Combustion Theory and Modelling, 2019
Vol. 23, No. 6, 1021–1053, https://doi.org/10.1080/13647830.2019.1606453

A combined PPAC-RCCE-ISAT methodology for efficient implementation of combustion chemistry

Ashish S. Newale*, Youwen Liang, Stephen B. Pope and Perrine Pepiot

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

(*Received 16 January 2019; accepted 2 April 2019*)

Probability density function (PDF) methods are now well established and can be used to accurately simulate flames with strong turbulence chemistry interactions. A pre-partitioned adaptive chemistry (PPAC) methodology (Liang et al., Combustion and Flame, 2015) has been proposed recently for the efficient implementation of combustion chemistry in particle PDF methods. PPAC generates a library of reduced kinetic models in an offline preprocessing stage. At runtime, PDF particles are dynamically assigned one reduced model and its corresponding reduced representation, leading to a significant decrease in both storage requirements and CPU time required for the integration of chemical source terms. In this work, we augment PPAC by combining it with storage retrieval and dimension reduction techniques. Specifically, this work combines PPAC with in-situ adaptive tabulation (ISAT) and rate-constrained chemical equilibrium (RCCE). Implementations for PPAC-ISAT, PPAC-RCCE, and PPAC-RCCE-ISAT are described and their performance is examined for a partially stirred reactor (PaSR) test case. The combined PPAC-RCCE-ISAT methodology shows a significant improvement over the stand-alone PPAC methodology by reducing the number of redundant direct integrations of similar compositions and reducing the number of variables that need to be retained at runtime.

Keywords: PPAC; RCCE; ISAT; adaptive chemistry; dimension reduction

Nomenclature

General

S	chemical source term
Δt	simulation time step
τ_{mix}	mixing timescale specification in PaSR
n_p	number of particles in PaSR
n_t	total number of time steps in a PaSR simulation
n_s	number of species in the detailed mechanism
$T^{(n)}$	temperature of particle n
$X_k^{(n)}$	value of quantity X for particle n on time step k before the mixing fractional step of a PaSR simulation
$X_{k+}^{(n)}$	value of quantity X for particle n on time step k after the reaction fractional step of a PaSR simulation

*Corresponding author. Email: asn79@cornell.edu

$X_k^{(n),D}$ value of quantity X for particle n on time step k of a PaSR simulation with the detailed mechanism
$X_k^{(n),A}$ value of quantity X for particle n on time step k of a PaSR simulation with adaptive chemistry
$n_{s,k}^{(n),A}$ number of species in the reduced model used for particle n on time step k
n_{rel} relative number of species
ε_X incurred error in quantity X
$\hat{\varepsilon}_X$ conservation error in quantity X

ISAT

ε error in the linear approximation to the reaction mapping
$A^{(n)}$ mapping gradient matrix for particle n
$f^{(l,n)}$ linear approximation to reaction mapping for particle n
$f^{(n)}$ reaction mapping for particle n
$x^{(n)}$ initial composition of particle n
x^q query composition

PPAC and combined techniques

$\mathbb{F}$ full composition space
$\mathbb{F}_J$ reduced composition space for model M_J
D database of compositions
M_D detailed kinetic model
M_J reduced kinetic model for region F_J
$\phi^{(n)}$ detailed representation of composition of particle n
$\phi^{(n),m}$ detailed representation of composition of particle n after mixing fractional step
$\phi^{(n),r}$ detailed representation of composition of particle n after reaction fractional step
$\phi_J^{(n)}$ reduced (skeletal) representation of composition of particle n based on M_J
$\phi_{J,r}^{(n)}$ reduced (RCCE) representation of composition of particle n based on M_J

RCCE

Φ vector of species specific moles and temperature
Φ^{DR} reaction mapping for a reconstructed composition
ε^r error in reaction mapping due to dimension reduction
$\{\phi_{\text{rep}}\}$ set of represented species
B constraint matrix for RCCE reduction
N number of compositions in database
r specific moles of represented species and elements
z species specific moles

1. Introduction

Turbulent combustion models can be broadly classified as topology-free and topology-based [1,2]. Topology-free models do not make any assumptions on the underlying flame structure, and are consequently more broadly applicable to a variety of reacting flows. However, this general applicability comes at the cost of added computational expense. Probability density function (PDF) methods [3,4] fall in the topology-free category and have been shown to accurately capture flames with strong turbulence-chemistry interactions in non-premixed [5], premixed [6], and partially premixed [7] modes of combustion.

PDF methods are known to be computationally more demanding than simpler topology-based approaches such as steady laminar flamelet models [8] in terms of CPU time and memory requirements [2,9]. The magnitude of difference between these turbulent combustion models for both CPU time and memory requirements is strongly dependent on the size of the chemical mechanism used in the computation.

Recent efforts towards reducing the computational expense of PDF methods have involved dimension reduction, storage retrieval, adaptive chemistry techniques, and their combinations. Dimension reduction techniques function by identifying low-dimensional manifolds in the high-dimensional composition space, thereby reducing the number of variables that need to be retained. A non-exhaustive list of dimension reduction techniques includes rate-constrained chemical equilibrium (RCCE) [10,11], computational singular perturbation (CSP) [12], intrinsic low-dimensional manifolds (ILDM) [13], and invariant constrained-equilibrium edge pre-image curve (ICE-PIC) [14]. On the other hand, storage retrieval techniques selectively store combustion chemistry computations. These stored values are used, where appropriate, to create accurate and efficient approximations for solutions to subsequent combustion chemistry computations. The suitable use of inexpensive approximations in lieu of directly performing combustion chemistry computations leads to a reduction in the overall cost of performing a simulation. Artificial neural networks (ANN) [15], in-situ adaptive tabulation (ISAT) [16], and piece-wise reusable implementation of solution mapping (PRISM) [17] are some examples of storage retrieval techniques explored in the context of reacting flows.

Adaptive chemistry approaches exploit the variety of compositions encountered in a turbulent combustion simulation by tailoring the chemistry representation for individual compositions. This leads to significant savings by avoiding the use of unnecessarily complicated chemical kinetic models for compositions that can be accurately described using only a few species. The corresponding reduced models can either be identified on-the-fly [18–20] or in an offline preprocessing stage [21–23]. The pre-partitioned adaptive chemistry (PPAC) approach proposed by Liang et al. [23] takes the latter approach: a set of reduced models is first automatically generated based on the expected behaviour of the LES/PDF simulation of interest. At runtime of the adaptive simulation, this set of reduced models is dynamically utilised to perform the chemistry integration. The use of reduced models instead of the detailed mechanism leads to a reduction in CPU cost. This approach also provides memory savings by reducing the number of variables that need to be retained at runtime. The reduction in memory requirement is a particularly important feature for particle PDF methods, which are known to have a significantly higher memory requirement compared to topology-based combustion models [2].

The aforementioned techniques are complementary to each other as they pertain to distinct aspects of chemistry modelling in a LES/PDF framework. Hence, they can be combined to provide added advantages compared to their standalone counterparts. This insight has recently led to a significant volume of work focused on combining a subset of these techniques. These efforts, classified according to the specific combination of techniques explored, are discussed below:

- **Coupled dimension reduction and storage retrieval**: Tang and Pope [24] developed a combined RCCE-ISAT method, which uses ISAT with compositions in their RCCE representation. Hiremath et al. [25] augmented RCCE-ISAT by combining it with the Greedy algorithm with local improvement (GALI) for automatic selection of near-optimal set of represented species for RCCE. Hiremath et al. [26] demonstrated the

use of RCCE-ISAT in a large scale LES/PDF computation of Sandia flame D [27]. The combined methodology was shown to reduce the wall clock time by 40% over standalone ISAT. Chatzopoulos and Rigopoulos [28] developed a combined RCCE-artificial neural network (ANN) approach. Compositions from laminar flamelet calculations in their RCCE representation were utilised to train an ANN. The methodology was tested for RANS/PDF simulations of DLR-A and DLR-B [29,30] and a speedup of the order of 100 was reported over a standalone RCCE computation. A subsequent work by Franke et al. [31] utilised this method in a LES/PDF simulation of Sydney flame L [32] with a similar speedup compared to standalone RCCE.

- **Coupled adaptive chemistry and storage retrieval**: Contino et al. [33] first proposed the use of dynamic adaptive chemistry (DAC) to perform the reaction integration required by ISAT. The proposed method was shown to outperform standalone ISAT and DAC for a simplified 2D homogeneous charge compression ignition simulations. Ren et al. [19] augmented DAC-ISAT by using a modified mapping gradient matrix and an improved version of ISAT. Recently, Xie et al. [34] proposed a dynamic adaptive acceleration method where either DAC or ISAT are utilised for chemistry integration based on a metric for composition inhomogeneity. The combined methodology showed a performance improvement of up to 50% compared to standalone ISAT for an internal combustion engine model simulations.

To the best of our knowledge, there is no existing work on a coupled dimension reduction and adaptive chemistry technique, or a coupled dimension reduction, storage-retrieval, and adaptive chemistry technique. Motivated by this review of the literature, the aim of the current work is to develop such an algorithm, which combines all three classes of techniques to attain their combined advantages in a single simulation tool. We systematically build toward this combination, by first developing a combined adaptive chemistry and storage retrieval technique, followed by a combined adaptive chemistry and dimension reduction technique. We use the PPAC methodology for adaptive chemistry, ISAT for storage-retrieval, and RCCE for dimension reduction. The performance of the combined methodologies is tested in a partially stirred reactor (PaSR) context. The subsequent sections of the paper include details of the PaSR implementation, a brief overview of ISAT, RCCE, PPAC, details of PPAC-ISAT, PPAC-RCCE, and the combined PPAC-RCCE-ISAT algorithms. An assessment of the performance of each combination is then provided in the same order.

2. Partially stirred reactor (PaSR)

A partially stirred reactor (PaSR) represents a particle PDF computation for an idealised statistically spatially homogeneous flow-field. PaSR is a computationally cheaper test case representative of PDF computations of turbulent flames [25]. Hence, the testing of various techniques described in subsequent sections is done using PaSR [16]. The PaSR implementation used in this work is described in detail in Liang et al. [23]. A summary of the PaSR features relevant to subsequent developments is provided below for reference.

The PaSR contains a constant, even number of particles, whose compositions evolve in time by increments of size Δt according to inflow/outflow events, and mixing and reaction fractional steps. In the inflow/outflow events, a certain number of particles are selected at random with equal probability, and their compositions are replaced by inflow stream compositions. The number of particles concerned and the inflow compositions are chosen

according to the user-specified mean residence time τ_{res} and streams mass flow rates, respectively. The mixing fractional step is performed using the pairwise mixing model [16], where particle compositions ($\phi^{(n)}$) evolve as follows (for odd values of n):

$$\begin{aligned}\frac{\mathrm{d}\phi^{(n)}}{\mathrm{d}t} &= -\frac{(\phi^{(n)} - \phi^{(n+1)})}{\tau_{\mathrm{mix}}},\\ \frac{\mathrm{d}\phi^{(n+1)}}{\mathrm{d}t} &= -\frac{(\phi^{(n+1)} - \phi^{(n)})}{\tau_{\mathrm{mix}}},\end{aligned} \tag{1}$$

where τ_{mix} is the user-defined mixing timescale. The particle composition during the reaction fractional step evolves as:

$$\frac{\mathrm{d}\phi^{(n)}}{\mathrm{d}t} = S(\phi^{(n)}), \tag{2}$$

where S represents the rate of change of $\phi^{(n)}$ due to chemical reactions. PaSRs initially go through a short transient before reaching a statistically stationary state.

3. Overview of individual techniques

This section reviews important details of ISAT, RCCE, and PPAC, which aid in understanding the subsequent development of combined techniques.

3.1. *In-situ adaptive tabulation (ISAT)*

ISAT [16,35] is a storage retrieval technique that provides efficient and accurate approximations of a function using previously tabulated function values. In the context of reacting flows, given an input composition (x), ISAT provides to within a user-specified error tolerance the reaction mapping (f) for that composition, that is, the solution, after Δt, of the chemical kinetic equations corresponding to the reaction fractional step (Equation (2)). ISAT uses a linear approximation to the reaction mapping ($f^{l,n}$), which is computed as follows:

$$f^{l,n}(x^q) = f^{(n)} + A^{(n)}(x^q - x^{(n)}), \quad \text{where } A_{ij}^{(n)} = \left(\frac{\partial f_i}{\partial x_j}\right)^{(n)}, \tag{3}$$

where x^q is the input, or query composition, and $f^{(n)}$ and $A^{(n)}$ are the stored reaction mapping and mapping gradient matrix for a previously resolved input composition $x^{(n)}$, respectively. The error in the linear approximation to the reaction mapping is defined as follows:

$$\varepsilon \equiv \left|\left|f^{l,n}(x) - f(x)\right|\right|_2. \tag{4}$$

Given an input composition, ISAT returns the linear approximation to its reaction mapping if the approximation error ε is within the user-specified error tolerance, with a high probability. The region of the composition space around $x^{(n)}$ where the approximation error is within the user-defined error tolerance is referred to as the region of accuracy (ROA). ISAT approximates the ROA using ellipsoids of accuracy (EOAs), which usually provide a conservative estimate to the ROA. At runtime, the EOAs can be modified or grown to more closely approximate the ROA as further information becomes available.

