A computationally-efficient, semi-implicit, iterative method for the time-integration of reacting flows with stiff chemistry

B. Savard, Y. Xuan, B. Bobbitt, G. Blanquart

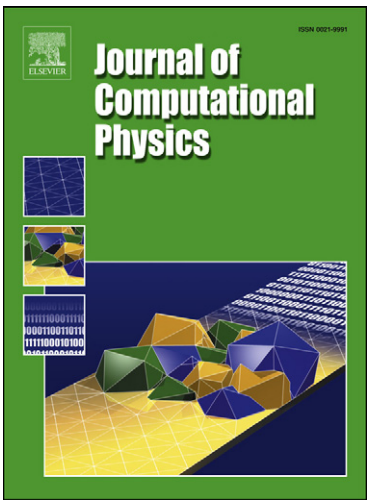

PII: S0021-9991(15)00264-8
DOI: http://dx.doi.org/10.1016/j.jcp.2015.04.018
Reference: YJCPH 5845

To appear in: *Journal of Computational Physics*

Received date: 10 October 2014
Revised date: 20 February 2015
Accepted date: 14 April 2015

Please cite this article in press as: B. Savard et al., A computationally-efficient, semi-implicit, iterative method for the time-integration of reacting flows with stiff chemistry, *J. Comput. Phys.* (2015), http://dx.doi.org/10.1016/j.jcp.2015.04.018

A computationally-efficient, semi-implicit, iterative method for the time-integration of reacting flows with stiff chemistry

B. Savard[a,1], Y. Xuan[b], B. Bobbitt[c], and G. Blanquart[c]

[a] *Graduate Aerospace Laboratories, California Institute of Technology, Pasadena, USA*

[b] *Department of Mechanical and Nuclear Engineering, Pennsylvania State University, University Park, USA*

[c] *Department of Mechanical Engineering, California Institute of Technology, Pasadena, USA*

Abstract

A semi-implicit preconditioned iterative method is proposed for the time-integration of the stiff chemistry in simulations of unsteady reacting flows, such as turbulent flames, using detailed chemical kinetic mechanisms. Emphasis is placed on the simultaneous treatment of convection, diffusion, and chemistry, without using operator splitting techniques. The preconditioner corresponds to an approximation of the diagonal of the chemical Jacobian. Upon convergence of the sub-iterations, the fully-implicit, second-order time-accurate, Crank-Nicolson formulation is recovered. Performance of the proposed method is tested theoretically and numerically on one-dimensional laminar and three-dimensional high Karlovitz turbulent premixed *n*-heptane/air flames. The species lifetimes contained in the diagonal preconditioner are found to capture all critical small chemical timescales, such that the largest stable time step size for the simulation of the turbulent flame with the proposed method is limited by the convective CFL, rather than chemistry. The theoretical and numerical stability limits are in good agreement and are independent of the number of sub-iterations. The results indicate that the overall procedure is second-order accurate in time, free of lagging errors, and the cost per iteration is similar to that of an explicit time integration. The theoretical analysis is extended to a wide range of flames (premixed and non-premixed), unburnt conditions, fuels, and chemical mechanisms. In all cases, the proposed method is found (theoretically) to be stable and to provide good convergence rate for the sub-iterations up to a time step size larger than 1 μs. This makes the proposed method ideal for the simulation of turbulent flames.

Key words: numerical integration, stiff chemistry, semi-implicit preconditioning, iterative method

1 Introduction

Simulations of reacting flow systems using detailed finite-rate chemistry are extremely challenging [1]. The expensive nature of the chemical source terms integration comes from four main challenges: 1) their high non-linearity in the Arrhenius form of the chemical reaction rate constants (*i.e.* high computational cost for each function evaluation) [2], 2) the typically large number of species involved, 3) the strong coupling between chemistry and transport processes (convection and diffusion) [3], and 4) their very large magnitude (or equivalently small timescales) [4,5]. As a result of all these challenges, detailed chemical mechanisms including a large number of species (above 50) and reactions (above 200) have been included in the numerical simulations of reacting flows only for relatively simple geometries (*e.g.* homogeneous or stratified reactors and statistically one-dimensional flames) [6–11]. The number of species (and number of reactions) included in the numerical simulations of two-dimensional and three-dimensional turbulent flames has been relatively limited [12–15,11,16]. Most of these simulations have been focused on investigating the combustion of relatively simple fuels (*e.g.* hydrogen, methane, and ethylene). Only very few studies have considered turbulent flames with large hydrocarbon fuels (*e.g.* propane and *n*-heptane) [11,17] due to the large inherent simulation cost. In order to perform numerical simulations with *detailed* chemical kinetics, robust, accurate, and efficient numerical algorithms are needed for solving the coupled, highly non-linear, multi-dimensional, Partial Differential Equations (PDE) governing the unsteady evolution of these complex reacting systems. There exist various methods to improve the efficiency of chemical source term integration in reacting flow problems. These are reviewed in the following paragraphs and are organized according to the challenges mentioned above.

First, the computational cost associated with the evaluation of exponential functions in the chemical source terms (challenge #1 mentioned above) can be reduced by using single precision calculations, or by tabulating the exponential functions [18]. However, the associated efficiency gain is not significant and the loss of accuracy might be problematic given the large range of timescales in a reacting flow simulation. Alternatively, the non-linear chemical source terms may be expanded using some type of low-order expansion [19–21]. However, it has been shown that these methods may be subject to severe time step size

[1] Corresponding author
Email address: bsavard@caltech.edu (B. Savard)

restrictions and stability issues when applied to combustion simulations [19], since they do not fully account for the non-linearity of the system [20].

Second, regardless of the chosen time-integration scheme, the cost of the chemical models could be alleviated by reducing the total number of species (challenge #2 mentioned above). This can be accomplished using Quasi-Steady-State (QSS) assumptions and Partial-Equilibrium (PE) approximations [22,23], or more advanced methods such as Directed Relation Graph (DRG) [24] and DRG with Error Propagation (DRGEP) [25] before being applied to the simulation [23,26]. Even with these techniques, it has been pointed out in previous work that the size of reduced mechanisms for practical hydrocarbon fuel surrogates is still too large to be used directly in Direct Numerical Simulations (DNS) with finite-rate chemistry [27]. Alternative chemistry reduction techniques based on separation of chemical timescales could be applied. Such techniques include the Computational Singular Perturbation (CSP) method [28] and the Intrinsic Low Dimensional Manifold (ILDM) [4] method. However, these methods require significant computational efforts to conduct chemical Jacobian decomposition and mode separation [27], which makes them not suited for the DNS of multi-dimensional reacting flows.

Third, to avoid the cost associated with the coupled reactive-transport system (challenge #3 mentioned above), most numerical frameworks rely on some variant of splitting techniques (*e.g.* Godunov [29] or Strang [30] splitting) followed by the chemical source term integration in a zero-dimensional setting. These techniques have been widely applied in the numerical simulations of turbulent reacting flows, for instance in the code developed at Lawrence Berkeley National Laboratory [31,32]. The lagging errors introduced by the operator splitting treatment are unimportant for steady-state configurations, but may be substantial in some circumstances, for instance in the proximity of unsteady premixed flame fronts [3,33], and become more severe when running with larger time step sizes. The effects of these errors have been the subject of many previous studies [33–37], and have been found to be case-dependent. Note that, with the application of these techniques, the resulting Ordinary Differential Equations (ODE) associated with the chemistry (instead of the coupled PDEs) remain stiff. To alleviate the high computational overhead associated with the integration of these stiff ODEs, methods relying on implicit numerical schemes based on Backward-Differentiation Formulas (BDF) have been developed [38,39] and implemented in packages such as VODE [40] and DASSL [41,42]. These packages integrate stiff chemical kinetics using BDFs with a modified iterative Newton procedure [41–44], and have been widely adopted in numerical simulations of chemically reacting flows [3,33,34]. Despite the significant computational efficiency gain brought by the stiff chemistry integration techniques discussed above, it is important to recall that these techniques are designed for time-dependent ODE systems (and not PDE), which arise from the application of time-splitting techniques to separate the reactive

(chemical kinetic) part of the PDE system (species transport equations) from the convective-diffusive part [45].

Forth, time-integration techniques designed for the coupled PDEs governing the unsteady evolution of complex reacting systems are discussed in the following. First and foremost, species chemical source terms can be integrated explicitly. For instance, an explicit time-integration scheme, along with QSS assumptions, is used in S3D [46], a massively parallel DNS solver developed at Sandia National Laboratories for the simulations of compressible, turbulent reacting flows. This code has been applied to the simulation of turbulent flames with relatively simple fuels, for instance hydrogen [47,48], methane [49], and ethylene [50]. Heavier fuels have been considered in ignition simulations of HCCI-like systems (*e.g.* *n*-heptane [6], iso-octane [8], and ethanol [9]). The application of explicit time-integration methods are commonly limited by prohibitively small time step sizes to resolve the smallest chemical timescales present in the system (challenge #4) [27,40]. This is the reason why the S3D code relies on "stiffness removal" techniques such as QSSA [22,49,23]. On the other hand, implicit time-integration methods generally yield better stability characteristics and allow for larger integration time step sizes than explicit methods [51]. Unfortunately, fully-implicit methods are generally prohibitively expensive [27], especially for unsteady problems [52], due to the large computational overhead for chemical Jacobian inversion within each time step. This makes fully-implicit time-integration methods prohibitive for simulations of reacting flows with large hydrocarbon fuels (*e.g.* *n*-heptane, kerosene, and diesel), where up to hundreds of species and reactions are typically required [53–55]. Further, Krylov-based iterative methods have been proposed [2,56,57] to reduce the computational burden associated with the construction, storage, and inversion of large, often non-sparse, Jacobian matrices [45,58]. Alternatively, chemical Jacobian diagonal-preconditioning has also been proposed for the time-integration of the PDE system [5,59,60].

While these simple diagonal preconditioners have been used for the simulation of steady-state chemically reacting flows [5,59,60], they were argued to be inappropriate for time-accurate simulations of unsteady flows [5]. That is why a large effort have been put in the development of iterative preconditioning methods for solving the ODEs describing 0D chemical systems. Examples of efficient methods can be found in Ref. [7,56]. However, these methods are typically tailored for very large chemical mechanisms (thousands of species) and are considerably more computationally expensive than explicit time-integration for mechanisms of small to medium sizes (tens to hundreds of species, with hundreds of reactions). In addition, these preconditioning methods rely on a decoupling of the chemistry and transport. In other words, the inversion of the sparse-chemical Jacobian is spatially local and not global. Consequently, computationally less expensive preconditioning iterative methods applied to the PDE system are desirable.

In view of the above discussion, the objective of the current work is to propose a time-integration method designed for the coupled, highly non-linear, PDEs governing the evolution of unsteady reacting flows such as highly turbulent flames. Emphasis is placed on the simultaneous treatment of convection, diffusion, and chemistry, without using operator splitting techniques. As such, a diagonal-preconditioned iterative (to account for the non-linearity of the system) method for the efficient integration of stiff chemistry in the numerical simulation of unsteady chemically reacting flows is proposed in a multi-dimensional setting.

The paper is organized as follows. The governing equations for chemically reacting flows under low Mach number approximation are presented in Section 2. In Section 3, the general iterative solving algorithm is first briefly described. Then, the time-marching step of the species transport equations is shown to be equivalent to a preconditioned Richardson iteration. A semi-implicit preconditioner is finally proposed to improve computational efficiency. In Section 4, the flow configurations used for the numerical tests are presented. Results on the performance of the method are presented in Section 5. Important properties such as convergence, stability, effects of the preconditioning matrix, and temporal accuracy are discussed. Finally, Section 6 includes an extension of the results to other reacting flows and a discussion of the advantages and the limitations of the proposed method.

2 Governing equations

The equations governing the unsteady evolution of the chemically reacting flows considered for the application of the proposed method are described in the following.

2.1 Fluid mechanics

The reacting mixture is assumed to contain a total number of N species and their chemistry is assumed to be given by a chemical kinetics mechanism involving K reactions, with forward and backward reactions counted separately. The chemically reacting flows of interest in the current study are of relatively low Mach number (M_a), typically below 0.3 [61,62,17]. Under this condition, the acoustic waves can be ignored and the pressure field can be decomposed into a spatially-invariant, but (potentially) time-dependent component, $P_0(t)$,

and a fluctuating hydrodynamic pressure, $p(\mathbf{x},t)$ [61–64], with

$$\frac{p(\mathbf{x},t)}{P_0(t)} = O\left(M_a^2\right). \tag{1}$$

To simplify the description of the numerical algorithm (yet without loss of generality), Soret and Dufour effects, body forces, and radiative heat transfer are ignored [3,33,62,65]. In addition, the species molecular diffusion is assumed to be described by the Fickian law [3,33,62–64]. Under these assumptions, the evolution of the system is governed by the following conservation equations of mass, momentum, energy, and species density [10,63,64]

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0 \tag{2}$$

$$\frac{\partial \rho \mathbf{u}}{\partial t} + \nabla \cdot (\rho \mathbf{u} \otimes \mathbf{u}) = -\nabla p + \nabla \cdot \tau \tag{3}$$

$$c_p \left[\frac{\partial \rho T}{\partial t} + \nabla \cdot (\rho \mathbf{u} T)\right] = \nabla \cdot (\rho c_p \alpha \nabla T) + \sum_i c_{p,i} \rho \left(\frac{\alpha}{Le_i} \nabla Y_i + Y_i \mathbf{V_{c,i}}\right) \cdot \nabla T + \dot{\omega}_T \tag{4}$$

$$\frac{\partial \rho Y_i}{\partial t} + \nabla \cdot (\rho \mathbf{u} Y_i) = \nabla \cdot \left(\rho \frac{\alpha}{Le_i} \nabla Y_i\right) + \nabla \cdot (\rho Y_i \mathbf{V_{c,i}}) + \dot{\omega}_i. \tag{5}$$

In the above equations, ρ is the density, $\mathbf{u}$ is the velocity vector, T denotes the temperature of the mixture, and Y_i is the mass fraction of species i. In the momentum equation (Eq. 3), τ is the deviatoric stress tensor, defined as

$$\tau = \mu \left[\nabla \mathbf{u} + (\nabla \mathbf{u})^T\right] - \frac{2}{3} \mu (\nabla \cdot \mathbf{u}) \mathbf{I}, \tag{6}$$

where $\mathbf{I}$ is the identity matrix and μ is the fluid viscosity. In the energy conservation equation (Eq. 4), $\dot{\omega}_T$ includes heat source terms due to chemical reactions, α is the thermal diffusivity, and c_p is the specific heat at constant pressure of the mixture, given by

$$c_p = \sum_{i=1}^{N} Y_i c_{p,i} \tag{7}$$

where $c_{p,i}$ is the specific heat at constant pressure of species i. In the species conservation equations (Eq. 5), $\dot{\omega}_i$ is the chemical source term of species i, and Le_i is the Lewis number of species i, defined as

$$Le_i = \frac{\alpha}{D_i}, \tag{8}$$

with D_i the mass diffusivity for species i. The correction velocity $\mathbf{V_{c,i}}$ in Eq. 5 accounts for gradients in the mixture molecular weight as well as ensures zero

net diffusion flux. It has the following expression [63,64]

$$\mathbf{V_{c,i}} = \frac{\alpha}{Le_i}\frac{\nabla W}{W} - \alpha\left(\sum_{j=1}^{N}\frac{\nabla Y_j}{Le_j}\right) - \alpha\frac{\nabla W}{W}\left(\sum_{j=1}^{N}\frac{Y_j}{Le_j}\right), \tag{9}$$

where

$$W = \left(\sum_{j=1}^{N}\frac{Y_j}{W_j}\right)^{-1} \tag{10}$$

represents the local mean molecular weight of the mixture, and W_j is the molecular weight of species j.

The above set of equations is complemented by the equation of thermodynamic state

$$\rho = \frac{P_0 W}{\widehat{R}T}, \tag{11}$$

where P_0 is the thermodynamic pressure (see Eq. 1) and $\widehat{R}$ is the universal gas constant.

2.2 Chemical model

The overall rate of change of species i, $\dot{\omega}_i$, in Eq. 5 can be split into a production term, $\dot{\omega}_i^+$, and a consumption term, $\dot{\omega}_i^-$, as

$$\dot{\omega}_i = \dot{\omega}_i^+ - \dot{\omega}_i^-. \tag{12}$$

It is important to note that both the production term $\dot{\omega}_i^+$ and the consumption term $\dot{\omega}_i^-$ are positive.

The production rate of species i, $\dot{\omega}_i^+$, is given by the sum of the contributions from all elementary chemical reactions leading to the formation of this species

$$\dot{\omega}_i^+ = W_i \sum_{\substack{j=1 \\ \nu_{ji}>0}}^{r}\left[k_j \prod_{s=1}^{N}\left(\frac{\rho Y_s}{W_s}\right)^{\nu_{js}}\right]. \tag{13}$$

where r is the total number of chemical reactions and ν_{js} is the stoichiometric coefficient of species s in reaction j. In the above expression, the rate constant of reaction j, k_j, is given by the Arrhenius form, $k_j\left(T\right) = A_j T^{b_j}\exp^{-T_{a,j}/T}$, where $T_{a,j}$ is the activation temperature of this reaction. Similarly, the consumption rate of species i, $\dot{\omega}_i^-$, is given by the sum of the contributions from all elementary chemical reactions leading to the destruction of this species

$$\dot{\omega}_i^- = W_i \sum_{\substack{j=1 \\ \nu_{ji}<0}}^{r}\left[k_j \prod_{s=1}^{N}\left(\frac{\rho Y_s}{W_s}\right)^{\nu_{js}}\right]. \tag{14}$$

The local heat release rate is given by

$$\dot{\omega}_T = -\sum_{j=1}^{N} h_i \dot{\omega}_i, \tag{15}$$

where

$$h_i = h_i^0 + \int_{T_0}^{T} c_{p,i} dT, \tag{16}$$

is the specific enthalpy of species i, and h_i^0 denotes its value under standard and reference conditions.

3 Numerical algorithm

As mentioned in the introduction, the objective of the current work is to propose an iterative diagonal preconditioning strategy for the efficient integration of the stiff system of equations introduced in the previous section. Towards this end, a brief description of the flow solver and the numerical algorithms used is given first. Second, the preconditioning method is introduced in the context of the solver previously described. Third, an extension of the preconditioner in a multi-dimensional setting is introduced. Finally, a summary of the characteristics of the method is provided.

