the numbers of grid points in the x , y , and z directions. This optimizes performance by reducing cache calls for both the species mass fractions and species diffusion coefficients.




Fig. 1. Dynamic algorithm for calculating multicomponent enthalpy and species diffusion source terms. Fluxes are located at cell faces while source terms are at cell centers. N is the number of species.

The algorithm is most efficient for a structured grid, but the proposed method is easily extendable to finite-volume discretizations on unstructured meshes with scalars located at the cell centers. In such schemes, the diffusion term is written as the sum of fluxes on each cell surface. In turn, these fluxes are written as differences of cell-averaged scalar values. The regrouping of the contributions of the diffusion term to each cell in Equations (30)–(32) would follow a similar approach.

3.4. Method stability

To evaluate the theoretical stability of the proposed treatment of the diffusion source terms, we will perform a one-dimensional von Neumann stability analysis. First, we decompose the vector of species mass fractions into the exact steady-state solution ($\mathbf{Y}^\circ$) and a small perturbation vector. Then, we expand this perturbation in a Fourier series by assuming a solution of the form

$$\mathbf{Y}(x, t) = \mathbf{Y}^\circ(x) + \mathbf{f}(t)e^{i\kappa x}, \quad (33)$$

where κ is the wavenumber and $\mathbf{f}(t)$ is the time-varying amplitude of the perturbation. Under small deviations from a steady-state solution, we can make the simplifying assumption that

$$\rho_k^{n+3/2} \approx \rho^{n+1/2} = \rho^\circ. \quad (34)$$

Similarly, all diffusion coefficients are evaluated from the steady-state solution.

From here, we rewrite Equation (21) in a point-wise form neglecting both the chemical source term—demonstrated to be stable by Savard et al. [23]—and the convective transport term, which is integrated explicitly in this stability analysis (i.e., not modified by sub-iterations). This transforms the set of N partial differential equations into a set of N ordinary differential equations, where N is the number of species. Equation (23) reduces to the form

$$\left(\mathbf{I} + \frac{\Delta t}{2} \mathbf{D}^{\text{MA}} \kappa'^2 \right) \left(\mathbf{f}_{k+1}^{n+3/2} - \mathbf{f}_k^{n+3/2} \right) = \mathbf{f}^{n+1/2} - \mathbf{f}_k^{n+3/2} - \frac{\Delta t}{2} \mathbf{D}^{\text{MC}} \kappa'^2 \left(\mathbf{f}_k^{n+3/2} + \mathbf{f}^{n+1/2} \right), \quad (35)$$

where κ'^2 is the modified wavenumber, and $\mathbf{D}^{\text{MA}}$ and $\mathbf{D}^{\text{MC}}$ are the mixture-averaged and multicomponent diffusion coefficient matrices calculated from Equations (8) and (13), respectively. For the second-order central differencing scheme used, κ'^2 takes the form

$$\kappa'^2 = \frac{2}{\Delta x^2} [1 - \cos(\kappa \Delta x)]. \quad (36)$$

While here we apply this to a second-order central difference scheme, the stability analysis holds for any spatial discretization of the diffusion terms in Equation (23). In the present case, the most unstable mode manifests as cell-to-cell oscillations corresponding to $\kappa = \pi/\Delta x$ and $\kappa'^2 = 4/\Delta x^2$.

Recall that $\mathbf{f}^{n+1/2}$ is the value at the previous time step as defined in step 0 of Section 3.1 and

$$\mathbf{f}_0^{n+3/2} \equiv \mathbf{f}^{n+1/2}. \quad (37)$$

Table 1

Eigenvalues for the multi-component (left) and mixture-averaged (center-left) diffusion matrices ($\mathbf{D}^{\text{MC}}$ and $\mathbf{D}^{\text{MA}}$) and absolute values of the eigenvalues for the amplification matrix ($\mathbf{B}$) for the implicit formulation (center-right) and explicit formulation (right) evaluated on the burned side of the lean hydrogen premixed flame (see section 4.1). $\dagger$ units are $10^{-3} \text{ m}^2/\text{s}$. A time-step size of $\Delta t = 10^{-5} \text{ s}$ was used for $\mathbf{B}$.

$\mathbf{D}^{\text{MC}\dagger}$	$\mathbf{D}^{\text{MA}\dagger}$	$\mathbf{B}^{\text{imp}}$	$\mathbf{B}^{\text{exp}}$
0.0000	0.2984	0.0095	0.0000
0.3033	0.3003	0.0096	25.629
0.3053	0.3014	0.0126	25.800
0.3070	0.3023	0.0131	25.945
0.3860	0.4069	0.0163	32.614
0.4644	0.4585	0.0265	39.239
0.4735	0.4672	0.0276	40.012
1.1163	1.0859	0.0450	94.327
1.8968	1.8145	0.9634	160.276

Dropping the superscripts for clarity, we can reduce Equation (35) to

$$\mathbf{f}_{k+1} = \mathbf{A}\mathbf{f}_0 + \mathbf{B}\mathbf{f}_k, \quad (38)$$

where

$$\mathbf{A} = \left(\mathbf{I} + \frac{\Delta t}{2} \kappa'^2 \mathbf{D}^{\text{MA}} \right)^{-1} \left(\mathbf{I} - \frac{\Delta t}{2} \kappa'^2 \mathbf{D}^{\text{MC}} \right) \quad (39)$$

and

$$\mathbf{B} = \left(\mathbf{I} + \frac{\Delta t}{2} \kappa'^2 \mathbf{D}^{\text{MA}} \right)^{-1} \left(\frac{\Delta t}{2} \kappa'^2 (\mathbf{D}^{\text{MA}} - \mathbf{D}^{\text{MC}}) \right). \quad (40)$$

Inspecting Equation (38), matrix $\mathbf{A}$ is multiplied by the constant value of the previous time step ($\mathbf{f}_0$) and therefore does not contribute to the stability of the sub-iterations. We focus on the properties of the $\mathbf{B}$ matrix, which acts as the amplification/growth factor. Theoretically, the stability of the sub-iterations is ensured if the spectral radius of matrix $\mathbf{B}$, defined as the largest absolute value of the eigenvalues, is less than one:

$$\rho(\mathbf{B}) \leq 1. \quad (41)$$

The matrix $\mathbf{B}$ has some interesting properties that deserve further discussion.

First, this matrix is proportional to the difference between the two diffusion matrices $\mathbf{D}^{\text{MA}}$ and $\mathbf{D}^{\text{MC}}$. Recall that the $\mathbf{D}^{\text{MA}}$ is a purely diagonal matrix. Table 1 compares the eigenvalues of these matrices on the burned side of the lean hydrogen premixed flames (see Section 4.1 for details on the flame). The burned side of the flame is characterized by the largest diffusion coefficients and is expected to be the most unstable location within a flame as far as diffusion is concerned. The two sets of eigenvalues are extremely close, which is expected as the mixture-averaged diffusion model approximates the multi-component diffusion model. As a result, the norm of the difference of the two matrices is expected to be much less than the norm of either matrix. In other words, we anticipate that $\rho(\mathbf{B}) \ll 1$ regardless of the time-step size (Δt) and grid spacing (Δx).

