



A pre-partitioned adaptive chemistry methodology for the efficient implementation of combustion chemistry in particle PDF methods



Youwen Liang*, Stephen B. Pope, Perrine Pepiot

Sibley School of Mechanical and Aerospace Engineering, Cornell University, NY 14853, USA

ARTICLE INFO

Article history:

Received 13 May 2015

Received in revised form 16 May 2015

Accepted 16 May 2015

Available online 9 June 2015

Keywords:

Adaptive chemistry

Reduced kinetics

LES/PDF

Partially Stirred Reactor

DRGEP

ABSTRACT

Large Eddy Simulation/particle Probability Density Function (LES/PDF) approaches are now well developed, and can be applied to turbulent combustion problems involving complex flows with strong turbulence–chemistry interactions. However, these methods are computationally expensive, restricting their use to simple fuels with relatively small detailed chemical mechanisms. To mitigate the cost in both CPU time and storage requirements, an adaptive strategy tailored for particle PDF methods is presented here, which provides for each particle a specialized reduced representation and kinetic model adjusted to its changing composition. Rather than performing chemical reduction at runtime to determine the optimal set of equations to use for a given particle, an analysis of the composition space likely to be accessed during the combustion simulation is performed in a pre-processing stage using simple Partially Stirred Reactor (PaSR) computations. In the pre-processing stage, the composition space is partitioned into a user-specified number of regions, over which suitable reduced chemical representations and kinetic models are generated automatically using the Directed Relation Graph with Error Propagation (DRGEP) reduction technique. A computational particle in the combustion simulation then carries only the variables present in the reduced representation and evolves according to the reduced kinetic model corresponding to the composition space region the particle belongs to. This region is identified efficiently using a low-dimensional binary-tree search algorithm, thereby keeping the run-time overhead associated with the adaptive strategy to a minimum. The performance of the algorithm is characterized for propane/air combustion in a PaSR with pairwise mixing. The results show that the reduction errors are well controlled by the specified error tolerance, and that the adaptive framework provides significant gains in cost and storage compared to traditional non-adaptive reduction approaches.

© 2015 The Combustion Institute. Published by Elsevier Inc. All rights reserved.

1. Introduction

The understanding of chemical kinetics for hydrocarbon fuel combustion has exploded over the past two decades, leading to the development of ever growing detailed kinetic schemes for a wide range of molecular species relevant to hydrocarbon combustion [1]. With more accurate rate rules and improved mechanistic considerations, some of the latest published mechanisms are approaching 10^4 species and typically five times as many reactions [2]. However, these advances in chemical kinetics and detailed model development have to be integrated with Computational Fluid Dynamics (CFD) tools to fully realize their potential in terms of improved understanding and optimization of practical combustion devices. The real challenge then is to maximize the level of chemical detail that can be afforded in CFD. Different modeling

approaches for turbulent flows impose different upper limits on the number of variables that can be used to describe the chemistry, ranging from around twenty in Direct Numerical Simulations [3], to a few dozens for LES/particle PDF methods [4], to a few hundreds for flamelet-based approaches [5,6].

Numerous techniques have been developed to drastically reduce the computation burden associated with a detailed description of the chemistry. A non-exhaustive list includes: (i) automatic elimination of negligible species and reactions to generate *skeletal* mechanisms [7–13]; (ii) dimension reduction techniques, which further decrease the number of species or degrees of freedom that need be considered, including chemical lumping [14,15], mathematical lumping [16], QSSA [17,18], RCCE [19], ILDM [20], REDIM [21] and ICE-PIC [22]; and (iii) storage/retrieval algorithms to reduce the computational cost of repetitive kinetics calculations, including ISAT [23,24] and PRISM [25]. (One component of the current methodology is DRGEP [9], which is in category (i).) Rather than being in competition, these general approaches can be used

* Corresponding author.

E-mail address: yli956@cornell.edu (Y. Liang).

Nomenclature

\mathbb{C}	classifying space, of dimension n_C	$\mathcal{R}_{F \rightarrow J}$	reduction operator
\mathcal{C}	mapping function between full \mathbb{F} and classifying \mathbb{C} spaces	R_J	J th region in the classifying space partition and in the composition space
\mathbf{C}	classification vector in \mathbb{C} , of dimension n_C	$\mathcal{R}_{J \rightarrow F}$	reconstruction operator
C_A	consumption rate of species A	\mathbf{R}^Φ	DRGEP coefficient vector for the single composition Φ
c_D	kinetic inhomogeneity cost of the set \mathcal{D}	$\tilde{\mathbf{R}}^\Phi$	scaled DRGEP coefficient vector for the single composition Φ
\mathcal{D}	database of composition vectors Φ	\mathbf{S}	chemical source term for the detailed chemical mechanism
D	distance between two reduced models	\mathbf{S}_J	chemical source term for the reduced kinetic model M_J
\mathbb{F}	full composition space, of dimension $n_s + 1$	\mathcal{T}	set of targets used in DRGEP
\mathbb{F}_J	reduced composition space for model M_J , of dimension $n_s^J + 1$	T	temperature (K)
H	hyperplane	t	time (s)
h	specific enthalpy	\mathbf{V}	cutting hyperplane normal vector
\mathcal{H}_D	kinetic homogeneity of the database \mathcal{D}	\mathbf{Y}	full species mass fraction vector, of dimension n_s
\mathcal{J}	classification function	\mathbf{y}_J	mass fraction vector in the representation ϕ_J , of dimension n_s^J
M_D	detailed chemical mechanism	$\hat{\mathbf{y}}$	mass fractions of locally-significant species
M_J	reduced kinetic model for region R_J	$\alpha_{T,\Phi}$	DRGEP scaling coefficient for target T and composition Φ
\dot{m}_k^{str}	normalized mass flow rate of the k th inflowing stream in the PaSR	Δt	simulation time step (s)
n_C	number of classifier variables; dimension of \mathbb{C}	ε	reaction-mapping error for a reduced kinetic model
n_D	number of compositions in the database \mathcal{D}	ε_c	specified error tolerance
n_e	number of elements in the detailed chemical mechanism	ε_{SA}^A	sensitivity coefficient of species A in Reduced Modeling
n_{in}	number of inflowing particles per time step in the PaSR	ε_{sig}	threshold for the mass fraction of an inactive species to be regionally or locally significant
n_p	number of particles in the PaSR	ε_Φ	reaction mapping error for composition Φ
n_{pair}	number of particles changing partners per time step in the PaSR	ε_X	incurred error in X in a PaSR simulation
N_R	number of regions in the partition of the classifying space	$\hat{\varepsilon}_X$	conservation error in X in a PaSR simulation
n_r	number of reactions in the detailed mechanism	θ_J	set of species in the reduced kinetic model M_J
n_r^J	number of reactions in the reduced kinetic model M_J	Θ	set of species in the detailed mechanism
n_s	number of species in the detailed chemical mechanism	$\hat{\theta}$	set of locally-significant species
n_s^J	number of species in the reduced kinetic model M_J	$\nu_{i,A}$	stoichiometric coefficient of species A in reaction i
n_t	number of time steps in the simulation	ξ_J	set of reactions in the reduced kinetic model M_J
n_{str}	number of inflowing streams in the PaSR	τ_{mix}	particle mixing time scale in the PaSR (s)
n_T	number of targets used in DRGEP	τ_{pair}	particle pairing time scale in the PaSR (s)
\mathbf{P}	cutting hyperplane anchor point	τ_{res}	residence time in the PaSR (s)
p	pressure	Υ	list of models available for a region, parameterized by the model error ε
P_A	production rate of species A	Φ	temperature-based representation of a composition in \mathbb{F} space
P_a	pseudo-production rate of atom a	Φ^D	temperature-based representation of a composition in \mathbb{F} space obtained using M_D
\mathbb{R}	DRGEP coefficient space, of dimension n_s	ϕ_J	reduced representation based on the reduced kinetic model M_J
R_{AB}	DRGEP path-dependent coefficient between species A and B	Φ^m	temperature-based representation of a composition in \mathbb{F} space after the mixing fractional step
r_{AB}	DRGEP direct interaction coefficient between species A and B	Φ'	enthalpy-based representation of a composition in \mathbb{F} space
$r_{AB,\{G\}}$	DRGEP direct interaction coefficient given a set of removed species G	Φ_k^{str}	composition of k th inflowing stream in the PaSR
r_{Ar_i}	DRGEP direct interaction coefficient between species A and reaction r_i	Ξ	set of reactions in the detailed mechanism
\mathbf{R}^D	DRGEP coefficient vector for all compositions in the database \mathcal{D}	ω_i	reaction rate of reaction i
$\tilde{\mathbf{R}}^D$	scaled DRGEP coefficient vector for all compositions in the database \mathcal{D}		
$\mathbf{R}^{D,r}$	DRGEP reaction coefficient vector for all compositions in the database \mathcal{D}		

in combination with a compounding of the benefits. Yet, reaching the level of reduction necessary to handle complex molecules or multi-component mixtures still appears as a major challenge.

One way forward is to observe that, for a given reactive flow configuration, thermodynamic conditions can vary widely in space and time, and in any small range of temperature and compositions, many species have negligible concentrations, and only a few are chemically active. Hence, the chemistry description in that

composition space region may require many fewer species and reactions than a more detailed model applicable over the entire composition space. This observation leads to the concept of *adaptive chemistry*, in which the chemical representation and kinetic equations are adapted to the local chemical activity and thermodynamic state.

Several attempts to devise computationally-viable adaptive chemistry strategies have been made recently [26–39]. In most

cases, on-the-fly reduction is used, in which at each time step of the simulation, the local composition is analyzed to determine the smallest set of important species to solve for, with the other unimportant species kept as “inert”. This methodology theoretically guarantees a maximal reduction in the time needed to integrate the chemistry equations, but is efficient only if the time taken by the reduction algorithm is comparatively small. Species elimination or dimension-reduction methods based on graph or element flux analysis methods have been shown to satisfy these requirements [27,29,34–38,40]. In particular, the Dynamic Adaptive Chemistry, or DAC, technique [27] has been applied on a variety of configurations, including methane/air combustion in partially-stirred reactor [34], 1D premixed, unsteady, freely-propagating methane/air flame [35], homogeneous ignition [37], and compression-ignition engine simulations [36]. Other on-the-fly reduction approaches considered HCCI engine simulations [39], and even large scale simulations of a turbulent non-premixed ethylene flame [30].

An alternative to on-the-fly reduction is to incorporate some amount of preprocessing calculations to reduce the overhead of the adaptive strategy at runtime. An early demonstration was made by Schwer et al. [26], who manually assembled a small library of reduced models addressing key chemistry regimes expected to be encountered in a flame. Criteria involving local species mass fractions were used to decide which model to use in each domain cell in the simulation of a turbulent hydrogen shear layer flame and an axisymmetric laminar partially premixed methane–air flames. Singer and Green [32] propose a snapshot method to tabulate *a priori* proper orthogonal decomposition projection matrices, which are then used adaptively at run time to project minor species mass fractions onto a reduced dimensional space, thereby accelerating flame calculations.

While major progress in the adaptive chemistry field has been made in reducing the cost associated with the chemistry calculations, most of the strategies listed above rely on the knowledge of the full composition vector at every given time and location. This avoids the conceptual problems of mechanism augmentation (*i.e.* re-introducing species) and transport-term evaluation. While storage is not an issue for relatively small kinetic schemes, this approach quickly becomes prohibitively expensive as the number of chemical variables is increased. Tosatto et al. [29] propose a solution to this problem by enabling species removal and re-introduction through a graph analysis modified to account for transport processes. While the integration of the resulting algorithm with flame simulation codes is not straightforward, very good performances are demonstrated in 2D axisymmetric co-flow flame calculations.

In this paper, we present and demonstrate a pre-partitioned adaptive chemistry strategy to handle large chemical mechanisms and tailored for particle Probability Density Function (PDF) methods. These methods are now well developed, and can be applied to turbulent combustion problems involving complex flows with strong turbulence–chemistry interactions [41,42]. In the particle PDF framework, the gas composition within the solution domain is represented by a large number (up to several millions) of notional particles, each carrying one realization of the fluid thermochemical state. One of the virtues of PDF methods is that reaction is in closed form, allowing for realistic chemistry to be incorporated. But that comes at a significant increase in CPU time compared to other approaches and models (*e.g.*, a factor of 3 or more [43]). As such, because the chemistry in PDF methods is described by a detailed (or reduced) mechanism and is treated in isolation in a fractional step, adaptive strategies appropriate for PDF methods to simulate turbulent combustion involve somewhat different issues than the Eulerian framework typically considered in the literature mentioned above, in which the energy and species

transport equations, discretized and solved on a grid, require closures for the reaction term that combine both chemical kinetics and turbulence/chemistry interaction assumptions.