3.1 Overview of the numerical solver

The simulations in this work are performed using the structured, multi-physics and multi-scale finite-difference code NGA [61]. The NGA code allows for accurate, robust, and flexible simulations of both laminar and turbulent reactive flows in complex geometries and has been applied in a wide range of test problems, including laminar and turbulent flows [10,66,67], constant and variable density flows [61,68,69], as well as Large-Eddy Simulations (LES) [66,70] and Direct Numerical Simulations (DNS) [12,69,71]. This numerical solver has been shown to conserve discretely mass, momentum, and kinetic energy, with arbitrarily high order spatial discretization [61].

The variable density flow solver in NGA uses both spatially and temporally staggered variables [61]. All scalar quantities (ρ, P, T, Y_i) are stored at the volume centers, and the velocity components are stored at their respective volume faces. The convective term in the species transport equations is discretized using the bounded quadratic upwind biased interpolative convective scheme (BQUICK) [72], and the diffusive term is discretized using a second-order centered scheme. The variables are advanced in time using the second-order semi-implicit Crank-Nicolson scheme of Pierce and Moin [73].

An iterative procedure is applied to fully cover the non-linearities in the Navier-Stokes equations. This iterative procedure has been found of critical importance for stability and accuracy considerations [61,62,73]. The numerical algorithmic sequence for one time step is described below, where a uniform time step Δt is employed. 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 level t^n to t^{n+1}. A total number of Q sub-iterations is assumed. Note that this algorithmic sequence is independent of the preconditioning strategy. As such, for clarity purposes, the chemical source terms are integrated explicitly in the following and the proposed preconditioning strategy is presented in the next section.

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

$$P_0^{n+3/2} = P^{n+1/2} \ , \ \mathbf{Y}_0^{n+3/2} = \mathbf{Y}^{n+1/2} \ , \text{ and } \ (\rho\mathbf{u})_0^{n+1} = (\rho\mathbf{u})^n, \tag{17}$$

where the subscript indicates the index of the sub-iteration. The Adams-Bashforth prediction is used for the initial density evaluation

$$\rho_0^{n+3/2} = 2\rho^{n+1/2} - \rho^{n-1/2}. \tag{18}$$

This ensures that the continuity equation is discretely satisfied at the beginning of the iterative procedure. The vector of chemical source terms is denoted by $\mathbf{\Omega} = (\dot{\omega}_1, ..., \dot{\omega}_N)$, and $\mathbf{\Omega}_0^{n+3/2}$ is evaluated using the thermo-chemical quantities obtained at the conclusion of the previous time step (explicit prediction).

For the sub-iteration $k = 1, \ldots, Q$

1. The scalar fields are advanced in time using the semi-implicit Crank-Nicolson method [61,73] for the convective and diffusive terms, and explicit integration for the chemical source terms

$$\mathbf{Y}_k^* = \frac{\mathbf{Y}^{n+1/2} + \mathbf{Y}_k^{n+3/2}}{2}, \tag{19}$$

$$\begin{aligned}\rho_k^{n+3/2}\mathbf{Y}_{k+1}^{n+3/2} = \rho^{n+1/2}\mathbf{Y}^{n+1/2} + \Delta t\left[\left(\mathbf{C}_k^{n+1} + \mathbf{D}_k^{n+1}\right)\cdot\mathbf{Y}_k^* + \mathbf{\Omega}_k^*\right] \\ + \frac{\Delta t}{2}\left(\frac{\partial \mathbf{C}}{\partial \mathbf{Y}} + \frac{\partial \mathbf{D}}{\partial \mathbf{Y}}\right)_k^{n+1} \cdot \left(\mathbf{Y}_{k+1}^{n+3/2} - \mathbf{Y}_k^{n+3/2}\right).\end{aligned} \tag{20}$$

To simplify the discrete notations for spatial differential operators, the operators corresponding to the convective and diffusive terms in the scalar

equations (Eq. 5) are written as $\mathbf{C}$ and $\mathbf{D}$, respectively. $\frac{\partial \mathbf{C}}{\partial \mathbf{Y}}$ and $\frac{\partial \mathbf{D}}{\partial \mathbf{Y}}$ are the Jacobian matrices corresponding to the convective and diffusive terms, respectively. $\mathbf{C}$ and $\frac{\partial \mathbf{C}}{\partial \mathbf{Y}}$ are functions of the density and the velocity, while $\mathbf{D}$ and $\frac{\partial \mathbf{D}}{\partial \mathbf{Y}}$ are functions of the density and the kinematic viscosity. They are consistently updated at each sub-iteration. Depending on the order of discretization, these operators are generally banded diagonal matrices (*e.g.* tridiagonal for 2nd order discretization and pentadiagonal for 3rd order discretization). It is important to note that the semi-implicit Crank-Nicolson method proposed by Pierce and Moin [73] is not applied to the time-integration of the species chemical source terms, Ω_k^*. As mentioned in the introduction, this is due to the extremely high computational cost associated with the calculation of the chemical Jacobian matrix, $\left(\frac{\partial \mathbf{\Omega}}{\partial \mathbf{Y}}\right)_k^{n+1}$, and the even more expensive inversion of this matrix.

The stiffness of chemically reacting flows is generally believed to be due to the stiff source terms in the species transport equations, but not due to the temperature transport equation [27,74,75]. An estimate for the temperature time scale in a n-C_7H_{16}/air premixed flame (test case presented in Section 4.1) gives $\tau_T \sim \left(T_b - T_u\right) / \left(\dot{\omega}_T / \left(\rho c_p\right)\right)_{\max} \sim 10^{-4}$ s, where T_b and T_u are the burnt and the unburnt temperatures respectively. This time scale is approximately an order of magnitude larger than the time step corresponding to a unity convective CFL in such a laminar flame. Therefore, the temperature equation (Eq. 4) is advanced in time in the exact same fashion as the species mass fractions (Eqs. 19 and 20) without any further implicit treatment. Cases for which the temperature time scale may not be considered large are discussed in Section 6.3. Since the focus is placed on the integration of the chemical source terms in the species equations, the discretized temperature equation is not shown for clarity.

2. The density field is predicted from thermodynamics using

$$\rho_{k+1}^{n+3/2} = \frac{P_0 \left(\sum_{i=1}^{N} \frac{Y_{i,k+1}^{n+3/2}}{W_i} \right)^{-1}}{\widehat{R} T_{k+1}^{n+3/2}}. \tag{21}$$

It is important to note that this density evaluation does not ensure conservation of the species densities, ρY_i, since no density rescaling such as the one proposed by Shunn *et al.* [62] is used. However, upon convergence of the sub-iterations, this formulation is equivalent to the density treatment proposed by Shunn *et al.*

3. The momentum equation is advanced in time using a similar semi-implicit Crank-Nicolson method as for the scalar fields

$$\mathbf{u}_k^* = \frac{\mathbf{u}^n + \mathbf{u}_k^{n+1}}{2}, \tag{22}$$

$$\frac{\rho^{n+1/2}+\rho_{k+1}^{n+3/2}}{2}\widehat{\mathbf{u}}_{k+1}^{n+1}=\frac{\rho^{n-1/2}+\rho^{n+1/2}}{2}\mathbf{u}^{n}+\Delta t\left[\left(\mathbf{C}_{\mathrm{u},k}^{n+1/2}+\mathbf{D}_{\mathrm{u},k}^{n+1/2}\right)\cdot\mathbf{u}_{k}^{*}+\nabla p_{k}^{n+3/2}\right]+\frac{\Delta t}{2}\left(\frac{\partial\mathbf{C}_{\mathrm{u}}}{\partial\mathbf{u}}+\frac{\partial\mathbf{D}_{\mathrm{u}}}{\partial\mathbf{u}}\right)_{k}^{n+1/2}\cdot\left(\widehat{\mathbf{u}}_{k+1}^{n+1}-\mathbf{u}_{k}^{n+1}\right),\quad(23)$$

where $\mathbf{C}_{\mathrm{u}}$ and $\mathbf{D}_{\mathrm{u}}$ are discrete operators associated with the convective and the viscous terms, respectively. $\widehat{\mathbf{u}}$ is the predicted velocity field used to compute the fluctuating hydrodynamic pressure (Step 4).

4. A Poisson equation is then solved for the fluctuating hydrodynamic pressure

$$\nabla^{2}\delta p_{k+1}^{n+3/2}=\frac{1}{\Delta t}\left[\nabla\cdot\left(\frac{\rho^{n+1/2}+\rho_{k+1}^{n+3/2}}{2}\widehat{\mathbf{u}}_{k+1}^{n+1}\right)+\frac{\rho_{k+1}^{n+3/2}-\rho^{n+1/2}}{\Delta t}\right]\quad(24)$$

The Poisson equation is solved using the high-fidelity HYPRE package [61,76]. The predicted velocity field is then updated through a projection step

$$\mathbf{u}_{k+1}^{n+1}=\widehat{\mathbf{u}}_{k+1}^{n+1}-\frac{2\Delta t}{\rho^{n+1/2}+\rho_{k+1}^{n+3/2}}\left(\nabla\delta p_{k}^{n+3/2}\right)\ \text{ and }\ p_{k+1}^{n+3/2}=p_{k}^{n+3/2}+\delta p_{k+1}^{n+3/2}.\quad(25)$$

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

$$\rho^{n+3/2}=\rho_{Q}^{n+3/2},\ p^{n+3/2}=p_{Q}^{n+3/2},\ \mathbf{u}^{n+1}=\mathbf{u}_{Q}^{n+1},\ \text{and}\ \mathbf{Y}^{\mathrm{n}+3/2}=\mathbf{Y}_{\mathrm{Q}}^{\mathrm{n}+3/2}.\quad(26)$$

It is important to note that the above formulation becomes equivalent to the fully-implicit Crank-Nicolson time-integration scheme upon convergence of the sub-iterations [73].

3.2 Preconditioning

Improvement of the above numerical procedure is based on modifying the time-marching step for species mass fraction fields only (step 1 in the procedure described in the previous section). All other intermediate steps are left unchanged.

3.2.1 Preconditioned iterative method

For simpler implementation, the set of equations (Eq. 20) is solved in practice in its residual form

$$\left[\rho_k^{n+3/2}\mathbf{I} - \frac{\Delta t}{2}\left(\frac{\partial \mathbf{C}}{\partial \mathbf{Y}} + \frac{\partial \mathbf{D}}{\partial \mathbf{Y}}\right)_k^{n+1}\right] \cdot \left(\mathbf{Y}_{k+1}^{n+3/2} - \mathbf{Y}_k^{n+3/2}\right)$$
$$= \rho^{n+1/2}\mathbf{Y}^{n+1/2} - \rho_k^{n+3/2}\mathbf{Y}_k^{n+3/2} + \Delta t\left[\left(\mathbf{C}_k^{n+1} + \mathbf{D}_k^{n+1}\right) \cdot \mathbf{Y}_k^* + \mathbf{\Omega}_k^*\right]. \quad (27)$$

The above equation is equivalent to

$$\mathbf{Y}_{k+1}^{n+3/2} = \mathbf{Y}_k^{n+3/2} - \Delta t \mathbf{J}^{-1} \cdot \mathbf{\Theta}_k, \quad (28)$$

where the matrix $\mathbf{J}$ is defined as

$$\mathbf{J} = \rho_k^{n+3/2}\mathbf{I} - \frac{\Delta t}{2}\left(\frac{\partial \mathbf{C}}{\partial \mathbf{Y}} + \frac{\partial \mathbf{D}}{\partial \mathbf{Y}}\right)_k^{n+1}, \quad (29)$$

and the vector

$$\mathbf{\Theta}_k = \frac{\rho_k^{n+3/2}\mathbf{Y}_k^{n+3/2} - \rho^{n+1/2}\mathbf{Y}^{n+1/2}}{\Delta t} - \left[\left(\mathbf{C}_k^{n+1} + \mathbf{D}_k^{n+1}\right) \cdot \mathbf{Y}_k^* + \mathbf{\Omega}_k^*\right] \quad (30)$$

is the error (residual) made on the species transport equation at the previous sub-iteration. When the sub-iterations are fully-converged, the residual, $\mathbf{\Theta}_k$, is zero.

Written in this form, the time-marching for species transport equations described above resembles the standard preconditioned Richardson-type iterative method [77], where the matrix $\mathbf{J}$ acts as a preconditioner. More precisely, the choice of the preconditioner, $\mathbf{J}$, can be arbitrary and does not modify the discrete form of the equations to solve (*i.e.* $\mathbf{\Theta}_k = \mathbf{0}$). It only changes the convergence characteristics of the iterative method. For instance, setting

$$\mathbf{J} = \rho_k^{n+3/2}\mathbf{I}, \quad (31)$$

is equivalent to the fully-explicit integration of the convective, diffusive, and chemical source terms in the species transport equations, while setting

$$\mathbf{J} = \rho_k^{n+3/2}\mathbf{I} - \frac{\Delta t}{2}\left(\frac{\partial \mathbf{C}}{\partial \mathbf{Y}} + \frac{\partial \mathbf{D}}{\partial \mathbf{Y}} + \frac{\partial \mathbf{\Omega}}{\partial \mathbf{Y}}\right)_k^{n+1} \quad (32)$$

corresponds to the full-implicit integration of the convective, diffusive, and chemical source terms in the species transport equations.

Clearly, there is a trade-off in the choice of the preconditioner. Since it is applied at each step of the iterative method, it is preferable to have a preconditioning matrix, $\mathbf{J}$, with low computing and inversion cost. The cheapest preconditioner would therefore be the one described by Eq. 31 (fully-explicit integration), which leads to poor convergence performance requiring extremely small time step sizes. On the other hand, the optimal preconditioner would be the one leading to the fully-implicit integration of the various terms (Eq. 32).

Unfortunately, since the chemical source terms of most species are generally dependent on a large number of other species, the chemical Jacobian matrix, $\frac{\partial \mathbf{\Omega}}{\partial \mathbf{Y}}$, is usually not sparse [7]. Therefore, its construction and inversion may become prohibitively expensive especially when a large number of species is considered.

3.2.2 Semi-implicit preconditioning for stiff chemistry

In an attempt to achieve better convergence characteristics while keeping the form of the preconditioner as simple as possible, a preconditioning method for stiff chemistry, which lies between the fully-explicit and fully-implicit extremes, is proposed.

The proposed preconditioner writes

$$\mathbf{J} = \rho_k^{n+3/2}\mathbf{I} - \frac{\Delta t}{2}\left(\frac{\partial \mathbf{C}}{\partial \mathbf{Y}} + \frac{\partial \mathbf{D}}{\partial \mathbf{Y}} - \mathbf{\Lambda}\right)_k^{n+1}, \tag{33}$$

where $\mathbf{\Lambda}$ is a diagonal matrix defined as

$$\mathbf{\Lambda}_{i,i} = \frac{\dot{\omega}_i^-}{Y_i}. \tag{34}$$

The matrix $\mathbf{\Lambda}$ may be regarded as a very good approximation of the diagonal of the chemical Jacobian:

$$\frac{\partial \dot{\omega}_i}{\partial Y_i} = \frac{\partial \dot{\omega}_i^+}{\partial Y_i} - \frac{\partial \dot{\omega}_i^-}{\partial Y_i} \approx 0 - \frac{\dot{\omega}_i^-}{Y_i} \tag{35}$$

as the production rate of species i (Eq. 13) is not a function of the species mass fraction and its consumption rate (Eq. 14) is often linear in Y_i. The i^{th} element of $\mathbf{\Lambda}$ represents an approximation of the inverse of the timescale corresponding to the chemical consumption of species i (approximation of the inverse of the consumption characteristic times as used in Ref. [5,27]). The proposed preconditioner aims to suppress the small timescales due to the fast consumption of the different species in the system with stiff chemistry.

As the matrix $\mathbf{\Lambda}$ is diagonal, the preconditioner corresponds to a species-wise relaxation, similar to the Jacobi preconditioner. The proposed method accounts for the non-linearities by coupling the transport equations through the iterative procedure (Eq. 30).

The proposed method is inspired by the work of Eberhardt and Imlay [60] who first introduced a diagonal preconditioning in a point-implicit algorithm for the simulation of steady-state reacting flows. The diagonal elements were found by computing the L_2-norm of the corresponding row of the chemical

Jacobian. However, this was found lo lead to an improper approximation of the chemical time scales and resulted in a lack of elemental conservation in time-marching algorithms [5,78,79]. In an effort to improve the accuracy of the method, Ju [5] replaced the jth element of the diagonal preconditioner by the maximum between the inverse of the consumption characteristic time of species j and the inverse of the production time of elementary reaction in which species j is the product. It was argued that this type of preconditioning is suited for the simulation of steady flows, but should fail to provide time-accurate solutions of unsteady flows [5]. An implicit correction with the diagonal approximation of the Jacobian would introduce errors that would accumulate over time.

In the present algorithm, the diagonal preconditioning is applied within an iterative procedure for each time step. This iterative method allows further reduction of the residuals and this is the reason why the method is suitable for the simulation of unsteady reacting flows.

3.3 Extension to multi-dimensions

In multi-dimensional numerical simulations, inverting the proposed preconditioner can be of high computational cost, despite its sparse nature. The method of Approximate Factorization (AF) is therefore used to convert the single, multi-dimensional problem into multiple, one-dimensional problems that can be solved efficiently [61,73]. The transport operator $\mathbf{F} = \frac{\partial \mathbf{C}}{\partial \mathbf{Y}} + \frac{\partial \mathbf{D}}{\partial \mathbf{Y}}$ can be split exactly into directional transport operators $\mathbf{F}_x$, $\mathbf{F}_y$, and $\mathbf{F}_z$, leading to

$$\begin{aligned}\mathbf{J} &= \rho_k^{n+3/2}\mathbf{I} - \frac{\Delta t}{2}\left(\mathbf{F}-\mathbf{\Lambda}\right)_k^{n+1} \\ &= \rho_k^{n+3/2}\mathbf{I} - \frac{\Delta t}{2}\mathbf{F}_{x,k}^{n+1} - \frac{\Delta t}{2}\mathbf{F}_{y,k}^{n+1} - \frac{\Delta t}{2}\mathbf{F}_{z,k}^{n+1} + \frac{\Delta t}{2}\mathbf{\Lambda}_k^{n+1}\end{aligned} \tag{36}$$

in a general three-dimensional orthogonal coordinate system. Accordingly, the following factorization is proposed

$$
\begin{aligned}
&\rho_k^{n+3/2}\mathbf{I} - \frac{\Delta t}{2}\mathbf{F}_{x,k}^{n+1} - \frac{\Delta t}{2}\mathbf{F}_{y,k}^{n+1} - \frac{\Delta t}{2}\mathbf{F}_{z,k}^{n+1} + \frac{\Delta t}{2}\mathbf{\Lambda}_k^{n+1} \\
&= \left(\rho_k^{n+3/2}\mathbf{I} + \frac{\Delta t}{2}\mathbf{\Lambda}_k^{n+1}\right) \\
&\cdot \left(\mathbf{I} - \frac{\Delta t}{2}\left(\rho_k^{n+3/2}\mathbf{I} + \frac{\Delta t}{2}\mathbf{\Lambda}_k^{n+1}\right)^{-1} \cdot \mathbf{F}_{x,k}^{n+1}\right) \\
&\cdot \left(\mathbf{I} - \frac{\Delta t}{2}\left(\rho_k^{n+3/2}\mathbf{I} + \frac{\Delta t}{2}\mathbf{\Lambda}_k^{n+1}\right)^{-1} \cdot \mathbf{F}_{y,k}^{n+1}\right) \\
&\cdot \left(\mathbf{I} - \frac{\Delta t}{2}\left(\rho_k^{n+3/2}\mathbf{I} + \frac{\Delta t}{2}\mathbf{\Lambda}_k^{n+1}\right)^{-1} \cdot \mathbf{F}_{z,k}^{n+1}\right) + O\left(\Delta t^2\right)
\end{aligned}
\tag{37}
$$

The proposed factorization does not degrade the temporal accuracy of the preconditioned time-integration scheme, since it introduces a second-order error in time, same order as the one introduced by the temporal discretization (Eqs. 19 and 20).