The following list briefly describes the sequence of operations ISAT follows to resolve a given input composition:

(1) **Retrieve**: If the input composition is covered by a stored EOA, the linear approximation to the reaction mapping is returned.
(2) **Grow**: If the composition is not part of any existing EOA, it is directly integrated and EOAs in close proximity of the input compositions are examined for grow attempts. If the linear approximation to the reaction mapping obtained using one of the EOAs in close proximity is within the user-specified error tolerance of the exact reaction mapping, the EOA is grown to include the new composition.
(3) **Add**: If all attempts at a grow fail and the table is not full, then a new entry corresponding to the input composition is added to the ISAT table.
(4) **Direct evaluation**: If the specified limit for the size of the table is reached and a retrieve or grow cannot be performed, the composition is integrated directly without saving the reaction mapping for subsequent use.

A more comprehensive description of ISAT can be found in Ref. [16]. The improved algorithm described by Lu and Pope [35] and implemented in ISAT5 is used in the current work.

3.2. *Rate-constrained chemical equilibrium (RCCE)*

RCCE [10,11] assumes that a non-equilibrium composition relaxes to chemical equilibrium through a series of states that lie on a low-dimensional constrained-equilibrium manifold. The constrained-equilibrium manifold provides a mapping between the specific moles of represented species and the elements in unrepresented species, and the specific moles of all the species. Hence at runtime, the variables that need to be retained include only the specific moles of represented species and the elements in unrepresented species. The number of represented species and elements are usually significantly fewer than the total number of species, which leads to dimension reduction.

To facilitate the coupling of RCCE and ISAT, the reaction integration in the current work is done in the full composition space. This corresponds to the RCCE-Trajectory in Full Space (TIFS) approach detailed in Hiremath and Pope [36]. This approach has been shown to be more accurate, efficient, and robust than other available approaches [36]. The choice of performing the reaction integration in the full space leads to the need for operators to switch from the full composition space to the RCCE space, and vice versa. The reconstruction from specific moles of the represented species and elements in the RCCE space to a representation in the full composition space, for a given pressure and enthalpy, is done using CEQ [37]. At any time t, the representation of a composition in the full space, $z(t)$, can be converted to its representation in the RCCE space, $r(t)$, as follows:

$$r(t) = B^T z(t), \tag{5}$$

where B is a constant matrix that can be determined from the elemental makeup of the unrepresented species [38]. As discussed previously, the representation of a composition in the RCCE space (r) is a vector of specific moles of represented species and elements in unrepresented species. The representation of a composition in the full space (z) is a vector of the specific moles of all species.

A critical step in using RCCE is the identification of the represented species. In the current work, the Greedy algorithm of Hiremath et al. [39] is used. The Greedy algorithm

sequentially adds to a set of represented species by making the locally optimal choice for a well-chosen database of N representative compositions. The local optimality is gauged by evaluating a reaction mapping error, ε^r, after a time step Δt:

$$\varepsilon^r(\Delta t, \{\phi_{\text{rep}}\}) = \frac{[\Phi^{DR(n)}(\Delta t, \{\phi_{\text{rep}}\}) - \Phi^{(n)}(\Delta t)]_{\text{rms}}}{[\Phi^{(n)}(\Delta t)]_{\text{rms}}}. \tag{6}$$

Here, Φ is defined to be a vector of the species specific moles and temperature: $\Phi \equiv \{z, T\}$. Specifically, $\Phi^{DR(n)}(\Delta t, \{\phi_{\text{rep}}\})$ is the reaction mapping for a reconstructed composition, for a chosen set of represented species ($\{\phi_{\text{rep}}\}$), obtained from its RCCE representation $r^{(n)}(0)$. $\Phi^{(n)}(\Delta t)$ is the reaction mapping for a full composition, $z^{(n)}(0)$, corresponding to the composition $r^{(n)}(0)$ in the RCCE space. $z^{(n)}(0)$ can be thought of as the exact reconstruction of $r^{(n)}(0)$. An example of the $[\cdot]_{\text{rms}}$ operator for $\Phi^{(n)}(\Delta t)$ is shown below:

$$[\Phi^{(n)}(\Delta t)]_{\text{rms}} = \sqrt{\frac{1}{N}\sum_{n=1}^{N} \|\Phi^{(n)}(\Delta t)\|_2^2}. \tag{7}$$

The reaction mapping error shown in (6) is computed using both species specific moles and temperature. We note that this is slightly different from the definition used in [39], where the reaction mapping error is computed using only species specific moles. This choice ensures a consistent error specification between PPAC and RCCE, and consequently facilitates their coupling. Further details on this coupling are provided in Section 4.2.

3.3. *Pre-partitioned adaptive chemistry methodology (PPAC)*

The last individual technique we detail that will subsequently be combined with the techniques discussed in the previous two sections is the pre-partitioned adaptive chemistry (PPAC) methodology [23]. As mentioned previously, the PPAC approach consists of an initial preprocessing stage, followed by a run time dynamic procedure. The offline preprocessing stage of PPAC entails the generation of a set of reduced kinetic models. Those reduced kinetic models are then selectively utilised at runtime to efficiently perform the reaction fractional step.

Figure 1 shows a schematic detailing both these stages. More specifically, the preprocessing stage consists of the following steps:

(1) **Database creation**: A database (D) of compositions is generated using the detailed kinetic mechanism (M_D). These compositions are assumed to be representative of those encountered at runtime.
(2) **Partitioning**: The full composition space $\mathbb{F}$ is partitioned into a user-specified number of regions ($\{\mathbb{F}_J\}$). The partition of the composition space is determined such that compositions belonging to each region are kinetically homogeneous. The kinetic homogeneity of a region is defined as the ability to generate a reduced mechanism of minimal size that can accurately describe the short time reaction mapping for any compositions belonging to it [23].
(3) **Reduced modelling**: Specialized reduced representations ($\{\phi_J\}$) and reduced kinetic models ($\{M_J\}$) are derived for each region identified in the previous step using the directed relation graph with error propagation (DRGEP, [40]) method according to user-specified targets.

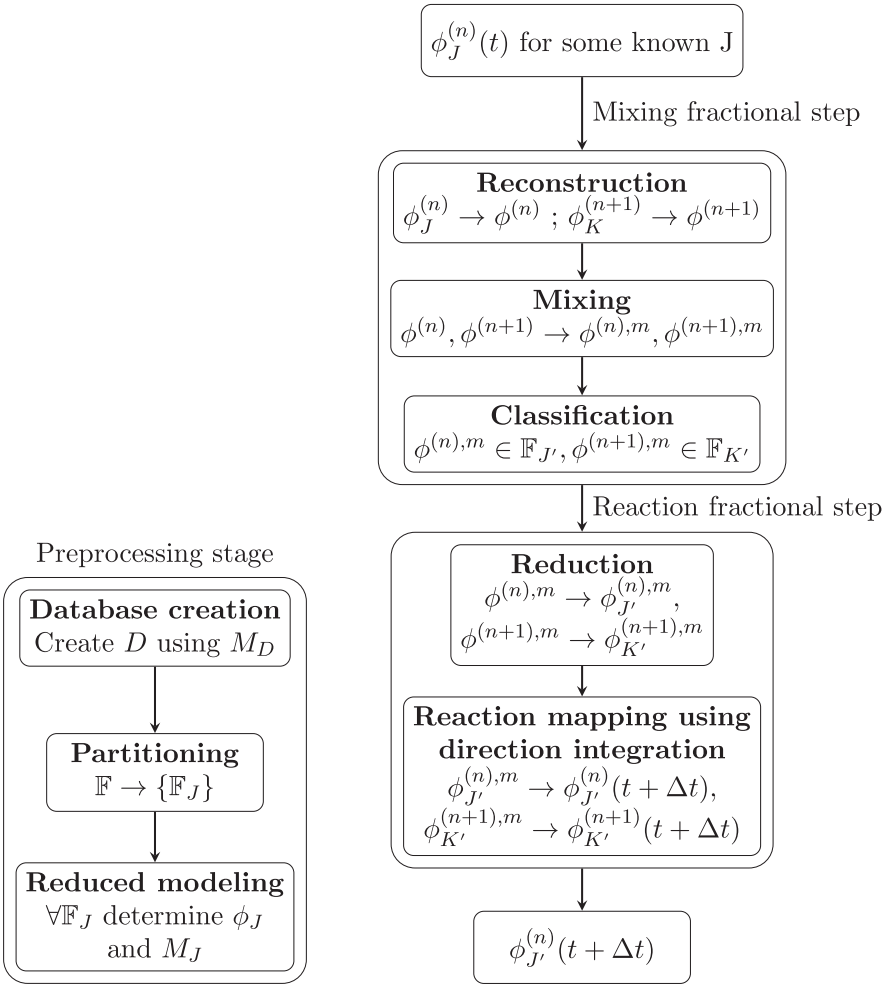

Figure 1. PPAC offline preprocessing stage and runtime adaptive algorithm, shown here for a PaSR simulation.

At the start of each time step in the online stage of PPAC, the particles carry a reduced skeletal representation ($\phi_J^{(n)}$) for a known region (J). The key steps through which a particle composition evolves in a time step are as follows:

(1) **Reconstruction**: At the start of the mixing fractional time step, the detailed representation of particle compositions ($\phi^{(n)}$) are reconstructed from their reduced skeletal representations.
(2) **Mixing**: The reconstructed composition are mixed (according to Equation (1)), yielding mixed compositions ($\phi^{(n),m}$) in the full composition space.
(3) **Classification**: The mixed compositions are classified into appropriate regions. After mixing, $\phi^{(n),m}$ has moved to region J'.
(4) **Reduction**: Particle compositions in the full composition space are then converted to the reduced skeletal representation corresponding to their respective regions ($\phi_{J'}^{(n),m}$).

(5) **Reaction mapping using direct integration**: The particles carrying the reduced skeletal representation are directly integrated using the reduced chemical kinetic models derived in the preprocessing stage corresponding to their respective regions.

The integrated values of particle composition, $\phi_{J'}^{(n),m}(t + \Delta t)$, are then used at the start of the next time step. For a comprehensive description of the PPAC methodology, we refer the reader to Liang et al. [23].

4. Development of combined techniques

We now develop coupled approaches based on the techniques discussed in the previous section. Specifically, we focus on coupled PPAC-ISAT, PPAC-RCCE, and PPAC-RCCE-ISAT. We start by investigating PPAC coupled with ISAT only, then PPAC coupled with RCCE only, to first assess the behaviour of each pair of techniques independently to one another. We then move on to the full PPAC-RCCE-ISAT combination.

4.1. *PPAC-ISAT*

The coupled PPAC-ISAT approach is examined first in this section, and a schematic for it is shown in Figure 2. The preprocessing stage of PPAC-ISAT being identical to that of PPAC, we do not include it in the schematic. A comparison between Figures 1 and 2 shows that the only difference in the online stage of PPAC and PPAC-ISAT is in the reaction fractional step: the reaction mapping is obtained using ISAT in PPAC-ISAT, whereas direct integration is systematically used in PPAC. The use of ISAT for resolving incoming compositions in their reduced skeletal representations implies that a distinct table is required for each individual region. The work required to retrieve from the table and storage needed per table entry scale as the square of the number of species. Additionally, the performance of ISAT has been shown to degrade beyond 50 species [25]. Hence, there is a significant benefit to building a distinct ISAT table for each reduced model with relatively few species, compared to a single table using a detailed mechanism with hundreds of species.

For brevity's sake, we do not repeat the steps in the online stage of PPAC-ISAT that are identical to that of PPAC, and focus instead on the reaction mapping using ISAT, which replaces item (5) in the PPAC procedure:

(5*) **Reaction mapping using ISAT**: The particles carrying the reduced skeletal representation are resolved by using the ISAT table corresponding to that specific region. If ISAT cannot resolve the input composition using a retrieve, this composition is integrated using the reduced model specific to that region as part of all other operations.