One noticeable difference between $\mathbf{D}^{\text{MC}}$ and $\mathbf{D}^{\text{MA}}$ is the presence of a null eigenvalue for multi-component diffusion. As described in Section 2.3, the multi-component matrix is singular, and its kernel is spanned by the species mass fraction matrix, here $\mathbf{Y}^\circ$. Consider the special case of $\mathbf{f}_0 = \mathbf{Y}^\circ$. Leveraging the fact that $\mathbf{f}_0$ lies in the kernel of $\mathbf{D}^{\text{MC}}$, the amplitude at the next sub-iteration will be

$$\mathbf{f}_1 = (\mathbf{A} + \mathbf{B}) \mathbf{f}_0 = \mathbf{f}_0. \quad (42)$$

By recursive reasoning, one can show the property holds for all sub-iterations. In other words, this mode is unaffected by the iterative process and remains the same between time steps. This time-invariant mode is nothing more than the steady-state solution $\mathbf{Y}^\circ$. To avoid “double counting” in Equation (33), the eigenvalue analysis of the matrix $\mathbf{B}$ should be performed on the linear space not including the vector $\mathbf{Y}^\circ$.

Table 1 provides an example of the eigenvalues of the amplification matrix $\mathbf{B}$ for a time-step size of $\Delta t = 10^{-5} \text{ s}$. Practically, the eigenvector associated with the largest eigenvalue of $\mathbf{B}$ forms a small angle ($\sim 0.08 \text{ deg}$) with the species mass fraction vector ($\mathbf{Y}^\circ$). Hence, most of it is in the kernel of $\mathbf{D}^{\text{MC}}$. Following the previous discussion, this eigenvalue (indicated in italics) should not be considered, and the overall stability is controlled by the second-largest eigenvalue, in

this case 0.0425. As expected, this eigenvalue is much less than unity, thus proving the stability of the iterative procedure. This should be compared to the stability of the explicit formulation obtained by setting $\mathbf{D}^{\text{MA}} = 0$ in Eq. (40). Under these conditions, the spectral radius of $\mathbf{B}$ becomes

$$\rho(\mathbf{B}) = 2 \frac{\Delta t}{\Delta x^2} \rho(\mathbf{D}^{\text{MC}}) \approx 2 \frac{\Delta t}{\Delta x^2} \max(\mathbf{D}^{\text{MA}}), \quad (43)$$

which resembles a Fourier number. As shown by the large eigenvalue of the explicit-method amplification matrix, $\mathbf{B}^{\text{exp}}$, in Table 1, solving the system of equations would not be stable at $\Delta t = 10^{-5}$ s without the proposed implicit formulation. Section 5.1 presents an in-depth comparison of this theoretical stability criterion against practical numerical convergence results for a one-dimensional freely propagating flame.

4. Test cases

We will evaluate the performance of the proposed iterative method and the relative cost of the implemented memory algorithm in Section 5. We base our evaluation on two flow configurations: a one-dimensional, unstretched, laminar flame and a three-dimensional, statistically stationary, turbulent flame; both are premixed hydrogen/air flames. All simulations used the same nine-species hydrogen mechanism of Hong et al. [47] with updated rate constants from the same group [48, 49]. This section describes the configuration and conditions used for the one- and three-dimensional simulations used for this study. Appendix B includes additional method verification.

4.1. One-dimensional premixed flame

To verify the implementation of the multicomponent mass-diffusion model and evaluate its accuracy, we performed one-dimensional, unstretched (flat), laminar flame simulations and compared these with similar mixture-averaged and multicomponent results computed using Cantera [50]. We selected the one-dimensional flat flame configuration because it restricts all transport to the streamwise direction. As a result, the spanwise fluxes are zero by definition for this geometry. This condition may not hold in a multidimensional flow simulation where the multicomponent diffusion fluxes may be misaligned with the species gradient vector. This simplified geometry allows us to directly compare the multicomponent mass diffusion model to the commonly used mixture-averaged diffusion model.

The simulations used an unburnt temperature of 298 K and pressure of 1 atm, with an equivalence ratio of $\phi = 0.4$ and inlet velocity equal to the laminar flame speed for all Cantera and NGA cases. The flame was centered in a computational domain comprised of 720 grid points where $\Delta x = 15.4 \mu\text{m}$. To ensure fidelity in the results, we selected the domain to have at least 20 points through the laminar flame, with the thickness defined using the maximum temperature gradient: $l_F = (T_{\text{max}} - T_{\text{min}})/|\nabla T|_{\text{max}}$. Schlup and Blanquart [6] used an identical configuration to investigate the impact of Soret and Dufour thermal diffusion effects.

We ran the Cantera simulations similarly using both mixture-averaged and multicomponent diffusion models with matching inlet conditions, equivalence ratio, and domain size. The freely-propagating adiabatic flat flame solver (FreeFlame) was used with grid refinement criteria for both slope and curvature set to 0.1 and a refinement ratio of 2.0 for 860 grid-points.

4.2. Three-dimensional flow configuration

We simulated a three-dimensional, turbulent, premixed, freely propagating flame as a test of the proposed algorithm for multicomponent mass diffusion and to assess the impact of diffusion model choice on global statistics such as the turbulent flame speed. The computational domain consists of inflow and convective outflow boundary conditions in the streamwise direction. The two spanwise directions use periodic boundaries. We set the inflow velocity to the mean turbulent flame speed, which keeps the flame statistically stationary such that turbulent statistics can be collected over an arbitrarily long run time. In the absence of mean shear, we use a linear turbulence-forcing method [33,51] to maintain the production of turbulent kinetic energy through the flame. We selected the Karlovitz number to fall near the boundary between the thin and distributed reaction zone regimes, while avoiding the broken reaction zone regime [52]. Moreover, the computational setup for this case is similar to those of Lapointe et al. [52], Burali et al. [5], and Schlup et al. [6], who studied differential-diffusion effects, local extinction, and flame broadening using the mixture-averaged model and constant non-unity Lewis number assumptions.