In the above factorization, the inversion of $\rho_k^{n+3/2}\mathbf{I}+\frac{\Delta t}{2}\mathbf{\Lambda}_k^{n+1}$ is computationally trivial since it is a diagonal matrix. As such, three simpler one-dimensional inversion problems in the x, y, and z directions can be solved sequentially using tridiagonal (2^{nd} order spatial discretization schemes) or pentadiagonal (3^{rd} order spatial discretization schemes) matrix inversion algorithms analytically, in a serial or parallel fashion. This is very important as it keeps the overall cost of any sub-iteration linear with the number of grid points and linear with the number of species.

3.4 Summary

A semi-implicit preconditioning strategy is proposed, in combination with an iterative method, for the time-integration of the stiff chemistry. The proposed method takes advantage of the iterative structure of the NGA code and, more precisely, its semi-implicit formulation of the transport terms. In addition, the method is compatible with the approximate factorization (see Section 3.3), which is necessary to maintain the performance of the code. Note that this preconditioning method could be applied as well within any other iterative algorithm.

The proposed semi-implicit preconditioning method is based on an approximation of the diagonal of the chemical Jacobian. The hypothesis, which will be tested in Section 5, is that the smallest chemical timescales are well approximated by this diagonal preconditioner, allowing the use of larger integration time step sizes. Another assumption that will be tested in Section 5 is that a good convergence rate of the sub-iterations can be obtained with sufficiently large time step sizes. This would allow the total number of operations per time

step to be very similar to that for the explicit time-integration of the chemistry. As mentioned in the introduction, other diagonal preconditioners were found to be associated with a lack of robustness and elemental conservation in non-iterative time-marching algorithms [78,79] or were argued to lack time-accuracy [5]. However, the proposed method is iterative and, upon convergence of the sub-iterations, the fully-implicit formulation is recovered. Therefore, any issues would be alleviated by the convergence of the sub-iterations.

4 Test cases

The performance of the proposed iterative method will be tested in Section 5 on two flow configurations: a one-dimensional unstretched premixed flame and a three-dimensional high Karlovitz turbulent premixed flame. These configurations are presented in the following.

4.1 One-dimensional premixed flame

The one-dimensional laminar unstretched premixed flame is selected as the first test case since it is the most canonical configuration (that includes convection, diffusion, and chemistry) and it is well suited for the quantitative evaluation of the stability and accuracy of a numerical scheme [3,27,33]. The condition of the present test case corresponds to a n-heptane/air flame with an equivalence ratio of 0.9 and an unburnt temperature and pressure of 298 K and 1 atm, respectively. N-heptane is used in this study as a representative of heavy hydrocarbons of relevance to transportation fuels. A reduced finite-rate chemistry model is used in the present work. The mechanism developed in Ref. [12] was reduced from 47 species and 290 reactions to 35 species and 217 reactions in an effort to alleviate the computational cost. Since the gas mixture is slightly lean, species that are only produced under rich conditions (and their associated reactions), namely C_5H_5, C_5H_6, and all aromatic species (benzene, naphthalene...) were removed from the mechanism. As the focus of the present work is placed on the time-integration, the Lewis numbers of all the species are set to unity for the present simulation (*i.e.* no differential diffusion). Under these conditions, the laminar flame speed is $S_L = 29$ cm/s, and the flame thickness is $l_F = 0.43$ mm, with $l_F = (T_b - T_u)/(\partial T/\partial x)_{\max}$. It is expected that using different fuels, mixture compositions, and unburnt conditions will lead to qualitatively similar results. This is discussed in more details in Section 6.

The computational domain is initialized with a fully-converged solution of a stationary flame. The flame front is initially located near the center of the

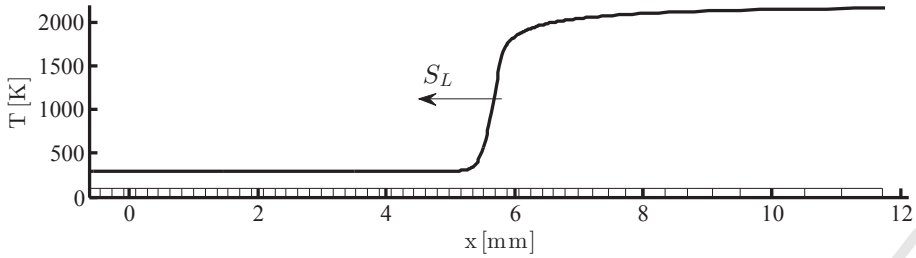

Fig. 1. Schematic of the computational domain and initial condition.

domain ($x_0 = 5.5$ mm) to reduce effects of the boundaries. The left boundary is set to be a wall, and the right boundary is an open flow. Once the simulation is started, the flame front moves towards the unburnt gases (left of the domain) at a speed which corresponds to the laminar flame speed. The length of the 1D computational domain is approximately 30 times the laminar flame thickness. The domain is discretized with a uniform grid cell spacing ($\Delta x = 18$ μm) except behind the initial flame location ($x > 5.5$ mm) where it is stretched with a factor of 1.1 (ratio of the grid cell size to its neighbor). A schematic of the configuration is shown in Fig. 1.

While arbitrarily high order (for the convective and viscous terms) is available in the NGA code, the current work relies on second-order spatial discretization of the viscous and convective terms of the Navier-Stokes equations. Grid convergence tests were performed and revealed a 2$^{\text{nd}}$ order accuracy in space (not shown). This spatial order of accuracy was found to be independent of the proposed time-integration scheme. These tests also determined that 24 points across the flame front (l_F) is sufficient to achieve satisfactory grid independence. In the following, this grid resolution is used for all numerical tests.

4.2 Three-dimensional turbulent premixed flame

The configuration chosen corresponds to the unity Lewis number flame presented in Ref. [17]. Figure 2 presents a schematic of the flow configuration. The left and right ends of the domain correspond to an inflow of unburnt gases and an outflow of burnt gases, respectively. The position of the flame is statistically steady as an inflow velocity equal to the turbulent flame speed is used. Turbulent forcing is employed to avoid a fast decay of turbulence due to viscous dissipation [68]. This forcing is not used at the inlet and the outlet to avoid negative velocities. Forcing starts at $0.5L$ after the inlet and is switched off at a distance of $3L$ from the end of the domain, with L being the domain width and height (Fig. 2). Such a distance is found sufficient for the turbulence to decay without forcing.

The chemical mechanism and the flame unburnt conditions are the same as

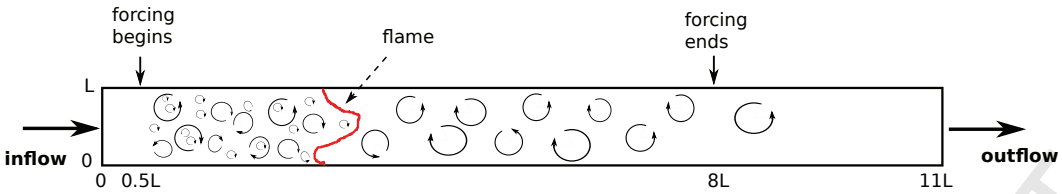

Fig. 2. Schematic of the three-dimensional turbulent premixed flame configuration.

for the one-dimensional premixed flame detailed in Section 4.1. Other important parameters of the simulation are listed in Table 1. Note that prior to this simulation, the flow field is established without any flame; homogeneous isotropic turbulence is obtained in the forced region. Then, the turbulent premixed flame is simulated with finite-rate chemistry for several eddy turnover times. The simulation is performed in parallel using 1920 processors on the cluster Hopper at the National Energy Research Scientific Computing Center (NERSC). Fig. 3 provides visual information about the flame simulated.

Domain width and height L [m]	2.33×10^{-3}
Domain size	$11L \times L \times L$
Grid size	$1408 \times 128 \times 128$
Spatial resolution Δx [m]	1.82×10^{-5}
Kolmogorov length scale η [m]	9×10^{-6}
Simulation time [s]	8.5×10^{-3}
Karlovitz number $Ka = \frac{u'^3 l_F}{S_L^3 l}$	98
Reynolds number $Re_t = \frac{u' l}{\nu_u}$	190

Table 1
Simulation parameters for the three-dimensional turbulent premixed flame. u' is the rms velocity fluctuation, l is the integral length scale, and ν_u is the kinematic viscosity in the unburnt gases.

The simulation is performed with a time step size of $\Delta t = 5.7 \times 10^{-7}$ s, which corresponds to a convective CFL condition of 0.8. With the proposed semi-implicit scheme, the stiffness of the chemical model was found not to impact the stability of the turbulent reacting flow simulation. More details on the stability of the numerical framework and the choice of time step size are provided in the following section.

5 Results

First, a theoretical analysis on the convergence of the sub-iterations for the species transport equations is presented in this section. Second, this analysis

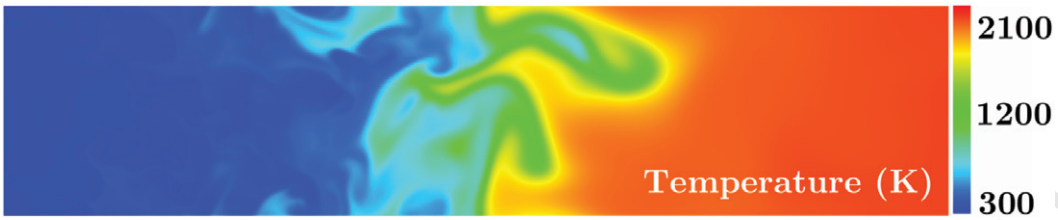

Fig. 3. Contours of temperature on a two-dimensional slice of the three-dimensional turbulent premixed flame.

is further discussed in light of the eigenvalue content of the proposed pre-condition matrix. Third, the convergence of the sub-iterations is evaluated numerically and compared to the theoretical analysis. Forth, the performance of the proposed method in terms of stability is presented both theoretically and numerically, using the test cases previously introduced. Fifth, since stability does not imply accuracy, the numerical accuracy of the proposed method is presented and its dependence on the time step size and the number of sub-iterations is discussed. Sixth, the performance of the proposed method in terms of elemental conservation is presented. In particular, the effects of the iterative procedure (and the number of sub-iterations used) are assessed. Finally, the computational efficiency of the method is discussed.

5.1 Theoretical analysis

To simplify the analysis, it is assumed transport is integrated explicitly (*i.e.* not modified by the sub-iterations). For the sub-iteration $k+1$ within a single iteration, Eq. 28 and 30 take the form

$$\mathbf{J}_k \cdot (\mathbf{Y}_{k+1} - \mathbf{Y}_k) = \rho_0 \mathbf{Y}_0 - \rho_k \mathbf{Y}_k + \Delta t \, (\mathbf{C} + \mathbf{D}) \cdot \mathbf{Y}_0 + \Delta t \boldsymbol{\Omega}_k^*, \tag{38}$$

where

$$\mathbf{J}_k = \rho_k \mathbf{I} + \frac{\Delta t}{2} \boldsymbol{\Lambda}_k. \tag{39}$$

The superscripts (n and $n+1$) have been dropped for clarity. The subscript 0 corresponds to the final solution of the previous iteration. The terms in the equation can be reorganized as

$$\begin{aligned} \left(\mathbf{I} + \frac{\Delta t}{2} \left(\frac{1}{\rho} \boldsymbol{\Lambda} \right)_k \right) \cdot (\mathbf{Y}_{k+1} - \mathbf{Y}_k) = & - \mathbf{Y}_k \\ & + \frac{1}{\rho_k} \left[\rho_0 \mathbf{Y}_0 + \Delta t \, (\mathbf{C} + \mathbf{D}) \cdot \mathbf{Y}_0 \right] \\ & + \frac{\Delta t}{\rho_k} \boldsymbol{\Omega}_k^*. \end{aligned} \tag{40}$$

Expanding $\left(\frac{1}{\rho}\mathbf{\Lambda}\right)_k$, $\left(\frac{1}{\rho}\mathbf{\Omega}^*\right)_k$, and $\frac{1}{\rho_k}$ around $\mathbf{Y}_0$ gives

$$\begin{aligned}
&\left(\mathbf{I}+\frac{\Delta t}{2}\left(\frac{1}{\rho}\mathbf{\Lambda}\right)_0\right)\cdot(\mathbf{Y}_{k+1}-\mathbf{Y}_k)=-\mathbf{Y}_k \\
&+\left[\frac{1}{\rho_0}+\left[\frac{\partial}{\partial\mathbf{Y}}\left(\frac{1}{\rho}\right)\right]_0\cdot(\mathbf{Y}_k-\mathbf{Y}_0)\right]\left[\rho_0\mathbf{Y}_0+\Delta t\,(\mathbf{C}+\mathbf{D})\cdot\mathbf{Y}_0+\Delta t\mathbf{\Omega}_0\right] \\
&+\frac{\Delta t}{2}\left(\frac{1}{\rho}\frac{\partial\mathbf{\Omega}}{\partial\mathbf{Y}}\right)_0\cdot(\mathbf{Y}_k-\mathbf{Y}_0)+\mathcal{O}\left(|\mathbf{Y}_k-\mathbf{Y}_0|^2\right).
\end{aligned} \tag{41}$$

Subtracting Eq. 41 evaluated at two consecutive sub-iterations, and neglecting the second-order terms yields

$$\begin{aligned}
&\left(\mathbf{I}+\frac{\Delta t}{2}\left(\frac{1}{\rho}\mathbf{\Lambda}\right)_0\right)\cdot(\mathbf{Y}_{k+1}-2\mathbf{Y}_k+\mathbf{Y}_{k-1})=-(\mathbf{Y}_k-\mathbf{Y}_{k-1}) \\
&+\left[(\rho\mathbf{Y}+\Delta t\,(\mathbf{C}+\mathbf{D})\cdot\mathbf{Y}+\Delta t\mathbf{\Omega})\otimes\frac{\partial}{\partial\mathbf{Y}}\left(\frac{1}{\rho}\right)\right]_0\cdot(\mathbf{Y}_k-\mathbf{Y}_{k-1}) \\
&+\frac{\Delta t}{2}\left(\frac{1}{\rho}\frac{\partial\mathbf{\Omega}}{\partial\mathbf{Y}}\right)_0\cdot(\mathbf{Y}_k-\mathbf{Y}_{k-1}).
\end{aligned} \tag{42}$$

A simpler expression reads

$$\mathbf{Y}_{k+1}-\mathbf{Y}_k=\mathbf{A}_0\cdot(\mathbf{Y}_k-\mathbf{Y}_{k-1}), \tag{43}$$

with

$$\mathbf{A}_0=\mathbf{I}-\left(\mathbf{I}+\frac{\Delta t}{2}\left(\frac{1}{\rho}\mathbf{\Lambda}\right)_0\right)^{-1}\left(\mathbf{I}-\frac{\Delta t}{2}\left(\frac{1}{\rho}\frac{\partial\mathbf{\Omega}}{\partial\mathbf{Y}}\right)_0-\mathbf{T}_0\right), \tag{44}$$

where

$$\mathbf{T}_0=\left[(\rho\mathbf{Y}+\Delta t\,(\mathbf{C}+\mathbf{D})\cdot\mathbf{Y}+\Delta t\mathbf{\Omega})\otimes\frac{\partial}{\partial\mathbf{Y}}\left(\frac{1}{\rho}\right)\right]_0. \tag{45}$$

The convergence of the sub-iterations is assured as long as the absolute values of all eigenvalues of $\mathbf{A}_0$ are less than unity, *i.e.* the spectral radius of $\mathbf{A}_0$ is less than unity. The opposite implies a divergence (in the linear sense) of the sub-iterations which is likely to be associated with an unstable simulation. Without surprise, in the limit of $\Delta t\to 0$, the spectral radius of $\mathbf{A}_0$ goes to zero and the convergence of the sub-iterations is ensured. For practical time step sizes, large eigenvalues of $\mathbf{A}_0$ can be due to the large magnitude of the chemical Jacobian $\left(\frac{1}{\rho}\frac{\partial\mathbf{\Omega}}{\partial\mathbf{Y}}\right)_0$ or the matrix $\mathbf{T}_0$. However, it can be easily shown that the only eigenvalue of $\mathbf{T}_0$ is a ratio of densities (density of the

explicit prediction *vs.* initial density) which is of order one and is negligible compared to the large eigenvalues of the chemical Jacobian (more details in Section 5.2). Consequently, the matrix $\mathbf{T}_0$ is neglected in the present theoretical analysis. This simplification will be further justified in the numerical tests in Section 5.4.2. Therefore, the rest of the theoretical analysis will consider the following matrix

$$\mathbf{A}'_0 = \mathbf{I} - \left(\mathbf{I} + \frac{\Delta t}{2}\left(\frac{1}{\rho}\mathbf{\Lambda}\right)_0\right)^{-1}\left(\mathbf{I} - \frac{\Delta t}{2}\left(\frac{1}{\rho}\frac{\partial \mathbf{\Omega}}{\partial \mathbf{Y}}\right)_0\right), \tag{46}$$

which is only a function of the full chemical Jacobian, and the approximate (diagonal) Jacobian. The convergence of the sub-iterations is then assured as long as the spectral radius of $\mathbf{A}'_0$ is less than unity.