4.2. *PPAC-RCCE*

While PPAC-ISAT is expected to improve CPU time compared to standalone PPAC, PPAC-RCCE, developed in this section, is expected to improve upon memory requirements. Figure 3 shows a schematic for the PPAC-RCCE method. Compared to PPAC, an additional step in the preprocessing stage is required for PPAC-RCCE to determine the represented species for each region and its corresponding reduced model. The represented species for each reduced model are determined with the Greedy algorithm [39] using compositions in the database D. The detailed compositions in D are classified and converted to their region-specific reduced skeletal representations. The Greedy algorithm is run for

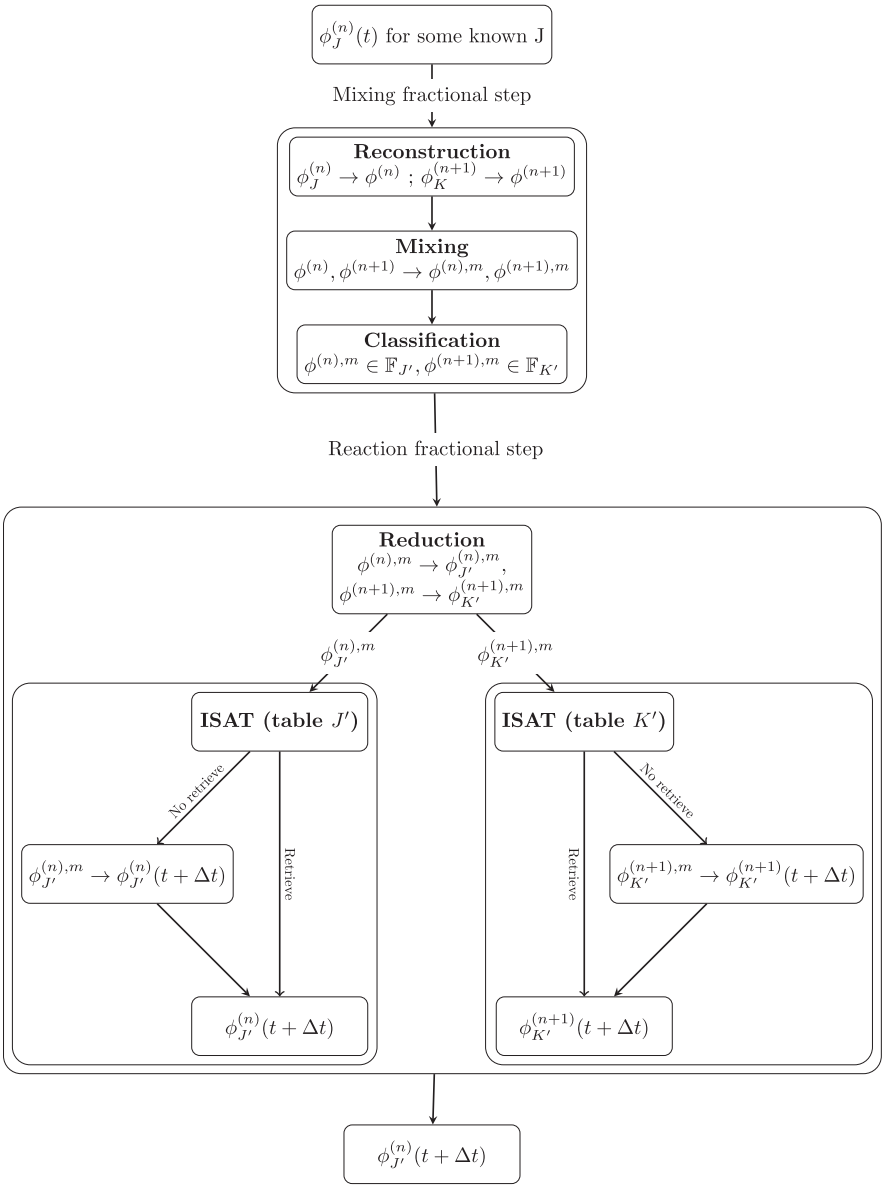

Figure 2. Coupled PPAC-ISAT methodology, online adaptive part only.

each region and its corresponding reduced model and compositions. The Greedy algorithm provides near-optimal sets of represented species with increasing cardinality that lead to non-increasing reaction mapping errors. The final choice of the represented species is made by the user based on the maximum reaction mapping error that is deemed acceptable. In the current work, this cutoff for determining the set of represented species is specified to be 10% of the DRGEP error cutoff. As noted previously (section 3.2), the reaction mapping error is computed using both specific moles of all species and temperature to ensure consistency with the DRGEP error specification.

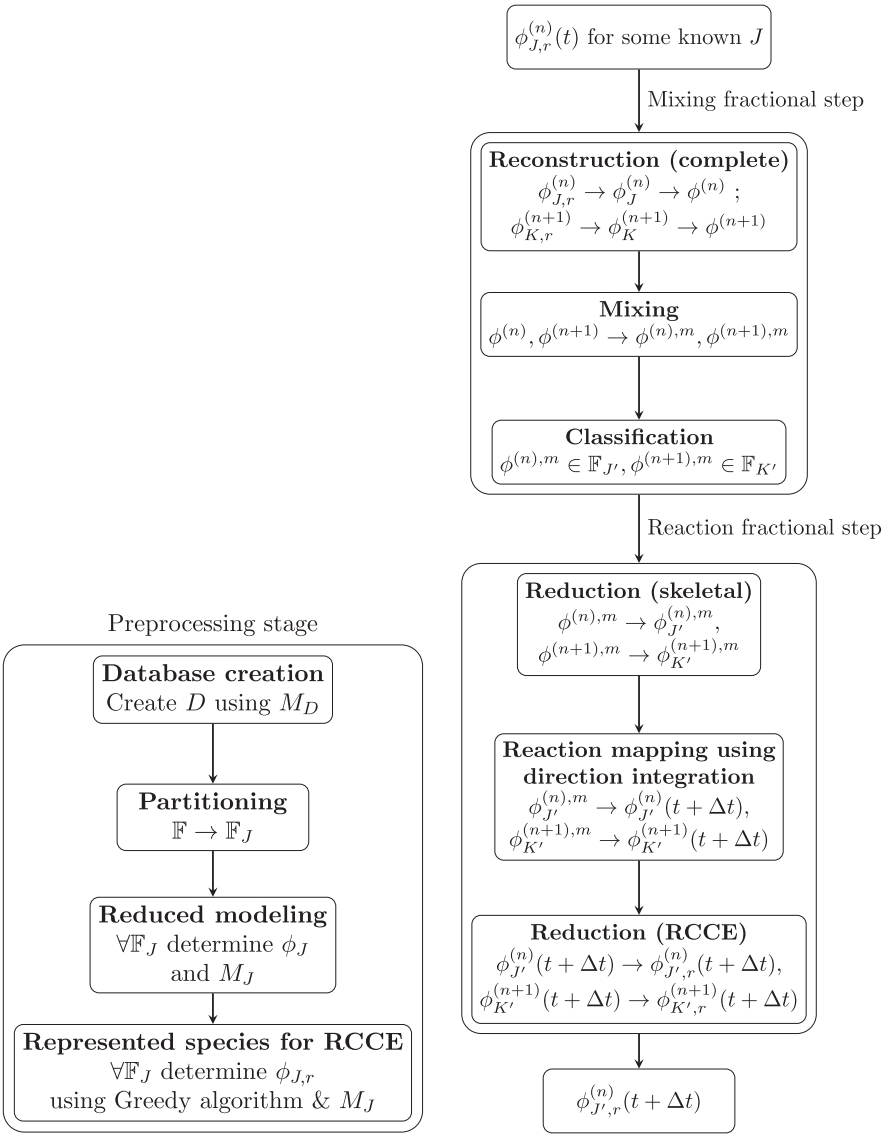

Figure 3. Offline preprocessing stage and online adaptive algorithm for the combined PPAC-R-CCE methodology.

At the start of each time step in the PPAC-RCCE simulation, the particles carry a reduced RCCE representation ($\boldsymbol{\phi}^{(n)}_{J,r}$) for their known region (J). The particle compositions evolve in time through the following key steps:

(1) **Reconstruction**: At the start of the mixing fractional time step, detailed representation of particle compositions are reconstructed from their reduced RCCE representations.
(2) **Mixing**: The reconstructed compositions are mixed in the full composition space.
(3) **Classification**: The mixed compositions are classified into appropriate regions.

(4) **Reduction (skeletal)**: The particle compositions in the full composition space are converted to the reduced skeletal representation corresponding to their respective regions.
(5) **Reaction mapping using direct integration**: The particles carrying the reduced skeletal representation are then integrated using the reduced chemical kinetic models corresponding to their respective regions.
(6) **Reduction (RCCE)**: The integrated compositions in reduced skeletal representation are finally converted to their reduced RCCE representation.

The particles, now carrying the region-specific, reduced RCCE representation $\phi_{J,r}^{(n)}(t + \Delta t)$, are ready to be used at the start of the next time step.

4.3. *PPAC-RCCE-ISAT*

Building on the algorithms now available for both PPAC-ISAT and PPAC-RCCE, we provide next a detailed description of the fully coupled PPAC-RCCE-ISAT methodology. The algorithm must satisfy a number of practical constraints, in particular: the particle compositions at the beginning and end of the time step must carry region-specific RCCE representations, mixing has to be done using the full representation, and ISAT must perform direct integrations (whenever necessary) using the region-specific reduced kinetic mechanisms. Due to these constraints, the order in which mixing and reaction proceed now does matter. We consider two different implementations, differentiated by the order in which the mixing and reaction fractional steps are performed: (MR) when mixing is done prior to reaction, and (RM) when the reaction fractional step is done first.

A visualisation of the PPAC-RCCE-ISAT (MR) algorithm is available in Figure 4, while the (RM) variant, with reaction fractional step coming first, is shown in Figure 5. The pre-processing stage for both options is identical to that of the PPAC-RCCE methodology, and is therefore omitted here. Figure 4 shows that for PPAC-RCCE-ISAT (MR), the conversion from the reduced RCCE space to the reduced skeletal composition space is performed prior to every mixing fractional step and during the reaction fractional step, whenever ISAT does not resolve the input composition as a retrieve. In contrast, PPAC-RCCE-ISAT (RM) only involves reconstructing the composition in the reduced skeletal space during the reaction fractional step if the input composition is not resolved as a retrieve. Reconstruction from the reduced RCCE space to the reduced skeletal space is a much more expensive operation than a reconstruction from the reduced skeletal space to the full composition space. Hence, we deem PPAC-RCCE-ISAT (RM) to be a superior option as it obviates the need for a reconstruction from the reduced RCCE space to reduced skeletal space in the mixing fractional step, and only consider the (RM) option in the remainder of this paper.

At the beginning of the time step, the particles carry the reduced RCCE representation of their composition ($\phi_{J,r}^{(n)}(t)$) corresponding to the region they belong to (J). The particle compositions then evolve in time through the following key steps:

(1) **Reaction mapping using ISAT**: The particle compositions are advanced in time according to their region-specific reduced kinetic model. Specifically, the table corresponding to the region the particle belongs to is used. If the input composition cannot be retrieved from the table, direct integration is required, which proceeds through the following sub-steps:
   (a) The input composition is reconstructed to its reduced skeletal representation

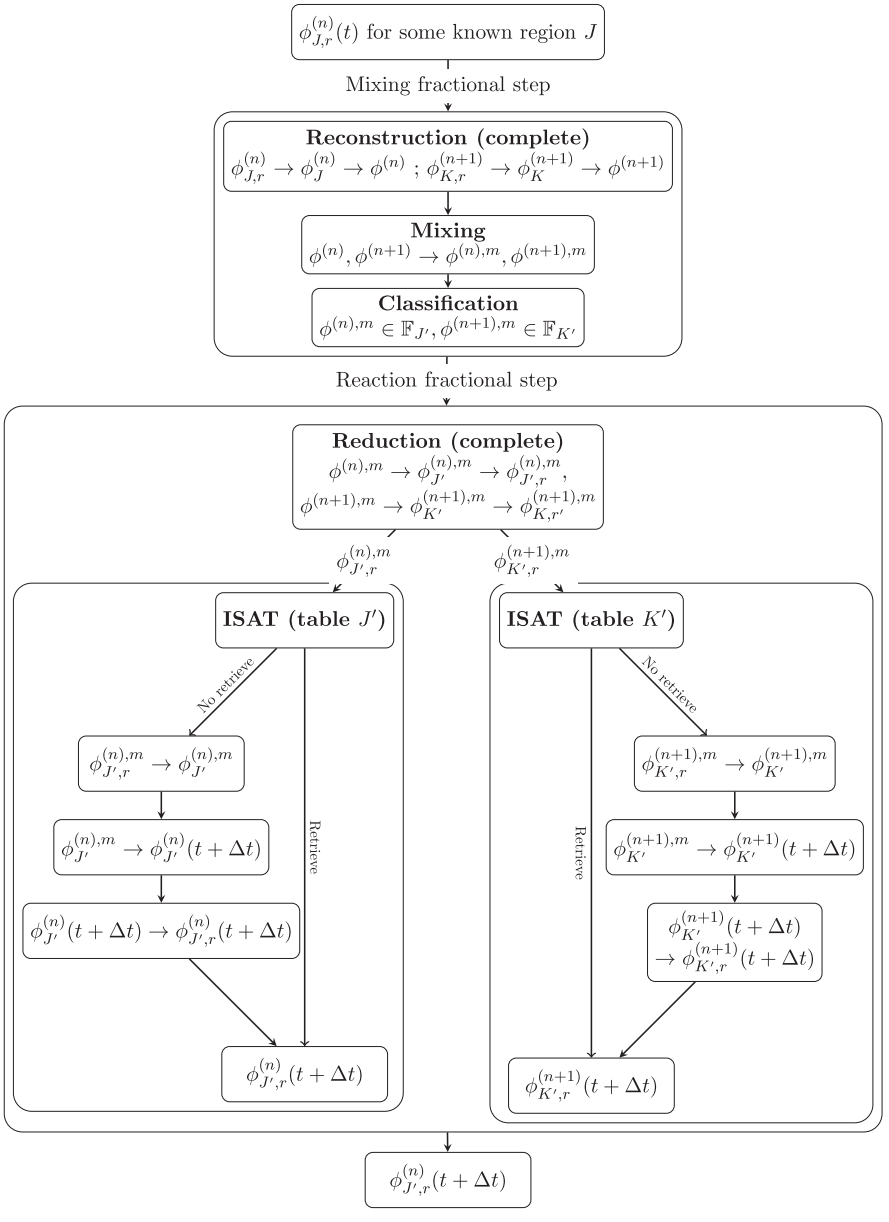

Figure 4. Adaptive algorithm for the PPAC-RCCE-ISAT (MR) methodology.