Table 2 provides further details of the computational domain, unburnt mixture, and inlet turbulence. The unburnt temperature and pressure are 298 K and 1 atm, respectively. The inlet equivalence ratio is $\phi = 0.4$, with an unburnt Karlovitz number $\text{Ka}_u = \tau_F/\tau_\eta = 149$, where $\tau_F = l_F/S_L$ is the flame time scale and $\tau_\eta = (\nu_u/\epsilon)^{1/2}$ is the Kolmogorov time scale of the incoming turbulence with unburnt kinematic viscosity ν_u and turbulent energy dissipation ϵ . The unburnt turbulent Reynolds number is $\text{Re}_t = u' l/\nu_u = 289$, where u' is the fluctuation of the mean velocity and l is the integral length scale. The mean inflow velocity at the inlet boundary condition approximately matches the turbulent flame speed so that the flame remains relatively centered in the domain and we can perform arbitrarily long simulations. Once the turbulence has fully developed, we run the simulations for 22 eddy turnover times, $\tau_{\text{eddy}} = k/\epsilon \approx 500 \mu\text{s}$.

Table 2

Three-dimensional simulations parameters. Δx is the grid spacing, η_u is Kolmogorov length scale of the unburnt gas, Δt is the simulation time-step size, ϕ is the equivalence ratio, P_0 is the thermodynamic pressure, T_u is the temperature of the unburnt mixture, T_{peak} is the temperature of peak fuel consumption rate in the one-dimensional laminar flame, S_L is the laminar flame speed, $l_F = (T_b - T_u) / |\nabla T|_{\text{max}}$ is the laminar flame thickness, $l = u'^3 / \epsilon$ is the integral length scale, u' is the turbulence fluctuations, ϵ is the turbulent energy dissipation rate, Ka_u is the Karlovitz number of the unburnt mixture, Re_t is the turbulent Reynolds number of the unburnt mixture, and ν_u is the unburnt kinematic viscosity.

Parameter	MA	MC
Domain	$8L \times L \times L$	
L	$190\Delta x$	
Grid	$1520 \times 190 \times 190$	
Δx [m]	4.24×10^{-5}	
η_u [m]	2.1×10^{-5}	
Δt [s]	6×10^{-7}	
ϕ	0.4	
P_0 [atm]	1	
T_u [K]	298	
T_{peak} [K]	1190	1180
S_L [m/s]	0.230	0.223
l_F [mm]	0.643	0.631
l/l_F	2	2.04
u'/S_L	18	18.6
$Ka_u = \tau_F/\tau_\eta$	149	151
$Re_t = (u'l)/\nu_u$	289	




Fig. 2. Two-dimensional schematic of the three-dimensional flame configuration. Adapted from Burali et al. and Schlup and Blanquart [5,6]. The red line indicates the approximate location of the flame. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

The domain has 1520 points in the streamwise direction and 190 points in both spanwise directions, with a uniform grid size of $\Delta x = l_F/16$. This domain is about $100l_F$ long and $12l_F$ in the spanwise directions. Given the prescribed turbulence intensity, this mesh has a grid spacing equivalent to $\Delta x \approx 2\eta_u$, where η_u is the Kolmogorov length scale for the unburnt region; this resolution improves in the burnt region of the flame. Lapointe et al. [52] previously confirmed the suitability of the selected grid spacing and resolution in the flame front using a mesh refinement study, which found no difference when using this grid spacing compared with half the size. Fig. 2 shows a two-dimensional schematic of the domain, including the locations of the flame and the forcing region. Fig. 3 shows a three-dimensional view of the iso-surface of T_{peak} defining the flame front, where T_{peak} is the temperature of peak fuel consumption rate in the one-dimensional laminar flame. The flame surface shows the complex behavior of the flame in the turbulent field.

5. Results and discussion

To start, we present a practical assessment of the method's convergence and stability, by comparing the numerical rate of convergence to the theoretical rate of convergence. Following this demonstration of the proposed method's stability, we verify the accuracy of the method through a posteriori assessment of one-dimensional, unstretched, premixed, laminar flame simulations. Finally, we present a preliminary evaluation of the relative differences between the mixture-averaged and multicomponent diffusion models for the three-dimensional turbulent premixed flame simulations.

5.1. Stability analysis results

We use the one-dimensional flame to numerically evaluate the convergence stability of the sub-iterations with respect to time-step size. The simulations for these tests were initialized from a mixture-averaged data file to provide a worst-case scenario for the initial iterative step in converging to the multicomponent solution. While the theoretical analysis was performed assuming explicit transport of the convective terms and constant density/diffusion coefficients, we performed




Fig. 3. Iso-surface of peak temperature colored by OH mass fraction for a three-dimensional turbulent hydrogen/air flame with multicomponent mass diffusion.




Fig. 4. Convergence of the density residual as a function of sub-iteration for the proposed semi-implicit method, for a smaller and larger time-step size. Dashed lines are the spectral radii shown in Fig. 5 and are determined by numerically fitting an exponential curve to the slope of the density residual

this test with semi-implicit transport and variable density/diffusion coefficients. This demonstrates the stability of the proposed preconditioner for the semi-implicit multicomponent diffusion transport in a practical numerical simulation. Savard et al. [23] previously showed the numerical stability of the chemical and convective terms, so we do not discuss these terms in detail in this analysis.

We focus on the maximum density residual over the whole domain, because its convergence is controlled by the convergence of all chemical species. Figs. 4a and 4b present the density residuals as a function of sub-iteration, starting from the initial time step, for a small and large time-step size, respectively. For the time-step sizes tested, converging (as opposed to diverging) sub-iterations implies a stable simulation, which agrees with behavior shown by Savard et al. [23]. In other words, unless the sub-iterations diverge, the simulation remains stable. As expected, the explicit method diverges quickly even at very small time-step sizes (Fig. 4a), while the semi-implicit method remains stable up to a time-step size of $\Delta t \leq 1 \times 10^{-5}$ s (Fig. 4b).

The rate of convergence of the sub-iterations for each of the source terms in Figs. 4a and 4b follows an exponential relationship, i.e., $\text{Res}_k \sim r^k$, where Res_k is the residual of the k th sub-iteration and r is the convergence rate. We compute the numerical convergence rate r by fitting an exponential curve to the slope of the density residuals; the convergence rate is represented by dashed lines in Fig. 4. Since density is a function of the species mass fractions, its convergence rate should tend towards that of the slowest-converging species mass fraction.

Fig. 5 compares the results of the theoretical and numerical stability analyses, showing the spectral radius of matrix **B** as a function of the time-step size for the one-dimensional test case. For the explicit scheme, the theoretical and numerical results agree well for the full range of time-step sizes. However, for the implicit scheme, the predicted spectral radius is much smaller than the measured one. The proposed implicit formulation thus yields a convergence rate that is not limited by diffusion, but rather constrained by other processes that were not considered in the stability analysis, such as chemistry. The predicted convergence rate can nonetheless be observed in the implicit case. As mentioned in Section 3.4,




Fig. 5. Theoretical convergence rate determined from diagonalizing matrix $\mathbf{B}$ corresponding to the worst-case modified wavenumber for the one-dimensional premixed flame, compared with the numerical convergence rates determined by fitting an exponential curve to the slope of the density residual.