Previous work on this problem includes that of Banerjee and Ierapetritou [33], who use elaborate clustering techniques to create a library of reduced kinetic models, and explicitly determine the highly non-convex regions in composition space over which those models are valid. They also detail how such a model library could be used in an adaptive simulation framework, providing proof-of-concept using the simpler Partially Stirred Reactor (PaSR) configuration. The present work also considers the PaSR as its development and validation framework due to its direct relevance to PDF simulations of turbulent combustion. However we adopt a much simpler conceptual approach, in which the composition space is partitioned into judiciously chosen regions *prior* to the derivation of reduced kinetic models. Computational particles are also constrained to carry only the minimum number of variables needed to represent their thermochemical state accurately, so that both CPU time and memory requirements associated with particle PDF methods are significantly reduced.

The two primary defining characteristics of the methodology are that the composition space is partitioned in a pre-processing step, and that DRGEP is used to generate the reduced kinetic models. The methodology can therefore be called Pre-Partitioned Adaptive Chemistry using Direction Relation Graphs with Error Propagation, or PPAC-DRGEP.

The remainder of this paper is organized as follows. In Section 2, we provide a mathematical description of how chemistry is represented in particles PDF computations, using a Partially Stirred Reactor configuration as the framework. In Section 3, an overview of the adaptive chemistry methodology is provided, followed by a more detailed description of the reduction technique employed (Section 4) and of each component of the algorithm (Section 5). The method is then applied to a propane/air non-premixed PaSR, described in Section 6, and its performance is analyzed in Section 7. Possible improvements and extensions to the methodology, and its application in LES/PDF simulations, are discussed in Section 8. A summary and conclusions are provided in Section 9.

2. Particle representation in a partially-stirred reactor (PaSR)

We consider a reacting mixture of ideal gases, consisting of n_s chemical species composed of n_e elements. For simplicity of exposition, the mixture is assumed to evolve at a fixed pressure p , so that (given p) the full thermochemical state, or composition, of the mixture is completely characterized by the n_s -vector of species mass fractions \mathbf{Y} and the mixture temperature T . Accordingly, we define the *composition* to be the $(n_s + 1)$ -vector: $\Phi \equiv \{\mathbf{Y}, T\}$. We denote by \mathbb{F} the full composition space, so that any composition vector Φ is an element of \mathbb{F} .

In some situations it is convenient to consider the enthalpy h as the energy variable, in place of T . We therefore define the alternative composition vector $\Phi' \equiv \{\mathbf{Y}, h\}$. These two representation of the composition contain the same information, and it is straightforward to transform from one to the other.

As mentioned above, a partially-stirred reactor (PaSR) is used as a reference configuration to develop, test, and validate the adaptive methodology. The reactor consists of a constant, even number n_p of particles, the n th particle having composition $\Phi^{(n)}$. Each particle is assigned a partner, and the particles are ordered so that particles n and $n + 1$ are partners, for odd n . The PaSR is continuously fed by a specified number n_{str} of inflow streams, the k th stream having prescribed composition Φ_k^{str} , and normalized mass flow rate \dot{m}_k^{str} . Also specified are the residence time, τ_{res} , a mixing time, τ_{mix} , and a pairing time, τ_{pair} . The particular PaSR test case considered is

Table 1

Specification of the parameters for the non-premixed propane/air piloted PaSR test case. Species are given as mass fractions. The initial conditions are the same as Stream 1.

Parameters	Name	Values	
<i>Reactor characteristics</i>			
Pressure	p	1 bar	
Number of particles	n_p	100	
Pairing time	τ_{pair}	1 ms	
Mixing time	τ_{mix}	1 ms	
Residence time	τ_{res}	10 ms	
Time step	Δt	0.1 ms	
Parameters	Stream 1	Stream 2	Stream 3
<i>Initial and Inflow conditions (Species data are mass fractions)</i>			
Normalized mass flow rates	0.05	0.893	0.057
Temperature	2272 K	300 K	300 K
C_3H_8	0.0	0.0	1.0
N_2	7.2085×10^{-1}	0.767	0.0
O_2	7.6607×10^{-3}	0.233	0.0
CO_2	1.6146×10^{-1}	0.0	0.0
H_2O	9.5159×10^{-2}	0.0	0.0
CO	1.2162×10^{-2}	0.0	0.0
OH	2.2616×10^{-3}	0.0	0.0
H_2	2.2880×10^{-4}	0.0	0.0
O	1.9634×10^{-4}	0.0	0.0
H	1.6928×10^{-5}	0.0	0.0
HO_2	8.0954×10^{-7}	0.0	0.0
H_2O_2	6.1000×10^{-8}	0.0	0.0
HCO	1.2371×10^{-9}	0.0	0.0
CH_2O	1.6891×10^{-11}	0.0	0.0

described in Section 6, where the values of these quantities are given in Table 1.

The compositions of the particles evolve in time in small time steps, Δt , in three fractional steps accounting for: inflow/outflow/pairing; mixing; and reaction. In the inflow/outflow/pairing fractional step, an even number, n_{in} , of particles are selected at random with equal probability and their compositions are replaced by inflow stream compositions. In expectation, that is, when considering its long-run average value, n_{in} is equal to $n_p \Delta t / \tau_{\text{res}}$, and the number from each inflow stream is proportional to its specified mass flow rate. A number, n_{pair} , pairs of particles are then selected for pairing. In expectation, n_{pair} is equal to $\frac{1}{2} n_p \Delta t / \tau_{\text{res}}$. These pairs include those involved in inflow/outflow and others selected at random with uniform probability. These particles are then randomly shuffled to reassign partners.

We denote by $\Phi^{(n)}(t)$ and $\Phi'^{(n),m}$ the enthalpy-based composition of the n th particle before and after the mixing fractional step. Then, mixing amounts to a small exchange of composition between partners:

$$\begin{aligned}\Phi'^{(n),m} &= \Phi^{(n)}(t) - \frac{\Delta t}{\tau_{\text{mix}}} (\Phi^{(n)}(t) - \Phi^{(n+1)}(t)), \\ \Phi'^{(n+1),m} &= \Phi^{(n+1)}(t) - \frac{\Delta t}{\tau_{\text{mix}}} (\Phi^{(n+1)}(t) - \Phi^{(n)}(t)),\end{aligned}\quad (1)$$

for all odd n .

At the start of the reaction fractional step, the compositions after mixing, $\Phi'^{(n),m}$, are converted to their temperature representations $\Phi^{(n),m}$. Starting from these initial conditions, the compositions then evolve due to reaction under isobaric, adiabatic conditions over a time Δt according to

$$\frac{d\Phi^{(n)}(t)}{dt} = \mathbf{S}(\Phi^{(n)}(t)),\quad (2)$$

where \mathbf{S} is the chemical source term defined by the user-provided detailed chemical mechanism. This mechanism, denoted by M_D ,

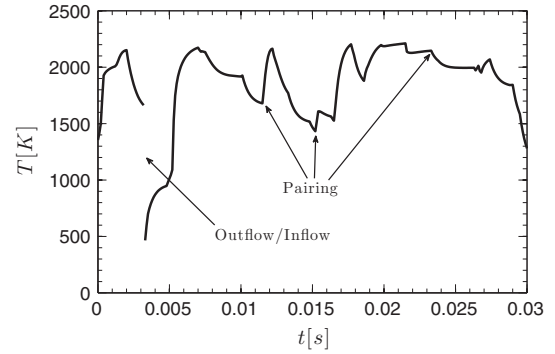


Fig. 1. Typical temperature evolution for an individual particle in a PaSR. Discontinuities in T correspond to inflow/outflow events, while discontinuities in dT/dt are associated with pairing events. Apart from those discrete, random events, the temperature evolves continuously due to mixing and reaction. (Only some of the pairing events are indicated.)

involves the set of n_s species Θ , and a set of n_r reactions Ξ . The reaction fractional step yields the final particle composition at $t + \Delta t$, i.e., $\Phi^{(n)}(t + \Delta t)$; and the mapping that Eq. (2) provides from $\Phi'^{(n),m}$ to $\Phi^{(n)}(t + \Delta t)$ is referred to as the *reaction mapping*.

Figure 1 shows the evolution of the temperature of a typical particle in the PaSR. Inflow/outflow events result in a discontinuity of the temperature; pairing results in a discontinuity in dT/dt ; and in between these discrete events, mixing and reaction result in a smooth evolution.

3. Overview of the pre-conditioned adaptive chemistry methodology

The overall pre-conditioned adaptive chemistry methodology combines an off-line, pre-processing stage (during which the composition space is partitioned and a set reduced kinetic models is created), and an online, dynamic procedure to identify, at time t , which reduced model to use for each particle in the combustion simulation. We refer to these two stages as the *pre-processing stage* and the *adaptive simulation*.

3.1. Pre-processing stage

The **pre-processing stage** consists of the following three tasks:

1. **Database creation:** Using the PaSR and the detailed chemical mechanism, generate a database \mathcal{D} consisting of a large number n_D of full compositions, representative of the compositions that occur during the simulation of interest.
2. **Partitioning:** Define a low-dimensional space \mathbb{C} , referred to hereafter as the classifying space, and a function mapping the composition space \mathbb{F} to the classifying space \mathbb{C} , $\mathcal{C}: \mathbb{F} \rightarrow \mathbb{C}$. Partition \mathbb{C} into a specified number N_R of regions, the j th region being denoted by R_j . This also implicitly partitions the composition space \mathbb{F} : corresponding to R_j , the j th region in \mathbb{F} is all points mapped by \mathcal{C} to R_j . Henceforth we use R_j to denote the j th region in either \mathbb{C} or \mathbb{F} , depending on the context.
3. **Reduced modeling:** For each region R_j , identify a reduced composition space \mathbb{F}_j (a sub-space of \mathbb{F}), a reduced representation of the composition, $\phi_j \equiv \{\mathbf{y}_j, T\} \in \mathbb{F}_j$, and an appropriate reduced kinetic model M_j . The reduced representation and kinetic model involve a reduced set of n'_s retained species θ_j and a reduced set of n'_r retained reactions ξ_j , with the expectation that n'_s and n'_r are significantly smaller than n_s and n_r , respectively. The n'_s -vector of the mass fraction of the retained species is denoted by \mathbf{y}_j . This procedure is called reduced modeling mapping for R_j :

$$M_J \equiv \{\phi_J, \theta_J, \xi_J\}. \tag{3}$$

3.2. Reduced representation of compositions

In the adaptive simulation, the reduced composition of the n th particle is denoted by $\phi^{(n)}$ and is defined as

$$\phi^{(n)} = \{J^{(n)}, \mathbf{y}_J^{(n)}, \hat{\theta}^{(n)}, \hat{\mathbf{y}}^{(n)}, T^{(n)}\}, \tag{4}$$

where J indicates that the representation is based on model M_J , and T is the temperature. The way the species are represented in Eq. (4) is schematically presented in Fig. 2, and further explanation is provided below.

As described in Section 5.3, the DRGEP procedure used to generate the model M_J identifies which species are kinetically active, and the rest are deemed to be inactive. Of the inactive species, some may be deemed to be regionally significant, that is, significant for a given region R_J in composition space, if their mass fraction has a significant probability to be above a specified threshold ϵ_{sig} (taken to be $\epsilon_{\text{sig}} = 10^{-3}$) among compositions in the region. This is further explained and made precise in Section 5.3. Both active and regionally-significant species are included in M_J , and their mass fractions are denoted by \mathbf{y}_J .

An inactive species that is not regionally significant is deemed to be locally significant, that is, significant for a specific composition, if its mass fraction in this composition exceeds ϵ_{sig} . In the reduced representation, Eq. (4), $\hat{\theta}$ denotes the set of locally-significant species, and $\hat{\mathbf{y}}$ denotes their mass fractions. It is found to be advantageous in terms of accuracy and efficiency to retain these locally-significant species in the particle composition, but not in the model M_J : neglecting these species leads to inaccuracy, while including them in the model leads to unduly large models and hence inefficiency. It is expected (and found in practice) that locally-significant species are seldom needed. It should be noted that to satisfy the normalization condition for mass fractions, the components of \mathbf{y}_J and $\hat{\mathbf{y}}$ together sum to unity.

Henceforth, we abbreviate the notation (Eq. (4)) for the reduced representation of composition to:

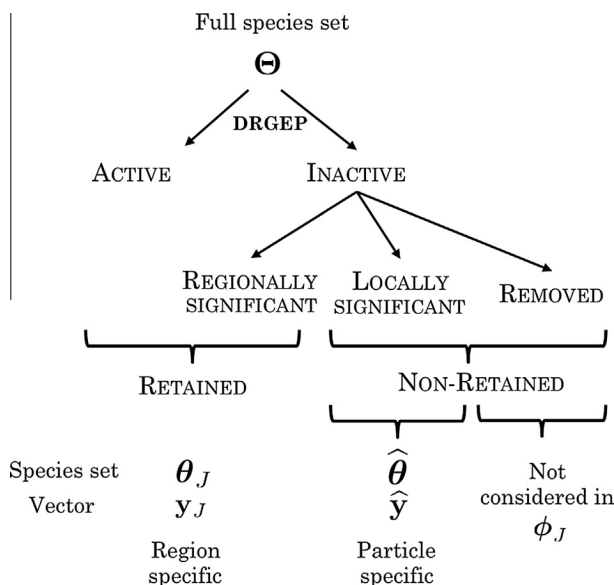


Fig. 2. Species nomenclature used in the reduced representation of compositions. Note that ξ_J , the set of retained reactions for region J , involves active species exclusively, with the exception of third-body expressions that may include regionally and locally significant species.

$$\phi_J^{(n)} = \{\mathbf{y}_J^{(n)}, \hat{\mathbf{y}}^{(n)}, T^{(n)}\}, \tag{5}$$

it being implied that $J^{(n)}$ and $\hat{\theta}^{(n)}$ are also known. We refer to this as the reduced composition based on model M_J , and note that it does not necessarily require that the composition belongs to region R_J .