(b) The reconstructed composition is then integrated using the reduced model corresponding to the region, resulting in the composition after reaction, $\boldsymbol{\phi}_J^{(n),r}$, still in its reduced skeletal representation.

(2) **Reconstruction (skeletal)**: Prior to the mixing fractional time step, detailed representation of particle compositions are reconstructed from their reduced skeletal representations.

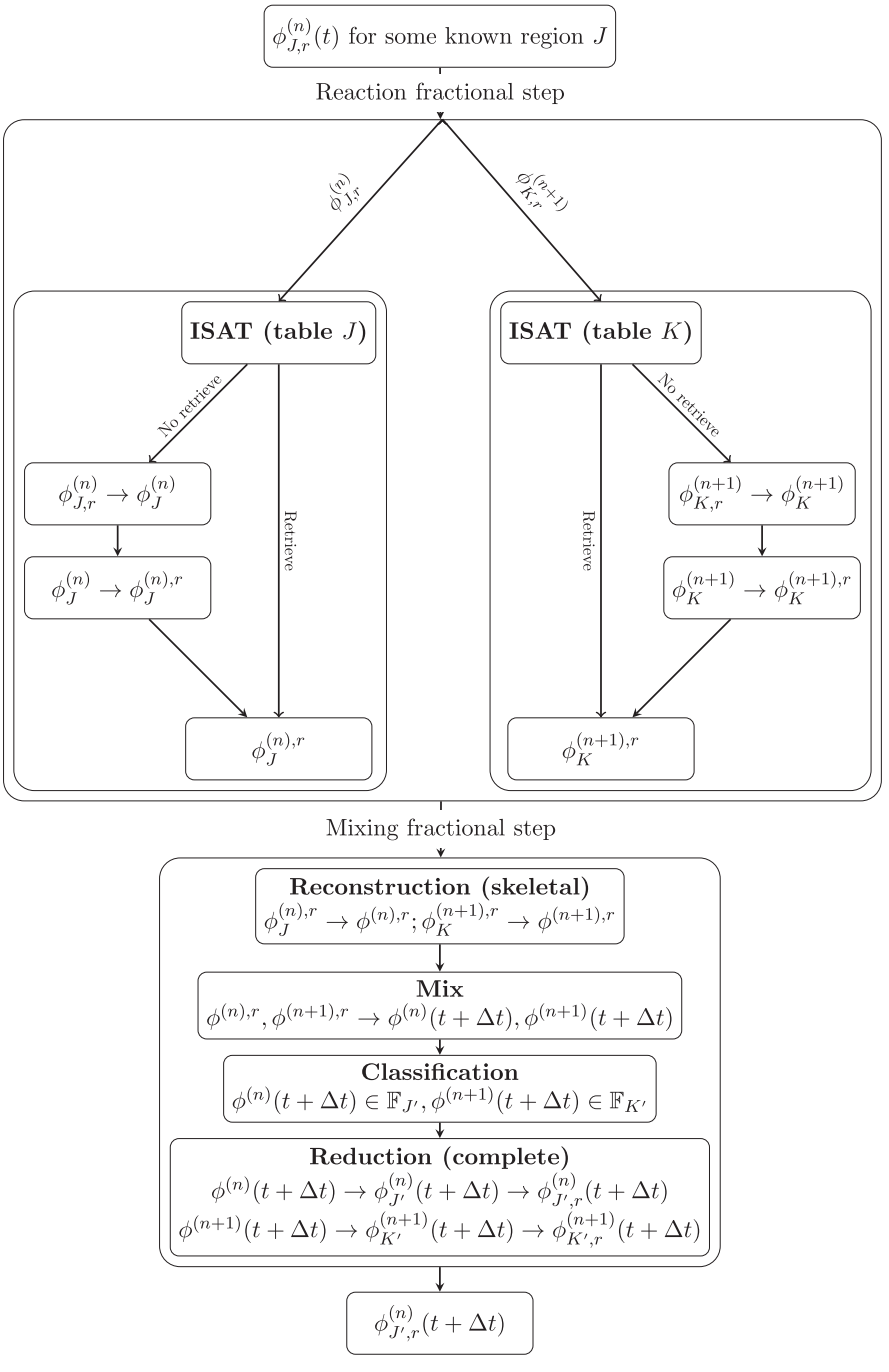

Figure 5. Adaptive algorithm for the PPAC-RCCE-ISAT (RM) methodology.

(3) **Mixing**: The reconstructed compositions are mixed, leading to mixed compositions in the full composition space.
(4) **Classification**: The mixed compositions are classified into appropriate regions.

(5) **Reduction (complete)**: The particle compositions in the full composition space are then converted to the reduced RCCE representation corresponding to their respective regions.

The compositions at the end of the mixing fractional step are now ready to be used for the next time step.

5. Results and discussion

5.1. *PaSR configuration*

A PaSR is used for testing the techniques and their combinations described in previous sections. The key parameters for the PaSR configuration used in the current work are listed in Table 1.

We use a non-premixed piloted PaSR as our test case and the composition of the three streams in terms of mass fractions is specified in Table 2. A detailed propane mechanism consisting of 115 species and 1308 reactions developed by Curran et al. [41] is used in this work. The current PaSR configuration is identical to the one used by Liang et al. [23] barring the number of particles: a higher number of particles is used here to reach

Table 1. PaSR parameters.

Parameter	Value
Pressure	1 bar
Number of particles	10,000
Time step	0.1 ms
Mixing time scale	1 ms
Pairing time scale	1 ms
Residence time	10 ms

Table 2. Stream specification for PaSR.

	Stream mass fraction		
Species	Pilot	Fuel	Oxidizer
C_3H_8	0.0	1.0	0.0
N_2	7.2085×10^{-1}	0.0	0.767
O_2	7.6607×10^{-3}	0.0	0.233
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

a larger number of queries for testing the performance of PPAC-ISAT and PPAC-RCCE-ISAT. The current configuration for PaSR is run for a physical time of 0.425 s, leading to a total of 4.25×10^7 input compositions, enough to gather proper statistics and assess ISAT performance for both PPAC-ISAT and PPAC-RCCE-ISAT methodologies.

5.2. *PPAC*

This section details the results for PPAC in the PaSR configuration described above. The results for standalone PPAC serve as a reference to assess the performance of the combined techniques. The initial composition database required for the preprocessing stage is generated using the same PaSR configuration, with a different random seed [23]. This ensures a justifiable test by avoiding identical particle evolution in time for the preprocessing and runtime stages. The number of regions used is 10. A larger number of regions is not used, as they lead to only minor improvements for the current configuration [23]. Additionally, this choice simplifies the presentation of subsequent analysis of the coupled techniques. The heat release rate, C_3H_8, O_2, OH, CO, and CO_2 are used as targets for reduction.

5.2.1. *Computational accuracy and memory requirements*

A metric to quantify the gain in terms of memory requirement is the relative number of species used in the simulation, defined as:

$$n_{\text{rel}} \equiv \frac{1}{n_t n_p n_s} \sum_{k=1}^{n_t} \sum_{n=1}^{n_p} n_{s,k}^{(n),A}, \tag{8}$$

where $n_{s,k}^{(n),A}$ is the number of species in the reduced model used for particle n on time step k, n_t is the total number of time steps in the simulation, n_p is the number of particles in PaSR, and n_s is number of species in the detailed mechanism. The incurred error for a quantity X is defined as follows:

$$\varepsilon_X \equiv \frac{\sum_{k=1}^{n_t} \sum_{n=1}^{n_p} \left| X_k^{(n),A} - X_k^{(n),D} \right|}{\sum_{k=1}^{n_t} \sum_{n=1}^{n_p} \left| X_k^{(n),D} \right|}, \tag{9}$$

where $X_k^{(n),A}$ and $X_k^{(n),D}$ are the values of the quantity X for particle n, at time step k, in the PaSR simulation using adaptive and detailed chemistry respectively. The PaSR test is performed several times, each with a different value of the specified DRGEP error tolerance, leading to different incurred errors, as defined by Equation (9).

Figure 6 shows the relative number of species as a function of the incurred error in temperature for PPAC. The performance of the degenerate, that is, non-adaptive, case is also shown in Figure 6. The degenerate case refers to the use of PPAC with a single region and is similar to the use of a conventional reduced mechanism obtained using DRGEP in a non-adaptive simulation. We observe that for approximately the same incurred error, PPAC leads to a much smaller relative number of species used in the chemical kinetics compared to the degenerate case. The reduction in the relative number of species shows the benefit of using specialised reduced models for different regions of the composition space.

The subsequent testing of PPAC-ISAT, PPAC-RCCE, and PPAC-RCCE-ISAT is not performed on the entire range of adaptive PaSR simulations used to gauge the performance of PPAC. We instead select a subset of four specified reduction error tolerances showing

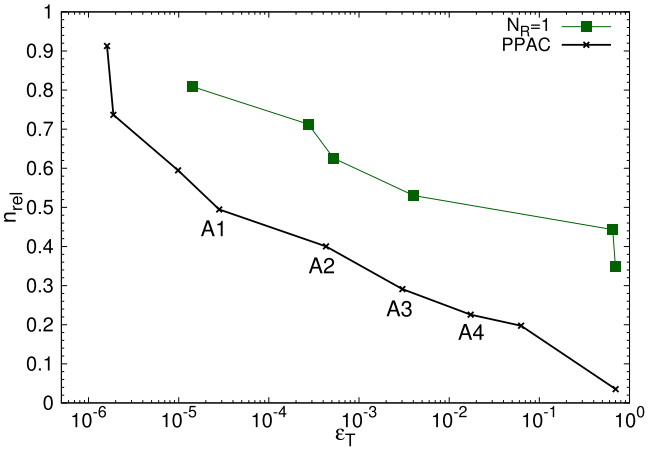

Figure 6. Relative number of species as a function of incurred error in temperature for the degenerate (non-adaptive) case ($N_R = 1$) and for PPAC($N_R = 10$).

the right balance between the relative number of species and the incurred error in temperature. We note that each of the four specified reduction error tolerances corresponds to a set of reduced kinetic models and this terminology is used interchangeably. For the purposes of the upcoming discussion, these cases are labelled as $A1$, $A2$, $A3$ and $A4$ as shown in Figure 6. The labels are chosen simply by indexing the adaptive PaSR runs used for testing PPAC in decreasing order of the relative number of species.

5.2.2. *Conservation errors*

PPAC entails a conversion from the full composition space to the reduced skeletal space which does not conserve elements or enthalpy. Hence, element and enthalpy conservation errors are an important check on the performance of PPAC. The conservation error for a quantity X is defined as follows:

$$\hat{\varepsilon}_X \equiv \frac{1}{n_p}\frac{1}{n_t}\sum_{n=1}^{n_t}\left|\frac{\sum_{n=1}^{n_p}\left(X_{k+}^{(n)} - X_k^{(n)}\right)}{\sum_{n=1}^{n_p} X_k^{(n)}}\right|, \tag{10}$$

where $X_k^{(n)}$ and $X_{k+}^{(n)}$ are the values of quantity X for particle n, at time step k, before and after the mixing and reaction fractional steps respectively. Figure 7 shows the conservation errors as a function of the incurred error in temperature. The conservation errors generally increase with increasing incurred error in temperature, and with decreasing relative number of species. Overall, the conservation errors remain bounded below 2×10^{-5} and are deemed acceptable.