Alternatively, one can consider the residuals in the relative species mass fractions instead of the residuals in their absolute values. These relative residuals are evaluated as follows

$$\mathbf{Y}^{\mathrm{rel}}_{k+1} - \mathbf{Y}^{\mathrm{rel}}_{k} = \mathbf{A}''_0 \cdot \left(\mathbf{Y}^{\mathrm{rel}}_{k} - \mathbf{Y}^{\mathrm{rel}}_{k-1}\right), \tag{47}$$

where $\mathbf{Y}^{\mathrm{rel}} = \mathbf{G}^{-1}\mathbf{Y}$, with $\mathbf{G} = \mathrm{diag}\left(Y_{0,1}, \ldots, Y_{0,N}\right)$. $Y_{0,i}$ is the mass fraction of species i obtained at the end of the previous iteration. The matrix $\mathbf{A}''_0$ reads

$$\mathbf{A}''_0 = \mathbf{G}^{-1}\mathbf{A}'_0\mathbf{G} = \mathbf{I} - \left(\mathbf{I} + \frac{\Delta t}{2}\left(\frac{1}{\rho}\mathbf{\Lambda}\right)_0\right)^{-1}\left(\mathbf{I} - \frac{\Delta t}{2}\mathbf{G}^{-1}\left(\frac{1}{\rho}\frac{\partial \mathbf{\Omega}}{\partial \mathbf{Y}}\right)_0\mathbf{G}\right). \tag{48}$$

Note that both the absolute and the relative value system of equations are analytically equivalent. In particular, the eigenvalues of $\mathbf{A}''_0$ are identical to those of $\mathbf{A}'_0$.

5.2 Eigenvalue analysis

It is clear from Eq. 46 that if the approximation of the diagonal of the chemical Jacobian (Eq. 34), further referred to as the precondition matrix, were to be equal to the full chemical Jacobian, $\frac{\partial \mathbf{\Omega}}{\partial \mathbf{Y}}$, the scheme would be, in the linear sense, unconditionally stable (recall that $\mathbf{T}_0 = 0$ is assumed). In all other cases, the first pertinent analysis to justify the choice of the preconditioner (Eq. 33) is to evaluate the eigenvalue content of the precondition matrix and compare it to that of the full chemical Jacobian. To ensure a good convergence rate of the sub-iterations, the eigenvalue content of the precondition matrix

should be close to that of the full Jacobian. This assumption is tested in the following.

More specifically, the eigenvalues of the chemical Jacobian $\frac{\partial \mathbf{\Omega}}{\partial \mathbf{Y}}$ correspond to the inverse of the chemical timescales (τ) present in the system. These timescales are associated with the rate of change of the species (or a combination of species) mass fractions in the system in the absence of transport. The idea of using a preconditioner is to allow the use of a time step larger than the smallest of these timescales. Consequently, given a time step Δt, it is important that the chemical timescales smaller than Δt be well represented by the precondition matrix, such that the source terms associated with these timescales are properly integrated over Δt. Note that the eigenvalues of the proposed diagonal precondition matrix correspond to an approximation of the inverse of the reciprocal species lifetimes.

The species lifetimes obtained for the precondition matrix (Eq. 34) and the chemical timescales obtained with the full Jacobian are compared in Fig. 4. These are evaluated with the mixture composition and the temperature ($T = 1615$ K) corresponding to the peak rate of heat release in a one-dimensional flame similar to that presented in Section 4.1. For this *a priori* analysis, the one-dimensional flame is computed with FlameMaster [80]. Note that, *a posteriori*, virtually identical results were obtained when the NGA solution was considered (Section 4.1). This is not surprising as the precondition matrix and the Jacobian only depend on the local mixture composition and temperature, which should not be dependent on the solver used. It is interesting to associate species to each of the chemical timescales shown in Fig. 4. However, it is important to note that not every chemical timescale corresponds to a species lifetime. The timescales are ordered from the smallest to the largest. They are then compared, entry by entry, between the preconditioned and the full Jacobian.

It is well known that because of elemental conservation, n_k of the eigenvalues of the full Jacobian are zero, with n_k the number of elements in the chemical system [4]. Only one of them is zero for the precondition matrix *i.e.* the one associated with N_2, as its source term is identically zero. This is to be expected from the definition of the precondition matrix. Therefore, only the 31 smallest timescales are shown in Fig. 4.

Although Fig. 4 does not provide direct information on the stability limit, it can be observed that all the timescales smaller than about 10^{-5} s are well approximated by the precondition matrix. In other words, the diagonal matrix represents accurately the smallest chemical timescales in the system. As mentioned previously, the proposed method is similar to the Jacobi method, which is guaranteed to converge in diagonal-dominant problems (but it can also converge in other cases). It is therefore interesting to assess if

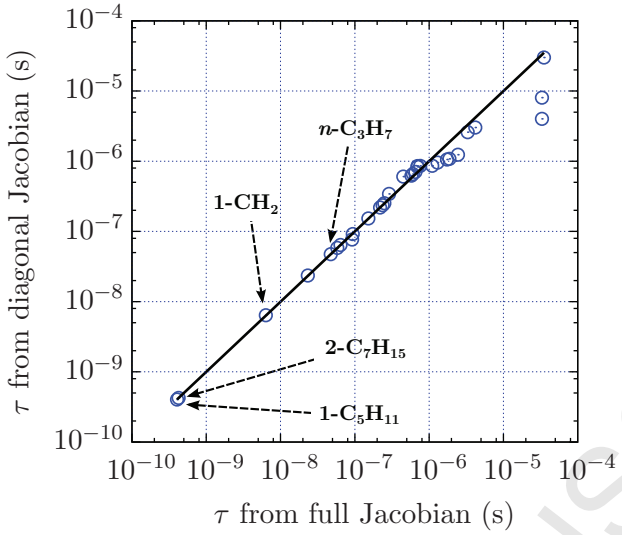

Fig. 4. Comparison of the chemical timescale (τ) of the full chemical Jacobian to the species lifetime of the preconditioned chemical Jacobian.

the matrix $\mathbf{I} - \frac{\Delta t}{2}\left(\frac{1}{\rho}\frac{\partial \mathbf{\Omega}}{\partial \mathbf{Y}}\right)_0$ (Eq. 46) is diagonal dominant for $\Delta t < 10^{-5}$ s. Figure 5(a) shows that the matrix is clearly not diagonal dominant for $\Delta t = 5 \times 10^{-6}$ s. However, the corresponding normalized matrix introduced in Eq. 48, $\mathbf{I} - \frac{\Delta t}{2}\mathbf{G}^{-1}\left(\frac{1}{\rho}\frac{\partial \mathbf{\Omega}}{\partial \mathbf{Y}}\right)_0 \mathbf{G}$ is close to be diagonal dominant for the same Δt, as shown in Fig. 5(b), *i.e.* all the terms of each row are smaller in magnitude than the corresponding term on the diagonal. Since both of these matrices have the same eigenvalue content and the same terms on the diagonal, it becomes clear why the precondition matrix approximates adequately the smaller timescales of the full Jacobian.

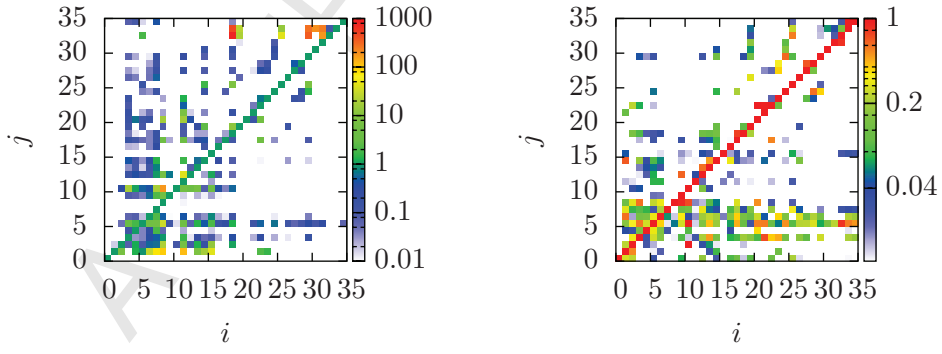

(a) $\mathbf{I} - \frac{\Delta t}{2}\left(\frac{1}{\rho}\frac{\partial \mathbf{\Omega}}{\partial \mathbf{Y}}\right)_0$ (b) $\mathbf{I} - \frac{\Delta t}{2}\mathbf{G}^{-1}\left(\frac{1}{\rho}\frac{\partial \mathbf{\Omega}}{\partial \mathbf{Y}}\right)_0 \mathbf{G}$

Fig. 5. Relative magnitude of the elements of each row compared to the element on the respective diagonals. $\Delta t = 5 \times 10^{-6}$ s.

The results from Fig. 5(b) suggest that the Jacobi method may be successfully

applied to such a matrix (it is guaranteed to converge if the matrix is diagonal dominant). This will be verified in the following sections as the stability and the convergence rate of the method is analyzed considering the spectral radius of $\mathbf{A}_0'$ (or, equivalently, $\mathbf{A}_0''$).

5.3 Convergence of the sub-iterations

The one-dimensional flame test case is used to evaluate numerically the convergence of the sub-iterations. The maximum density residual over the whole domain is investigated as its convergence is controlled by the convergence of all the chemical species. Figure 6 displays the residual through two complete time steps for four different time step sizes. As the time step size decreases, the residuals decrease faster. The rate of convergence of the sub-iterations is observed to follow an exponential relationship, *i.e.* $\mathrm{Res_k} \sim \mathrm{r^k}$ with r the convergence rate. The numerical convergence rate r is computed by fitting an exponential curve to the density residuals. Since density is an analytical function of the species mass fractions, its convergence rate should tend towards that of the slowest converging species mass fraction. In other words, this convergence rate should be close to the spectral radius of $\mathbf{A}_0'$ (Eq. 46).

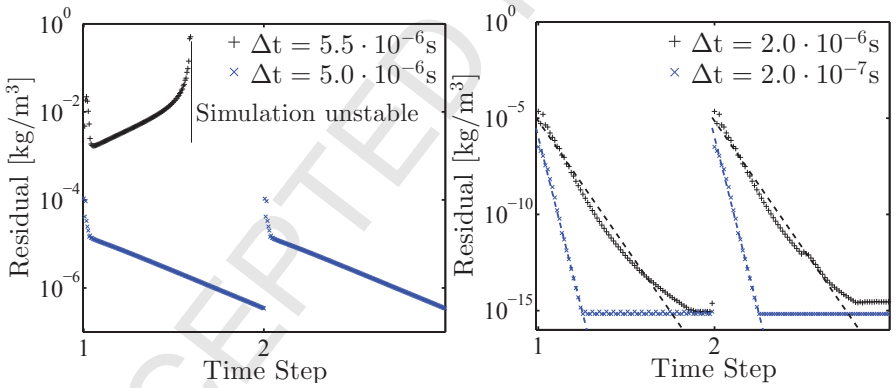

(a) At the stability limit (b) Far from the stability limit

Fig. 6. Evolution of the density residual as a function of sub-iterations over two time steps for the proposed semi-implicit time-integration scheme. The dashed lines correspond to fitted exponential curves averaged over several time steps.

The numerical and theoretical (spectral radius) convergence rates are compared in Fig. 7. The numerical convergence rates are in relatively good agreement with the theoretical values. This further justifies neglecting the variable density matrix, $\mathbf{T}_0$, in the present theoretical analysis (Section 5.1). Using the spectral radius (*i.e.* the largest eigenvalue) as a measure of the convergence rate is a worst case scenario, as the projection of the density residuals on the

associated eigenvector might be identically zero (to machine precision). This might explain partially better numerical convergence rates than theoretically predicted.

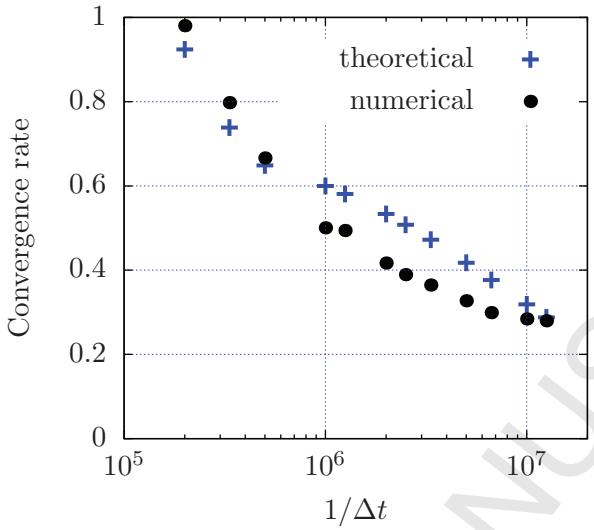

Fig. 7. Rate of convergence of the sub-iterations: comparison between theoretical (largest spectral radius of $\mathbf{A}'_0$) and numerical (rate of convergence of the largest density residuals) results for the one-dimensional flame.

5.4 Stability

The stability of the preconditioned iterative method relies on the stability of the sub-iterations. By stability, it is meant that the solution remains bounded in time. Note that stability does not require the sub-iterations to be fully converged, nor does it imply accuracy. Due to the high non-linearity and the complexity of the governing equations, the stability condition for the proposed preconditioned iterative method and its relation to sub-iteration convergence rate are investigated theoretically (simplified system) and numerically in the present section.

5.4.1 Theoretical stability

Assuming the residuals of the species mass fractions at each sub-iteration remain sufficiently small for the linear analysis presented in Section 5.1 to be valid, then the sub-iterations do not diverge if the spectral radius of $\mathbf{A}'_0$ (Eq. 46), is less than 1. In other words, if this spectral radius is less than 1, stability is ensured, independently of the number of sub-iterations used. However, the accuracy of the solution will be affected by the number of sub-iterations (discussed in Section 5.5).

The spectral radius of $\mathbf{A}'_0$ is plotted as a function of temperature in Fig. 8(a) for different values of Δt. The matrices $\left(\frac{1}{\rho}\mathbf{\Lambda}\right)_0$ and $\left(\frac{1}{\rho}\frac{\partial \mathbf{\Omega}}{\partial \mathbf{Y}}\right)_0$ are evaluated from the one-dimensional flame solution, computed with FlameMaster [80]. Note that the maximum of each curve corresponds to a point in Fig. 7. The spectral radius is shown to be a strong function of temperature. It is not surprising that the chemical system is more sensitive to perturbations at higher temperatures (where the Arrhenius rate constants are larger). It is important to remember that the stability limit in a reacting flow also depends on the nature of the transport terms, which have been ignored for the present analysis, as well as the coupling between the scalar transport and the Navier-Stokes equations. In particular, the most unstable location may not systematically occur at the highest temperature in the domain.

For comparison, the matrix corresponding to $\mathbf{A}'_0$ using an explicit time-integration of the chemical source term is introduced:

$$\mathbf{A}'_{0,\mathrm{exp}} = \frac{\Delta t}{2}\left(\frac{1}{\rho}\frac{\partial \mathbf{\Omega}}{\partial \mathbf{Y}}\right)_0 . \tag{49}$$

The spectral radius of $\mathbf{A}'_{0,\mathrm{exp}}$ is plotted against temperature in Fig. 8(b) for different values of Δt. The (theoretical) stability limit for the explicit scheme is $\Delta t = 5.6 \times 10^{-11}$ s; whereas it is $\Delta t = 6.1 \times 10^{-6}$ s for the semi-implicit scheme. Under the present conditions, it can be clearly observed that the proposed method has the potential to increase the stability limit by several orders of magnitude.

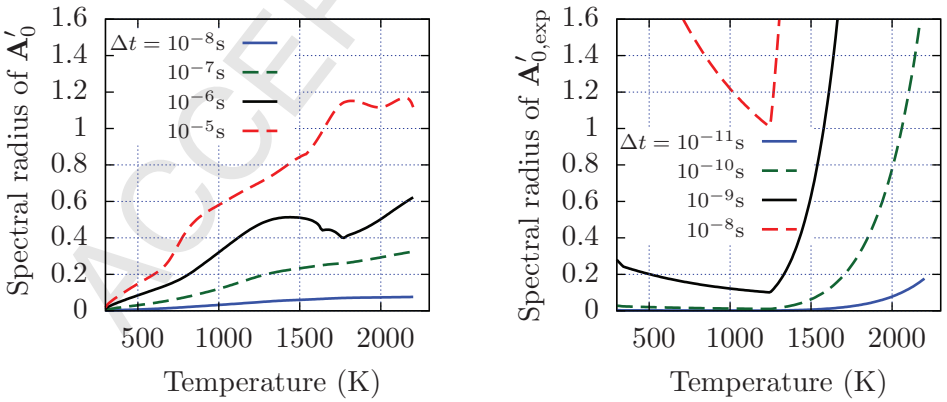

(a) Proposed semi-implicit

(b) Explicit

Fig. 8. Spectral radius of $\mathbf{A}'_0$ and $\mathbf{A}'_{0,\mathrm{exp}}$ as a function of temperature in the one-dimensional premixed flame.

5.4.2 Numerical stability

While the previous section provided a stability criterion through theoretical analysis of a system of reduced complexity, the stability of the scheme is now analyzed numerically in this section. Note that, while the theoretical analysis was done assuming explicit transport, the test cases are performed with semi-implicit transport (as described in Section 3). It is important to demonstrate the performance of the proposed preconditioning method within the algorithmic setting used in practical simulations. For the present tests, the same stability limits were found for the transport terms treated explicitly and implicitly. This is not surprising because the convective CFL number is less than unity for all test cases. It also confirms that the transport terms may be neglected in the previous theoretical analysis (Section 5.1).

The 1D test case is considered first. For all time step sizes tested, it was found that converging (as opposed to *converged*) sub-iterations implied a stable simulation. In other words, unless the sub-iterations diverge, the simulation remains stable. As shown in Fig. 6(a), the largest time step size that can be used for the simulation to be stable is found to be $\Delta t = 5 \times 10^{-6}$ (with $\Delta t = 5.5 \times 10^{-6}$ leading to unstable results). This value is very close to the theoretical stability limit of $\Delta t = 6.1 \times 10^{-6}$ s (Section 5.4.1). Note that the largest numerically stable time step size using an explicit time-integration of the chemical source term is $\Delta t = 2 \times 10^{-10}$ s, also close to the theoretical stability limit ($\Delta t = 5.6 \times 10^{-11}$ s). These stability limits are summarized in Table 2.

	Numerical	Theoretical
Proposed scheme	5×10^{-6} s	6.1×10^{-6} s
Explicit scheme	2×10^{-10} s	5.6×10^{-11} s

Table 2
Largest stable time step size for the proposed semi-implicit scheme and the explicit time-integration of the chemical source terms for the 1D flame test case. Numerical and theoretical results (see Section 5.4.1) are compared.