3.3. Adaptive simulation

Once the classifying space partition and a corresponding set of reduced kinetic models $M_{J=1 \dots N_R}$ have been constructed, the time evolution of a particle composition in the PaSR can be computed using reduced representations and kinetic models dynamically chosen among this set. This adaptive algorithm, illustrated in Fig. 3 for $N_R = 4$, proceeds in the following sub-steps:

1. At the beginning of each time step, the composition of particle n is expressed using the reduced representation based on model M_J (for some known J): $\phi_J^{(n)}(t) = \{\mathbf{y}_J^{(n)}(t), \hat{\mathbf{y}}^{(n)}(t), T^{(n)}(t)\}$.
2. In the mixing fractional step involving two paired particles n and $n + 1$ with reduced compositions $\phi_J^{(n)}$ and $\phi_{K'}^{(n+1)}$ pertaining to models M_J and $M_{K'}$, respectively, the full representations $\Phi^{(n)}$ and $\Phi^{(n+1)}$ are first reconstructed. These reconstructed compositions are then mixed according to Eq. (1) to yield the full compositions after mixing, denoted by $\Phi^{(n),m}$ and $\Phi^{(n+1),m}$.
3. The classification algorithm is used to identify which regions R_J and $R_{K'}$ that $\Phi^{(n),m}$ and $\Phi^{(n+1),m}$ belong to.
4. The full compositions $\Phi^{(n),m}$ and $\Phi^{(n+1),m}$ are then reduced to the reduced representations $\phi_J^{(n),m}$ and $\phi_{K'}^{(n+1),m}$, based on models M_J and $M_{K'}$.
5. Finally, in the reaction fractional step, the reduced representation of particle n located in region R_J , $\phi_J^{(n)}$, is integrated over the time increment Δt according to the reduced kinetic model M_J . These values of J and $\phi_J^{(n)}$ are then used in sub-step 1 of the next time step. How the locally-significant species are treated in the reaction fractional step is described in Section 5.3.




Fig. 3. Overview of the adaptive strategy. The classifying space \mathbb{C} is partitioned into $N_R = 4$ regions, R_1 to R_4 , for which there are corresponding reduced kinetic models M_1 to M_4 . At the beginning of the time step, particle n , represented by the black circle, has the reduced representation based on model M_1 . In the mixing fractional step, the reduced representation $\phi_{J=1}^{(n)}$ is first reconstructed to a full representation $\Phi^{(n)}$ (1), and mixed using the pairwise mixing model of Eq. (1) (2). Because of the mixing process, the full representation after mixing, $\Phi^{(n),m}$, has moved to region R_4 , as determined by classification (3). $\Phi^{(n),m}$ is reduced to $\phi_{J=4}^{(n),m}$ (4), then integrated in time according to the reduced model M_4 (5) to yield the particle composition at the end of the time step: $\phi_{J=4}^{(n)}(t + \Delta t)$. (Note that the final composition, although based on model M_4 is not necessarily in region R_4 .)

3.4. Required operations

The adaptive procedure outlined above requires the definition of the following generic operations, applicable to any particle with reduced composition ϕ_j . Detailed descriptions and implementations are provided in Section 5.

1. **Reconstruction:** The full (F) representation $\Phi \in \mathbb{F}$, reconstructed from the reduced representation ϕ_j , is provided by:

$$\Phi = \mathcal{R}_{j \rightarrow F}(\phi_j). \quad (6)$$

2. **Classification:** Given a full representation Φ of an individual particle composition, the classification function \mathcal{J} that returns the index of the region it belongs to is defined as:

$$J = \mathcal{J}(\Phi). \quad (7)$$

3. **Reduction:** The full representation Φ of a given particle composition can be converted to the reduced representation based on model M_j by the *reduction operation*:

$$\phi_j = \mathcal{R}_{F \rightarrow j}(\Phi). \quad (8)$$

Note that information is lost in the reduction, so a subsequent reconstruction does not (in general) recover the original composition, i.e., $\Phi \neq \mathcal{R}_{j \rightarrow F}(\mathcal{R}_{F \rightarrow j}(\Phi))$. This loss of information in the reduction process leads to an error in the adaptive simulation, which is referred to as *conversion error*. On the other hand, reconstruction is done consistently to satisfy $\phi_j = \mathcal{R}_{F \rightarrow j}(\mathcal{R}_{j \rightarrow F}(\phi_j))$.

In the following sections, we first provide (in Section 4) an overview of the Directed Relation Graph with Error Propagation method, which lies at the core of this adaptive chemistry methodology. We then describe in detail (in Section 5) the algorithms required at each stage, namely our approach to:

- In the pre-processing stage: generating a database \mathcal{D} of representative compositions; identifying a suitable classifying space \mathbb{C} ; subsequently partitioning \mathbb{C} into N_R regions (**Partitioning**); and, developing accurate reduced kinetic models for each of these regions (**Reduced Modeling**);
- In the adaptive simulation stage: Reconstructing the full representation of a given particle from its reduced representation (**Reconstruction**), determining the region a particle belongs to based on its current representation (**Classification**), and finally converting a particle composition from a full reconstructed representation to its appropriate reduced representation (**Reduction**).

The reader less interested in the details of the implementation may want to skip to Section 6.

4. Directed Relation Graph with Error Propagation (DRGEP)

By analyzing production rates derived from an ensemble of composition states, the Directed Relation Graph with Error Propagation (DRGEP) method [9] quantifies the coupling between the species and reactions included in a chemical mechanism and some user-specified targets. The main output from this analysis is a set of importance coefficients, which are referred to as DRGEP coefficients, allowing us to rank species and reactions from most important to least important for the prediction of the chosen targets. DRGEP coefficients are used at two different levels in the pre-processing stage: (i) in the *Partitioning* stage, to identify regions in composition space that share common chemical characteristics; and (ii) in the *Reduced Modeling* stage, to derive region-specific reduced kinetic models. The main ingredients of the DRGEP method are recalled below, the first one being the

determination of a set of n_T targets \mathcal{T} , most often particular species or heat release, on which we want to quantify the influence of individual species and reactions.

4.1. Direct interaction coefficients

Direct interaction coefficients are defined as the measure of the coupling between two species that are **directly** related through an elementary reaction, that is, two species that appear concurrently in the same reaction. For a given composition Φ and kinetic model M , the coupling coefficient between two such species A and B is expressed as:

$$r_{AB} \equiv \frac{\left| \sum_{i=1, n_R} v_{i,A} \omega_i \delta_B^i \right|}{\max(P_A, C_A)}, \quad (9)$$

where the production and consumption of species A are defined as:

$$P_A = \sum_{i=1, n_R} \max(0, v_{i,A} \omega_i), \quad (10)$$

$$C_A = \sum_{i=1, n_R} \max(0, -v_{i,A} \omega_i). \quad (11)$$

In the above equations, ω_i are the net reaction rate of the i th reaction evaluated from composition Φ and kinetic model M , $v_{i,A}$ is the stoichiometric coefficient of species A in reaction i , and

$$\delta_B^i = \begin{cases} 1 & \text{if the } i\text{th reaction involves species } B, \\ 0 & \text{otherwise.} \end{cases} \quad (12)$$

Eq. (9) can be extended to better quantify the impact of removing a species *in addition to* a previously removed set of species [9]:

$$r_{AB, \{G\}} \equiv \frac{\left| \sum_{i=1, n_R} v_{i,A} \omega_i \delta_{B, \{G\}}^i \right|}{\max(P_A, C_A)}, \quad (13)$$

where $\{G\}$ is the set of species already removed. $\delta_{B, \{G\}}^i$ is unity, if the i th reaction involves B or any species in subset $\{G\}$, and is zero otherwise.

4.2. Path-dependent coefficients

To go beyond direct interactions, DRGEP defines path-dependent coefficients that quantify the coupling between any directly or indirectly related species A and B . Assuming geometric damping, and again, for a given composition Φ and kinetic model M , the coupling between A and B through a reaction path p is written as:

$$r_{AB,p} = \prod_{i=1}^{n-1} r_{S_i S_{i+1}}, \quad (14)$$

with $S_1 = A$, $S_n = B$, S_i is on the reaction path p that links A and B , n is the number of reactions involved in path p . Since many paths can exist linking A to B , only the most important one is retained:

$$R_{AB} \equiv \max_{\text{all paths } p} r_{AB,p}. \quad (15)$$

R_{AB} is calculated following the implementation used in [9]. Briefly, starting from target A , all possible paths of degree n , n being the number of edges separating A from a given species B are explored successively, starting at $n = 1$, and R_{AB} tracks the maximum $r_{AB,p}$ encountered on those paths. The algorithm stops when paths of degree $n + 1$ return coefficients below a very small cut-off number, guaranteeing that none of the higher-degree paths would be yield $r_{AB,p}$ values larger than the current maximum. Additional details on graph search algorithms in this specific context can be found in [44]. R_{AB} can be interpreted as the magnitude of the error made in the prediction of species A if species B is removed [9].

4.3. Target-specific coefficients

In a similar way, the overall importance of B to the target set \mathcal{T} for a given composition Φ and kinetic model M is defined by:

$$R_{B,\Phi} \equiv \max_{T \in \mathcal{T}} R_{TB}. \quad (16)$$

4.4. Extension to multiple composition states

The above results can be extended to an ensemble of compositions \mathcal{D} to yield the DRGEP coefficient of species B over \mathcal{D} :

$$R_B^{\mathcal{D}} = \max_{\Phi \in \mathcal{D}} R_{B,\Phi}. \quad (17)$$

Whether referring to a single composition Φ or a set of compositions \mathcal{D} , the DRGEP coefficients can be viewed as an n_s -vector, \mathbf{R}^{Φ} or $\mathbf{R}^{\mathcal{D}}$, respectively, in an n_s -dimensional real space, which we term the DRGEP coefficient space and denote by \mathbb{R} .

4.5. Scaling factors

By construction, the DRGEP coefficients are for a large part independent of the mass fractions of target species present in the compositions used to derive them. An alternative definition with scaling factors is therefore introduced in [9] to account for the relative importance of individual targets among the compositions contained in \mathcal{D} . Defined for any composition Φ in \mathcal{D} , these scaling coefficients $\alpha_{T,\Phi}$ are written:

$$\alpha_{T,\Phi} = \max_{\text{all atoms } a} \left(\frac{\alpha_{T,\Phi}^a}{\max_{\Phi \in \mathcal{D}} \alpha_{T,\Phi}^a} \right), \quad (18)$$

where

$$\alpha_{T,\Phi}^a = \frac{N_{a,T} |P_T - C_T|}{P_a}. \quad (19)$$

In the above equations, a refers to each element present in the system (C, H, O and N for hydrocarbon combustion), $N_{a,T}$ is the number of atoms a in target species T , and P_T and C_T are the production and consumption rates, respectively, of any target species T (Eqs. (10) and (11)). P_a is the pseudo-production rate of atom a , defined as

$$P_a = \sum_{\text{all species } A} N_{a,A} \max(0, P_A - C_A). \quad (20)$$

This scaling coefficient is unity for compositions in which the target contributes to its maximum (over all compositions in \mathcal{D}) to the exchange of atoms between species, and zero when, for instance, the target species has been consumed entirely. Eq. (16) is re-defined to include the scaling factor:

$$\tilde{R}_{B,\Phi} \equiv \max_{T \in \mathcal{T}} \alpha_{T,\Phi} R_{TB}, \quad (21)$$

the scaled DRGEP coefficients still being defined as in Eq. (17):

$$\tilde{R}_B^{\mathcal{D}} = \max_{\Phi \in \mathcal{D}} \tilde{R}_{B,\Phi}. \quad (22)$$

The scaled coefficients vector $\tilde{\mathbf{R}}^{\mathcal{D}}$ are used in the following, and the tilde is dropped.

4.6. DRGEP reaction coefficients

Simply replacing Eq. (9) by

$$r_{Ar_i} \equiv \frac{|v_{iA} \omega_i|}{\max(P_A, C_A)}, \quad (23)$$

in the above procedure quantifies the importance of individual reactions on the prediction of the targets, resulting in a n_r -vector $\mathbf{R}^{\mathcal{D},r}$, used during the reduced modeling stage.