5.3. *PPAC-ISAT*

We now examine the performance of PPAC-ISAT for the PaSR configuration discussed in 5.1, for the set of reduced models A1–A4. The performance of PPAC-ISAT is explored for ISAT error tolerance specifications of $10^{-4}, 5 \times 10^{-5}$, and 10^{-5}. We reiterate that ISAT is used to resolve incoming compositions in their reduced skeletal representations. Consequently, at runtime, multiple ISAT tables are simultaneously built corresponding to the

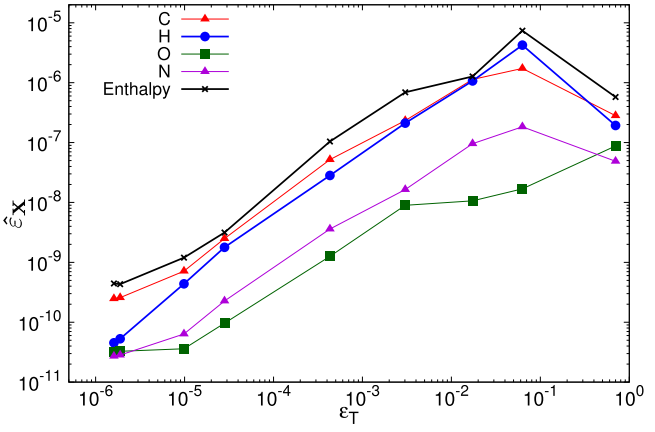

Figure 7. Element and enthalpy conservation errors as a function of incurred error in temperature for PPAC.

user-specified number of regions. The maximum size for each ISAT table is specified to be 500 MB.

5.3.1. *Computational accuracy and memory requirements*

Figure 8 shows the relative number of species as a function of the incurred error in temperature for PPAC-ISAT with different ISAT error tolerance specifications. The curves for PPAC and the degenerate case are included as reference.The curves for PPAC-ISAT connect four adaptive PaSR runs using the sets of reduced models A1–A4 and a common ISAT error tolerance specification. We recall that ISAT is a storage retrieval technique designed for reducing CPU time, and does not lead to dimension reduction. Hence, the relative number of species for PPAC-ISAT runs are essentially identical to PPAC runs.

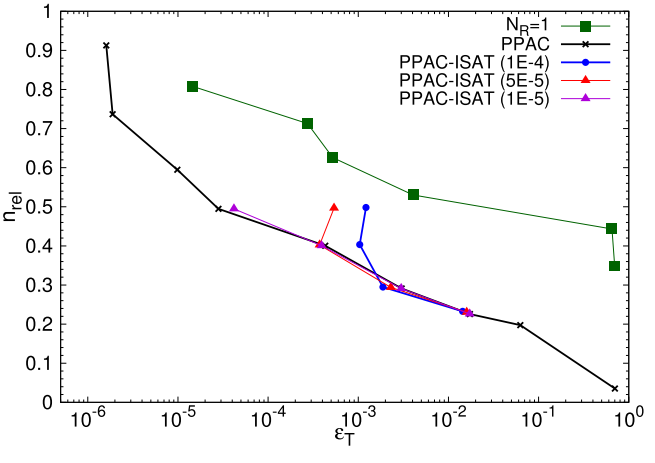

Figure 8. Relative number of species as a function of incurred error in temperature for the degenerate case, PPAC, and PPAC-ISAT, showing that the choice of error tolerance in ISAT, if not tight enough, can affect the performance of the PPAC-ISAT methodology.

PPAC-ISAT incurs errors due to reduction and tabulation. Here, reduction error refers to the errors incurred due to the use of reduced chemical kinetic models at runtime. Consequently, PPAC and PPAC-ISAT runs that use the same set of reduced models incur approximately the same reduction error. Hence, the relative contribution of the reduction and tabulation errors can be inferred by comparing the overall incurred error for adaptive PaSR runs for PPAC and PPAC-ISAT using the same set of reduced models. Cases *A*3 and *A*4, associated with larger reduction ratios, and therefore larger incurred error, are not affected by the introduction of ISAT: we infer that for those cases, the reduction error dominates the tabulation error for all ISAT error tolerance specifications. However, for the lower reduction error tolerance cases *A*1 and *A*2, there exists a cutoff for the ISAT error tolerance above which the PPAC-ISAT errors are dominated by the tabulation error. From Figure 8, we can conclude that the ISAT error tolerance should be set to a value at least an order of magnitude smaller than the desired incurred error.

5.3.2. *Conservation errors*

PPAC-ISAT, in a similar fashion as PPAC, incurs errors in element and enthalpy conservation due to the conversion operation from the full space to the skeletal space. Figure 9 shows the conservation errors for the three different ISAT error tolerance specifications as a function of the incurred error in temperature. The conservation errors similar to PPAC remain bounded below 2×10^{-5} and are deemed to be satisfactory.

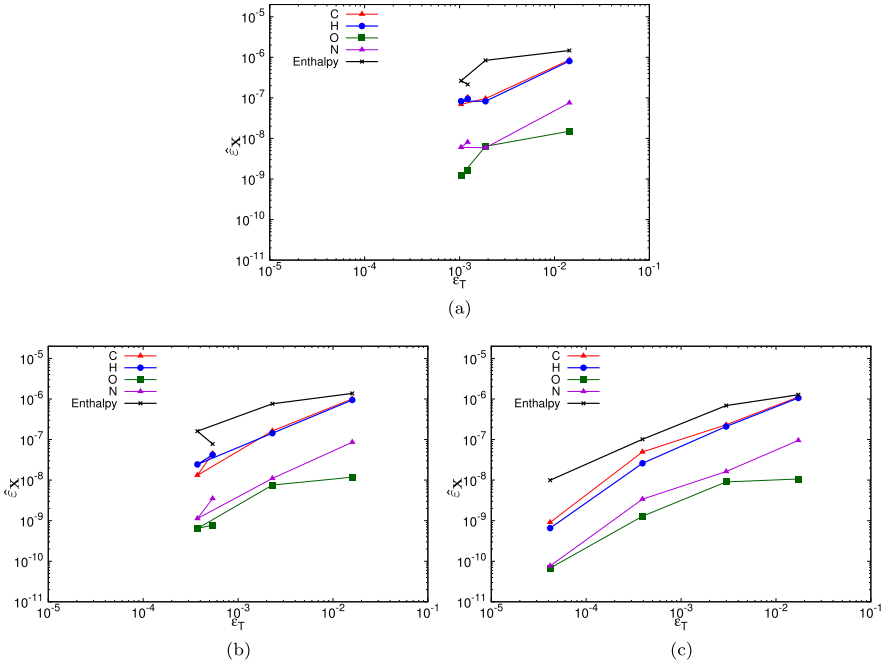

Figure 9. Element and enthalpy conservation errors as a function of incurred error in temperature for PPAC-ISAT runs with three different ISAT error tolerance specification. (a) ISAT error tolerance: 1×10^{-4}, (b) ISAT error tolerance: 5×10^{-5} and (c) ISAT error tolerance: 1×10^{-5}.

5.3.3. *Computational efficiency*

An important metric to gauge the performance of ISAT is the fraction of queries that are resolved by ISAT as a retrieve. For these queries, ISAT returns the computationally efficient linear approximation to the reaction mapping as opposed to directly solving an expensive set of stiff ODEs. Table 3 shows the percent of total queries resolved as a retrieve for the various PPAC-ISAT runs. Overall, we observe that a large ($\approx 70-90\%$) fraction of the queries are resolved as a retrieve, with an expected trend of decreasing fraction of retrieves with increasingly stringent error tolerance specification for ISAT.

The total CPU time of an adaptive simulation divided by the total CPU time of a non-adaptive simulation with the same configuration using a detailed mechanism is defined as the relative time [23]. Table 4 reports the relative time for both PPAC and PPAC-ISAT runs. We note that the relative time for PPAC-ISAT is closely related to the retrieve percentage reported in Table 3. A higher retrieve percentage leads to a lower relative time, as it involves a more frequent use of the computationally efficient linear approximation. For the same set of reduced models, we observe an increasing relative time as the ISAT error tolerance specification is decreased due to the decreasing retrieve percentage. For the set of reduced models *A*3, with an ISAT error tolerance specification of 10^{-4}, we observe that for an incurred error in temperature of 1.1×10^{-3}, about 94% of the queries are resolved as a retrieve leading to a relative time of 0.075.

5.3.4. *Performance analysis*

This section details a more in-depth investigation into the performance of PPAC-ISAT. We recall that PPAC-ISAT resolves queries using multiple tables, each one corresponding to one of the regions in the partitioned composition space. On the other hand, standalone ISAT resolves queries with a single table using the detailed mechanism. The performance of PPAC-ISAT is expected to improve upon the performance of standalone ISAT due to the following:

Table 3. Total retrieve percentage for various PPAC-ISAT runs.

	ISAT error tolerance specification		
Run number	10^{-4}	5×10^{-5}	10^{-5}
A1	91.628	85.608	73.103
A2	92.707	86.914	74.572
A3	94.392	89.205	74.859
A4	94.662	90.145	75.059

Table 4. Relative time for various PPAC and PPAC-ISAT runs.

		ISAT error tolerance specification		
Run number	PPAC	10^{-4}	5×10^{-5}	10^{-5}
A1	0.399	0.154	0.219	0.303
A2	0.299	0.119	0.170	0.232
A3	0.189	0.075	0.109	0.159
A4	0.116	0.048	0.068	0.105

- Partitioning, which leads to individual ISAT tables resolving queries that are kinetically homogeneous, and therefore, more likely to be retrieved.
- The use of reduced kinetic models, which reduces the dimension of the composition space where lookup is performed and the CPU time required for reaction integration.

To quantitatively compare the performance of ISAT and PPAC-ISAT and specifically determine the advantages of partitioning and the use of reduced models, we examine the three variants shown in Figure 10 and described in detail below.

(a) A single ISAT table is used with the detailed mechanism to resolve all queries. This variant corresponds to standalone ISAT.
(b) Multiple ISAT tables corresponding to the user-specified number of regions are used with the detailed mechanism to resolve the queries.
(c) Multiple ISAT tables corresponding to the user-specified number of regions are used with the corresponding reduced mechanisms to resolve all queries. This variant corresponds to PPAC-ISAT.

To test the aforementioned three variants, a database of detailed compositions is used. The database is obtained from a non-adaptive PaSR run using the same configuration as before, with the detailed mechanism. An ISAT error tolerance specification of 10^{-4} is used for all variants. The maximum table size is specified to be 5 GB for variant (a) and for each table in variant (b). The variant (c) is tested using a maximum size specification of 500 MB for each table. The detailed compositions are classified and converted to their reduced skeletal representation for variant (c) before calling ISAT. Variant (c) is tested using the sets of reduced models $A2$ and $A3$.

The metrics used to characterise the performance of ISAT are the retrieve percentage, relative build time, and relative query time. The build time is defined to be the total time spent on queries resolved as add or grows [25]. The query time is defined to be the average time spent on resolving a query once the table is built. The build time for individual tables is estimated using the method detailed in Hiremath et al. [25]. For variants (b) and (c), which maintain multiple ISAT tables, we use the mean build time and a weighted mean for the query time. The weighing factor utilised is simply the number of queries used for determining the query time of individual reduced models. The build and query times are reported relative to the standalone ISAT variant.

Table 5 shows the relative build time, relative query time, and retrieve percentage for the three variants. A significant reduction in the relative build time is observed for variant (b) compared to variant (a). This reduction is attributed to the partitioning, which leads to individual ISAT tables resolving queries more likely to be similar from the kinetics point of view. The resolution of kinetically homogeneous queries implies that a smaller subset of the accessed region is tabulated, reducing the number of queries required for building the table and consequently reducing the build time. An additional reduction in the build time is observed for variant (c) compared to (b), explained by computationally cheaper adds and grows due to the use of reduced kinetic models in variant (c). Additionally, the use of reduced models causes a reduction in the relative query time, since the table lookup is performed in a smaller number of dimensions.

To explore the performance of variant (c) in more depth, we show the relative build and query time of individual reduced models in Figure 11. The fraction of species for the region-specific reduced model (skeletal n_{rel}), i.e. the number of species in a reduced model divided by the number of species in the detailed mechanism, is also shown in Figure 11.