These results suggest that the stability limit can be well approximated by the theory, *i.e.* the maximum Δt such that the spectral radius of $\mathbf{A}_0'$ (Eq. 46) is less than unity. As such, this should also correspond to the numerical stability limit for the 3D turbulent premixed flame. However, $\Delta t = 5 \times 10^{-6}$ s corresponds to a convective CFL number of approximately 7, which is too large for the overall spatio-temporal scheme to be stable. In other words, for the turbulent case, the largest stable time step size is constrained not by the time-integration of the chemical source terms but by the convective CFL condition. This allowed the simulation to be performed with a convective CFL of 0.8 (see Section 4.2), which corresponds to $\Delta t = 5.7 \times 10^{-7}$ s.

To understand why such a large time step size could be used for the turbulent flame simulation, the spectral radius of $\mathbf{A}'_0$ (Eq. 45) is plotted against temperature throughout the whole domain of the 3D simulation in Fig. 9. For reference, the one-dimensional equivalent is added on top of the scatter plot. This plot suggests that the stability of the chemical system alone (Eq. 38) is only slightly altered by turbulence.

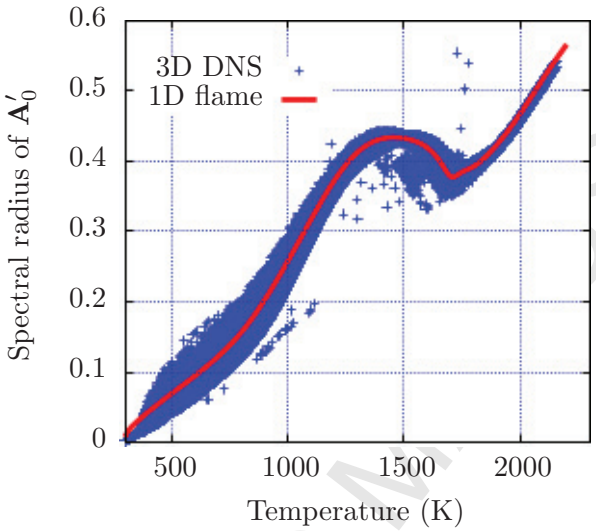

Fig. 9. Scatter plot of the spectral radius of $\mathbf{A}'_0$, with $\Delta t = 5.7 \times 10^{-7}$ s, as a function of temperature in the three-dimensional turbulent flame. The one-dimensional profile is added for comparison.

5.4.3 Summary

It was shown in the present section that the stability limit of the proposed method is well approximated by the largest Δt such that the spectral radius of $\mathbf{A}'_0$ is less than 1. This stability limit is independent of the number of sub-iterations used. On the other hand, the number of sub-iterations does affect accuracy, and it will be discussed in the next section.

An important result is the fact that the stability of the three-dimensional turbulent flame was constrained by the convective CFL limit, rather the time-integration of the chemical source terms. This means that the proposed method has a great potential, in terms of stability, for such flow simulations. Theoretical estimates of the stability limits for other types of flows, fuels, conditions, and chemical mechanisms will be discussed in Section 6.1.

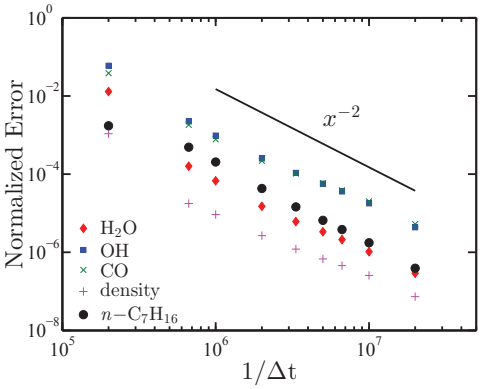

Fig. 10. Temporal accuracy of the method as a function of the time step size for the one-dimensional, propagating flame. The errors for the various species mass fractions are evaluated as the absolute difference of their integrated value in temperature space compared with a reference solution obtained with $\Delta t = 2 \times 10^{-8}$ s.

5.5 Accuracy

In the previous subsection, we established the stability limit(s) of the proposed scheme. We now investigate the accuracy for a given stable simulation.

5.5.1 Order of accuracy

The order of accuracy (*i.e.* the power-dependence of the error as the time step size is reduced) of the proposed method is determined for the 1D (freely propagating) flame test case. In the absence of chemistry, the algorithm used in NGA can be formally shown to be second-order accurate if two or more sub-iterations are used [73]. In practice, four sub-iterations are typically used (to improve stability and to achieve adequate accuracy of the fractional-step) [61,68,69]. For the present test case, four sub-iterations are also used to evaluate the order of accuracy. The impact of the number of sub-iterations on the absolute magnitude of these errors is discussed in the next sub-section.

Simulations with different time steps are performed and results are presented in Fig. 10. The errors for various quantities are evaluated as the absolute difference of their integrated value in temperature space compared with a reference solution obtained with $\Delta t = 2 \times 10^{-8}$ s. The species n-C_7H_{16}, OH, CO, and H_2O are chosen as representatives of reactants, radicals, intermediates, and products. All of these quantities are found to demonstrate second-order accuracy in time, as shown in Fig. 10. It is interesting to note that the expected order of accuracy is already recovered with only four sub-iterations.

5.5.2 Magnitude of errors

It is important to distinguish order of accuracy (as the time step size goes to zero) and absolute magnitude of errors. In order to illustrate the quality of the solution using various time step sizes, Table 3 compares the laminar flame speeds, Fig. 11(a) presents the temperature profiles in physical space (shifted to coincide at $T = 400$ K), and Fig. 11(b) shows the intermediate species n-C_3H_7 mass fraction *vs.* temperature profiles. The first two quantities are chosen as they correspond to the two most important quantities associated with a laminar flame. The n-C_3H_7 mass fraction *vs.* temperature profile is chosen as it is the quantity the most sensitive to the time step size (small reciprocal lifetime where its mass fraction is non-zero). It is clear that, up to a time step size of $\Delta t = 2 \times 10^{-6}$ s, errors are negligible even with only four sub-iterations (0.3% error in laminar flame speed and virtually no difference in the temperature and species profiles). At the stability limit ($\Delta t = 5 \times 10^{-6}$ s), the solution is deteriorated when only four sub-iterations are used. However, using a large number of sub-iterations, the solution reaches similar level of accuracy.

Δt (s)	Q	$\rho_{\max}(\mathbf{A}'_0)$	S_L (cm/s)
2×10^{-7}	4	0.42	28.65
2×10^{-6}	4	0.65	28.57
5×10^{-6}	4	0.92	26.93
5×10^{-6}	200	0.92	28.64

Table 3
Laminar flame speed obtained from simulations with various time step sizes and number of sub-iterations. $\rho_{\max}(\mathbf{A}'_0)$ is the theoretical maximum spectral radius of $\mathbf{A}'_0$ in the flame.

These results demonstrate that, for time step sizes smaller than $\Delta t = 2 \times 10^{-6}$ s, sufficiently accurate solutions for the 1D flame are obtained with as little as four sub-iterations. For time step sizes between $\Delta t = 2 \times 10^{-6}$ s and the stability limit, more sub-iterations are needed to reach sufficient accuracy. This is a direct consequence of the decreasing convergence rates with increasing time step sizes, as presented in Fig. 7. There should exist a pair of time step size/number of sub-iterations such that performance is optimized for a given level of accuracy. However, since the method is meant to be used with turbulent flames, and not 1D flames, such optimization would be of limited interest since the largest time step size allowed by the convective CFL condition in the turbulent flame is smaller than $\Delta t = 2 \times 10^{-6}$ s. Because the spectral radius in the turbulent flame is similar to that of the 1D flame (see Fig. 9), four sub-iterations should be sufficient to obtain accurate solutions. This is further tested below.

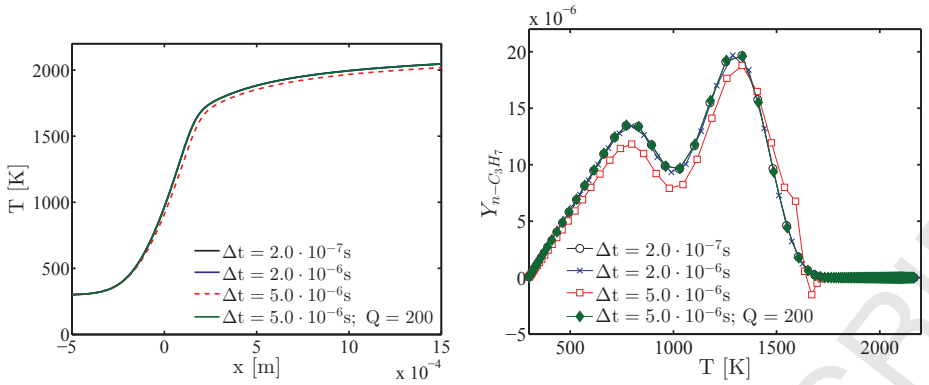

(a) Temperature *vs.* distance

(b) n-C_3H_7 mass fraction *vs.* temperature

Fig. 11. Impact of time step size and number of sub-iteration on the accuracy of 1D propagating flames. When not mentioned, four sub-iterations are used ($Q = 4$).

As mentioned in the previous section, the turbulent premixed flame simulation was performed with a convective CFL number of 0.8 ($\Delta t = 5.7 \times 10^{-7}$ s). In order to evaluate the quality of the solution, a comparative simulation was performed with a smaller time step size, $\Delta t = 8 \times 10^{-8}$ s (four sub-iterations are used for both simulations). Note that this comparative simulation was also run until statistically steady state was reached. While the time step size varies by a factor of 7 between the two simulations, virtually no differences could be identified between the simulations. Figure 12 shows a representative comparison of joint probability density functions of the species mass fraction *vs.* temperature. This type of joint probability density functions is used to evaluate the impact of turbulence on the chemistry [17] and it is important to make sure that any deviation away from a laminar flame is not due to numerical artifacts. The results obtained using both time step sizes are virtually identical, which is consistent with the observations made with the one-dimensional flame test case.

5.6 Mass conservation

As it can be noted from the presentation of the method in Section 3, the sum of the species mass fraction is not implicitly recovered. In other words, with the proposed diagonal preconditioner, the sum of mass fraction is not guaranteed to remain equal to unity. While element conservation is not ensured with the proposed scheme using only one sub-iteration, in practice, with the number of sub-iterations used, elemental mass fractions were found to be adequately conserved. This can be observed in Fig. 13 for the 3D turbulent premixed flame, at a CFL of 0.8 with only four sub-iterations being used. Under the

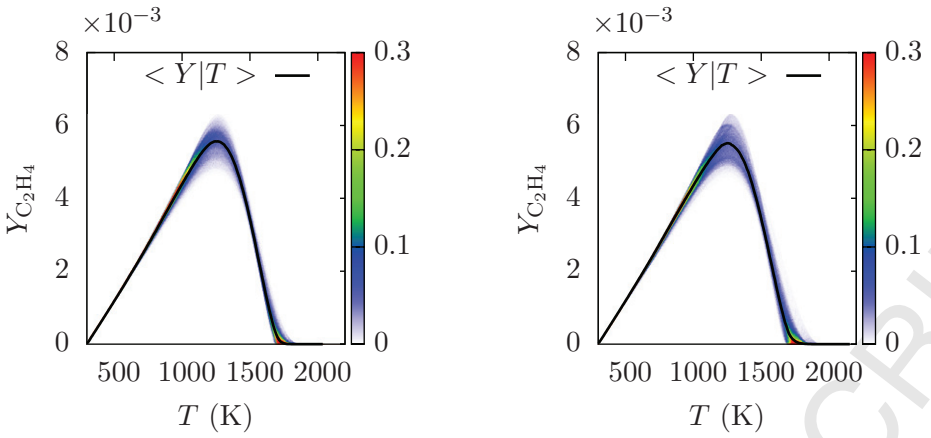

(a) $\Delta t = 5.7 \times 10^{-7}$ s (CFL= 0.8) (b) $\Delta t = 8 \times 10^{-8}$ s

Fig. 12. Comparison between the joint probability density function of the C_2H_4 mass fractions *vs.* temperature obtained with different time step sizes.

assumption of unity Lewis number transport, the elemental mass fractions (Y_C, Y_H, Y_O) should remain perfectly at their inlet values. Any deviations are evidence of mass conservation errors. The maximum deviations in the domain are only about 1% of the inlet values. These maximum deviations are found to occur in the oxidation layer of the turbulent flame.

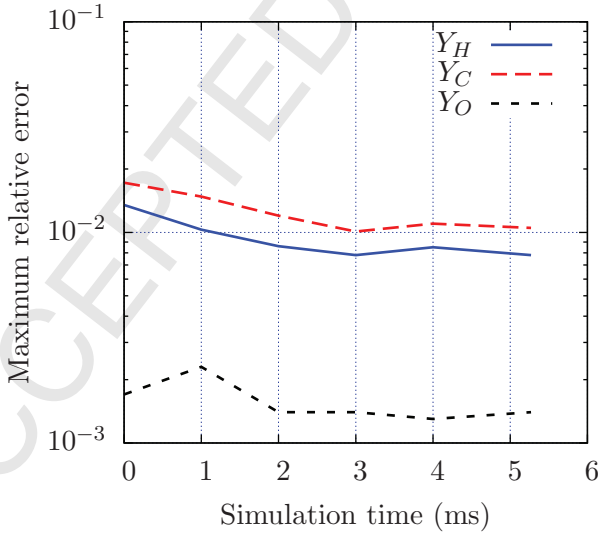

Fig. 13. Maximum deviation in the domain from the inlet elemental mass fractions *vs.* simulation time. The results shown are for the three-dimensional turbulent flame with $\Delta t = 5.7 \times 10^{-7}$ s and four sub-iterations.

Using more sub-iterations leads to better elemental conservation, as shown in Fig. 14 for the 1D flame test case with $\Delta t = 2 \times 10^{-6}$ s. In particular, the rate of convergence of these elemental mass fraction residuals follows the theoretical

spectral radius of 0.65. Note that each of these simulations were performed with a different number of sub-iterations Q until a constant propagation speed (flame speed) and a constant flame structure was reached.

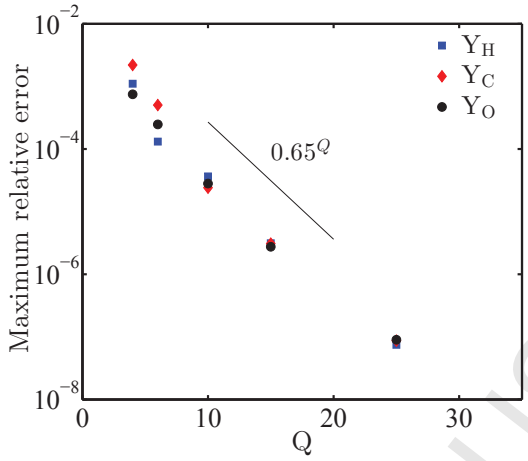

Fig. 14. Maximum deviation in the domain from the inlet elemental mass fractions *vs.* number of sub-iterations, Q. The results shown are for the one-dimensional, propagating flame using $\Delta t = 2 \times 10^{-6}$ s. The theoretical convergence rate is added for comparison.

5.7 Computational efficiency

As presented in Section 3, the cost per sub-iteration with the proposed semi-implicit scheme is virtually identical to that using an explicit time-integration of the chemical source terms. In addition, with the proposed scheme, a number of four sub-iterations was found to be sufficient in practice to achieve adequate accuracy and elemental conservation. As previously mentioned, this number of sub-iterations is typically used for the simulation of non-reacting flows [61,68,69] and reacting flows with explicit time-integration. Therefore, the cost per iteration with the proposed scheme is (in practice) similar to that using an explicit time-integration of the chemical source terms.

Since the proposed scheme does not alter the cost per iteration compared to an explicit time-integration of the chemical source terms, the increase in efficiency (speed-up) is equal to the increase in largest stable time step size. For unsteady flames, the optimal time step size corresponds to the convective CFL limit (the theoretical limit is unity, but 0.8 is the target within our numerical framework). While the time step size in the 1D flame simulations was limited by the chemistry, the convective CFL number was 0.2 ($\Delta t = 2 \times 10^{-6}$ s). This means that for turbulent flames with $u_{\max}/S_{L,b}$ larger than 4, with $u_{\max}$ the maximum velocity in the domain and $S_{L,b}$ the laminar flame speed in the burnt gas ($u_{\max}/S_{L,b} = 1$ for a 1D flame), the time step size will be limited by

the convective CFL limit. This was the case with the 3D turbulent flame test case in which $u_{\max}/S_{L,b} \approx 14$. Consequently, the proposed time-integration method is optimally efficient for such turbulent flames, in the context of the numerical framework of NGA.

6 Discussion

The proposed preconditioner was shown to exhibit very good performance for the test cases analyzed, and in particular it allowed the use of a time step size limited by the convective CFL (3D turbulent flame). In this section, in light of the results presented previously, the theoretical analysis of the stability limit using the proposed scheme is extended to various flames, fuels, unburnt conditions, and kinetic mechanisms. An ignition case is also considered. Additional validation of the numerical stability is provided for a few selected cases. A quantification of the accuracy of the solutions obtained with the proposed method is also presented for these cases. Then, the advantages of the proposed preconditioner over alternative methods are highlighted. Subsequently, the limitations of the present method are discussed. The objective of this section is to help one decide if the proposed method is well suited for a specific unsteady reacting flow simulation.

6.1 Extension

The theoretical analysis presented in Section 5.1 is general and does not depend on the fuel, the chemical mechanism, the flow configuration, or the unburnt conditions. In particular, the eigenvalue analysis (Section 5.2) and the theoretical stability conditions (Section 5.4.1) can be applied to any reacting flow. These analyses, which consider a dependence on the local mixture composition and the temperature only, were further validated in a one-dimensional flame configuration (Section 5.4.2). The results were also argued to be independent of the transport terms. Finally, it was shown that, even in a highly turbulent three-dimensional flame, the departure from a one-dimensional flame solution was not sufficient to significantly influence the stability (both theoretical and numerical). In summary, the theoretical analysis (without transport) is sufficient to determine the stability and the convergence rate of any laminar or turbulent reacting flow.

As such, the theoretical results are extended in this section by considering a wide range of one-dimensional flame and zero-dimensional ignition solutions (each computed with FlameMaster [80]). First, for premixed flames, the effects of the fuel, the unburnt conditions, and the chemical mechanism on the

theoretical stability limits are investigated. Then, the analysis is performed for non-premixed flamelets with different scalar dissipation rates and a 2D coflow diffusion flame. Finally, a homogeneous ignition case at constant pressure is considered.

6.1.1 Premixed flames

The theoretical performance of the proposed semi-implicit method is tested on a series of one-dimensional unstretched premixed flames.