5. Implementation details

In the following, we assume that a database \mathcal{D} has been assembled, which contains a large number of full compositions Φ that densely sample the portion of the composition space relevant for the combustion simulation of interest. The number of compositions in this database is $n_{\mathcal{D}}$. Details on how to obtain such a database using a suitably chosen PaSR are provided in Section 6.

5.1. Classifying space \mathbb{C}

We define the classifying space \mathbb{C} , of specified dimension n_c , by a subset $\mathcal{I} = \{I_1 \dots I_{n_c}\}$ of the composition variables included in Φ (species in Θ and temperature), and the injective mapping function \mathcal{C} :

$$\mathcal{C}: \mathbb{F} \rightarrow \mathbb{C}, \quad \mathcal{C}(\Phi) = \mathbf{C}, \quad (24)$$

where the i th component of \mathbf{C} is defined by

$$C_i = \frac{\Phi_{I_i} - \Phi_{I_i,\min}^{\mathcal{D}}}{\Phi_{I_i,\max}^{\mathcal{D}} - \Phi_{I_i,\min}^{\mathcal{D}}}. \quad (25)$$

In Eq. (25), $\Phi_{I_i,\min/\max}^{\mathcal{D}}$ refers to the minimum/maximum values of Φ_{I_i} observed across the entire database \mathcal{D} . The subset \mathcal{I} is chosen so that the corresponding variables span the n_c -dimensional subspace of \mathbb{F} closest to the n_c -dimensional principal component space formed by the first n_c principal components of the compositions in the database \mathcal{D} with each variable rescaled between 0 and 1 following Eq. (25). More information on this procedure can be found in Yang et al. [45]. This definition of \mathcal{C} satisfies two important constraints: (i) \mathbf{C} is cheap to evaluate, which is essential for a computationally-efficient classification method, and (ii), the function \mathcal{C} retains most of the compositional variance found in \mathcal{D} , which we expect to translate into good selectivity when classifying compositions. In addition, each component of \mathbf{C} is bounded between zero and one.

5.2. Partitioning

The objective of the *Partitioning* algorithm is to partition the classifying space \mathbb{C} into N_R regions $R_{J=1,\dots,N_R}$, determined so that for each index J , compositions $\Phi \in \mathbb{F}$ that satisfy

$$\mathcal{J}(\Phi) = J, \quad (26)$$

share similar chemical kinetic characteristics. This has two practical implications: (i) the compositions in R_J can be represented using a much smaller representation ϕ_J ; and (ii) the short-term time evolution (over time interval Δt) of each of these compositions in R_J can be described accurately by a common specialized kinetic model M_J much smaller than the original detailed model $M_{\mathcal{D}}$. Instead of considering all compositions $\Phi \in \mathbb{F}$ satisfying Eq. (26) to delineate the regions, we first restrict the analysis on the $n_{\mathcal{D}}$ discrete sample compositions contained in the database \mathcal{D} mentioned above. The *Partitioning* procedure is described in three stages: the first one presents the overall, iterative structure of the algorithm, the second provides the steps required at each iteration, and the third describes how the optimal set of N_R regions is chosen.

- *Algorithm structure.* Cutting hyperplanes H (by definition, of dimension $n_c - 1$) are iteratively introduced in \mathbb{C} to isolate groups of sample compositions with similar short-term kinetics. The partition structure is recorded in a binary tree, whose nodes contain the hyperplanes equations (fully determined by their normal vector \mathbf{V} and one anchor point \mathbf{P} , both n_c -vectors) and whose leaves are associated with increasingly refined regions and the subset of sample compositions they contain. A

two-dimensional illustration of the iterative procedure used to create and store the list of hyperplanes that define each region in \mathbb{C} is provided in Fig. 4 for a database \mathcal{D} containing 6 compositions.

- *Iteration step.* The same task has to be accomplished during each iteration of the partitioning procedure: given an initial region that contains a subset $\mathcal{D}_m \subset \mathcal{D}$ of $n_{\mathcal{D}_m}$ compositions (determined by a previous iteration, or initially $\mathcal{D}_m = \mathcal{D}$), find the hyperplane H in \mathbb{C} -space that optimally divides the region, and therefore the set of compositions \mathcal{D}_m , into two sub-regions/subsets with more homogeneous kinetic characteristics. The concept of kinetic homogeneity of a region is defined here as the potential to derive a specialized kinetic model of minimal size able to describe accurately the short-term time evolution of any composition belonging to this region.

We use the positive DRGEP coefficients $\mathbf{R}^{\mathcal{D}_m}$ (Eq. (22)) defined based on \mathcal{D}_m to form an estimate of the kinetic homogeneity of the ensemble of compositions \mathcal{D}_m as:

$$\mathcal{H}_{\mathcal{D}_m} = \|\mathbf{R}^{\mathcal{D}_m}\|_1 = \sum_{i=1}^{n_s} |R_i^{\mathcal{D}_m}| = \sum_{i=1}^{n_s} R_i^{\mathcal{D}_m}. \quad (27)$$

The smaller that \mathcal{H} is, the more specialized and reduced the corresponding kinetic model can be, as is briefly justified below. By construction of the DRGEP coefficients, and especially Eq. (22), this definition satisfies the following observations: (i) the largest value of \mathcal{H} is obtained for the full composition database \mathcal{D} (no sub-division); (ii) the value of \mathcal{H} of any subset of \mathcal{D}_m is necessarily equal to or smaller than $\mathcal{H}_{\mathcal{D}_m}$; and (iii) if all compositions in a set are identical, the homogeneity of any subset is equal to that of the set itself.

Reduced kinetic models tailored for a given set of sample compositions are obtained by removing from the mechanism all species with a DRGEP coefficient less than a user-defined threshold. This threshold can be interpreted as an approximation of the actual prediction error introduced by the reduction [9]. As a consequence, a low value of \mathcal{H} for a set is associated with low DRGEP coefficients and more species removed from the model for a given reduction error threshold. The objective of each iteration is therefore to identify a hyperplane H_m that divides \mathcal{D}_m into two complementary (left and right) subsets, $\mathcal{D}_{m,L}$ and $\mathcal{D}_{m,R}$, contained in subregions $R_{m,L}$ and $R_{m,R}$, respectively, that minimizes the kinetic inhomogeneity, or “cost”:

$$c_{\mathcal{D}_m} = \mathcal{H}_{\mathcal{D}_{m,L}} + \mathcal{H}_{\mathcal{D}_{m,R}} = \|\mathbf{R}^{\mathcal{D}_{m,L}}\|_1 + \|\mathbf{R}^{\mathcal{D}_{m,R}}\|_1. \quad (28)$$

A global minimum, if it exists, would be computationally prohibitive to obtain given the large number of compositions typically included in the initial database \mathcal{D} . Therefore, a simplified approach is adopted here that achieves a significant decrease in the value of \mathcal{H} of each subregion, while remaining computationally affordable. The algorithm is described in Fig. 5.

- *Consolidation into a partition with N_R regions.* Without any constraints on the number of regions, the iterative partitioning proceeds until \mathbb{C} is partitioned into $n_{\mathcal{D}}$ regions, each of them containing a single composition Φ from \mathcal{D} . To restrict the number of regions to the specified value N_R , the complete tree is analyzed to identify the upper tree with N_R leaves that minimizes

$$c_{\text{tree}} = \sum_{J=1}^{N_R} \mathcal{H}_{\mathcal{D}_J}, \quad (29)$$

where the databases \mathcal{D}_J are those formed by all compositions Φ contained in region J . This is done in two steps, efficiently implemented using dynamic programming and recursion: in a first pass, the values of \mathcal{H} associated with each leaf and node in the

complete tree are computed. By construction, leaves of the complete tree represent regions containing a single composition (denoted by Φ_{leaf}), hence

$$\mathcal{H}_{\text{leaf}} = \|\mathbf{R}^{\Phi_{\text{leaf}}}\|_1. \quad (30)$$

On the other hand, the value of \mathcal{H} of a node is defined as

$$\mathcal{H}_{\text{node}} = \|\mathbf{R}^{\mathcal{D}_{\text{node}}}\|_1, \quad (31)$$

where $\mathcal{D}_{\text{node}}$ is the set of compositions in the leaves originating from this node. Once the optimal upper tree is identified, all hyperplanes below the terminal elements of this upper tree are removed, effectively consolidating regions into the specified number N_R . Figure 6 illustrates the consolidation process on the same partition problem as before, bringing the number of regions from a maximum of 6 to the specified value of $N_R = 4$.

The final result of this partitioning is a binary tree with a hyperplane $H = \{\mathbf{V}, \mathbf{P}\}$ at each node, and the index J of a distinct region at each leaf. Given a value of the classifying variable \mathbf{C} , the tree can be traversed to determine the corresponding region J . The tree is traversed by moving to a node’s left or right child, depending on the sign of $(\mathbf{C} - \mathbf{P}) \cdot \mathbf{V}$.

5.3. Reduced modeling

The *Partitioning* step establishes the structure of a partition of \mathbb{C} and \mathbb{F} into N_R regions, each region R_j containing a subset \mathcal{D}_j of size $n_{\mathcal{D}_j}$ of the original composition database \mathcal{D} . The objective of the *Reduced Modeling* step is to identify a reduced kinetic model M_j able to accurately describe the short-term time evolution of any composition belonging to region R_j . We assume that $n_{\mathcal{D}_j}$ is large enough to characterize the chemistry relevant to region R_j , and therefore it is sufficient to obtain a model M_j which is accurate for all compositions in \mathcal{D}_j .

5.3.1. Definition of error

In order to quantify the accuracy of a model M_j , we define a normalized error in composition incurred over a reaction sub-step. For a composition in \mathcal{D} , here denoted by $\Phi^D(0)$, in region J (i.e., $\mathcal{J}(\Phi^D(0)) = J$), the composition after a reaction sub-step is

$$\Phi^D(\Delta t) = \int_0^{\Delta t} \mathbf{S}(\Phi^D(t)) dt, \quad (32)$$

where \mathbf{S} is the chemical source term given by the detailed mechanism. Using instead the model M_j , the corresponding reconstructed composition after the reaction sub-step is denoted by $\Phi(\Delta t)$. Based on these two compositions after reaction, we define the model error (for model M_j and the initial condition $\Phi^D(0)$) to be

$$\varepsilon_{\Phi^D(0)} \equiv \frac{\|\Phi^D(\Delta t) - \Phi(\Delta t)\|_2}{\|\Phi^D(\Delta t)\|_2}. \quad (33)$$

5.3.2. Reaction sub-step with reduced modeling

If the reduced representation contains no locally-important species, so that all non-zero species are present in the model M_j , then the reaction sub-step using the reduced kinetic model consists of the following steps:

1. $\phi_j(0) = \mathcal{R}_{F \rightarrow J}(\Phi^D(0)),$ (34)

2. $\phi_j(\Delta t) = \int_0^{\Delta t} \mathbf{S}_j(\phi_j(t)) dt,$ (35)

3. $\Phi(\Delta t) = \mathcal{R}_{J \rightarrow F}(\phi_j(\Delta t)),$ (36)

where \mathbf{S}_j is the chemical source term given by the reduced kinetic model M_j .