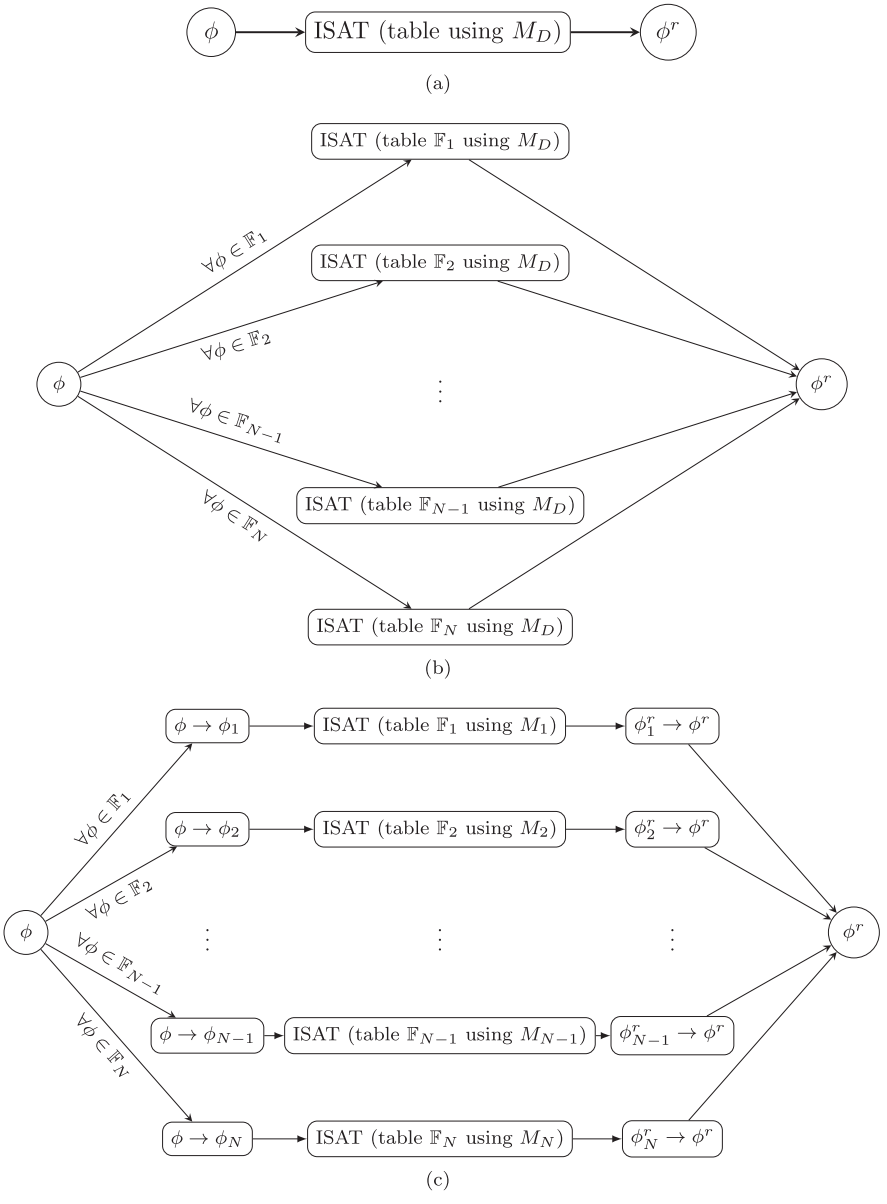

Figure 10. Three variants of query resolution with ISAT used to examine the benefits of partitioning and the use of reduced models for reaction integration.

The final bar in Figure 11 shows the fraction of queries classified as belonging to a specific region (FQCR). The partitioning of the composition space is the same for both *A*2 and *A*3. Hence, the fraction of classified compositions is identical for *A*2 and *A*3 and is reported only once. We observe that more than 50% of the queries are resolved by models M1–M3, which contain between 20% and 27%, and 9% and 15% of the species in the detailed mechanism for *A*2 and *A*3, respectively. The regions corresponding to these

Table 5. ISAT performance metrics for the three PPAC-ISAT variants under investigation.

Variant	Relative build time	Relative query time	Retrieve percentage
(a)	1	1	93.2
(b)	0.181	0.84	96.5
(c) ($\sim A2$)	0.029	0.59	93.0
(c) ($\sim A3$)	0.027	0.36	94.5

reduced models have relative query times of about 10^{-3}. The resolution of more than half the queries at a fraction of the query time of standalone ISAT shows the advantages of using PPAC-ISAT. The compositions belonging to these regions are essentially non-reactive and consequently can be described effectively by reduced models with such a small number of species. The fraction of non-reactive compositions in a turbulent combustion computation is expected to be significantly higher than that of a PaSR computation. For instance, a significant fraction of the domain used for the LES/PDF computation of Sandia flame D reported by Hiremath et al. [9] is filled with particles having the unmixed coflow composition. Due to the expected increase in the fraction of non-reactive compositions, we expect further improvement in the performance of PPAC-ISAT over standalone ISAT in a turbulent combustion simulation.

5.4. *PPAC-RCCE*

This next section explores the performance of RCCE coupled with PPAC. The methodology is the one described in Section 4.2, and as above, we utilise the set of reduced models $A1 - A4$ and the PaSR configuration detailed in 5.1.

5.4.1. *Computational accuracy and memory requirements*

A key differentiator between PPAC-RCCE and PPAC is that particles in PPAC-RCCE carry reduced RCCE representations, whereas they carry reduced skeletal representations in PPAC. Consequently, we expect PPAC-RCCE to lead to a significant reduction in the average number of species carried by the particles. This hypothesis can be tested by looking at the relative number of species n_{rel}, defined in Equation (8), as a function of the incurred error in temperature, as shown in Figure 12. A dimension reduction of $\approx 30-40\%$ over standalone PPAC is indeed observed for PPAC-RCCE runs in the current study. Despite this dimension reduction, the incurred error in temperature is approximately the same as that for the corresponding PPAC runs. We recall that the represented species for each reduced model are chosen such that the error in reaction mapping due to dimension reduction is at most 10% of the DRGEP error cutoff. Additionally, PPAC-RCCE uses the same sets of reduced models as PPAC, and PPAC incurs only reduction errors. Hence, comparing PPAC and PPAC-RCCE, we can conclude that the selection criteria discussed above ensures that the incurred error is dictated by the reduction error.

5.4.2. *Conservation errors*

PPAC-RCCE does not conserve elements or enthalpy due to the conversion from the full space to the reduced skeletal space. Figure 13 shows the element and enthalpy conservation

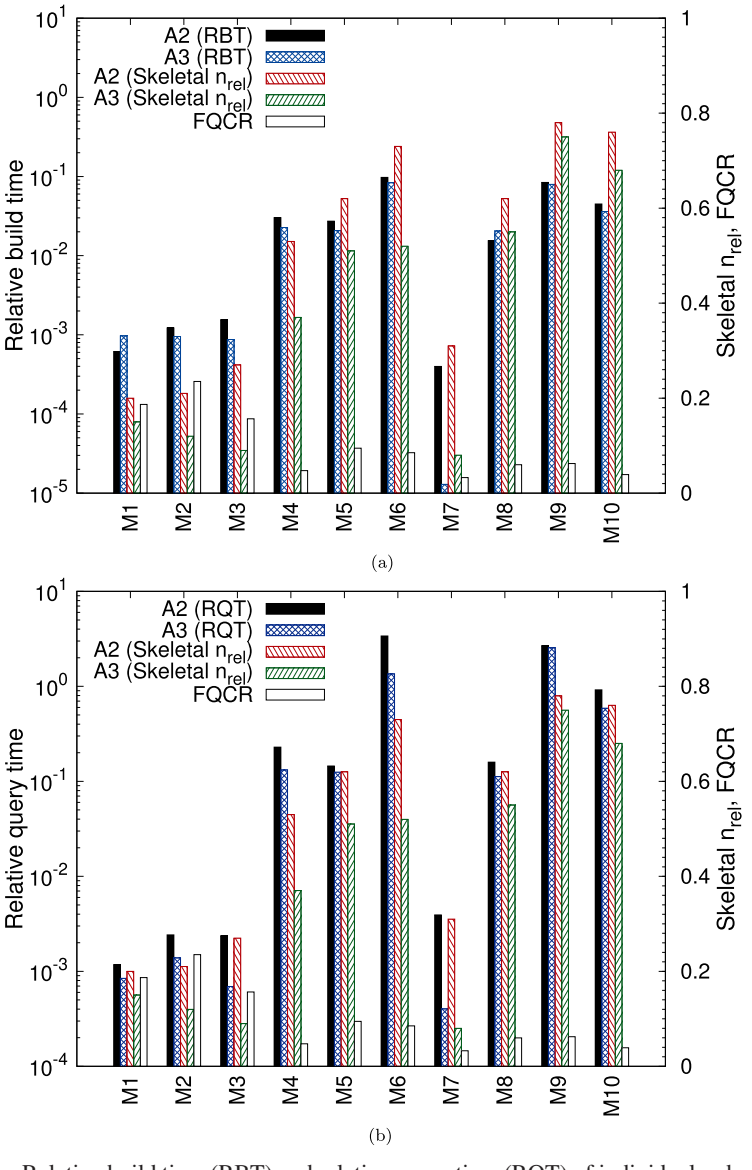

Figure 11. Relative build time (RBT) and relative query time (RQT) of individual reduced models for variant (c) using *A*2 and *A*3. The fraction of species for each region-specific reduced model (Skeletal n_{rel}) are also reported for reference. The final bar shows the fraction of queries classified as belonging to a specific region (FQCR). (a) Relative build time and (b) Relative query time.

errors for PPAC-RCCE as a function of the incurred error in temperature. The conservation errors for PPAC-RCCE are bounded below 2×10^{-5} and are deemed to be acceptable.

5.5. *PPAC-RCCE-ISAT*

The last two sections showed results for PPAC coupled with a storage-retrieval (ISAT) and a dimension reduction (RCCE) technique. A natural extension to PPAC-ISAT and

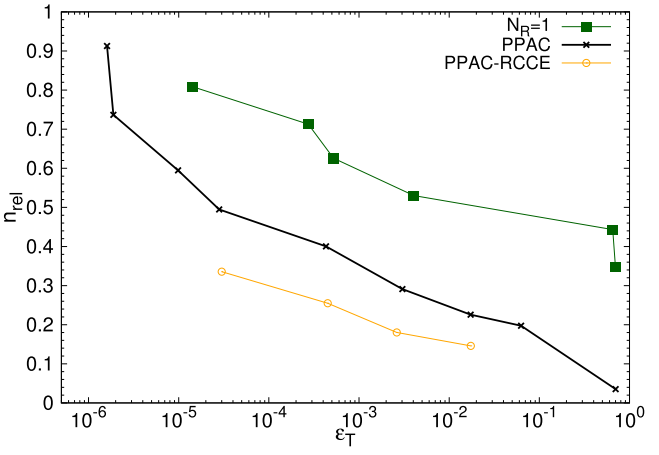

Figure 12. Relative number of species as a function of incurred error in temperature for the degenerate case, PPAC and PPAC-RCCE.

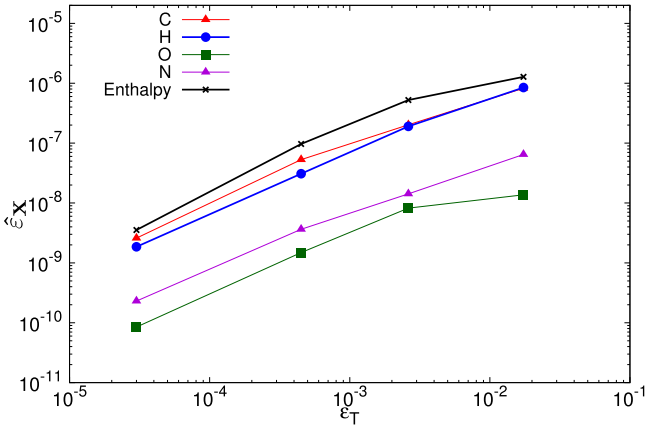

Figure 13. Element and enthalpy conservation errors as a function of incurred error in temperature for PPAC-RCCE.

PPAC-RCCE is the combined PPAC-RCCE-ISAT. The results for its performance in the previously described PaSR configuration, using the set of reduced models $A1 - A4$ are presented in this section. We test PPAC-RCCE-ISAT for ISAT error tolerance specifications of 10^{-4}, 5×10^{-5}, and 10^{-5}. PPAC-RCCE-ISAT resolves queries which are in the reduced RCCE space. Hence, PPAC-RCCE-ISAT, akin to PPAC-ISAT, builds multiple tables at runtime, one for each of the regions in the PPAC composition space partition. The maximum ISAT table size is specified as 500 MB for each table. We recall that the preprocessing stage for PPAC-RCCE-ISAT is identical to that of PPAC-RCCE. Hence, we use the same sets of represented species for PPAC-RCCE-ISAT as for PPAC-RCCE.