First, the unburnt conditions are kept fixed and various fuels are considered. Four additional fuels are tested: H_2 combined with the chemical model presented in Ref. [81] (9 species 52 reactions), CH_4 with GRI-3.0 [82] (36 species 422 reactions), *i*-C_8H_{18} with CaltechMech v2.1 [83] (171 species 1835 reactions), and *n*-$C_{12}H_{26}$, also with CaltechMech. The analysis performed in Section 5.2 is repeated with these fuels. The species lifetimes at the location of the peak heat release obtained from the semi-implicit precondition matrix (Eq. 34) are compared to the chemical timescales obtained from the full chemical Jacobian in Fig. 15. Same as observed in Section 5.2, for all fuels (and mechanisms) tested, the diagonal Jacobian (precondition matrix) approximates very well almost all the timescales smaller than 10^{-5} s. Similarly, as presented in Table 4, the stability limits using the proposed scheme are very close to the one found in Section 5.4.1.

Second, the same *n*-C_7H_{16} premixed flame with the unburnt conditions presented in Section 4.1 is computed using CaltechMech. Again, the stability limit (Table 4) is only marginally affected by the chemical mechanism.

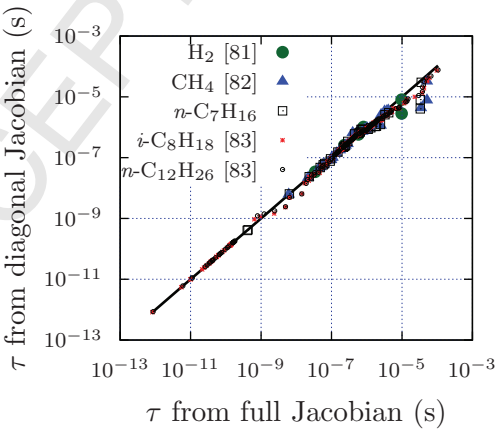

Fig. 15. Comparison of the chemical timescale (τ) of the full chemical Jacobian to the species lifetime of the preconditioned chemical Jacobian at the peak rate of heat release, considering various fuels and chemical mechanisms.

Third, a series of unburnt conditions are used for the *n*-C_7H_{16} premixed flame (using the mechanism introduced in Section 4.1). As the unburnt conditions

encountered in practical combustion devices correspond typically to higher temperature and pressure, both these quantities are increased in this series of tests. Then, the equivalence ratio, which can significantly vary in a combustion device, is modified in the test cases to cover a wide range centered around stoichiometry. Table 4 presents the theoretical stability limit for all these cases. Again, the largest (theoretically) stable time step size vary only slightly throughout all cases.

Fuel	ϕ	P_0 (atm)	T_u (K)	$\Delta t_{\max}$ (s) explicit	$\Delta t_{\max}$ (s) semi-implicit	Equivalent convective CFL
H_2 [81]	0.9	1	298	5.2×10^{-8}	1.6×10^{-6}	1.2
CH_4 [82]	0.9	1	298	2.7×10^{-9}	5.0×10^{-6}	0.47
n-C_7H_{16}	0.9	1	298	5.6×10^{-11}	6.1×10^{-6}	0.76
i-C_8H_{18} [83]	0.9	1	298	2.3×10^{-14}	4.7×10^{-6}	0.45
n-$C_{12}H_{26}$ [83]	0.9	1	298	2.2×10^{-14}	4.6×10^{-6}	0.55
n-C_7H_{16} [83]	0.9	1	298	4.5×10^{-14}	4.6×10^{-6}	0.58
n-C_7H_{16}	0.9	1	400	4.3×10^{-11}	5.2×10^{-6}	0.89
n-C_7H_{16}	0.9	1	600	2.7×10^{-11}	3.8×10^{-6}	1.1
n-C_7H_{16}	0.9	2	298	5.2×10^{-11}	6.2×10^{-6}	1.3
n-C_7H_{16}	0.9	10	298	5.0×10^{-11}	6.5×10^{-6}	3.1
n-C_7H_{16}	0.7	1	298	2.1×10^{-10}	9.5×10^{-6}	0.58
n-C_7H_{16}	1.1	1	298	3.9×10^{-11}	4.3×10^{-6}	0.62
n-C_7H_{16}	1.3	1	298	7.4×10^{-11}	4.3×10^{-6}	0.43

Table 4
Theoretical largest stable time step size for the proposed semi-implicit scheme and the explicit time-integration of the chemical source terms with various unstretched one-dimensional premixed flames.

An important conclusion can be drawn from the results shown in Table 4: the stability limit, using the proposed iterative semi-implicit preconditioning method is only marginally sensitive to the fuel, the unburnt condition, and the mechanism used. As mentioned in Section 5.7, the target time step size is the convective CFL limit. An effective CFL number is therefore computed for each of these 1D premixed flame (see Table 4), assuming 24 grid points per flame thickness (see Section 4.1) are necessary for accurate simulation of a 1D laminar flame ($u_{\max} = S_{L,b}$). The effective CFL number would be larger for turbulent flames and would increase with the turbulent intensity. From these results, it is obvious that even for moderately turbulent flames, the time step

size would be restricted by the convective CFL, rather than the chemistry.

However, these results relate to the stability limit only, and do not suggest anything about accuracy. As shown in Section 5.5, the accuracy is a function of the spectral radius, which is a function of the time step size, and the number of sub-iterations used. Depending on the flow configuration considered (turbulent, laminar, strained,...), four sub-iterations may be sufficient for the solution to be accurate, even at the stability limit. However, the opposite is also possible. Under such circumstance, either the time step size has to be decreased or the number of sub-iterations has to be increased. This choice depends on the spectral radius *vs.* time step size profile (previously shown in Fig. 7).

Figure 16 presents such profiles for flames corresponding to each of the chemical mechanisms used in this section (and presented in Table 4). The first three mechanisms, although used for different flame conditions, exhibit very similar profiles. This means that, with these mechanisms, a moderate decrease in time step size from the stability limit translates in an appreciable decrease in spectral radius. For the 1D test case analyzed in Section 4.1, a time step size three times smaller than the theoretical stability limit was shown to provide sufficient accuracy with only four sub-iterations.

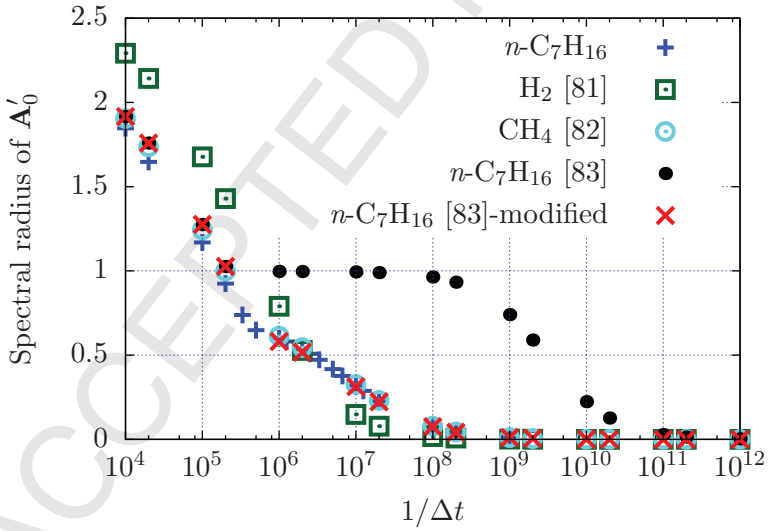

Fig. 16. Spectral radius of $\mathbf{A}'_0$ *vs.* the inverse of the time step size for various cases presented in Table 4. The red cross symbols correspond to the profile for the n-C_7H_16/air flame with the stiff reaction (Eq. 50) removed from CaltechMech.

Interestingly, the spectral radius profile exhibits a plateau just below the stability limit over a wide range of time step size with CaltechMech. This means that, in order to obtain a minimal level of convergence of the sub-iterations, either a very large number of sub-iterations or a very small time step size

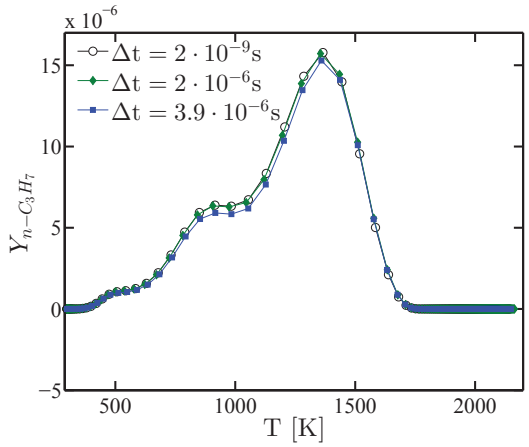

Fig. 17. Profiles of n-C_3H_7 mass fraction *vs.* temperature in the one-dimensional, n-C_7H_{16} propagating flame (similar to the test case of Section 4.1) with CaltechMech. The solutions from using three different time step sizes, each with 4 sub-iterations, are compared.

would be needed. However, this is not what is observed numerically, as shown in Fig. 17, since accurate solutions are obtained with time step sizes as large as half the numerical stability limit, still with only four sub-iterations.

This result is better understood by considering the density residuals *vs.* sub-iterations for the n-C_7H_{16} flame with CaltechMech, with $\Delta t = 2 \times 10^{-6}$ s, presented in Fig. 18 (similar to Fig. 6). Their convergence rate is found to be much closer to the third largest eigenvalue of $\mathbf{A}'_0$ (0.689), as opposed to its largest one (*i.e.* spectral radius), corresponding to 0.997. In other words, the projection of the species mass fractions residuals on the eigenvectors associated with the two largest eigenvalues is negligible. This is consistent with the fact that these two largest eigenvalues are only due to the presence in the chemical mechanism of the following fast reversible reaction

$$A_1C_2H_2{-}C_8H_7 \rightleftarrows A_1C_2H_3^{\star}{-}C_8H_7, \tag{50}$$

which involves species that have negligible mass fractions in the flame simulated (these species are soot precursors and should not be present in the lean flames considered in this paper). The high pressure limit rate constant (the only one available in the literature) was prescribed for this reaction. Such rate constant is obviously too large, especially for the present atmospheric flames. The spectral radius *vs.* time step size profile obtained when this reaction is removed from the chemical mechanism is shown in Fig 16 (red cross symbols). The profile is virtually identical to the one obtained with the 35-species mechanism introduced in Section 4.1. This mean that the proposed method is also efficient with a mechanism as large as CaltechMech, which is far larger than any other mechanism used for the simulation of three-dimensional turbulent premixed flames [11,16,17,49,84–87].

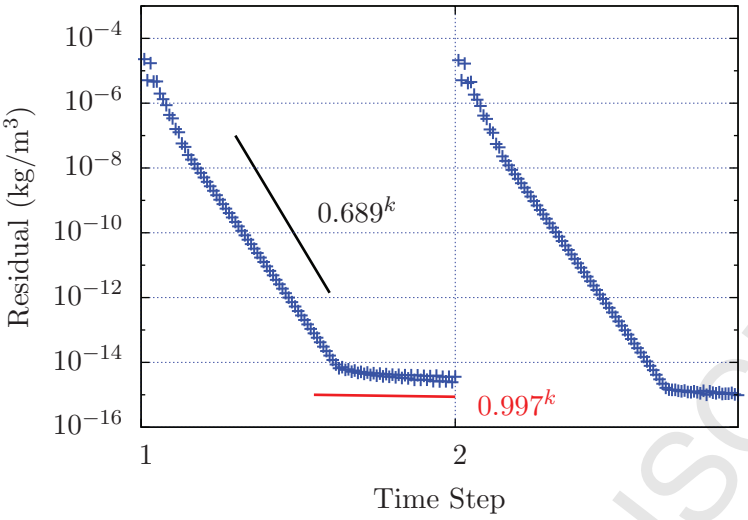

Fig. 18. Evolution of the density residual as a function of sub-iterations over two time steps for the n-C_7H_{16} flame with CaltechMech using the proposed time-integration scheme with $\Delta t = 2 \times 10^{-6}$ s. The convergence rates given by the first (red) and the third (black) largest eigenvalues are shown for comparison. k references to the sub-iteration index.

6.1.2 Non-premixed flames

The theoretical stability limit using the proposed scheme is now evaluated for a series of (unity Lewis number) non-premixed flamelets [88]. These one-dimensional flamelets correspond to solutions close to the axis of symmetry of counter-flow diffusion flames and to local solutions close to the stoichiometric isosurface of mixture fraction of turbulent flames.

First, two n-C_7H_{16}/air flamelets are considered: one with a small scalar dissipation rate (typically found in laminar co-flow diffusion flames and turbulent diffusion flames, at moderate Reynolds number) and one with a large scalar dissipation rate, corresponding to half the dissipation rate leading to extinction. The results are shown in Table 5. Once again, the stability limits using the proposed scheme are very similar to the values found for the series of premixed flames (previous section). Additionally, the dissipation rate does not seem to have a strong effect on the stability of the scheme.

Second, two C_2H_4/air flamelets are considered: again, one with a small scalar dissipation rate, and one with a large dissipation rate. These two flamelets are used to estimate theoretically the stability limit of a 2D-coflow diffusion flame. The 2D flame corresponds to an International Sooting Flam Workshop target flame (more details in Ref. [89,90]). The ethylene fuel (17.6% by mass) is diluted with nitrogen (82.4% by mass) (in both the flamelets and the 2D flame). The steady-state solution for the temperature field is shown in Fig. 19. The two dissipation rates considered in Table 5 correspond to the

Fuel	χ_{st} (1/s)	P_0 (atm)	T_f (K)	T_o (K)	$\Delta t_{\max}$ (s) explicit	$\Delta t_{\max}$ (s) semi-implicit
n-C_7H_{16} [12]	1	1	400	800	1.4×10^{-11}	3.7×10^{-6}
n-C_7H_{16} [12]	320	1	400	800	8.3×10^{-11}	2.9×10^{-6}
C_2H_4 [12]	0.025	4	298	298	4.4×10^{-11}	4.9×10^{-6}
C_2H_4 [12]	138	4	298	298	1.4×10^{-10}	2.7×10^{-6}

Table 5
Theoretical largest stable time step size for the proposed semi-implicit scheme and the explicit time-integration of the chemical source terms with non-premixed flamelets. χ_{st} is the scalar dissipation rate at stoichiometry, T_f the temperature on the fuel side, and T_o the temperature on the oxidizer side. The oxidizer is air and the chemical mechanism considers 47 species and 290 reactions [12].

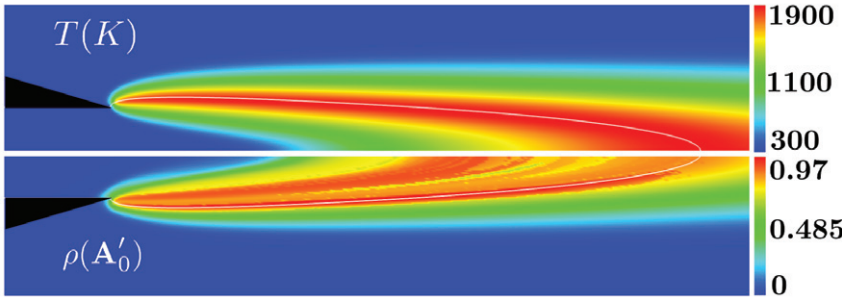

Fig. 19. Contours of temperature (top) and spectral radius of $\mathbf{A}_0'$ (bottom) from the two-dimensional coflow laminar flame, obtained with a time step of 4.0×10^{-6} s.

maximum and minimum values found in the 2D simulation. As expected, the more restrictive time step size is encountered at the largest dissipation rate. In contrast, the 2D numerical simulation was found to be stable up to a time step size of $\Delta t = 4.0 \times 10^{-6}$ s, which is larger than the theoretical prediction of $\Delta t = 2.7 \times 10^{-6}$ s. With this "practical" time step, the maximum spectral radius of $\mathbf{A}_0'$ found in the 2D domain (see Fig. 19) is about 0.97. This difference can be partially explained by the fact that, in the region of largest dissipation rates (at the burner exit), the flame is extinguished and does not compare well with a flamelet.

6.1.3 0D ignition

Although the proposed time-integration scheme was developed primarily for the simulation of multi-dimensional turbulent flames, it could potentially be applied to the simulation of flows with ignition events. In order to partially assess the potential of the method for such flows, a canonical 0D, constant pressure ignition case is considered. The initial conditions as well as the the-

oretical stability limits (computed from a FlameMaster solution) are listed in Table 6. These conditions are meant to be representative of ignition events in HCCI-like engines [6,91]. All simulations are performed with the CaltechMech mechanism.

Fuel	ϕ	P_0 (atm)	T_0 (K)	$\Delta t_{\max}$ (s) explicit	$\Delta t_{\max}$ (s) semi-implicit
n-C_7H_{16} [83]	0.7	30	850	1.1×10^{-12}	3.7×10^{-7}

Table 6
Theoretical largest stable time step size for the proposed semi-implicit scheme and the explicit time-integration with a 0D isobaric ignition case. T_0 is the initial temperature.

The same numerical simulation is performed with the present semi-implicit scheme in NGA to evaluate the practical stability limit. The stability limit identified numerically is $\Delta t = 5.2 \times 10^{-7}$ s, which is close to the theoretical limit. Unfortunately, at such large time step size, using four sub-iterations, the solution is deteriorated: the ignition delay time is over-predicted and the burnt temperature is under-predicted. This is also observed at the theoretical stability limit, as shown in Table 7. Interestingly, with only four sub-iterations and using a time step size of $\Delta t = 2 \times 10^{-7}$ s, the solution is very close to the FlameMaster prediction, as seen in Table 7. While only a limited analysis, the present results show the applicability of the proposed semi-implicit scheme for ignition events.

Framework	Δt (s)	t_{ign} (ms)	T_b (K)
FlameMaster		1.560	2129
NGA	3.7×10^{-7}	1.703	2002
NGA	2.0×10^{-7}	1.566	2126

Table 7
Comparison of the ignition delay time t_{ign} and the burnt temperature obtained with FlameMaster, and with the proposed framework using two different time step sizes. Four sub-iterations are used.

6.2 Advantages over other methods

The performance of the proposed preconditioner is compared to that of the fully-implicit preconditioner, operator-splitting methods, and stiffness removal through QSSA in the following.

6.2.1 Fully-implicit method

As mentioned in the introduction, the use of fully-implicit time-integration of the chemical source terms is known to be prohibitively expensive for the simulation of turbulent reacting flows [27], as the inversion of the full chemical Jacobian, $\frac{\partial \mathbf{\Omega}}{\partial \mathbf{Y}}$, at every point of the domain and at every time step becomes very expensive when more than 10-20 species are considered. Another problem that would arise using a fully-implicit preconditioner with the numerical framework presented in Section 3 is that the extension to multi-dimensions using the approximate factorization introduced in Section 3.3 could no longer be applied, as the fully-implicit chemical Jacobian is not a diagonal matrix. Such a factorization is necessary for the efficiency of the overall procedure.