Fig. 6. A posteriori comparisons of species mass fractions relative to mixture local temperature in a hydrogen/air flame with $\phi = 0.4$ using NGA and Cantera.

the eigenvector associated with the largest eigenvalue of $\mathbf{B}^{\text{imp}}$ in Table 1 forms a small angle with $\mathbf{Y}^\circ$. Hence, a fraction of the error, albeit tiny, is associated with this eigenvector, which will slowly converge. In Fig. 4b, the convergence rate for the last sub-iterations of the implicit case closely matches that eigenvalue.

Overall, these results suggest that the theory well-approximates actual stability and provides a practical limit for the numerical stability of the proposed algorithm.

5.2. Method verification

To verify the multicomponent model, we compare a posteriori the one-dimensional unstretched species profiles and laminar flame speeds. Fig. 6 compares the nine species profiles for the steady-state one-dimensional flat flame solutions relative to local mixture temperature for the multicomponent and mixture-averaged models from both NGA and Cantera. The profiles all agree within 1% at all points, with the exception of N_2 . The laminar flame speeds (S_L^0) for these simulations




Fig. 7. Relative accuracy of the method as a function of time step size for the one-dimensional, freely propagating flame test case with seven sub-iterations. Errors are defined as the absolute difference of their integrated value in temperature space compared with a reference solution obtained for $\Delta t = 1 \times 10^7$ s and seven sub-iterations. Black dashed line corresponds to $y = x^{-2}$.

are approximately 23.0 cm/s and 22.3 cm/s for the mixture-averaged and multicomponent diffusion NGA cases, respectively; the laminar flame speeds for both cases agree with those from Cantera within 1%. The unstretched laminar flame speed is

$$S_L^0 = - \frac{\int \rho \dot{\omega}_{H_2} dx}{\rho_u Y_{H_2,u}}, \quad (44)$$

where ρ_u is the unburnt mixture density and $Y_{H_2,u}$ is the unburnt fuel mass fraction. We attribute the larger difference in the species profile for N_2 to the correction velocity term associated with the mixture-averaged diffusion model, which is weighted by mass fraction and thus can be heavily impacted by differences in N_2 due to its high concentration throughout the flame. The minor differences between the multicomponent species profiles are less than 1% at all points. The strong agreement between the other eight species profiles for both the NGA and Cantera results verifies the multicomponent model's functionality.

5.3. Accuracy

With the proposed algorithm's stability limits and functionality verified, we now examine the accuracy for a given stable simulation. We determine the order of accuracy of the method based on the 1D freely propagating flame case by examining the power-law dependence of the error as a function of the time-step size.

Fig. 7 shows the normalized error for the 1D freely propagating flame case for various time steps. We initialize the simulation using an input flame profile corresponding to a fully converged statistically stationary flame, generated with a time-step size of $\Delta t = 1 \times 10^7$ s and seven sub-iterations. A wall is then set at the simulation inlet, allowing the flame to propagate upstream in the domain. We then let the reference flame propagate for two flame pass-through times to ensure a fully converged freely propagating flame profile free of any initial transients due to the transition from the input stationary flame profile. This reference file then serves as the input for a set of freely propagating flames with time-step sizes ranging 10^{-5} – 10^{-7} s and for seven sub-iterations. Finally, we allow these test flames to propagate for an additional flame pass-through time to ensure statistical independence from the initial reference-flame input file.

With the freely propagating flame tests completed, we interpolated the species and density fields to a constant temperature space corresponding to the temperature distribution in the flame region. This interpolation ensures a direct comparison of the species error independent of variation in the temperature space over the range of time-step sizes. We then calculate error as the L^2 -norm of the species and density profiles in temperature space, relative to the reference flame profile with $\Delta t = 1 \times 10^{-7}$ s and seven sub-iterations:

$$\text{error} = \sqrt{\frac{\int (Y_i - Y_{i,\text{ref}})^2 dT}{\int Y_{i,\text{ref}}^2 dT}} \quad (45)$$

and

$$\text{error} = \sqrt{\frac{\int (\rho - \rho_{\text{ref}})^2 dT}{\int \rho_{\text{ref}}^2 dT}}. \quad (46)$$

We selected the species H_2 , H_2O , OH , and H to evaluate the accuracy of the method because they represent the reactants, intermediate species, and products present in hydrogen combustion. Density (ρ) is also included to globally assess error,




Fig. 8. Impact of time-step size and number of sub-iterations on the accuracy of one-dimensional freely propagating flames.




Fig. 9. Turbulent flame speed history for three-dimensional, freely propagating, premixed, turbulent hydrogen/air flame with $\phi = 0.4$.

since it depends on all species. As shown in Fig. 7a, all quantities exhibit second-order accuracy in time with seven sub-iterations. The errors corresponding to the L^1 - and L^∞ -norms are similar in magnitude and also demonstrate second-order accuracy in time with seven sub-iterations.

While the method is fully second-order accurate for seven sub-iterations and above, the solution transitions to first-order accuracy as the number of sub-iterations decreases. Fig. 7b shows that the solution exhibits first-order accuracy when using four sub-iterations. Between four and seven sub-iterations the solution is second-order accurate for large time-step sizes but transitions to first-order accuracy as the time step size decreases. The range of time-step sizes that achieve second-order accuracy grows until the solution becomes fully second-order accurate at seven sub-iterations for all time-step sizes considered.

To evaluate the absolute magnitude of error associated with the proposed method, as opposed to the order of accuracy (as time-step size approaches zero), Figs. 8a and 8b present the temperature as a function of distance and fuel mass fraction as a function of temperature, respectively, for a range of freely propagating flames with several time-step sizes and sub-iterations. The solutions exhibit negligible error in both temperature and fuel mass fractions for the time-step sizes considered, and even when using as few as four sub-iterations; these tests demonstrate the high accuracy and robustness of the proposed method.

5.4. Three-dimensional assessment of diffusion flux models

In this section we assess a posteriori the species mass diffusion fluxes in the doubly periodic three-dimensional flames [4–6]. Differential diffusion effects cause the instabilities found in lean hydrogen/air flames, and at high Karlovitz numbers the turbulence time scales match the order of diffusion time scales.

To assess the impact of the mixture-averaged and multicomponent mass diffusion models on flame chemistry, we compare a posteriori the turbulent and chemistry statistics. We allow the flames to develop in a turbulent flow field, and compute the statistics after the transients from the initial flow and scalar fields have advected through the domain. As an initial assessment, we calculate the effective turbulent flame propagation speeds:

$$S_T = - \frac{\int_V \rho \dot{\omega}_{H_2} dV}{\rho_u Y_{H_2,u} L^2}. \quad (47)$$

Fig. 9 shows the time history of the turbulent flame speed over twenty-two eddy turn-over times (τ_{eddy}). The average normalized flames speeds from the mixture-averaged and multicomponent models differ by 15%: $S_T^{MA}/S_L = 29.6$ and $S_T^{MC}/S_L^0 = 34.7$, respectively. Further study is needed on whether the mixture-averaged diffusion model fully captures the fundamental physics of multicomponent diffusion.