5.5.1. *Computational accuracy and memory requirements*

The relative number of species as a function of the incurred error in temperature is shown in Figure 14. The PPAC-RCCE-ISAT curves, similar to PPAC-ISAT, connect four adaptive

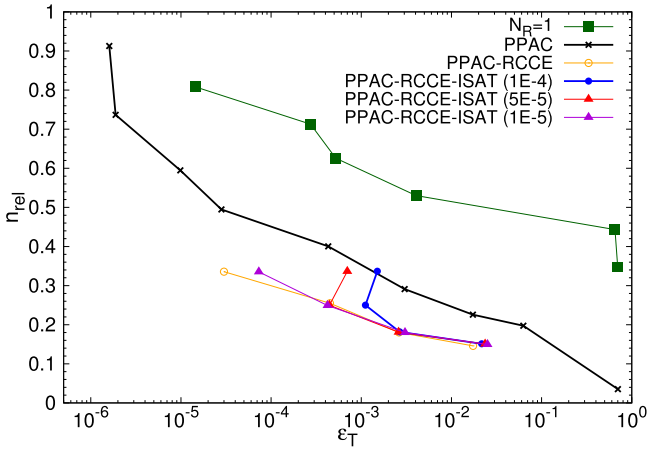

Figure 14. Relative number of species as a function of incurred in temperature for the degenerate case, PPAC, PPAC-RCCE, and PPAC-RCCE-ISAT with various ISAT error tolerances.

PaSR runs using the sets of reduced models A1–A4 and a common ISAT error tolerance specification. RCCE-ISAT is used to resolve queries in the reduced RCCE representation. Hence, we expect and observe the relative number of species for PPAC-RCCE-ISAT to be approximately equal to that of PPAC-RCCE for points that share the same set of reduced models. PPAC-RCCE-ISAT incurs reduction, dimension reduction, and tabulation errors.

As has been discussed previously, the represented species are chosen such that the reaction mapping error due to dimension reduction is 10% of the DRGEP error cutoff. In the previous section, we concluded that this leads to reduction errors dominating the dimension reduction errors for PPAC-RCCE. We expect this relation between errors to hold for PPAC-RCCE-ISAT, since PPAC-RCCE and PPAC-RCCE-ISAT utilise the same set of reduced models and represented species. Consequently, the behaviour of the incurred error in temperature for the various PPAC-RCCE-ISAT runs can be understood in terms of varying contributions from the reduction and tabulation errors.

We reiterate that PPAC-RCCE uses approximately the same relative number of species and incurs approximately the same reduction and dimension reduction errors as PPAC-RCCE-ISAT. Additionally, PPAC-RCCE does not incur any tabulation errors. Hence, we can determine the relative contribution of the reduction and tabulation errors for PPAC-RCCE-ISAT by directly comparing its performance with that of PPAC-RCCE for cases using the same sets of reduced models. For $A3$ and $A4$, we infer that reduction error dominates for the all three ISAT error tolerance specification. For $A1$ and $A2$, we observe that the tabulation error dominates for error tolerance specification of 10^{-4}. At more stringent error tolerance specifications, we infer that the reduction error dominates for run $A2$ whereas the tabulation error is higher for run $A1$.

5.5.2. *Conservation errors*

PPAC-RCCE-ISAT similar to previously discussed coupled techniques does not conserve elements or enthalpy. Figure 15 shows the variation of the conservation errors for PPAC-RCCE-ISAT for the three different ISAT error tolerance specifications. The conservation errors are bounded below 2×10^{-5}. Hence, they are deemed to be acceptable.

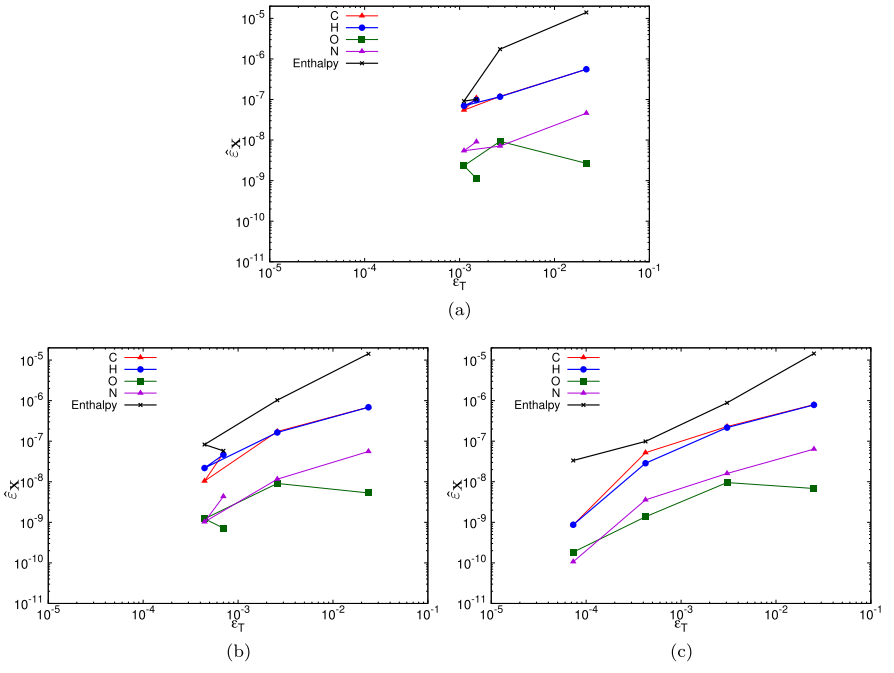

Figure 15. Element and enthalpy conservation errors as a function of incurred error in temperature for PPAC-RCCE-ISAT runs with three different ISAT error tolerance specification. (a) ISAT error tolerance: 1×10^{-4}, (b) ISAT error tolerance: 5×10^{-5} and (c) ISAT error tolerance: 1×10^{-5}.

5.5.3. *Computational efficiency*

Table 6 shows the percent of total queries resolved by RCCE-ISAT as a retrieve for the various PPAC-RCCE-ISAT runs. We observe the expected trend of decreasing retrieve fraction for increasingly stringent ISAT error tolerance specification. Table 7 shows the variation of relative time for PPAC and PPAC-RCCE-ISAT for different ISAT error tolerance specifications. For the same set of reduced models, we observe an increasing relative time for decreasing ISAT error tolerance specification. This can be directly attributed to the decreasing fraction of queries resolved as retrieves. For the same ISAT error tolerance specification, we observe a decreasing relative time, with a decreasing relative number of species.

Table 6. Total retrieve percentage for various PPAC-RCCE-ISAT runs.

	ISAT error tolerance specification		
Run number	10^{-4}	5×10^{-5}	10^{-5}
A1	93.273	87.959	75.894
A1	93.273	87.959	75.894
A2	95.452	91.286	77.411
A3	95.970	92.747	80.870
A4	96.498	93.419	78.408

Table 7. Relative time for various PPAC and PPAC-RCCE-ISAT runs.

		ISAT error tolerance specification		
Run number	PPAC	10^{-4}	5×10^{-5}	10^{-5}
A1	0.390	0.135	0.194	0.288
A2	0.289	0.083	0.132	0.231
A3	0.188	0.063	0.087	0.155
A4	0.115	0.042	0.059	0.111

Finally, we consider $A3$ with PPAC-RCCE-ISAT and an ISAT error tolerance specification of 10^{-4}. For this point, an incurred error in temperature of 2×10^{-3} is obtained while retaining only a 0.2 fraction of species in the detailed mechanism and with $\approx 96\%$ retrieves leading to a relative time of 0.063. These results show the combined benefits attained by coupling PPAC, RCCE and ISAT techniques. They additionally illustrate the potential of the combined PPAC-RCCE-ISAT to handle detailed mechanisms very efficiently for particle PDF methods.

5.5.4. *Performance analysis*

We analyse the performance of a variant corresponding to PPAC-RCCE-ISAT in comparison to variants (a)–(c) discussed in the PPAC-ISAT section. Specifically, we examine the following variant:

(d) Multiple ISAT tables corresponding to the user-specified number of regions are used with the corresponding reduced mechanisms with incoming queries in their reduced RCCE representations.

We utilise the same database of compositions used for testing the variants (a)–(c). The detailed compositions are converted to their reduced RCCE representations before calling ISAT. As before, we use an ISAT error tolerance specification of 10^{-4}, an individual maximum table size specification of 500 MB, and consider the set of reduced models $A2$ and $A3$.

The relative build and query times, and the retrieve percentage are shown in Table 8. The values for the variants (a)–(c) are repeated for reference. Comparing variants (c) and (d), we observe approximately the same relative build time and a reduction in the relative query time of $\approx 21\%$ and $\approx 9\%$ respectively. The reduction in query time can be attributed to the use of ISAT to resolve queries in their reduced RCCE representation for variant (d)

Table 8. Performance metrics for three variants for query resolution using ISAT.

Variant	Relative build time	Relative query time	Retrieve percentage
(a)	1	1	93.2
(b)	0.181	0.84	96.5
(c) ($\sim A2$)	0.029	0.59	93.0
(c) ($\sim A3$)	0.027	0.36	94.5
(d) ($\sim A2$)	0.036	0.38	95.7
(d) ($\sim A3$)	0.034	0.27	96.3

as opposed to their reduced skeletal representation for variant (c). We also observe a slight increase in the retrieve % for variant (d) compared variant (c), which can also be attributed to the tabulation being performed in a lesser number of dimensions for variant (d).

For a further examination of the performance of variant (d), we show the relative build and query time of individual reduced models in Figure 16. Additionally, we show the fraction of represented species (RCCE n_{rel}) for each region-specific reduced model and

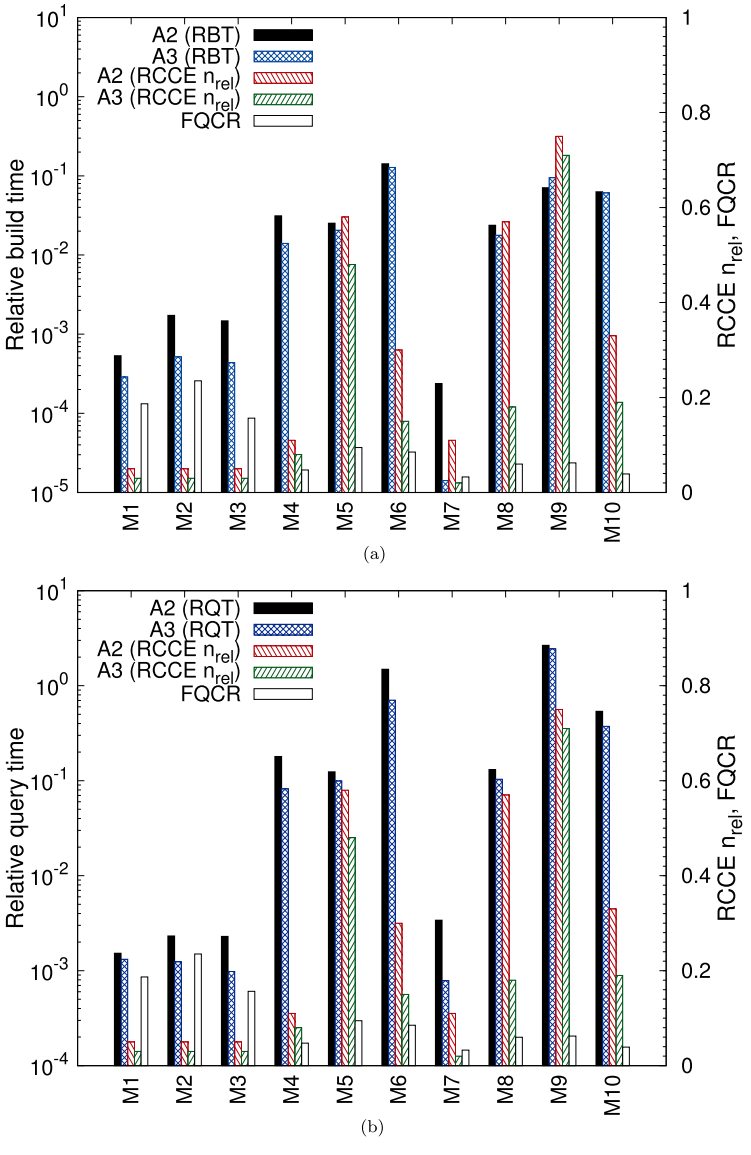

Figure 16. Relative build time (RBT) and relative query time (RQT) of individual reduced models for variant (d) using *A*2 and *A*3. The fraction of represented species for each region-specific reduced model (RCCE n_{rel}) are also reported for reference. The final bar shows the fraction of queries classified as belonging to a specific region (FQCR). (a) Relative build time and (b) Relative query time.

the fraction of queries classified as belonging to the region for reference. The fraction of represented species is computed as the number of represented species corresponding to a reduced model divided by the number of species in the detailed mechanism. We observe that 58% of the queries for $A2$ and 61% of the queries for $A3$ are resolved with reduced models represented by 5% or less of the species in the detailed mechanism. These queries are resolved at approximately a 10^{-3} fraction of the query time for standalone ISAT, demonstrating the advantage of using PPAC-RCCE-ISAT. As noted previously, for a turbulent combustion computation, a significantly larger fraction of non-reactive compositions will be encountered compared to the current PaSR test case. Hence, we expect a further improvement in performance of the combined PPAC-RCCE-ISAT technique in a turbulent combustion computation.