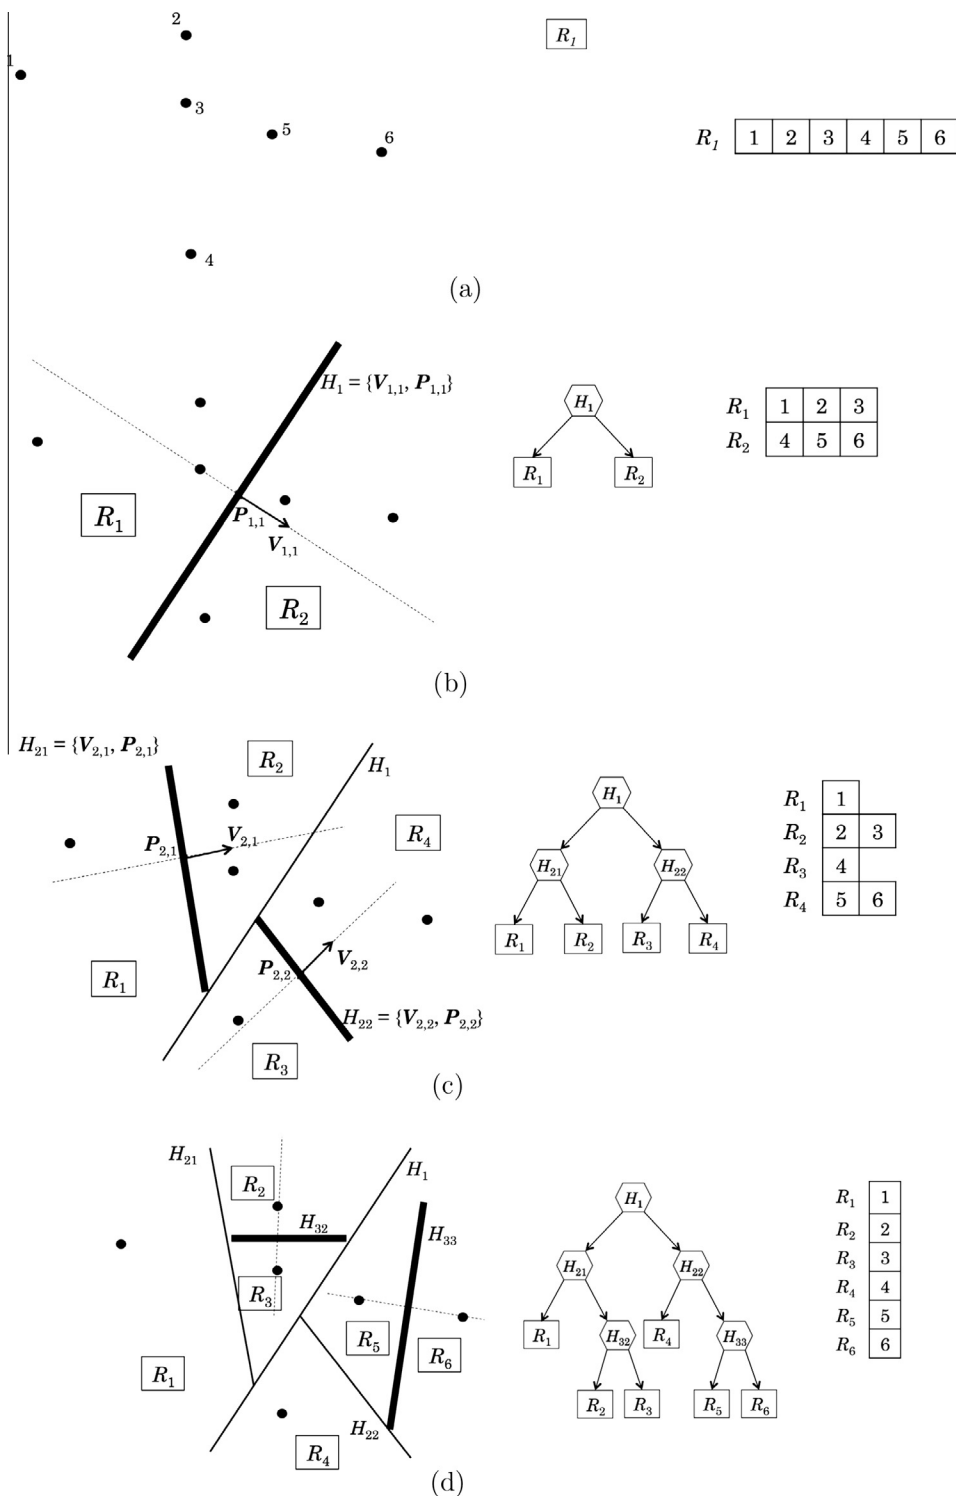


Fig. 4. Illustration of the iterative partitioning algorithm applied to a two-dimensional classifying space \mathbb{C} and a sample composition database containing 6 elements. (a) *Initial state.* Numbered dots indicate the location in \mathbb{C} space of each composition in the database, obtained through $\mathbf{C}^{(n)} = C(\Phi^{(n)})$, $n = 1, \dots, 6$. All compositions initially belong to a single region R_1 . (b) *First iteration.* A hyperplane H_1 , defined by its normal vector $V_{1,1}$ and anchor point $P_{1,1}$, is placed in \mathbb{C} , creating two distinct regions R_1 and R_2 . (Note that in each iteration the regions' indices may be re-defined.) The binary tree recording the partition structure now counts one node, containing H_1 , and two leaves, corresponding to R_1 and R_2 . All points \mathbf{P} in R_1 satisfy $(\mathbf{P} - \mathbf{P}_{1,1}) \cdot \mathbf{V}_{1,1} < 0$, while all points in R_2 satisfy $(\mathbf{P} - \mathbf{P}_{1,1}) \cdot \mathbf{V}_{1,1} \leq 0$. In this example, compositions are split into $\{1, 2, 3\}$ in R_1 , and $\{4, 5, 6\}$ in R_2 . (c) *Second iteration.* Each previously determined region containing two or more compositions is split further into two subregions using two additional cutting hyperplanes, H_{21} and H_{22} . A total of four regions are obtained. The partition binary tree is updated correspondingly. (d) *Final iteration.* Each region still containing more than one composition is split further. Once each region contains just a single composition, no further refinement is possible. The final result is a binary tree with n_D terminal leaves and $n_D - 1$ hyperplane nodes. Compositions belonging to regions far from one another in the tree are expected to have different kinetic characteristics.

If instead the reduced composition contains locally-significant species (with mass fractions \bar{y}), the treatment is more complicated, because these species are not included in the reduced model. (This

is especially relevant when considering PPAC-DRGEP and ISAT combinations, as discussed in Section 8.1.) Locally-significant species are kinetically inactive (i.e., essentially inert): the only role

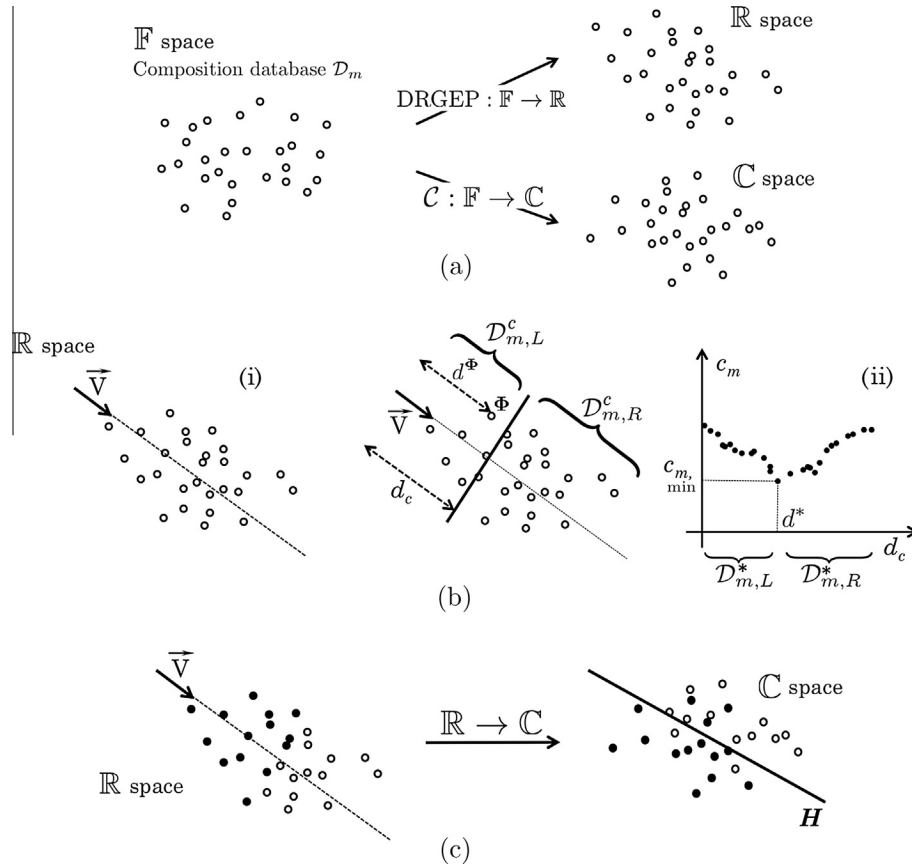


Fig. 5. Tasks performed at each step of the *Partitioning* iterative procedure. (a) For each composition Φ in database \mathcal{D}_m , the DRGEP coefficients \mathbf{R}^Φ (\mathbb{R} space) and the classifying composition \mathbf{C} (\mathbb{C} space) are computed. (b) In \mathbb{R} space, the first principal component vector \mathbf{V} of the ensemble of normalized DRGEP coefficients $\hat{\mathbf{R}}^\Phi$, defined as $\hat{\mathbf{R}}^\Phi = \mathbf{R}^\Phi / \frac{1}{n_D} \sum_{\Phi \in \mathcal{D}_m} \mathbf{R}^\Phi$ for $i = 1, \dots, n_S$, is computed (left). Compositions in \mathcal{D}_m are then projected orthogonally onto \mathbf{V} and ranked according to their distance d^Φ along the direction of \mathbf{V} . Sub-databases $\mathcal{D}_{m,L}^c$ and $\mathcal{D}_{m,R}^c$ are defined for any value d_c , by separating compositions with $d^\Phi < d_c$ from those with $d^\Phi \geq d_c$ (middle). The value of d_c, d^* , leading to minimum of c_m determines $\mathcal{D}_{m,L}^*$ and $\mathcal{D}_{m,R}^*$ used in the final step (right). (c) Each composition is colored based on its belonging to $\mathcal{D}_{m,L}^*$ or $\mathcal{D}_{m,R}^*$. The final step aims at identifying the hyperplane in \mathbb{C} -space that best reproduces this coloring, that is, the hyperplane that creates two subdatabases $\mathcal{D}_{m,L}$ and $\mathcal{D}_{m,R}$ as close as possible to $\mathcal{D}_{m,L}^*$ and $\mathcal{D}_{m,R}^*$. (Note that one-to-one correspondence cannot be achieved in general due to the non-linearity of the relationship between Φ and \mathbf{R}^Φ .) This is done using covariance ellipsoids, as described in [Appendix](#).

they play is through their thermal capacity and as acting as third bodies in reactions. We approximate these effects by treating all locally-significant species as being equivalent to an equal mass of a specified inert, which we take to be N_2 .

To implement this treatment in the reaction fractional step described above, after the reduction process (Eq. (34)), the mass fraction of N_2 is increased by $\hat{y}_{\text{tot}} \equiv \sum \hat{y}_i$, and the value of $\hat{\mathbf{y}}$ in $\phi_J(0)$ is set to zero. Then after reaction (Eq. (35)), the mass fraction of N_2 is decreased by \hat{y}_{tot} and the original value of $\hat{\mathbf{y}}$ is restored.

It may be noted that this treatment requires that a specified inert must be one of the species represented in each model, although in principle different models could have different species as the designated inert. For the PaSR case considered here, N_2 is present in every model, so this condition is satisfied.

5.3.3. Generation of reduced kinetic models

For each region J , the DRGEP methodology is used to generate a set of models of different sizes and resulting in different errors. Such a pair of model and error is denoted by $\{M_J|_\varepsilon, \varepsilon\}$ and has the following properties:

1. For every composition Φ in \mathcal{D}_J , the model error ε_Φ is at most ε .
2. All *regionally-significant* species (as defined below) are retained in the model.

3. *Locally-significant* species (as defined below) are not retained in the model, but are treated according to the procedure above when evaluating ε_Φ .

4. Subject to 1, 2, and 3: the number of species included in $M_J|_\varepsilon$ is minimum, as is the number of reactions.

A species is deemed to be *regionally significant* if its mass fraction exceeds $\varepsilon_{\text{sig}} = 10^{-3}$ for at least 10% of the samples in \mathcal{D}_J , that is, 10% of the sample compositions contained in region R_j of the composition space. It is deemed *locally significant* for a composition Φ if it is not *regionally significant*, but its mass fraction in Φ still exceeds $\varepsilon_{\text{sig}} = 10^{-3}$. The reason for imposing condition 2 is that the DRGEP methodology eliminates essentially inert species, even if they have large mass fractions, and eliminating such species from the model results in large conversion errors. Condition 3 is a fail-safe mechanism that prevents any one composition with unusually large amount (compared to other compositions in the region) of an inactive species to dominate the model error.

Algorithm 1 details how the DRGEP methodology is used to generate a list of reduced models. Each successive model in the list has smaller size and (usually) larger error.

Algorithm 1. Species and reaction elimination algorithm applied to the database \mathcal{D}_J to produce a list Υ of k pairs $\{M_J|_\varepsilon, \varepsilon\}$ of reduced models $M_J|_\varepsilon$ with errors ε .




Fig. 6. Consolidation of the complete tree obtained in Fig. 4 into an optimal upper tree with $N_R = 4$ regions. (a) The consolidation procedure first computes the \mathcal{H} values of each element in the tree: the 6 terminal leaves R_1 to R_6 , and the 5 nodes H_1 , H_{21} , H_{22} , H_{32} , and H_{33} . The optimal upper tree is chosen among the 3 possible 4-leaf upper trees of this structure, as that minimizing c_{tree} (Eq. (29)). In this case, $c_{\text{tree}}^{\min} = \mathcal{H}_{(\Phi_1)} + \mathcal{H}_{(\Phi_2)} + \mathcal{H}_{(\Phi_3)} + \mathcal{H}_{(\Phi_4, \Phi_5, \Phi_6)}$. (b) Accordingly, the hyperplane H_{22} and H_{33} are removed, thereby merging regions R_4 , R_5 , and R_6 . The final *Partitioning* binary tree is trimmed to the optimal 4-leaf upper tree identified above. Sample compositions are now labeled as belonging to one, and only one, region out of the 4 contained in the partition.