Fig. 10. Conditional means on temperature for the three-dimensional, freely propagating, premixed, turbulent hydrogen/air flame with $\phi = 0.4$.




Fig. 11. Computational time per grid point for computing diffusion coefficients and diffusion mass fluxes using kinetic models with 9, 35, and 172 species; black dashed lines correspond to linear ($y = x$) and quadratic ($y = x^2$) scaling trends respectively. MC and MA stand for multicomponent and mixture-averaged, respectively.

To further assess any differences between the mixture-averaged and multicomponent mass diffusion models, Fig. 10 presents the means of fuel mass fraction and its source term conditioned on temperature for the full time domain. The differences in the calculated conditional means are small: less than 5.5%. This agreement also extends into super-adiabatic regions for the hydrogen/air flame; these regions, also called “hot spots”, result from differential diffusion and have been predicted both in theoretical studies [53] and numerical analyses of lean hydrogen/air mixtures [54–56]. However, these small differences in global flame statistics do not explain the 15% difference observed in the turbulent flame speeds between the mixture-averaged and multicomponent diffusion models. These results raise questions on the appropriateness of the mixture-averaged diffusion assumption for direct numerical simulation and warrants further investigation.

5.5. Computational cost

This section discusses the relative cost for implementing the full multicomponent mass diffusion to provide context for its use. The presented timing comparisons examine how the method scales with both number of chemical species and spatial dimension.

We tested three chemical kinetic models (containing 9 [5], 35 [36,37], and 172 species [57–59]) in a one-dimensional flat flame simulation to determine the cost of multicomponent mass diffusion over a wide range of model sizes. Fig. 11 shows the computational time per grid point for computing the diffusion mass fluxes on a desktop workstation using an Intel Xeon-X5660 CPU with a 2.80 GHz clock speed. The presented timings include calculation of both the diffusion coefficients and mass diffusion fluxes for all aspects of the code.

For the tested chemical kinetic models, the mixture-averaged model scales linearly while the multicomponent model scales quadratically with the number of species. The multicomponent model is more expensive and does take more time per-point for all three test cases. For the largest kinetic model (with 172 species) the multicomponent case is noticeably more expensive than the mixture-averaged model. The increased cost for the multicomponent simulations comes primarily from the CHEMKIN II [28] routine used to determine the ordinary multicomponent diffusion coefficient matrix.

The relevant cost for the proposed method can be split into three primary categories: the costs of calculating the multicomponent diffusion coefficients, calculating the multicomponent diffusion fluxes, and the semi-implicit integration scheme.




Fig. 12. Computational time per grid point for each of the three flame configurations: one dimensional (blue), two dimensional (red), and three dimensional (yellow).




Fig. 13. Comparison of numerical costs for the three-dimensional hydrogen flame simulations.

Since the proposed method for implementing full multicomponent mass diffusion focuses on efficient low-memory calculation of the diffusion fluxes, rather than the multicomponent diffusion coefficients, the cost of CHEMKIN should be considered independently of the proposed algorithm. Moreover, the semi-implicit scheme is the same for the mixture-averaged and multicomponent cases, because both cases use the mixture-averaged diffusion coefficient matrix to approximate the Jacobian for the diffusion source terms. As a result, the two methods have similar implementation and computational expense, with the exception of using CHEMKIN II [28].

To evaluate how the multicomponent model scales with increasing spatial dimension, and evaluate the relative cost of using CHEMKIN II [28], we acquired timings for one- (720 grid points), two- (1888×472 grid points), and three-dimensional ($1520 \times 190 \times 190$ grid points) configurations covering the cases presented in this work, with the additional two-dimensional case matching similar timing tests by Schlup et al. [6]. These timing tests represent an average cost per point and are determined by averaging the timings taken for the 20 time steps, skipping the first and last integrations. Fig. 12 presents the computational timings for each part of the code for both diffusion models, where “Scalar” includes scalar field calculation; “Diffusion” includes the flux calculation and $\mathbf{D}^{\text{MA}}$ calculation for the implicit solver; “Chemistry”, “Velocity”, and “Pressure” are as named; and “Rest” account for any remaining computations. “Scalar” includes the semi-implicit solver for integrating the diffusion source terms, while the semi-implicit solvers for chemistry and velocity are included in their named categories. To facilitate comparison between the two models, Fig. 13 presents the total computational time per grid point for both three-dimensional hydrogen simulations as a stacked bar chart broken down by each section of code. We performed these computations on the National Energy Research Scientific Computing Center (NERSC) high-performance computing cluster Cori (Cray XC40) [60].

While much of the code exhibits a similar cost per grid point, regardless of the dimensionality of the problem, the chemistry is more expensive for the one- and two-dimensional cases. This cost increase is due to NGA’s structure. NGA was written and optimized for three-dimensional configurations, thus the one- and two-dimensional cases are artificially more expensive, especially in the chemistry calculations [6]. In addition, for three dimensions the cost of the pressure solver increases due to using the HYPRE package [42]. The one- and two-dimensional cases both implement an exact FFT-tridiagonal solver, while HYPRE—used for the three-dimensional cases—is iterative and thus more expensive. Despite the minor increase in cost for the pressure solver in three dimensions, the cost is negligible when considering larger kinetic models (i.e., more than 35 species).

Consistent with Fig. 11, the cost of calculating “Diffusion” increases with model complexity; recall that $\mathbf{D}^{\text{MA}}$ is calculated for both the mixture-averaged and multicomponent solvers. However, the multicomponent diffusion mass flux calculation

represents only 21% of the total simulation time for the three-dimensional case. As expected, the cost of calling CHEMKIN II for the diffusion coefficients is large and accounts for roughly 23% of the three-dimensional simulation time. Interestingly, the cost of diffusion increases only slightly moving from one dimension to two dimensions. This results from the high efficiency of the dynamic memory-allocation algorithm used to implement this model (see Section 3.3). Moreover, the multicomponent diffusion implementation is less expensive than the mixture-averaged model for the one-dimensional case and equivalent in cost for the two-dimensional case. Overall, by reducing memory requirements and optimizing calls to memory, the memory algorithm implemented for the multicomponent model maintains low computational expense.

These results indicate that, for hydrogen-air combustion, the multicomponent model is more expensive than the mixture-averaged model; however, the differences in “Diffusion” costs between the two models are due to the use of CHEMKIN II [28]. Thus, the slowdown could be minimized by implementing a more-efficient package for calculating the mass-diffusion coefficients such as EGLIB [9,19,20]; however, the total cost of computing mass diffusion fluxes remains notable, even for the mixture-averaged case.