6. Conclusions

We have augmented the recently proposed PPAC methodology, by combining it with storage retrieval (ISAT) and dimension reduction (RCCE) techniques. Specifically, we have developed coupled implementations for PPAC-ISAT, PPAC-RCCE, and PPAC-RCCE-ISAT. The techniques examined in this work constitute a hierarchy as shown in Figure 17. Each coupled technique provides a performance improvement over the methodology below it in the hierarchy. Quantitative values of these improvements for a representative test case in a PaSR context are reported on the connecting arrows in the Figure 17. PPAC-RCCE-ISAT for this representative test case, with an ISAT error tolerance specification of 10^{-4}, leads to an incurred error in temperature of 2×10^{-3} by retaining only 20% of the species in the detailed mechanism and resolving approximately 96% of the queries as retrieves. The

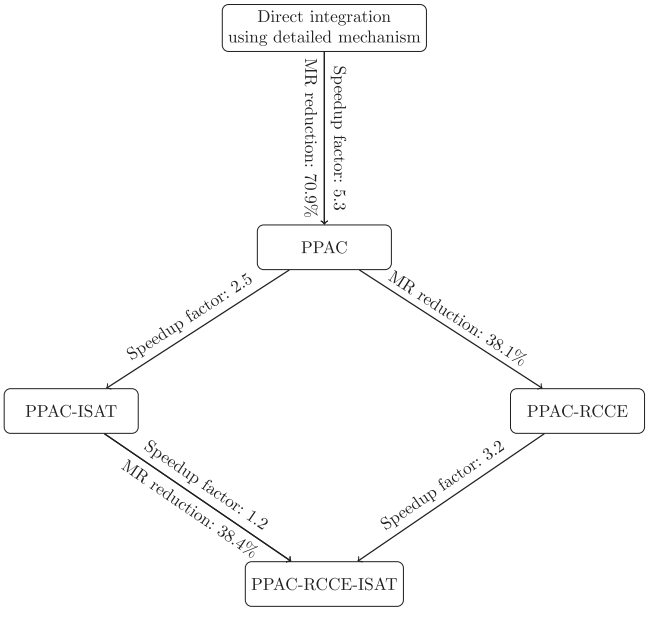

Figure 17. Hierarchy of methods considered in this work. The advantages are reported on the connecting arrows for the representative test case $A3$. The combinations that include ISAT utilised an error tolerance specification of 10^{-4}. MR refers to memory requirement.

high percentage of queries resolved as retrieves produces a timing fraction of 0.063 relative to the use of the detailed mechanism with direct integration, i.e. a speedup of a factor of 15. The coupled PPAC-RCCE-ISAT approach leads to a reduction in the CPU cost of 67% and a reduction in the memory requirement of 38% over a standalone PPAC simulation using the same set of reduced models. Thus, PPAC-RCCE-ISAT provides promising gains in performance over standalone PPAC. We anticipate that the coupled PPAC-RCCE-ISAT approach will provide improved gains in a LES/PDF context.

Acknowledgments

ASN is immensely grateful to Dr. Varun Hiremath for discussions regarding some of the subtleties of ISAT and RCCE-ISAT implementations in ISAT-CK7.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

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

ORCID

Perrine Pepiot http://orcid.org/0000-0001-9870-2809

References

[1] H. Wu, Y.C. See, Q. Wang, and M. Ihme, *A Pareto-efficient combustion framework with submodel assignment for predicting complex flame configurations*, Combust. Flame162(11) (2015), pp. 4208–4230.
[2] S.B. Pope, *Small scales, many species and the manifold challenges of turbulent combustion*, Proc. Combust. Inst. 34(1) (2013), pp. 1–31.
[3] S.B. Pope, *PDF methods for turbulent reactive flows*, Prog. Energy Combust. Sci. 11(2) (1985), pp. 119–192.
[4] D.C. Haworth, *Progress in probability density function methods for turbulent reacting flows*, Prog. Energy Combust. Sci. 36(2) (2010), pp. 168–259.
[5] H. Wang and S.B. Pope, *Large eddy simulation/probability density function modeling of a turbulent $CH_4/H_2/N_2$ jet flame*, Proc. Combust. Inst. 33(1) (2011), pp. 1319–1330.
[6] R.R. Tirunagari and S.B. Pope, *An investigation of turbulent premixed counterflow flames using large-eddy simulations and probability density function methods*, Combust. Flame 166 (2016), pp. 229–242.
[7] W. Han, V. Raman, and Z. Chen, *LES/PDF modeling of autoignition in a lifted turbulent flame: Analysis of flame sensitivity to differential diffusion and scalar mixing time-scale*, Combust. Flame 171 (2016), pp. 69–86.
[8] N. Peters, *Laminar diffusion flamelet models in non-premixed turbulent combustion*, Prog. Energy Combust. Sci. 10(3) (1984), pp. 319–339.
[9] V. Hiremath, S.R. Lantz, H. Wang, and S.B. Pope, *Computationally-efficient and scalable parallel implementation of chemistry in simulations of turbulent combustion*, Combust. Flame 159(10) (2012), pp. 3096–3109.
[10] J.C. Keck and D. Gillespie, *Rate-controlled partial-equilibrium method for treating reacting gas mixtures*, Combust. Flame 17(2) (1971), pp. 237–241.
[11] J.C. Keck, *Rate-controlled constrained-equilibrium theory of chemical reactions in complex systems*, Prog. Energy Combust. Sci. 16(2) (1990), pp. 125–154.

[12] S.H. Lam and D.A. Goussis, *The CSP method for simplifying kinetics*, Int. J. Chem. Kinet. 26(4) (1994), pp. 461–486.
[13] U. Maas and S.B. Pope, *Implementation of simplified chemical kinetics based on intrinsic low-dimensional manifolds*, in *Symposium (International) on Combustion*, Vol. 24, Elsevier, 1992, pp. 103–112.
[14] Z. Ren, S.B. Pope, A. Vladimirsky, and J.M. Guckenheimer, *The invariant constrained equilibrium edge preimage curve method for the dimension reduction of chemical kinetics*, J. Chem. Phys. 124(11) (2006), pp. 114111.
[15] F.C. Christo, A.R. Masri, E.M. Nebot, and S.B. Pope, *An integrated PDF/neural network approach for simulating turbulent reacting systems*, in *Symposium (International) on Combustion*, Vol. 26, Elsevier, 1996, pp. 43–48.
[16] S.B. Pope, *Computationally efficient implementation of combustion chemistry using in situ adaptive tabulation*, 1997.
[17] S.R. Tonse, N.W. Moriarty, N.J. Brown, and M. Frenklach, *PRISM: Piecewise reusable implementation of solution mapping. An economical strategy for chemical kinetics*, Isr. J. Chem. 39(1) (1999), pp. 97–106.
[18] L. Liang, J.G. Stevens, and J.T. Farrell, *A dynamic adaptive chemistry scheme for reactive flow computations*, Proc. Combust. Inst. 32(1) (2009), pp. 527–534.
[19] Z. Ren, Y. Liu, T. Lu, L. Lu, O.O. Oluwole, and G.M. Goldin, *The use of dynamic adaptive chemistry and tabulation in reactive flow simulations*, Combust. Flame 161(1) (2014), pp. 127–137.
[20] S. Yang, R. Ranjan, V. Yang, S. Menon, and W. Sun, *Parallel on-the-fly adaptive kinetics in direct numerical simulation of turbulent premixed flame*, in *Proceedings of the Combustion Institute*, 2016.
[21] I. Banerjee and M.G. Ierapetritou, *An adaptive reduction scheme to model reactive flow*, Combust. Flame 144(3) (2006), pp. 619–633.
[22] D.A. Schwer, P. Lu, and W.H. Green, *An adaptive chemistry approach to modeling complex kinetics in reacting flows*, Combust. Flame 133(4) (2003), pp. 451–465.
[23] Y. Liang, S.B. Pope, and P. Pepiot, *A pre-partitioned adaptive chemistry methodology for the efficient implementation of combustion chemistry in particle PDF methods*. Combust. Flame 162(9) (2015), pp. 3236–3253.
[24] Q. Tang and S.B. Pope, *Implementation of combustion chemistry by in situ adaptive tabulation of rate-controlled constrained equilibrium manifolds*, Proc. Combust. Inst. 29(1) (2002), pp. 1411–1417.
[25] V. Hiremath, Z. Ren, and S.B. Pope, *Combined dimension reduction and tabulation strategy using ISAT-RCCE-GALI for the efficient implementation of combustion chemistry*. Combust. Flame 158(11) (2011), pp. 2113–2127.
[26] V. Hiremath, S.R. Lantz, H. Wang, and S.B. Pope, *Large-scale parallel simulations of turbulent combustion using combined dimension reduction and tabulation of chemistry*, Proc. Combust. Inst.34(1) (2013), pp. 205–215.
[27] R.S. Barlow and J.H. Frank, *Effects of turbulence on species mass fractions in methane/air jet flames*, in *Symposium (International) on Combustion*, Vol. 27, Elsevier, 1998, pp. 1087–1095.
[28] A.K. Chatzopoulos and S. Rigopoulos, *A chemistry tabulation approach via rate-controlled constrained equilibrium (RCCE) and artificial neural networks (ANNs), with application to turbulent non-premixed $CH_4/H_2/N_2$ flames*, Proc. Combust. Inst. 34(1) (2013), pp. 1465–1473.
[29] W. Meier, R.S. Barlow, Y.L. Chen, and J.Y. Chen, *Raman/Rayleigh/LIF measurements in a turbulent CH4/H2/N2 jet diffusion flame: Experimental techniques and turbulence–chemistry interaction*. Combust. Flame 123(3) (2000), pp. 326–343.
[30] C. Schneider, A. Dreizler, J. Janicka, and E.P. Hassel, *Flow field measurements of stable and locally extinguishing hydrocarbon-fuelled jet flames*, Combust. Flame 135(1–2) (2003), pp. 185–190.
[31] L.L.C. Franke, A.K. Chatzopoulos, and S. Rigopoulos, *Tabulation of combustion chemistry via artificial neural networks (ANNs): Methodology and application to LES-PDF simulation of Sydney flame L*. Combust. Flame 185 (2017), pp. 245–260.
[32] A.R. Masri, R.W. Bilger, and R.W. Dibble, *The local structure of turbulent nonpremixed flames near extinction*, Combust. Flame 81(3–4) (1990), pp. 260–276.

[33] F. Contino, H. Jeanmart, T. Lucchini, and G. D'Errico, *Coupling of in situ adaptive tabulation and dynamic adaptive chemistry: An effective method for solving combustion in engine simulations*, Proc. Combust. Inst. 33(2) (2011), pp. 3057–3064.
[34] W. Xie, Z. Lu, Z. Ren, and G.M. Goldin, *Dynamic adaptive acceleration of chemical kinetics with consistent error control*, Combust. Flame 197 (2018), pp. 389–399.
[35] L. Lu and S.B. Pope, *An improved algorithm for in situ adaptive tabulation*, J. Comput. Phys. 228(2) (2009), pp. 361–386.
[36] V. Hiremath and S.B. Pope, *A study of the rate-controlled constrained-equilibrium dimension reduction method and its different implementations*, Combust. Theory Model. 17(2) (2013), pp. 260–293.
[37] S.B. Pope, *Gibbs function continuation for the stable computation of chemical equilibrium*, Combust. Flame 139(3) (2004), pp. 222–226.
[38] Z. Ren, G.M. Goldin, V. Hiremath, and S.B. Pope, *Reduced description of reactive flows with tabulation of chemistry*, Combust. Theory Model. 15(6) (2011), pp. 827–848.
[39] V. Hiremath, Z. Ren, and S.B. Pope, *A greedy algorithm for species selection in dimension reduction of combustion chemistry*, Combust. Theory Model. 14(5) (2010), pp. 619–652.
[40] P. Pepiot-Desjardins and H. Pitsch, *An efficient error-propagation-based reduction method for large chemical kinetic mechanisms*, Combust. Flame 154(1–2) (2008), pp. 67–81.
[41] H.J. Curran, T.M. Jayaweera, W.J. Pitz, and C.K. Westbrook, *A detailed modeling study of propane oxidation*, in *Western States Section of the Combustion Institute*, 2004, pp. 1–6.