- 1: Set all species as RETAINED and ACTIVE.
- 2: Set all reactions to RETAINED.
- 3: $k = 0$
- 4: $\delta = 10^{-2}$
- 5: **for** $i = 1$ to n_s **do**
- 6: Compute \mathbf{R}^{D_i} , account for NON-RETAINED species using Eq. (13).
- 7: Sort the elements of \mathbf{R}^{D_i} by increasing value.
- 8: Set the status of species i with smallest, non-zero $R_i^{D_i}$ coefficient to INACTIVE
- 9: Set reactions involving species i to REMOVED.
- 10: If species i is not *regionally* or *locally significant*, set its status to REMOVED.
- 11: Create reduced model M_j from all RETAINED species and reactions.
- 12: Evaluate the model error $\varepsilon_S = \max_{\Phi \in \mathcal{D}_j}(\varepsilon_{\Phi})$ using M_j .
- 13: Compute $\mathbf{R}^{D_j, r}$, account for NON-RETAINED species and REMOVED reactions (Eq. (13)).
- 14: Sort the elements of $\mathbf{R}^{D_j, r}$ by increasing value.
- 15: $\varepsilon_R \leftarrow \varepsilon_S$.
- 16: **while** $\varepsilon_R \leq (1 + \delta)\varepsilon_S$ **do**
- 17: Set status of reaction j with smallest $R_j^{D_j, r}$ coefficient to REMOVED.
- 18: Update reduced model M_j from all RETAINED species and reactions.

- 19: Evaluate $\varepsilon_R = \max_{\Phi \in \mathcal{D}_j}(\varepsilon_{\Phi})$ using M_j .
- 20: $\varepsilon \leftarrow \varepsilon_R, M_j|_{\varepsilon} \leftarrow M_j$
- 21: $k \leftarrow k + 1$
- 22: Add the pair $\{M_j|_{\varepsilon}, \varepsilon\}$ as the k th element of the model list Υ .

To improve the ranking ability of DRGEP, sensitivity analysis coefficients can be used to replace \mathbf{R}^{D_j} when ε becomes larger than a pre-defined value. The sensitivity coefficient of species A is set to 0 if A is INACTIVE, and set to $\varepsilon_{SA}^A = \max_{\Phi \in \mathcal{D}_j} \varepsilon_{\Phi}^A$, where ε_{Φ}^A is evaluated using the k th model in Υ , further reduced by removing species A and all corresponding reactions.

When an adaptive simulation is performed with a specified error tolerance ε_c , for every region J , the model M_j used is the smallest one in the list Υ satisfying $\varepsilon \leq \varepsilon_c$.

5.4. Classification

This and the next two sub-sections describe the operators needed in the adaptive simulation stage, namely *Classification*, *Reconstruction*, and *Reduction*. Note that both *Reduction* and *Reconstruction* are also used to evaluate ε_{Φ} in the Reduced Modeling algorithm.

The *Classification* operator \mathcal{J} identifies the region J a composition Φ belongs to through a simple binary tree search done in

the tree formed in the *Partitioning* stage. First, the value of the classifier is determined as $\mathbf{C} = \mathcal{C}(\Phi)$, (Eq. (24)). Each node in the tree contains the equation of a cutting hyperplane $H = \{\mathbf{V}, \mathbf{P}\}$, where \mathbf{V} and \mathbf{P} are the normal vector and an anchor point of this hyperplane in the classifying space. Starting from the root of the tree, the scalar product $(\mathbf{C} - \mathbf{P}) \cdot \mathbf{V}$ is computed, and its sign decides which node (or leaf) to access next. The procedure is repeated until a leaf is reached. The index of the region associated with this leaf is returned as J .

5.5. Reconstruction

The $\mathcal{R}_{J \rightarrow F}$ *Reconstruction* operator reconstructs a full composition vector Φ from a reduced representation ϕ_J that involves the reduced set of species θ_J and locally-significant species with mass fractions $\hat{\mathbf{y}}$. While various reconstruction methods have been tested involving various levels of complexity and assumptions, we conclude that the most robust operation is simply to set the state represented by Φ to be identical to that represented by ϕ_J . This is achieved by: setting all species not represented in ϕ_J to zero in Φ ; setting all other species mass fraction in Φ to the corresponding mass fractions in ϕ_J ; and setting the temperature in Φ to be the same as that in ϕ_J . This operation conserves enthalpy and species (and therefore also elements and the normalization condition on mass fractions).

5.6. Reduction

The last operator that needs to be defined is the $\mathcal{R}_{F \rightarrow J}$ *Reduction* operator, which performs the reverse of the *Reconstruction* operator, i.e. taking a full representation and converting it into the reduced representation ϕ_J . A straightforward reduction is adopted here which conceptually involves the following steps:

1. Consider the full representation $\Phi = \{\mathbf{Y}, T\}$.
2. Set $Y_i = 0$ for every species i that is neither included in the model M_J nor is locally significant (which is determined by comparing Y_i to ε_{sig}).
3. Rescale the mass fractions \mathbf{Y} so that they sum to unity.
4. Set the mass fractions \mathbf{y}_J and $\hat{\mathbf{y}}$ to the corresponding values in \mathbf{Y} , and set T to the same temperature as in the full composition to obtain the reduced composition $\phi_J = \{\mathbf{y}_J, \hat{\mathbf{y}}, T\}$.

Because of steps 2 and 3, neither species nor elements are conserved in the reduction. In addition, enthalpy is not conserved (as it could be with a different specification of T in the reduced representation). The errors resulting from discarding and rescaling species are referred to as *conversion errors*. We explored variants with improved conservation properties, but none proved superior overall in accuracy and efficiency compared to the simple treatment described above.

5.7. Summary of inputs to the adaptive chemistry methodology

We summarize here the inputs to the adaptive chemistry methodology, and give some of their values used in the tests. The required inputs are:

- The detailed chemical mechanism M_D .
- The targets \mathcal{T} specified for DRGEP, which we take to be fuel, O_2 , OH , CO , CO_2 , and heat release.
- The PaSR parameters, including the time step Δt .
- The number n_D of samples in \mathcal{D} , which we take to be $n_D = 10,000$.

- The dimension, n_C , of the classifying space, which we take to be $n_C = 10$.
- The number of regions, N_R , which we vary between 1 and 100.
- The error tolerance, ε_c , which we vary over a wide range.
- The mass-fraction threshold for significant species ε_{sig} , which we take to be $\varepsilon_{\text{sig}} = 10^{-3}$.

6. PaSR test case and composition database

The configuration chosen to develop and test the adaptive methodology presented above is a non-premixed propane/air piloted PaSR. We consider the detailed mechanism for propane oxidation developed by Curran et al. [46], which consists of $n_s = 115$ species involved in $n_r = 1300$ reactions. The PaSR has three inflow streams: pure fuel, pure oxidizer (air), and a pilot stream defined as the burnt products of a stoichiometric propane/air mixture in chemical equilibrium. The pressure is 1 bar throughout. The fuel and air streams are at a temperature of 300 K, while the pilot stream is at the adiabatic flame temperature of 2272 K. Initially, all particles in the reactor are set to the pilot composition. The number of particles, and hence the total mass, in the reactor remains constant. The particle compositions change in time through reaction and mixing, as described in Section 2. All relevant parameters are included in Table 1.

The PaSR is run with the above parameters for 10 residence times. We observe a transient state, which depends on the initial conditions, for about the first three residence times, after which the PaSR operates in a statistically stationary state. A sample composition database \mathcal{D} containing $n_D = 10,000$ distinct compositions is assembled by randomly sampling the particles in the reactor over the entire duration of the simulation. (The performances of the adaptive chemistry approach is found to be rather insensitive to the number of compositions n_D included in \mathcal{D} .)

7. Results

In this section we present a quantitative assessment of both the accuracy and efficiency of the adaptive chemistry strategy, as applied to the PaSR simulation.

7.1. PaSR test runs

The parameters used in the PaSR simulations are provided in Section 6; some of the parameters used in the adaptive methodology are given in Section 5.7; and the remainder are given below. A series of PaSR simulations are performed with a broad range of values of the error tolerance, ε_c , and for five values of N_R , namely $N_R = 1, 10, 30, 60$, and 100. Consequently, in the pre-processing stage, for each value of N_R , the partitioning algorithm is applied to determine the N_R regions, and for each region the reduced modeling is used to produce the list of models $M_J|_{\varepsilon}$. Then, prior to a simulation with a specified value of ε_c , the smallest models are selected from the list, subject to the condition $\varepsilon \leq \varepsilon_c$.

The adaptive PaSR simulations are then carried out according to the algorithm described in Section 3 and illustrated in Fig. 3. A (non-adaptive) PaSR simulation is also performed using the detailed chemical mechanism M_D and the full representation to serve as reference for error measurements. Quantities extracted from the tests with the adaptive strategy and with detailed chemistry are indicated with the superscripts “A” and “D”, respectively.

It is found that the results obtained for the degenerate case $N_R = 1$ are qualitatively different from those obtained with larger values of N_R . This case is degenerate in that the adaptive methodology is in fact not adaptive. Henceforth we therefore refer to this as the *non-adaptive case*.

In order to create a fair test of the adaptive method, the seed that generates the pseudo-random numbers used in the PaSR simulations is changed from that used to create the sample database \mathcal{D} . Therefore the PaSR particles follow trajectories in the same areas of the composition space, but they are not strictly identical to those obtained in the pre-processing stage. Particle evolutions are recorded over a total time τ equal to 10 residence times, corresponding to $n_t = 1000$ time steps.

7.2. Definition of error

Errors are measured for species mass fractions and temperature. Rather than comparing reactor statistics, such as mean values, we opt for a more stringent test of our approach by comparing individual particle trajectories obtained in the detailed and adaptive PaSR simulations. We choose the following integrated error measure, written here for any quantity X (species mass fraction or temperature):

$$\varepsilon_X = \frac{\sum_{k=1}^{n_t} \sum_{n=1}^{n_p} |X_k^{(n),A} - X_k^{(n),D}|}{\sum_{k=1}^{n_t} \sum_{n=1}^{n_p} |X_k^{(n),D}|} \quad (37)$$

In the above equation, $X_k^{(n),A}$ and $X_k^{(n),D}$ denote the value of quantity X for particle n on time step k of a PaSR simulation with adaptive chemistry and detailed chemistry, respectively.

7.3. Control of errors

Figure 7 shows the incurred error in temperature, ε_T , as a function of the specified error tolerance ε_c for each of the five values of N_R . As may be seen, the adaptive strategy is successful in controlling the error, in that, for each value of N_R , the incurred error ε_T decreases monotonically as the error tolerance ε_c is decreased, until a floor is reached at the very low value of 10^{-6} . In a turbulent combustion calculation, an error of 1% in species mass fractions and 5 K in temperature (i.e., $\varepsilon_T \approx 3 \times 10^{-3}$) is generally regarded as more than acceptable.

We recall that there are two sources of error. The first is in the reduced kinetic model, and this is directly controlled by ε_c . The second is the conversion error. This is controlled indirectly by ε_c , since the smaller ε_c , the larger the reduced models, and hence the smaller the number (and magnitude) of species omitted from the models. As may be seen from Fig. 7, for a fixed value of the error




Fig. 7. The incurred error in temperature, ε_T , vs. the specified error tolerance, ε_c , in PaSR simulations. Comparison between the non-adaptive case, $N_R = 1$ (\star) and adaptive cases with $N_R = 10$ (\diamond), 30 (\square), 60 (\bullet), and 100 (\blacktriangle). The unit slope line (thin black line) is added for reference.

tolerance ε_c , the incurred error ε_T increases with increasing N_R . This is likely due to a combination of two effects. First, the larger N_R , the more crossings there are likely to be from one region to another, and hence more conversions are performed. Second, the larger N_R , the smaller (on average) the regions, and hence fewer species are needed to represent the more specialized chemistry, and hence there are more species omitted from the models.

It is the incurred error (not the error tolerance) that is of consequence, and in comparing two approaches or variants, this should be done at fixed incurred error, not at fixed error tolerance. Consequently, henceforth, we use ε_T to parameterize the error, and we plot quantities of interest against ε_T rather than ε_c .

Figure 8 shows the incurred error for different species mass fractions plotted against the incurred error in temperature for the adaptive case with $N_R = 30$. As may be seen, for CO_2 , C_3H_8 and H_2O , the species errors are very close to the temperature error, whereas the error for CO is about 5 times that for T . Results for different values of N_R and for different species are similar to those shown in Fig. 8. The important observation is that ε_T provides a good characterization of all of the errors.

To help relate the error levels as defined in Eq. (37) to actual temperature changes, in Fig. 9 (top), the mean reactor temperature is plotted over three residence times for the detailed and adaptive simulations with $\varepsilon_T = 10^{-3}$ and $N_R = 30$. As may be seen, the two curves are essentially indistinguishable, and so we plot their difference in Fig. 9 (bottom). Overall, the maximum temperature discrepancies are around 2.5 K out of 1650 K, and more importantly, these discrepancies are found not to increase in any significant manner over time, an essential property for the adaptive strategy to be viable in large-scale simulations. This behavior has been observed consistently for all quantities discussed below.