6. Summary and future work

This article presents an efficient and stable scheme for implementing multicomponent mass diffusion in reacting-flow DNS with minimal memory expense. The proposed scheme exhibits reasonable computational cost for chemical kinetic models of up to 100 species; this performance could be further improved by implementing a more-efficient method for calculating the multicomponent diffusion coefficient matrix.

The results presented for hydrogen flames suggest that the mixture-averaged mass diffusion model may suffice for DNS of three-dimensional, premixed turbulent flames in the regimes and configurations considered. However, we observed a 15% difference in the turbulent flame speeds between the two models, though the differences in the conditional means of the fuel source term and mass fraction were negligible. The difference observed in turbulent flame speeds raises questions about using the mixture-averaged model in DNS of turbulent reacting flows. Moreover, the algorithm proposed in this study provides a fast, efficient, method for implementing multicomponent mass diffusion in reacting-flow simulations, which may eliminate the need for the mixture-averaged assumption. However, despite these results, we do not have sufficient data to draw firm conclusions on the accuracy and appropriateness of mixture-averaged assumptions for all flames (i.e., all fuels, configurations, and regimes). Additional data are needed from studies of different fuels—namely large hydrocarbons—and kinetic models with more species. Thus, future work should focus on extending these comparisons to other fuels and flame configurations.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant Nos. 1314109-DGE and CBET-1832548. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility operated under Contract No. DE-AC02-05CH11231.

Appendix A. Availability of material

The figures in this article, as well as the data and plotting scripts necessary to reproduce them, are available openly under the CC-BY license [61]. Furthermore, the full simulation results from NGA are available for the three-dimensional multicomponent [62] and mixture-averaged [63] hydrogen/air flames.

Appendix B. Method verification

To verify the method implementation, we generated an artificial species profile where the direction and relative magnitudes of the flux vectors could be predicted a priori to remain independent of any differential diffusion effects that may exist in a physical system. Specifically, we created a two-dimensional V-shaped species profile with a central angle of 45° and projected it into three dimensions as shown in Fig. B.14a.

Such a profile results in flux vectors that are constant in the y -direction, are of equal magnitude and opposite sign in the z -direction reflected over the x - y -plane, and vary in magnitude but remain constant in sign matching the initial input profile in the x -direction. These predictions should be consistent independent of chemical species or other scalar value for the artificial input profile. We ran the algorithm for one “complete” set of sub-iterations to convergence and normalized the resulting diffusion flux vectors to ensure the relative magnitudes and direction were consistent with our expectations.

Fig. B.14 shows the results of this artificial test case. The resulting normalized flux vectors agree with expectation and have equal magnitudes in the x - and z -directions corresponding to the 45° artificial flame angle. This result indicates proper functionality of the proposed method.




Fig. B.14. Normalized flux vectors resulting from an artificial species profile after one full iteration of semi-implicit multicomponent diffusion calculation.

References

- [1] R.B. Bird, W.E. Stewart, E.N. Lightfoot, *Transport Phenomena*, John Wiley & Sons, Inc., New York, 1960.
- [2] J. Warnatz, Calculation of the structure of laminar flat flames I: flame velocity of freely propagating ozone decomposition flames, *Ber. Bunsenges. Phys. Chem.* 82 (2) (1978) 193–200, <https://doi.org/10.1002/bbpc.197800010>.
- [3] M.E. Coltrin, R.J. Kee, J.A. Miller, A mathematical model of silicon chemical vapor deposition, *J. Electrochem. Soc.* 133 (6) (1986) 1206, <https://doi.org/10.1149/1.2108820>.
- [4] S. Lapointe, G. Blanquart, Fuel and chemistry effects in high Karlovitz premixed turbulent flames, *Combust. Flame* 167 (2016) 294–307, <https://doi.org/10.1016/j.combustflame.2016.01.035>.
- [5] N. Burali, S. Lapointe, B. Bobbitt, G. Blanquart, Y. Xuan, Assessment of the constant non-unity Lewis number assumption in chemically-reacting flows, *Combust. Theory Model.* 20 (4) (2016) 632–657, <https://doi.org/10.1080/13647830.2016.1164344>.
- [6] J. Schlup, G. Blanquart, Validation of a mixture-averaged thermal diffusion model for premixed lean hydrogen flames, *Combust. Theory Model.* 22 (2) (2018) 264–290, <https://doi.org/10.1080/13647830.2017.1398350>.
- [7] T.P. Coffee, J.M. Heimerl, Transport algorithms for premixed, laminar steady-state flames, *Combust. Flame* 43 (1981) 273–289, [https://doi.org/10.1016/0010-2180\(81\)90027-4](https://doi.org/10.1016/0010-2180(81)90027-4).
- [8] J. Warnatz, Influence of transport models and boundary conditions on flame structure, in: N. Peters, J. Warnatz (Eds.), *Numerical Methods in Laminar Flame Propagation*, in: *Notes on Numerical Fluid Mechanics*, vol. 6, Vieweg+Teubner Verlag, Wiesbaden, 1982, pp. 87–111.
- [9] A. Ern, V. Giovangigli, Impact of detailed multicomponent transport on planar and counterflow hydrogen/air and methane/air flames, *Combust. Sci. Technol.* 149 (1–6) (1999) 157–181, <https://doi.org/10.1080/00102209908952104>.
- [10] H. Bongers, L.P.H. De Goeij, The effect of simplified transport modeling on the burning velocity of laminar premixed flames, *Combust. Sci. Technol.* 175 (10) (2003) 1915–1928, <https://doi.org/10.1080/713713111>.
- [11] Y. Xin, W. Liang, W. Liu, T. Lu, C.K. Law, A reduced multicomponent diffusion model, *Combust. Flame* 162 (1) (2015) 68–74, <https://doi.org/10.1016/j.combustflame.2014.07.019>.
- [12] M. Faghih, W. Han, Z. Chen, Effects of Soret diffusion on premixed flame propagation under engine-relevant conditions, *Combust. Flame* 194 (2018) 175–179, <https://doi.org/10.1016/j.combustflame.2018.04.031>.