The major advantage of fully-implicit time-integration of the chemical source terms over the proposed semi-implicit scheme is the use of a time step size not restricted by the chemistry. For steady-state problems, this might be justified/useful. However, for turbulent reacting flows (such as the one presented in Section 4.2), the characteristic hydrodynamic timescales of the turbulent flow (relative to the turbulence and the chemistry) need to be resolved. These are often sufficiently small that the cost increase for the fully-implicit method is not justified anymore.

6.2.2 Operator-splitting methods

The preconditioned iterative method integrates simultaneously the chemical, diffusive, and convective terms at the *same* effective time level. This guarantees that the numerical scheme used is free of lagging errors. These errors are of particular importance in unsteady reacting flows, where chemistry, diffusion, and convection are closely coupled, especially close to the thin flame fronts [3].

Using operator-split formulations, the chemical source terms are decoupled from the diffusive and the convective terms in order to be integrated using stiff ODE solvers. Therefore, the application of these methods for the simulation of reacting flow problems leads typically to integration accuracy degradation [19,92,93]. This is demonstrated in Fig. 20 for the mass fraction of n-C_3H_7. Using Godunov splitting, large numerical errors due to operator-splitting are observed with $\Delta t > 5 \times 10^{-7}$ s (the ODE solver used is DVODE [40] with 10^{-8} and 10^{-20} for the relative and absolute tolerances, respectively [56]). At this point, the numerical time step size surpasses the diffusion timescales. In contrast, the proposed preconditioned iterative method does not suffer from these errors, since the convection, diffusion, and chemistry are all integrated simultaneously. Note that, while Strang splitting is known to perform generally better than Godunov splitting, its extension to a low Mach number code based on spatial and temporal staggering is not trivial. This is the reason

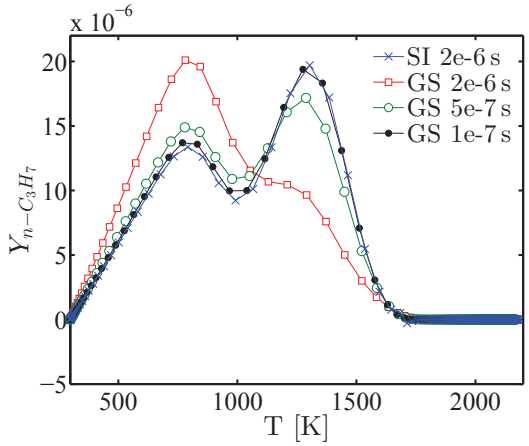

Fig. 20. n-C_3H_7 mass fraction of from the 1D (freely propagating) flame solution. Solutions from using Godunov splitting (GS) are compared, for different time step sizes, to the solution using the proposed preconditioning method.

why Godunov splitting, easily implementable on a staggered grid, was used for comparison.

Figure. 21 presents the computational cost per grid point per simulation time using 1) the proposed preconditioned iterative method, 2) Godunov splitting, and 3) explicit time-integration of the chemical source terms. For all cases, four sub-iterations are used. Without surprise, the cost using small time step sizes ($\Delta t < 5 \times 10^{-8}$ s) is similar for all methods. For Godunov splitting, it is computationally as cheap as the explicit method when the chemical source terms are not stiff, which is the case at small time step sizes. However, using large time step sizes, the chemical source terms become stiff and the cost associated with solving the stiff ODEs increases. At large time step sizes, the cost associated with Godunov splitting increases up to twice larger than that associated with the proposed method. On the other hand, for the proposed preconditioned iterative method, as mentioned earlier, the total number of operations per time step is virtually the same as that associated with an explicit time-integration of the chemical source terms. In particular, the cost per iteration does not vary with time step size. As such, the computational cost per grid point per simulation time is proportional to the inverse of the time step size. In summary, for the present 1D flame test case, with a time step of $\Delta t = 2 \times 10^{-6}$ s, the computational cost associated with the proposed method is smaller than that associated with Godunov splitting, while being free of lagging errors (see Fig. 20).

In the simulation of turbulent flames, the chemical source terms are zero almost everywhere (unburnt/burnt regions in a premixed flame; fuel/oxidizer streams in a non-premixed flame) except at the flame front. This means that, if the domain is partitioned in the direction perpendicular to the flame, the cost of a single time step, using an operator-splitting method, will vary be-

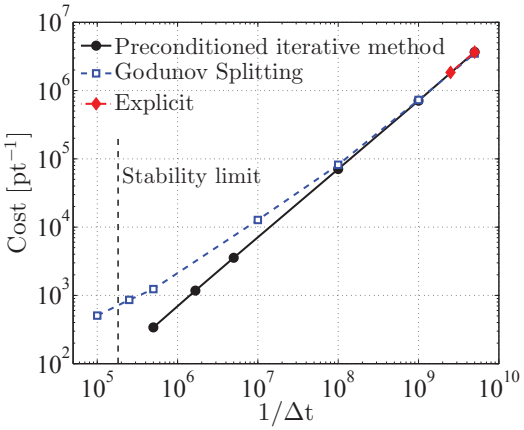

Fig. 21. Computational cost of 1D stationary flame simulation for different chemical integration methods. Cost is calculated as cpu time per point (s/pt), per second of simulation time (s). All cases use four sub-iterations.

tween the different partitions. Unfortunately, a partition cannot advance faster in time than the others. Therefore, the computational time is dictated by the slowest partition. For the three-dimensional simulation test case, this made the simulation impracticable using Godunov splitting. This could be partially alleviated by considering load balancing [94] at the cost of making the code more complicated.

6.2.3 Stiffness removal through QSSA

As mentioned in the introduction, a way to remove the stiffness of the species transport equations (and reduce the number of transport equations) is to put the species with small chemical timescales in Quasi Steady State (QSS) [22]. Application of this method is particularly interesting for compressible codes, for which the stability limit is controlled by either chemistry or acoustics [46]. The acoustics timescale is smaller than the convective timescale (subsonic flows) and may be relatively close to the smallest chemical timescales. As seen in Fig. 4, it is very likely that only a few species (and their associated reactions) are responsible for the small chemical timescales. After removal of these species (and their associated reactions), using QSSA, the stability of the solver would be limited by the acoustics only.

However, putting the species with the smallest associated timescales in quasi-steady state may not always be justified. With the quasi steady state assumption, algebraic expressions can be found for these species. In Fig. 22, these expressions are compared to their true values in the 3D turbulent premixed flame (see Section 4.2). Several species typically placed in QSS in previous studies [49,95] are considered. It is obvious that the QSSA is valid for 1-CH_2, but not for n-C_3H_7 nor 2-C_7H_{15} in the present turbulent premixed flame. This can be explained by the fact that the timescales corresponding to these species,

although very small at high temperature, are very large at low temperatures, as can be seen in Fig. 23. A way to counter this behavior that has been used in the literature [13] is to preheat the unburnt mixture to make the flame more "robust", *i.e.* to make sure that the species responsible for the stiffness of the system can be put in quasi-steady state. However, this obviously modifies the nature of the flame simulated. Note that a recently developed dynamic stiffness removal relies on local, rather than global, QSSA [23]. However, to the best of the authors' knowledge, the method has only been applied to the simulation of ignition problems [6,23].

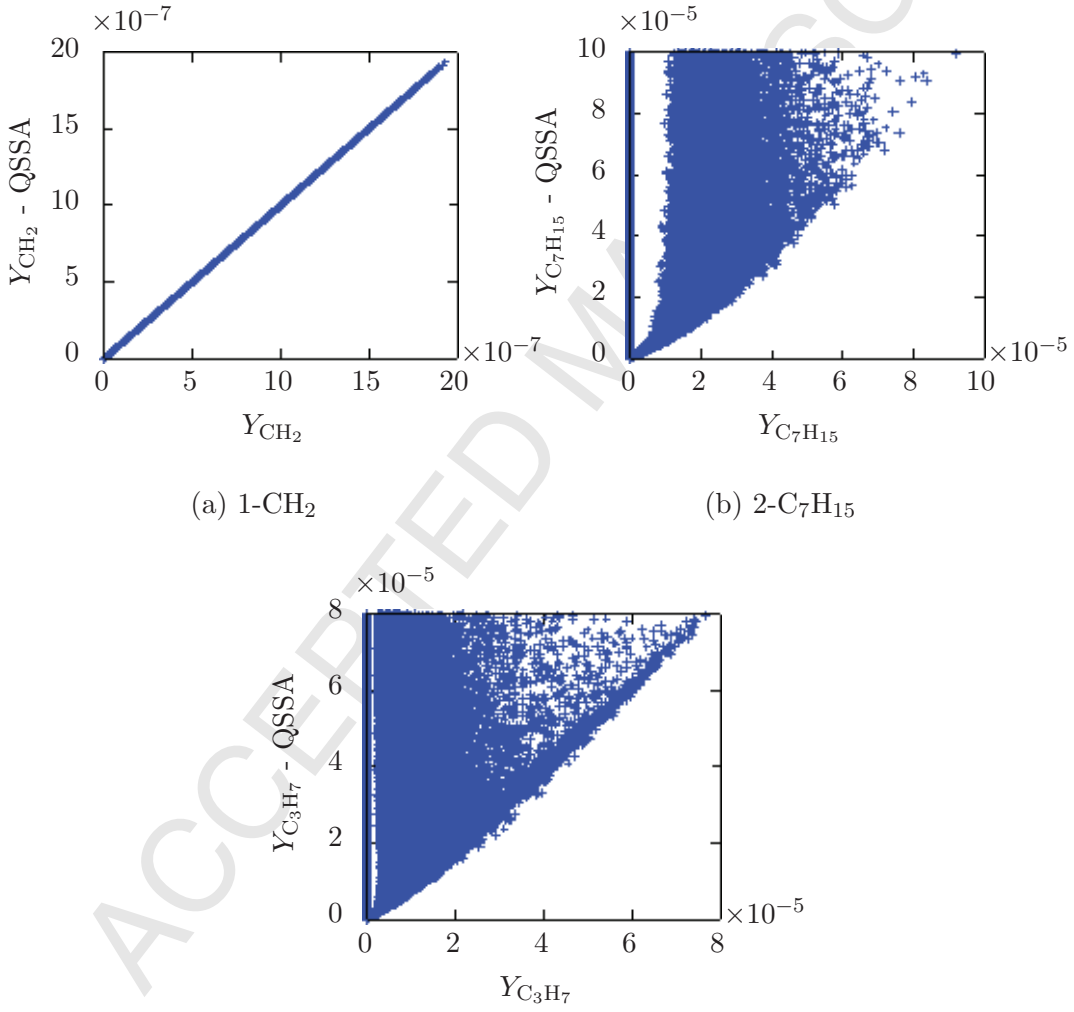

(a) 1-CH_2

(b) 2-C_7H_{15}

(c) *n*-C_3H_7

Fig. 22. Scatter plots of species mass fractions computed from the algebraic expression assuming QSS *vs.* their actual value in the three-dimensional turbulent flame. A straight line ($y = x$) is expected for perfect QSS species.

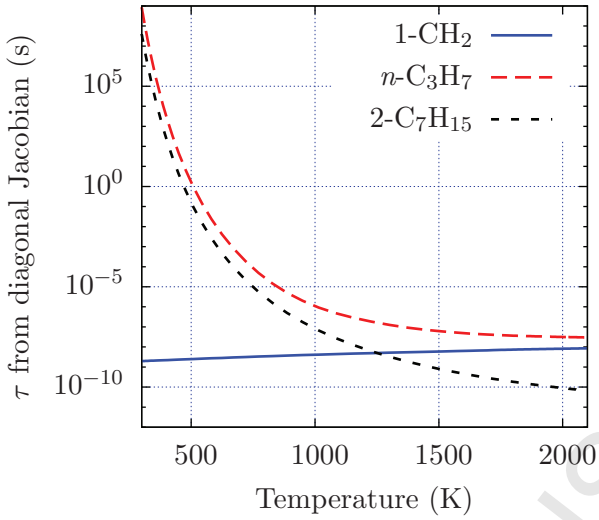

Fig. 23. Chemical consumption timescale associated with 1-CH_2, 2-C_7H_{15}, and n-C_3H_7 *vs.* temperature in the one-dimensional flame.

6.3 Limitations

A first limitation of the method is that it is only efficient for the simulation of unsteady reacting flows, in which the interplay between the flow field and the chemistry has to be captured through adequate temporal resolution (small time step size). While one might decide to use the present method to reach the solution of a steady-state problem, as shown in Section 6.1.2 for the 2D-coflow diffusion flame, the associated cost would be large. In such a case, the use of a large time step size is desirable to reach the time-independent solution. For all the examples provided, the largest stable time step size is of the order of 1×10^{-6} s, which makes the method inefficient to reach a steady-state flow solution.

Second, the method behaves poorly when a fast reversible reaction is present in the chemical mechanism, since this leads to a spectral radius of $\mathbf{A}_0'$ close to unity over a wide range of time step sizes (as discussed in Section 6.1.1). Although such fast reaction was found to be unphysical in CaltechMech (a better reaction rate should be implemented), it is not clear if it is always the case. More importantly, for very large mechanisms as those developed at the Lawrence-Livermore National Laboratories (LLNL) such reactions are present. It is impractical to identify each of these reactions and assess if they can or cannot be removed for the specific reacting flow being simulated (as was done in Section 6.1.1 with CaltechMech). A time-integration method used with such mechanisms has to be efficient even in the presence of such reactions (an example can be found in Ref. [56]). This is achieved at the cost of making the preconditioner more complex (and non-diagonal). Therefore, the proposed

method is not expected to be efficient for the simulation of reacting flows with very large chemical mechanisms such as those developed at LLNL. However, these mechanisms are mainly used for 0D ignition calculations and are too large to be used for the simulation of turbulent flames.

Third, the temperature equation is integrated explicitly, which may limit the largest stable time step size in some reacting flow configurations. For the laminar *n*-heptane/air flame (Section 4.1), the temperature time scale (related to heat release) is of the order of 10^{-4} s. This temperature time scales goes down to 10^{-5} s for the 30 atm 0D ignition case (section 6.1.3). For turbulent reacting flows, the convective CFL limit is generally more restrictive than any of these time scales. Therefore, explicit treatment of the chemistry in the temperature equation should not affect the performance of the proposed scheme for the applications it is intended for, *i.e.* unsteady reacting flows such as turbulent flames. If the method were to be used to simulate flows in which the temperature time scale is smaller than the convective CFL limit, then a similar implicit treatment of the temperature equation as the one proposed for the species equations may be desirable.

Forth, the proposed approximation of the diagonal of the chemical Jacobian (Eq. 35) may, in some cases, introduce non-negligible deviation from the exact diagonal. For species whose consumption rate is mostly due to recombination reactions, the corresponding term in the approximate diagonal may be up to twice smaller (in magnitude). For instance, for the *n*-heptane/air flame tested with the 35-species mechanism, the reaction $OH + OH \rightarrow O + H_2O$ accounts for most of the consumption rate of OH. As a consequence, the exact term in the diagonal of the chemical Jacobian corresponding to OH is about 1.7 times larger than its approximation. However, when the approximation is replaced by the exact diagonal of the Jacobian, the increase in efficiency of the method was found to be negligible. More specifically, the stability limit increases by only 25%, and the convergence rate is unaffected for time step sizes smaller than 2×10^{-6} s. Since computing the exact diagonal requires additional operations, the proposed implementation is marginally more efficient. Under other circumstances, for instance in wall/flame interactions, where the importance of the H recombination reaction has been shown [96], or in hypersonic flows, replacing the approximate diagonal by the exact diagonal may lead to better efficiency.

In summary, use of the proposed preconditioner is particularly relevant to moderately to highly turbulent (premixed or non-premixed) flames (high Karlovitz numbers for premixed flames) in which the convective CFL limit is more restrictive than the largest stable time step size (due to the chemistry) with the proposed time-integration method.

7 Conclusion

A semi-implicit preconditioning strategy, applied to an iterative method, is proposed for the time-integration of the stiff chemistry in the simulation of unsteady reacting flows, such as turbulent flames. The preconditioner consists of an approximation of the diagonal of the chemical Jacobian. It is integrated into the iterative procedure already implemented in the NGA code, in order to account for the non-linearities of the governing equations. Upon convergence of the sub-iterations, the fully-implicit Crank-Nicolson method is recovered. Therefore, the stability of the scheme is dictated by the stability of the sub-iterations.

The performance of the proposed method was numerically tested on two flow configurations: a one-dimensional unstretched premixed flame and a three-dimensional turbulent premixed flame, both with an unburnt mixture of air and *n*-heptane. First, the species lifetimes evaluated from the preconditioned chemical Jacobian represent appropriately the smallest chemical timescales. Second, a theoretical approximation of the rate of convergence of the sub-iterations was derived and shown to be in good agreement with numerical results. Third, the stability limit was found to be well approximated by the theoretical analysis. It was also shown that the stability limit does not depend on the number of sub-iterations. Forth, the method was shown to be second-order accurate in time, even with only four sub-iterations. Increasing the number of sub-iterations led to a reduction of the magnitude of the errors. With a time step size as large as a third of the stability limit, four sub-iterations were shown to be sufficient to achieve acceptable accuracy. Fifth, while other methods using diagonal preconditioned chemical Jacobians have been shown to lack elemental conservation or were argued to not be time-accurate [5,78,79], the proposed method was shown to conserve properly elements over time thanks to the sub-iterations. Sixth, the computational cost of a single iteration with the proposed method is similar to that of an explicit time-integration scheme (since the same number of sub-iterations are used). Therefore, the simulation speed-up achieved with the proposed method corresponds to the increase in the largest stable time step size. For the three-dimensional turbulent premixed flame, the simulation could be performed with a convective CFL of 0.8 (optimal, with or without chemistry).

The theoretical analysis for stability and convergence rate is general and is not limited by the type of fuel, chemical mechanism or flow configuration. Therefore, it was repeated, in the context of one-dimensional premixed flames, with several fuels, unburnt conditions, and chemical mechanism. It was also performed with non-premixed flamelets using different scalar dissipation rates. The method provided good convergence rates of the sub-iterations close to the stability limit for all the chemical mechanisms considered. Consequently, the

proposed preconditioning method showed great potential for the efficient time-integration of turbulent flames. Although not a primary target, the method was also shown to work for a homogeneous ignition case.

The proposed semi-implicit preconditioning, in combination with the iterative method, was argued to be far less computationally expensive than a fully-implicit method and was shown to be as inexpensive or less expensive than operator-splitting methods, while being more accurate. It was also observed that the QSS assumption may not be used for conventional species in the turbulent flame presently studied. As such the proposed method is more suited than alternative methods for the type of flow studied, *i.e.* high Karlovitz flames.