7.4. Computer time

The PaSR simulations are performed in parallel on a Beowulf cluster based on Nehalem X series processors, QDR Infiniband, and 2 GB of RAM per core. For the simulation using detailed chemistry, the wall clock time needed to reach $t = 10\tau_{\text{res}}$ on 12 cores is 4020 s. For a given number of regions N_R , the partitioning, reduction, and model evaluation over the range of incurred errors shown in Fig. 8 require on average 10^3 CPU hours. This is but a small fraction of the typical cost of LES/PDF simulations, which can quickly reach upward of 10^5 CPU hours even for simple flames. This cost therefore appears very reasonable, especially since the computational overhead due to the adaptive treatment of the chemistry at run time in the PPAC approach is negligible.




Fig. 8. Incurred error in species ε_{Y_i} vs. incurred error in temperature ε_T for $N_R = 30$. Species are: CO_2 (\diamond), CO (\square), C_3H_8 (\blacktriangle), H_2O (\star).




Fig. 9. Top: Evolution of the mean temperature in the PaSR for the detailed ($\langle T^D \rangle$, red solid line) and adaptive ($\langle T^A \rangle$, blue dashed line) simulations using $\varepsilon_T = 10^{-3}$ and $N_R = 30$. (The two lines are indistinguishable.) Bottom: the absolute difference $\Delta T = |\langle T^A \rangle - \langle T^D \rangle|$ in the mean temperature between the two simulations. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

We define the *relative time* t_{rel} of an adaptive simulation as its CPU time divided by that of the simulation with detailed chemistry. Figure 10 shows the relative time t_{rel} as a function of the incurred temperature error for the different values of N_R . It may be observed that, as expected, all the simulations using reduced models require less CPU time than those using the detailed mechanism (i.e., t_{rel} is less than unity). Furthermore, the adaptive cases ($N_R > 1$) require, typically, just half the time of the non-adaptive case ($N_R = 1$). At an incurred error of $\varepsilon_T = 3 \times 10^{-3}$, the CPU time required for the adaptive cases is about 15% of that for the detailed mechanism, and about 55% of that for the non-adaptive case. In general, the CPU time decreases weakly with N_R , for $N_R \geq 10$.

The above comparisons are much in favor of the adaptive approach, even in a PaSR with pair-wise mixing. It is expected that its benefits will be even more pronounced in a turbulent combustion calculation, since the portion of particles carrying inerts or burnt compositions, described efficiently with few variables, is much higher than in a PaSR.




Fig. 10. Relative CPU time (adaptive chemistry to detailed chemistry) vs. the incurred temperature error. Comparison between the non-adaptive case, $N_R = 1$, (\star) and adaptive cases with $N_R = 10$ (\blacklozenge), 30 (\blacktriangle), 60 (\bullet), and 100 (\blacktriangle).

The results presented so far use DRGEP combined species and reaction elimination, as described in Section 5.3. While already effective, the standard DRGEP approach can be complemented by sensitivity analysis, to refine the ranking of species identified as of intermediate importance by DRGEP. Some incremental gain in CPU cost is observed when switching to sensitivity coefficients in lieu of DRGEP coefficients for ε_c larger than 10^{-8} , as shown in Fig. 11. In view of these results, the cost vs. benefit of using expensive sensitivity analysis methods to further decrease the complexity of the region-specific reduced representations and models for larger reference mechanisms warrants further investigation.

7.5. Computer memory

Another quantity of interest relates to the storage requirements of the PaSR simulations. This can be measured by the *relative number of species*, n_{rel} , defined as the number of species in the reduced models used (averaged over all particles and time steps) divided by the number of species in the detailed mechanism, n_s :

$$n_{rel} \equiv \frac{1}{n_t n_p n_s} \sum_{k=1}^{n_t} \sum_{n=1}^{n_p} n_{s,k}^{(n)A}, \quad (38)$$

where $n_{s,k}^{(n)A}$ is the number of species in the reduced model used for particle n on time step k .

Figure 12 shows n_{rel} as a function of the incurred error ε_T for each value of N_R . This shows, for any specified incurred temperature error in the adaptive PaSR simulations, the degree of reduction that can be achieved by the method. All curves are monotonically decreasing with increasing ε_T , from values very close to unity, corresponding to very low error levels at one extreme (in which the reduced models are very close to the detailed model), to large errors and very small reduced models at the other extreme. A clear difference is observed between the non-adaptive case ($N_R = 1$) and the adaptive cases, with the number of species retained to achieve a specified error tolerance being up to a factor of two lower for the adaptive cases (compared to the non-adaptive case). Consider $\varepsilon_T = 10^{-3}$ as a typical desired error level: in this case of the order of 30% of the species are needed (on average) in the adaptive simulations, compared to roughly 60% for the non-adaptive case ($N_R = 1$). Some minor reduction in n_{rel} is observed when increasing the number of regions N_R above 10. However, in the range of error




Fig. 11. Relative CPU cost vs. the incurred temperature error ε_T for $N_R = 30$. Comparison between reduced modeling performed using standard DRGEP coefficients (\blacktriangle), and reduced modeling performed using standard DRGEP coefficients up to $\varepsilon_c = 10^{-8}$, and sensitivity coefficients thereafter (\bullet). The non-adaptive results of Fig. 10, which are not using sensitivity analysis, are added here for reference.




Fig. 12. Relative number of species n_{rel} (adaptive to detailed) vs. the incurred error in temperature ε_T . Comparison between the non-adaptive case, $N_R = 1$ (\star) and adaptive cases with $N_R = 10$ (\blacklozenge), 30 (\blacksquare), 60 (\bullet), and 100 (\blacktriangle).

typically of interest (10^{-4} – 10^{-2}), cases with $N_R \geq 30$ can be considered equivalent. Unless stated otherwise, $N_R = 30$ is therefore used in all cases below.

7.6. Conservation properties

In the detailed PaSR simulation, over the mixing and reaction fractional steps, when summed (or averaged) over all particles, both enthalpy and the mass fractions of elements are conserved. However, in the adaptive simulations, the reduction operator $\mathcal{R}_{F \rightarrow J}$ (see Section 5.6) is not conservative. For elements, this is because some species mass fractions are set to zero, and then the mass fractions are rescaled to sum to unity. For the enthalpy, it is because the reduction is performed as constant temperature rather than at constant enthalpy. However, it is expected that the resulting conservation errors are small, since if a species is not part of a representation, its mass fraction is small by construction (certainly less than ε_{sig}). To test this hypothesis, the total enthalpy and element mass fractions are followed over time in adaptive PaSR simulations using $N_R = 30$ and a range of error tolerances. The normalized conservation error $\hat{\varepsilon}_X$ for a conserved quantity X is defined as:

$$\hat{\varepsilon}_X = \frac{1}{n_p} \frac{1}{n_t} \sum_{k=1}^{n_t} \left| \frac{\sum_{n=1}^{n_p} (X_{k+}^{(n)} - X_k^{(n)})}{\sum_{n=1}^{n_p} X_k^{(n)}} \right|, \quad (39)$$

where, for particle n on time step k , $X_k^{(n)}$ and $X_{k+}^{(n)}$ denote the value of X before and after the mixing and reaction fractional steps, respectively. Conservation errors for enthalpy and element mass fractions are plotted in Fig. 13 as functions of the temperature error ε_T . Note that as mentioned above, those conservation errors are integrated over the full simulation, and no significant increase of the instantaneous values is observed over time. When the number of species retained decreases (leading to higher conversion errors) and temperature error increases, errors for the conserved quantities remain below 10^{-5} for all cases and quantities of interest, indicating satisfactory treatment of particle conversion between various representations.

7.7. Regions accessed and model utilization

A major assumption at the core of the adaptive strategy is that the sample composition database \mathcal{D} is representative of the particle

compositions that occur in the adaptive simulations. To confirm the validity of this assumption *a posteriori*, the distribution of compositions in \mathcal{D} across regions is compared to the frequency at which particles are classified into those same regions during the adaptive PaSR simulation. Recall that, since pairing in the PaSR is a random process, we do not recover the database compositions during the adaptive simulation. Results are shown in Fig. 14, demonstrating similar region access frequency between pre-processing and adaptive simulation stages. Figure 15 illustrates the trajectory of a particular particle in \mathbb{C} space, by providing side-to-side the temperature evolution and the index of the region the particle belongs to at time t . Comparing both plots shows that changes in pairing partners are often followed by a change in region (*i.e.* model switches often occurs a few time steps after discontinuities in temperature gradients) but particles can also switch to a different representation outside of those discrete events.

While the above result brings confidence that \mathcal{D} contains the needed relevant kinetic information, being able to predict model utilization also has important implications for large-scale simulations, especially in terms of load balancing and coupling with tabulation methods such as ISAT.

8. Discussion

In this section we discuss some possible improvements and extensions to the adaptive methodology, and then consider future applications of the method in LES/PDF simulations.




Fig. 13. Conservation errors as function of the temperature error in the adaptive PaSR simulation: Enthalpy (\star , solid line), mass fraction of carbon atoms (\blacksquare), hydrogen atoms (\bullet), oxygen atoms (\blacklozenge), and nitrogen atoms (\blacktriangle).




Fig. 14. Fraction of compositions in \mathcal{D} included in Region J (\bullet) compared to the overall probability P_J of a particle to be classified as belonging to region J (bars) during the adaptive simulation stage. Regions are sorted according to the fraction of database compositions they contain.




Fig. 15. Temporal evolution of temperature T (left) and region index J (right) for one particle in an adaptive simulation with $N_R = 30$ in a non-premixed propane/air PaSR. Jumps in T are due to the inflow/outflow events.

8.1. Improvements and extensions

It is evident that the number of regions N_R is not a computational limitation, in that of order tens are sufficient, whereas of order hundreds could be handled without difficulty in computations. One can therefore consider *local improvements* to the partitioning: if a region is found to require a large reduced kinetic mechanism, it can (recursively) be split into two smaller regions, provided that (as expected) these two regions have smaller models.

The most obvious and beneficial extension to the methodology is to combine it with *in situ* adaptive simulation (ISAT) [23,24]. For each region J , a separate ISAT table can be used to perform the reaction fractional step much more efficiently than by integrating the kinetic equations. In ISAT, both the work to retrieve from the table, and the storage needed per table entry scale as the square of the number of species. Furthermore, the performance of ISAT is found to degrade significantly beyond about 50 species. Hence, there is great benefit in having N_R tables for reduced representation with relatively few species (e.g., $n_s^j < 50$), compared to one table involving hundreds or thousands of species.

It is straightforward to combine the adaptive methodology with ISAT in serial computations (or to use *purely local processing* [43] in parallel simulations); and the fact that there are many tables opens up new possibilities for efficient parallel implementations.

8.2. Use in LES/PDF simulations

For the PaSR test case used here, it is guaranteed that the database \mathcal{D} of compositions generated in the pre-processing stage are representative of those in the adaptive simulations, since both stages use statistically-identical PaSR simulations. For an adaptive LES/PDF simulation, we propose two sub-stages in the pre-processing stage in order to generate a suitably representative database. In the first sub-stage, one or more PaSR simulations are performed with conditions representative of the LES/PDF simulation to generate the first database \mathcal{D}_1 , and the reduced models are generated with a reasonably small error tolerance. Then, using the adaptive method, a *coarse* LES/PDF simulation is performed, for example, with double the grid spacing in each direction, and half the number of particles and simulation time compared to the full simulation, resulting is approximately 1/64 of the cost of the full simulation. This coarse LES/PDF is used to generate the representative database \mathcal{D} used subsequently in the adaptive strategy. (A small error tolerance is used in the first sub-stage to ensure that

all needed species are included in \mathcal{D}_1 .) When considering large detailed mechanisms (with more than, say, 300 species) it is impracticable to perform even a coarse LES/PDF with the detailed mechanism. This is why a PaSR is used in the first sub-stage. Note that further investigation is required to evaluate the relative costs involved in such a process, especially when compared to on-the-fly adaptive techniques. Yet, this provides a path forward to systematically handle simulations with very large kinetic mechanisms, which would not be otherwise feasible.

In contrast to on-the-fly adaptive chemistry methodologies, in the present method, the errors are not controlled during the simulation, but are instead controlled as part of the pre-processing stage. Their control during the adaptive simulation depends on the database \mathcal{D} used in the pre-processing stage being completely representative. It is, therefore, desirable to be able to measure the errors during the adaptive simulation, and hence to verify that they are indeed controlled (and possibly to invoke a correction strategy if they are not). This can be done with negligible computational penalty (as it is in ISAT [24]) by randomly sampling a small fraction (e.g., 10^{-4}) of the particle compositions prior to the reaction fractional step, and comparing the reaction mappings obtained with the reduced model and the detailed mechanism.