- [13] J. De Charentenay, A. Ern, Multicomponent transport impact on turbulent premixed H_2/O_2 flames, *Combust. Theory Model.* 6 (3) (2002) 439–462, <https://doi.org/10.1088/1364-7830/6/3/304>.
- [14] V. Giovangigli, Multicomponent transport in laminar flames, *Proc. Combust. Inst.* 35 (1) (2015) 625–637, <https://doi.org/10.1016/j.proci.2014.08.011>.
- [15] S. Dworkin, M. Smooke, V. Giovangigli, The impact of detailed multicomponent transport and thermal diffusion effects on soot formation in ethylene/air flames, *Proc. Combust. Inst.* 32 (1) (2009) 1165–1172, <https://doi.org/10.1016/j.proci.2008.05.061>.
- [16] T. Lu, C.K. Law, Toward accommodating realistic fuel chemistry in large-scale computations, *Prog. Energy Combust. Sci.* 35 (2) (2009) 192–215, <https://doi.org/10.1016/j.pecs.2008.10.002>.
- [17] G. Borghesi, J. Bellan, A priori and a posteriori investigations for developing large eddy simulations of multi-species turbulent mixing under high-pressure conditions, *Phys. Fluids* 27 (3) (2015) 035117, <https://doi.org/10.1063/1.4916284>.
- [18] E. Masi, J. Bellan, K.G. Harstad, N.A. Okong'o, Multi-species turbulent mixing under supercritical-pressure conditions: modelling, direct numerical simulation and analysis revealing species spinodal decomposition, *J. Fluid Mech.* 721 (2013) 578–626, <https://doi.org/10.1017/jfm.2013.70>.
- [19] A. Ern, V. Giovangigli, Fast and accurate multicomponent transport property evaluation, *J. Comput. Phys.* 120 (1) (1995) 105–116, <https://doi.org/10.1006/JCPH.1995.1151>.
- [20] A. Ern, V. Giovangigli, Thermal diffusion effects in hydrogen-air and methane-air flames, *Combust. Theory Model.* 2 (4) (1998) 349–372, <https://doi.org/10.1088/1364-7830/2/4/001>.
- [21] S. Ambikasaran, K. Narayanaswamy, An accurate, fast, mathematically robust, universal, non-iterative algorithm for computing multi-component diffusion velocities, *Proc. Combust. Inst.* 36 (1) (2017) 507–515, <https://doi.org/10.1016/j.proci.2016.05.055>.
- [22] O. Desjardins, G. Blanquart, G. Balarac, H. Pitsch, High order conservative finite difference scheme for variable density low Mach number turbulent flows, *J. Comput. Phys.* 227 (15) (2008) 7125–7159, <https://doi.org/10.1016/j.jcp.2008.03.027>.
- [23] B. Savard, Y. Xuan, B. Bobbitt, G. Blanquart, A computationally-efficient, semi-implicit, iterative method for the time-integration of reacting flows with stiff chemistry, *J. Comput. Phys.* 295 (2015) 740–769, <https://doi.org/10.1016/j.jcp.2015.04.018>.
- [24] J.O. Hirschfelder, C.F. Curtiss, R.B. Bird, *Molecular Theory of Gases and Liquids*, Wiley, New York, 1954.
- [25] C.F. Curtiss, J.O. Hirschfelder, Transport properties of multicomponent gas mixtures, *J. Chem. Phys.* 17 (6) (1949) 550–555, <https://doi.org/10.1063/1.1747319>.
- [26] J.F. Grcar, J.B. Bell, M.S. Day, The Soret effect in naturally propagating, premixed, lean, hydrogen-air flames, *Proc. Combust. Inst.* 32 (2009) 1173–1180, <https://doi.org/10.1016/j.proci.2008.06.075>.
- [27] J. Schlup, G. Blanquart, A reduced thermal diffusion model for H and H_2 , *Combust. Flame* 191 (2018) 1–8, <https://doi.org/10.1016/j.combustflame.2017.12.022>.
- [28] R. Kee, F. Rupley, J. Miller, Chemkin-II: a Fortran chemical kinetics package for the analysis of gas-phase chemical kinetics, Sandia National Laboratories Report SAND89-8009, 1989.
- [29] G. Dixon-Lewis, Flame structure and flame reaction kinetics. II. Transport phenomena in multicomponent systems, *Proc. R. Soc., Math. Phys. Eng. Sci.* 307 (1488) (1968) 111–135, <https://doi.org/10.1098/rspa.1968.0178>.
- [30] Y. Xuan, G. Blanquart, M.E. Mueller, Modeling curvature effects in diffusion flames using a laminar flamelet model, *Combust. Flame* 161 (5) (2014) 1294–1309, <https://doi.org/10.1016/j.combustflame.2013.10.028>.
- [31] Y. Xuan, G. Blanquart, Effects of aromatic chemistry-turbulence interactions on soot formation in a turbulent non-premixed flame, *Proc. Combust. Inst.* 35 (2) (2015) 1911–1919, <https://doi.org/10.1016/j.proci.2014.06.138>.
- [32] Y. Xuan, G. Blanquart, Two-dimensional flow effects on soot formation in laminar premixed flames, *Combust. Flame* 166 (2016) 113–124, <https://doi.org/10.1016/j.combustflame.2016.01.007>.
- [33] P.L. Carroll, G. Blanquart, A proposed modification to Lundgren's physical space velocity forcing method for isotropic turbulence, *Phys. Fluids* 25 (10) (2013) 105114, <https://doi.org/10.1063/1.4826315>.
- [34] S. Verma, Y. Xuan, G. Blanquart, An improved bounded semi-Lagrangian scheme for the turbulent transport of passive scalars, *J. Comput. Phys.* 272 (2014) 1–22, <https://doi.org/10.1016/j.jcp.2014.03.062>.
- [35] M.E. Mueller, H. Pitsch, LES model for sooting turbulent nonpremixed flames, *Combust. Flame* 159 (6) (2012) 2166–2180, <https://doi.org/10.1016/j.combustflame.2012.02.001>.
- [36] F. Bisetti, G. Blanquart, M.E. Mueller, H. Pitsch, On the formation and early evolution of soot in turbulent nonpremixed flames, *Combust. Flame* 159 (1) (2012) 317–335, <https://doi.org/10.1016/j.combustflame.2011.05.021>.
- [37] B. Savard, G. Blanquart, Broken reaction zone and differential diffusion effects in high Karlovitz $n-C_7H_{16}$ premixed turbulent flames, *Combust. Flame* 162 (5) (2015) 2020–2033, <https://doi.org/10.1016/j.combustflame.2014.12.020>.
- [38] M. Herrmann, G. Blanquart, V. Raman, Flux corrected finite volume scheme for preserving scalar boundedness in reacting large-eddy simulations, *AIAA J.* 44 (12) (2006) 2879–2886, <https://doi.org/10.2514/1.18235>.
- [39] C. Pierce, *Progress-variable approach for large-eddy simulation of turbulent combustion*, Ph.D. thesis, Stanford University, 2001.
- [40] L. Shunn, F. Ham, P. Moin, Verification of variable-density flow solvers using manufactured solutions, *J. Comput. Phys.* 231 (9) (2012) 3801–3827, <https://doi.org/10.1016/j.jcp.2012.01.027>.
- [41] O. Desjardins, G. Blanquart, G. Balarac, H. Pitsch, High order conservative finite difference scheme for variable density low Mach number turbulent flows, *J. Comput. Phys.* 227 (15) (2008) 7125–7159, <https://doi.org/10.1016/j.jcp.2008.03.027>.
- [42] R.D. Falgout, U.M. Yang, *hypre: a library of high performance preconditioners*, in: P.M.A. Soot, A.G. Hoekstra, C.J.K. Tan, J.J. Dongarra (Eds.), *Computational Science – ICCS 2002*, Springer, Berlin, Heidelberg, 2002, pp. 632–641.
- [43] H.A. van der Vorst, Bi-CGSTAB: a fast and smoothly converging variant of bi-CG for the solution of nonsymmetric linear systems, *SIAM J. Sci. Statist. Comput.* 13 (2) (1992) 631–644, <https://doi.org/10.1137/0913035>.
- [44] M. Frigo, S. Johnson, The design and implementation of FFTW3, *Proc. IEEE* 93 (2) (2005) 216–231, <https://doi.org/10.1109/jproc.2004.840301>.
- [45] L.F. Richardson, The approximate arithmetical solution by finite differences of physical problems involving differential equations, with an application to the stresses in a masonry dam, *Philos. Trans. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci.* 210 (459–470) (1911) 307–357, <https://doi.org/10.1098/rsta.1911.0009>.
- [46] F. Perini, E. Galligani, R.D. Reitz, A study of direct and Krylov iterative sparse solver techniques to approach linear scaling of the integration of chemical kinetics with detailed combustion mechanisms, *Combust. Flame* 161 (5) (2014) 1180–1195, <https://doi.org/10.1016/j.combustflame.2013.11.017>.
- [47] Z. Hong, D.F. Davidson, R.K. Hanson, An improved H_2/O_2 mechanism based on recent shock tube/laser absorption measurements, *Combust. Flame* 158 (4) (2011) 633–644, <https://doi.org/10.1016/j.combustflame.2010.10.002>.
- [48] K.-Y. Lam, D.F. Davidson, R.K. Hanson, A shock tube study of $H_2 + OH \rightarrow H_2O + H$ using OH laser absorption, *Int. J. Chem. Kinet.* 45 (6) (2013) 363–373, <https://doi.org/10.1002/kin.20771>.
- [49] Z. Hong, K.-Y. Lam, R. Sur, S. Wang, D.F. Davidson, R.K. Hanson, On the rate constants of $OH + HO_2$ and $HO_2 + HO_2$: a comprehensive study of H_2O_2 thermal decomposition using multi-species laser absorption, *Proc. Combust. Inst.* 34 (1) (2013) 565–571, <https://doi.org/10.1016/j.proci.2012.06.108>.
- [50] D.G. Goodwin, H.K. Moffat, R.L. Speth, Cantera: An Object-oriented Software Toolkit for Chemical Kinetics, Thermodynamics, and Transport Processes, Version 2.3.0, <https://doi.org/10.5281/zenodo.170284>, <https://www.cantera.org>, 2017.
- [51] C. Rosales, C. Meneveau, Linear forcing in numerical simulations of isotropic turbulence: physical space implementations and convergence properties, *Phys. Fluids* 17 (9) (2005) 095106, <https://doi.org/10.1063/1.2047568>.