By extension, it was suggested that the proposed method is suited for reacting flows in which the convective timescales are of the order of 10^{-6} s or less. These correspond to moderately to highly turbulent (non-premixed or premixed) flames (high Karlovitz for premixed flames).

Acknowledgements

The authors gratefully acknowledge funding from the U.S. Department of Energy-Basic Energy Sciences (DE-SC006591) under the supervision of Dr. Wade Sisk, the Air Force Office of Scientific Research (Award FA9550-12-1-0144) under the supervision of Dr. Chiping Li, the Natural Sciences and Engineering Research Council of Canada (NSERC Postgraduate Scholarship D), and the Los Angeles chapter of the Achievement Rewards for College Scientists Foundation. This research used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. The authors would also like to thank Nicholas Burali who performed the 2D coflow diffusion flame simulation.

References

[1] M. Smooke, R. Mitchell, D. Keyes, Numerical-solution of 2-dimensional axisymmetric laminar diffusion flames, Combust. Sci. Technol 67 (4–6) (1989) 85–122.

[2] F. Bisetti, Integration of large chemical kinetic mechanisms via exponential methods with Krylov approximations to Jacobian matrix functions, Combust. Theory Model. 16 (3) (2012) 387–418.

[3] H. N. Najm, P. S. Wyckoff, O. M. Knio, A semi-implicit numerical scheme for reacting flow: I. Stiff chemistry, J. Comput. Phys. 143 (2) (1998) 381–402.

[4] U. Maas, S. B. Pope, Simplifying chemical kinetics: intrinsic low-dimensional manifolds in composition space, Combust. Flame 88 (3) (1992) 239–264.

[5] Y. Ju, Lower-upper scheme for chemically reacting flow with finite rate chemistry, AIAA J. 33 (8) (1995) 1418–1425.

[6] C. S. Yoo, T. Lu, J. H. Chen, C. K. Law, Direct numerical simulations of ignition of a lean *n*-heptane/air mixture with temperature inhomogeneities at constant volume: Parametric study, Combust. Flame 158 (9) (2011) 1727–1741.

[7] 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) 1180–1195.

[8] C. S. Yoo, Z. Luo, H. Kim, J. H. Chen, A DNS study of ignition characteristics of a lean *iso*-octane/air mixture under HCCI and SACI conditions, Proc. Comb. Inst. 34 (2013) 2985–2993.

[9] A. Bhagatwala, J. H. Chen, T. Lu, Direct numerical simulations of HCCI/SACI with ethanol, Combust. Flame 161 (2014) 1826–1841.

[10] Y. Xuan, G. Blanquart, Numerical modeling of sooting tendencies in a laminar co-flow diffusion flame, Combust. Flame 160 (9) (2013) 1657–1666.

[11] A. Aspden, M. Day, J. Bell, Lewis number effects in distributed flames, Proc. Comb. Inst. 33 (2011) 1473–1480.

[12] 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.

[13] D. O. Lignell, J. H. Chen, P. J. Smith, T. Lu, C. K. Law, The effect of flame structure on soot formation and transport in turbulent nonpremixed flames using direct numerical simulation, Combust. Flame 151 (1–2) (2007) 2–28.

[14] D. O. Lignell, J. H. Chen, P. J. Smith, Three-dimensional direct numerical simulation of soot formation and transport in a temporally evolving nonpremixed ethylene jet flame, Combust. Flame 155 (1–2) (2008) 316–333.

[15] J. H. Chen, Petascale direct numerical simulation of turbulent combustion - fundamental insights towards predictive models, Proc. Comb. Inst. 33 (1) (2011) 99–123.

[16] A. Aspden, M. Day, J. Bell, Turbulence-flame interactions in lean premixed hydrogen: transition to the distributed burning regime, J. Fluid Mech. 680 (2011) 287–320.

[17] B. Savard, B. Bobbitt, G. Blanquart, Structure of a high Karlovitz *n*-C_7H_{16} premixed turbulent flame, Proc. Comb. Inst. 35 (2015) 1377–1384.

[18] J.-M. Muller, Elementary functions: algorithms and implementation, 2nd Edition, Birkhäuser, 2005.

[19] Y. D'Angelo, B. Larrouturou, Comparison and analysis of some numerical schemes for stiff complex chemistry problems, Modélisation mathématique et analyse numérique 29 (3) (1995) 259–301.

[20] D. L. Ropp, J. N. Shadid, C. C. Ober, Studies of the accuracy of time integration methods for reaction-diffusion equations, J. Comput. Phys. 194 (2) (2004) 544–574.

[21] C. C. Ober, J. N. Shadid, Studies on the accuracy of time-integration methods for the radiation-diffusion equations, J. Comput. Phys. 195 (2) (2004) 743–772.

[22] T. Lu, C. K. Law, A criterion based on computational singular perturbation for the identification of quasi steady state species: A reduced mechanism for methane oxidation with NO chemistry, Combust. Flame 154 (4) (2008) 761–774.

[23] T. Lu, C. K. Law, C. S. Yoo, J. H. Chen, Dynamic stiffness removal for direct numerical simulations, Combust. Flame 156 (8) (2009) 1542–1551.

[24] T. Lu, C. K. Law, A directed relation graph method for mechanism reduction, Proc. Comb. Inst. 30 (1) (2005) 1333–1341.

[25] P. Pepiot-Desjardins, H. Pitsch, An efficient error-propagation-based reduction method for large chemical kinetic mechanisms, Combust. Flame 154 (1) (2008) 67–81.

[26] N. Peters, Reducing mechanisms, in: M. D. Smooke (Ed.), Reduced kinetic mechanisms and asymptotic approximations for methane-air flames, Vol. 384 of Lecture Notes in Physics, Springer Berlin Heidelberg, 1991, pp. 48–67.

[27] X. Gou, W. Sun, Z. Chen, Y. Ju, A dynamic multi-timescale method for combustion modeling with detailed and reduced chemical kinetic mechanisms, Combust. Flame 157 (6) (2010) 1111–1121.

[28] S. Lam, D. Coussis, Understanding complex chemical kinetics with computational singular perturbation, in: Symposium (International) on Combustion, Vol. 22, Elsevier, 1989, pp. 931–941.

[29] K. Bagrinovskii, S. Godunov, Difference schemes for multidimensional problems (in Russian), Dokl. Akad. Nauk. USSR 115 (1957) 431–433.

[30] G. Strang, On the construction and comparison of difference schemes, SIAM J. Numer. Anal. 5 (3) (1968) 506–517.

[31] M. Day, J. Bell, Numerical simulation of laminar reacting flows with complex chemistry, Combust. Theory Model. 4 (4) (1999) 535–556.

[32] A. Almgren, J. Bell, P. Colella, L. Howell, M. Welcome, A conservative adaptive projection method for the variable density incompressible Navier–Stokes equations, J. Comput. Phys. 142 (1998) 1–46.

[33] O. M. Knio, H. N. Najm, P. S. Wyckoff, A semi-implicit numerical scheme for reacting flow: II. Stiff, operator-split formulation, J. Comput. Phys. 154 (2) (1999) 428–467.

[34] R. Yu, J. Yu, X.-S. Bai, An improved high-order scheme for DNS of low Mach number turbulent reacting flows based on stiff chemistry solver, J. Comput. Phys. 231 (16) (2012) 5504–5521.

[35] H. Yoshida, Construction of higher order symplectic integrators, Phys. Lett. A 150 (5) (1990) 262–268.

[36] C. A. Kennedy, M. H. Carpenter, Additive Runge-Kutta schemes for convection-diffusion-reaction equations, Appl. Numer. Math. 44 (1–2) (2003) 139–181.

[37] A.-K. Kassam, L. N. Trefethen, Fourth-order time-stepping for stiff PDEs, SIAM J. Sci. Comput. 26 (4) (2005) 1214–1233.

[38] A. Sandu, J. Verwer, M. V. Loon, G. Carmichael, F. Potra, D. Dabdub, J. Seinfeld, Benchmarking stiff ode solvers for atmospheric chemistry problems-I. implicit vs explicit, Atmos. Environ. 31 (19) (1997) 3151–3166, EUMAC: European Modelling of Atmospheric Constituents.

[39] A. Sandu, J. Verwer, J. Blom, E. Spee, G. Carmichael, F. Potra, Benchmarking stiff ode solvers for atmospheric chemistry problems II: Rosenbrock solvers, Atmos. Environ. 31 (20) (1997) 3459–3472.

[40] P. N. Brown, G. D. Byrne, A. C. Hindmarsh, VODE: A variable-coefficient ODE solver, SIAM J. Sci. Stat. Comput. 10 (5) (1989) 1038–1051.

[41] L. R. Petzold, A Description of DASSL: a differential/algebraic system solver, Sandia National Laboratories Report, SAND282-8637.

[42] K. E. Brenan, S. L. Campbell, A Description of DASSL: a differential/algebraic equation solver, in: Scientific Computing, North-Holland, Amsterdam, The Netherlands, 1983, pp. 65–68.

[43] H. Robertson, The solution of a set of reaction rate equations, Numerical analysis: an introduction (1966) 178–182.

[44] R. D. Saylor, G. D. Ford, On the comparison of numerical methods for the integration of kinetic equations in atmospheric chemistry and transport models, Atmos. Environ. 29 (19) (1995) 2585–2593.

[45] C. J. Aro, G. H. Rodrigue, Preconditioned time differencing for stiff ODEs in diurnal atmospheric kinetics, Comput. Phys. Commun. 92 (1) (1995) 27–53.

[46] E. Hawkes, R. Sankaran, C. Sutherland, J. Chen, Direct numerical simulation of turbulent combustion: fundamental insights towards predictive models, J. Phys. Conf. Ser. 16 (2005) 65–79.

[47] M. A. Mueller, T. J. Kim, R. A. Yetter, F. L. Dryer, Flow reactor studies and kinetic modeling of the H_2/O_2 reaction, Int. J. Chem. Kinet. 31 (2) (1999) 113–125.

[48] J. Li, Z. Zhao, A. Kazakov, F. L. Dryer, An updated comprehensive kinetic model of hydrogen combustion, Int. J. Chem. Kinet. 36 (10) (2004) 566–575.

[49] R. Sankaran, E. Hawkes, J. Chen, T. Lu, C. K. Law, Structure of a spatially developing turbulent lean methane-air Bunsen flame, Proc. Comb. Inst. 31 (2007) 1291–1298.

[50] C. S. Yoo, E. Richardson, R. Sankaran, J. H. Chen, A DNS study on the stabilization mechanism of a turbulent lifted ethylene jet flame in highly-heated coflow, Proc. Comb. Inst. 33 (1) (2011) 1619–1627.

[51] P. N. Brown, D. E. Shumaker, C. S. Woodward, Fully implicit solution of large-scale non–equilibrium radiation diffusion with high order time integration, J. Comput. Phys. 204 (2) (2005) 760–783.

[52] Z. Chen, M. P. Burke, Y. Ju, Effects of Lewis number and ignition energy on the determination of laminar flame speed using propagating spherical flames, Proc. Comb. Inst. 32 (1) (2009) 1253–1260.

[53] H. Wang, M. Frenklach, A detailed kinetic modeling study of aromatics formation in laminar premixed acetylene and ethylene flames, Combust. Flame 110 (1–2) (1997) 173–221.

[54] T. Lu, C. K. Law, Toward accommodating realistic fuel chemistry in large-scale computations, Prog. Energy Combust. Sci. 35 (2) (2009) 192–215.

[55] 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.

[56] M. J. McNenly, R. A. Whitesides, D. L. Flowers, Faster solvers for large kinetic mechanisms using adaptive preconditioners, Proc. Comb. Inst. 35 (1) (2015) 581–587.

[57] P. Tranquilli, A. Sandu, Rosenbrock-Krylov methods for large systems of differential equations, SIAM J. Sci. Comput. 36 (3) (2014) A1313–A1338.

[58] J. E. Dennis Jr, R. B. Schnabel, Numerical methods for unconstrained optimization and nonlinear equations, Vol. 16, SIAM, 1996.

[59] C. Park, S. Yoon, Fully coupled implicit method for thermochemical nonequilibrium air at suborbital flight speeds, AIAA J. 28 (1) (1991) 31–39.

[60] S. Eberhardt, S. Imlay, Diagonal implicit scheme for computing flows with finite rate chemistry, J. Thermophys. Heat Tr. 6 (2) (1992) 208–216.

[61] 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.

[62] L. Shunn, F. Ham, P. Moin, Verification of variable-density flow solvers using manufactured solutions, J. Comput. Phys. 231 (9) (2012) 3801–3827.

[63] F. A. Williams, Combustion Theory 2nd Edition, Addison-Wesley, 1985.

[64] N. Peters, Turbulent Combustion, Cambridge University Press, UK, 2000.

[65] S. Dworkin, M. Smooke, V. Giovangigli, The impact of detailed multicomponent transport and thermal diffusion effects on soot formation in ethylene/air flames, Proc. Comb. Inst. 32 (1) (2009) 1165–1172.

[66] Y. Xuan, G. Blanquart, Effects of aromatic chemistry-turbulence interactions on soot formation in a turbulent non-premixed flame, Proc. Comb. Inst. 35 (2) (2015) 1911–1919.

[67] 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.

[68] 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.

[69] S. Verma, Y. Xuan, G. Blanquart, An improved bounded semi-lagrangian scheme for the turbulent transport of passive scalars, J. Comput. Phys. 272 (0) (2014) 1 – 22.

[70] M. E. Mueller, H. Pitsch, LES model for sooting turbulent nonpremixed flames, Combust. Flame 159 (6) (2012) 2166–2180.

[71] P. L. Carroll, G. Blanquart, The effect of velocity field forcing techniques on the Karman-Howarth equation, J. Turbul. 15 (7) (2014) 429–448.

[72] 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.

[73] C. D. Pierce, Progress-variable approach for large-eddy simulation of turbulent combustion, Ph.D. thesis, Stanford University (2001).

[74] U. Maas, S. Pope, Simplifying chemical kinetics: Intrinsic low-dimensional manifolds in composition space, Combust. Flame 88 (3–4) (1992) 239–264.

[75] S. Lam, Singular perturbation for stiff equations using numerical methods, in: Recent Advances in the Aerospace Sciences, Springer, 1985, pp. 3–19.

[76] R. D. Falgout, U. M. Yang, HYPRE: A library of high performance preconditioners, in: Computational Science—ICCS 2002, Springer, 2002, pp. 632–641.

[77] 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. A 210 (1911) 307–357.

[78] G. Candler, P. Subbareddy, I. Nompelis, Decoupled implicit method for aerothermodynamics and reacting flows, AIAA J. 51 (5) (2013) 1245–1254.

[79] G. Candler, D. Olynick, Hypersonic flow simulations using a diagonal implicit method, in: R. Glowinski (Ed.), Computing Methods in Applied Sciences and Engineering, Nova Science Publishers, 1991, pp. 29–47.

[80] H. Pitsch, FlameMaster: A C++ computer program for 0D combustion and 1D laminar flame calculations. Available at http://www.stanford.edu/∼hpitsch/. (1998).

[81] Z. Hong, D. Davidson, R. Hanson, An improved H-2/O-2 mechanism based on recent shock tube/laser absorption measurements, Combust. Flame 158 (4) (2011) 633–644.

[82] G. P. Smith, D. M. Golden, M. Frenklach, N. W. Moriarty, B. Eiteneer, M. Goldenberg, C. T. Bowman, R. K. Hanson, S. Song, W. C. Gardiner, V. V. Lissianski, Z. Qin, GRI-Mech 3.0. Available at http://www.me.berkeley.edu/gri_mech/.

[83] G. Blanquart, CaltechMech v2.1. Available at http://theforce.caltech.edu/resources/.

[84] W. Wang, K. Luo, J. Fan, Direct numerical simulation and conditional statistics of hydrogen/air turbulent premixed flames, Energ. Fuel 27 (2013) 549–560.

[85] E. Hawkes, O. Chatakonda, H. Kolla, A. Kerstein, J. Chen, A petascale direct numerical simulation study of the modelling of flame wrinkling for large-eddy simulations in intense turbulence, Combust. Flame 159 (2012) 2690–2703.

[86] M. Day, J. Bell, P.-T. Bremer, V. Pascucci, V. Beckner, M. Lijewski, Turbulence effects on cellular burning structures in lean premixed hydrogen flames, Combust. Flame 156 (2009) 1035–1045.

[87] A. Aspden, M. Day, J. Bell, Turbulence-chemistry interaction in lean premixed hydrogen combustion, Proc. Comb. Inst. 35 (2) (2015) 1321–1329.

[88] N. Peters, Local quenching due to flame stretch and non-premixed turbulent combustion, Combust. Sci. Technol 30 (1) (1983) 1–17.

[89] R. K. A. Kailasanathan, T. Yelverton, T. Fang, W. Roberts, Effect of diluents on soot precursor formation and temperature in ethylene laminar diffusion flames, Combust. Flame 160 (3) (2013) 656–670.

[90] R. K. A. Kailasanathan, E. Book, T. Fang, W. Roberts, Hydrocarbon species concentrations in nitrogen diluted ethylene-air laminar jet diffusion flames at elevated pressures, Proc. Comb. Inst. 34 (1) (2013) 1035–1043.

[91] Y. Ju, W. Sun, M. Burke, X. Gou, Z. Chen, Multi-timescale modeling of ignition and ame regimes of *n*-heptane-air mixtures near spark assisted homogeneous charge compression ignition conditions, Proc. Comb. Inst. 33 (1) (2011) 1245–1251.

[92] G. Goyal, P. Paul, H. Mukunda, S. Deshpande, Time dependent operator-split and unsplit schemes for one dimensional premixed flames, Comb. Sci. Tech. 60 (1-3) (1988) 167–189.

[93] X. Zhong, Additive semi-implicit Runge–Kutta methods for computing high-speed nonequilibrium reactive flows, J. Comput. Phys. 128 (1) (1996) 19–31.

[94] V. Hiremath, S. Lantz, H. Wang, S. Pope, Large-scale parallel simulations of turbulent combustion using combined dimension reduction and tabulation of chemistry, Proc. Comb. Inst. 34 (2013) 205–215.

[95] T. Lu, C. K. Law, Strategies for mechanism reduction for large hydrocarbons: *n*-heptane, Combust. Flame 154 (1–2) (2008) 153–163.

[96] A. Gruber, E. R. Sankaran, E. R. Hawkes, J. Chen, Turbulent flame-wall interaction: a direct numerical simulation study, J. Fluid Mech. 658 (2010) 5–32.