9. Conclusions

A novel adaptive strategy specifically tailored for particle PDF methods to simulate turbulent combustion systems has been presented.

Rather than reducing the chemistry on-the-fly during the combustion simulation, a set of reduced representations and reduced kinetic models is assembled *a priori*, which are assigned adaptively to each individual particle during the simulation. The approach includes the following major characteristics:

- In the pre-processing stage: a database \mathcal{D} of representative compositions is generated; a classifying space is defined and suitably partitioned into N_R regions, for each of which DRGEP is used to generate a locally-accurate reduced kinetic mechanism M_j , much smaller than the detailed mechanism; and the particle composition in region R_j is represented by a reduced representation ϕ_j .
- In the adaptive simulation stage, for each PDF particle: a classification algorithm based on a low-dimensional binary tree is used to identify the region J the particle's composition belongs to; the reduced kinetic model M_j is used to advance the particle

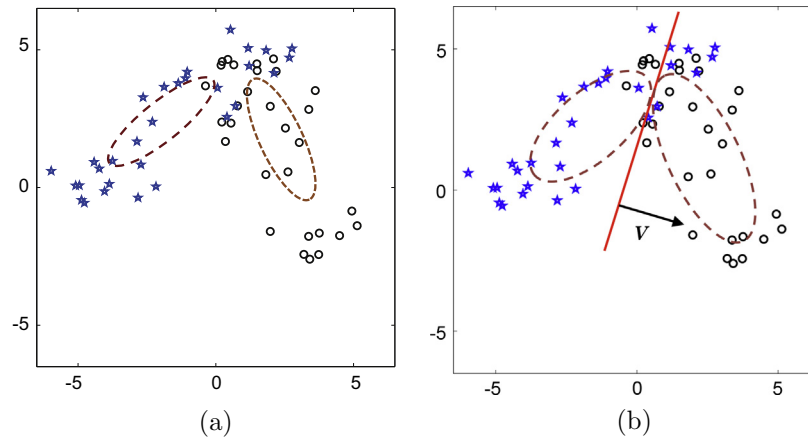


Fig. 16. Two-dimensional illustration of the use of ellipsoids to identify the cutting hyperplane H in \mathbb{C} -space. (a) Initial subsets $\mathcal{D}_{m,L}^*$ (blue stars) and $\mathcal{D}_{m,R}^*$ (open black circles), along with ellipsoids E_L and E_R . (b) The size s of E_L and E_R are adjusted so that they intersect at one point. The hyperplane normal vector is chosen as that normal to the tangent plane. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

composition in the reaction fractional step; and, the reduced representation of a particle composition can be adaptively chosen and converted as it moves from one region to another.

Based on these characteristics, the methodology is referred to as Pre-Partitioned Adaptive Chemistry using DRGEP, or PPAC-DRGEP.

The performance of the adaptive methodology has been demonstrated for a propane/air non-premixed PaSR. The results show that the incurred errors are well controlled by the specified error tolerance, and can be reduced well below acceptable levels. Even though the methodology does not exactly conserve elements or enthalpy, nevertheless, the error in these conserved quantities is very small. Typically, the CPU time using the adaptive strategy is 15% of that using the detailed mechanism; and 50% of that using a single reduced mechanism. On average, only 30% of the species are included in the reduced models. These computational gains are expected to be much larger in PDF simulations and with larger detailed mechanisms.

As discussed in Section 8, in future work, the adaptive methodology will be combined with ISAT and used in LES/PDF simulations of turbulent flames.

Acknowledgment

This material is based upon work supported by the U.S. Department of Energy Office of Science, Office of Basic Energy Sciences under Award Number DE-FG02-90ER14128.

Appendix

This appendix details the algorithm used to accomplish the partitioning step depicted in Fig. 5c. We consider a set of compositions \mathcal{D}_m , divided into two complementary subsets $\mathcal{D}_{m,L}^*$ and $\mathcal{D}_{m,R}^*$, of size n_L and n_R respectively. $\mathcal{D}_{m,L}^*$ and $\mathcal{D}_{m,R}^*$ have been obtained by minimization in \mathbb{R} space of c_m along the first principal component direction of the normalized DRGEP coefficient vectors \mathbf{R}^Φ , for all $\Phi \in \mathcal{D}$. We want to identify an hyperplane $H = \{\mathbf{V}, \mathbf{P}\}$ in \mathbb{C} such that

$$\begin{aligned} (\mathbf{C}_\Phi - \mathbf{P}) \cdot \mathbf{V} &> 0 \quad \text{for } \Phi \in \mathcal{D}_{m,L}^*, \quad \text{and} \\ (\mathbf{C}_\Phi - \mathbf{P}) \cdot \mathbf{V} &< 0 \quad \text{for } \Phi \in \mathcal{D}_{m,R}^*, \end{aligned} \quad (40)$$

where $\mathbf{C}_\Phi = \mathcal{C}(\Phi)$ (Eq. (24)). Due to the non-linear implicit transformation between \mathbb{C} and \mathbb{R} spaces, the existence of an hyperplane in \mathbb{C} that satisfies those inequalities for all $\Phi \in \mathcal{D}_m$ is not guaranteed.

Instead, a pragmatic approach based on covariance ellipsoids is used. The steps are as follows (all operations are made in \mathbb{C} -space):

1. Form the means \mathbf{m}_L and \mathbf{m}_R of \mathbf{C}_Φ vectors from compositions contained in $\mathcal{D}_{m,L}^*$ and $\mathcal{D}_{m,R}^*$, respectively.
2. Form the ellipsoid E_L defined by

$$E_L \equiv \{\mathbf{C} | (\mathbf{C} - \mathbf{m}_L)^T \mathbf{C}_L^{-1} (\mathbf{C} - \mathbf{m}_L) \leq s^2\}, \quad (41)$$

where $\mathbf{C} \in \mathbb{C}$, \mathbf{C}_L is the covariance matrix of all \mathbf{C}_Φ from $\mathcal{D}_{m,L}^*$, and s determines the size of the ellipsoid. E_R is similarly defined. This step is illustrated in Fig. 16a.

3. Determine the smallest value of s at which E_L and E_R intersect. This is achieved using the ELL_LIB library developed by Pope [47].
4. The normal vector \mathbf{V} of the cutting hyperplane H is taken to be that of the plane tangent to both ellipsoids at the point of intersection (Fig. 16b).
5. The anchor point \mathbf{P} is chosen along the direction of \mathbf{V} to minimize the kinetic inhomogeneity c_m of the resulting $\mathcal{D}_{m,L}$ and $\mathcal{D}_{m,R}$, following the same procedure as that in Fig. 5b(ii).

References

- [1] T. Lu, C.K. Law, *Prog. Energy Combust. Sci.* 35 (2009) 192–215.
- [2] O. Herbinet, W.J. Pitz, C.K. Westbrook, *Combust. Flame* 154 (2008) 507–528.
- [3] R.W. Grout, A. Gruber, C.S. Yoo, J.H. Chen, *Proc. Combust. Inst.* 158 (2011) 2113–2127.
- [4] V. Hiremath, S.R. Lantz, H. Wang, S.B. Pope, *Proc. Combust. Inst.* 34 (1) (2013) 205–215.
- [5] N. Peters, *Turbulent Combustion*, Cambridge University Press, 2000.
- [6] C.D. Pierce, P. Moin, *Progress-variable Approach for Large Eddy Simulation of Turbulent Combustion*, Tech. Rep. TF80, Flow Physics and Computation Division, Dept. Mech. Eng., Stanford University, 2001.
- [7] T. Lu, C.K. Law, *Combust. Flame* 144 (2006) 24–36.
- [8] T. Lu, C.K. Law, *Combust. Flame* 146 (2006) 472–483.
- [9] P. Pepiot, H. Pitsch, *Combust. Flame* 154 (2008) 67–81.
- [10] K.E. Niemeyer, C.J. Sung, M.P. Raju, *Combust. Flame* 157 (9) (2010) 1760–1770.
- [11] M. Valorani, F. Creta, D.A. Goussis, J.C. Lee, H.N. Najm, *Combust. Flame* 146 (1–2) (2006) 29–51.
- [12] T. Nagy, T. Turanyi, *Combust. Flame* 156 (2009) 417–428.
- [13] W. Sun, Z. Chen, X. Gou, Y. Ju, *Combust. Flame* 157 (2010) 1298–1307.
- [14] P. Pepiot, H. Pitsch, *Combust. Theory Model.* 12 (2008) 1089–1108.
- [15] S.S. Ahmed, F. Mauss, G. Moréac, T. Zeuch, *Phys. Chem. Chem. Phys.* 9 (2007) 1107–1126.
- [16] H. Huang, M. Fairweather, J.F. Griffiths, A.S. Tomlin, R.B. Brad, *Proc. Combust. Inst.* 30 (2005) 1309–1316.
- [17] M. Bodenstein, S.C. Lind, *Z. Phys. Chem.* 57 (1906) 168.
- [18] J.Y. Chen, *Combust. Sci. Technol.* 57 (1–3) (1988) 89–94.
- [19] J.C. Keck, D. Gillespie, *Combust. Flame* 17 (1971) 237.
- [20] U. Maas, S.B. Pope, *Combust. Flame* 88 (1992) 239–264.

- [21] V. Bykov, U. Maas, *Combust. Theory Model.* 11 (2007) 839–862.
- [22] Z. Ren, S.B. Pope, A. Vladimirov, J.M. Guckenheimer, *J. Chem. Phys.* 124 (2006) 114111.
- [23] S.B. Pope, *Combust. Theory Model.* 1 (1997) 41–63.
- [24] L. Lu, S.R. Lantz, Z. Ren, S.B. Pope, *J. Comput. Phys.* 228 (2009) 5490–5525.
- [25] S.R. Toner, N.W. Moriarty, N.J. Brown, M. Frenklach, *Israel J. Chem.* 39 (1999) 97–106.
- [26] D.A. Schwer, P. Lu, W.H. Green, *Combust. Flame* 133 (4) (2003) 451–465.
- [27] L. Liang, J.G. Stevens, J.T. Farrell, *Proc. Combust. Inst.* 32 (1) (2009) 527–534.
- [28] F. Contino, H. Jeanmart, T. Lucchini, G. D'Errico, *Proc. Combust. Inst.* 33 (2) (2011) 3057–3064.
- [29] L. Tosatto, B.A.V. Bennett, M.D. Smooke, *Combust. Flame* 158 (5) (2011) 820–835.
- [30] T. Løvas, S. Navarro-Martinez, S. Rigopoulos, *Proc. Combust. Inst.* 33 (1) (2011) 1339–1346.
- [31] B.J. Debusschere, Y.M. Marzouk, H.N. Najm, B. Rhoads, D.A. Goussis, M. Valorani, *Combust. Theory Model.* 16 (1) (2012) 173–198.
- [32] M. Singer, W. Green, *Appl. Numer. Math.* 59 (2009) 272–279.
- [33] I. Banerjee, M. Ierapetritou, *Combust. Flame* 144 (3) (2006) 619–633.
- [34] H. Yang, Z. Ren, T. Lu, G.M. Goldin, *Combust. Theory Model.* 17 (1) (2013) 167–183.
- [35] Z. Ren, C. Xu, T. Lu, M.A. Singer, *J. Comput. Phys.* 263 (2014) 19–36.
- [36] A. Viggiano, V. Magi, *Appl. Energy* 113 (2014) 848–863.
- [37] X. Gou, Z. Chen, W. Sun, Y. Ju, *Combust. Flame* 160 (2) (2013) 225–231.
- [38] Z. Ren, Y. Liu, T. Lu, L. Lu, O.O. Oluwole, G.M. Goldin, *Combust. Flame* 161 (1) (2014) 127–137.
- [39] K. He, I.P. Androulakis, M.G. Ierapetritou, *Energy Fuels* 25 (8) (2011) 3369–3376.
- [40] N.J. Curtis, K.E. Niemeyer, C.J. Sung, *Combust. Flame* 162 (4) (2015) 1358–1374.
- [41] H. Wang, S.B. Pope, *Proc. Combust. Inst.* 33 (1) (2011) 1319–1330.
- [42] Y. Yang, H. Wang, S.B. Pope, J.H. Chen, *Proc. Combust. Inst.* 34 (1) (2013) 1241–1249.
- [43] V. Hiremath, S.R. Lantz, H. Wang, S.B. Pope, *Combust. Flame* 159 (2012) 3096–3109.
- [44] K.E. Niemeyer, C.-J. Sung, *Combust. Flame* 158 (8) (2011) 1439–1443.
- [45] Y. Yang, S.B. Pope, J. Chen, *Combust. Flame* 160 (10) (2013) 1967–1980.
- [46] H. Curran, T. Jayaweera, W. Pitz, A detailed modeling study of propane oxidation, Western States Section of the Combustion Institute, 2004.
- [47] S.B. Pope, ELL_LIB: A Fortran Library for Ellipsoids. <http://tcg.mae.cornell.edu/ELL_LIB>.