- [52] S. Lapointe, B. Savard, G. Blanquart, Differential diffusion effects, distributed burning, and local extinctions in high Karlovitz premixed flames, *Combust. Flame* 162 (9) (2015) 3341–3355, <https://doi.org/10.1016/j.combustflame.2015.06.001>.
- [53] F.A. Williams, *Combustion Theory*, Benjamin/Cummings, 1985.
- [54] M. Day, J. Bell, P. Bremer, V. Pascucci, V. Beckner, M. Lijewski, Turbulence effects on cellular burning structures in lean premixed hydrogen flames, *Combust. Flame* 156 (5) (2009) 1035–1045, <https://doi.org/10.1016/j.combustflame.2008.10.029>.
- [55] A.J. Aspden, M.S. Day, J.B. Bell, Turbulence-flame interactions in lean premixed hydrogen: transition to the distributed burning regime, *J. Fluid Mech.* 680 (2011) 287–320, <https://doi.org/10.1017/jfm.2011.164>.
- [56] A. Aspden, M. Day, J. Bell, Turbulence-chemistry interaction in lean premixed hydrogen combustion, *Proc. Combust. Inst.* 35 (2) (2015) 1321–1329, <https://doi.org/10.1016/j.proci.2014.08.012>.
- [57] G. Blanquart, H. Pitsch, Thermochemical properties of polycyclic aromatic hydrocarbons (PAH) from G3MP2B3 calculations, *J. Phys. Chem. A* 111 (28) (2007) 6510–6520, <https://doi.org/10.1021/jp068579W>.
- [58] G. Blanquart, P. Pepiot-Desjardins, H. Pitsch, Chemical mechanism for high temperature combustion of engine relevant fuels with emphasis on soot precursors, *Combust. Flame* 156 (3) (2009) 588–607, <https://doi.org/10.1016/j.combustflame.2008.12.007>.
- [59] K. Narayanaswamy, G. Blanquart, H. Pitsch, A consistent chemical mechanism for oxidation of substituted aromatic species, *Combust. Flame* 157 (10) (2010) 1879–1898, <https://doi.org/10.1016/j.combustflame.2010.07.009>.
- [60] Y. He, B. Cook, J. Deslippe, B. Friesen, R. Gerber, R. Hartman-Baker, A. Koniges, T. Kurth, S. Leak, W.-S. Yang, Z. Zhao, E. Baron, P. Hauschildt, Preparing NERSC users for Cori, a Cray XC40 system with Intel many integrated cores, *Concurr. Comput. Pract. Exp.* 30 (1) (2017) e4291, <https://doi.org/10.1002/cpe.4291>.
- [61] A.J. Fillo, J. Schlup, G. Beardsell, G. Blanquart, K.E. Niemeyer, Figures, plotting scripts, and data for “A fast, low-cost, and stable memory algorithm for implementing multicomponent transport in direct numerical simulations” [dataset], Zenodo, 2019, <https://doi.org/10.5281/zenodo.3519910>.
- [62] A.J. Fillo, J. Schlup, G. Beardsell, G. Blanquart, K.E. Niemeyer, Direct numerical simulation results for turbulent hydrogen/air flame: multicomponent diffusion model [dataset], Oregon State University, 2019, <https://doi.org/10.7267/2f75rf80s>.
- [63] A.J. Fillo, J. Schlup, G. Beardsell, G. Blanquart, K.E. Niemeyer, Direct numerical simulation results for turbulent hydrogen/air flame: mixture-averaged diffusion model [dataset], Oregon State University, 2019, <https://doi.org/10.7267/c247f0072>.