

9<sup>th</sup> U. S. National Combustion Meeting  
Organized by the Central States Section of the Combustion Institute  
May 17–20, 2015  
Cincinnati, Ohio

# An adaptive methodology for the efficient implementation of detailed chemistry in simulations of turbulent non-premixed combustion

*Y. Liang, S. B. Pope, P. Pepiot.*

*Sibley School of Mechanical and Aerospace Engineering, Cornell University*

## Abstract

Large Eddy Simulation/particle Probability Density Function (LES/PDF) approaches are well developed and have been applied to turbulent combustion problems involving complex flows with strong turbulence-chemistry interactions. However, the level of details that can be used to describe the chemistry in those simulations is severely restricted by the available computational resources. To reduce both the CPU time and memory cost, a pre-partitioned adaptive chemistry (PPAC) methodology tailored to LES/PDF simulation has been developed in previous work, in which each particle is assigned a specialized reduced representation and chemical model tailored to their individual composition. Instead of performing chemical reduction at runtime to determine the optimal set of equations to use for a given particle, an analysis of the composition space region likely accessed during the turbulent flow simulation is performed using simple Partially Stirred Reactor (PaSR) computations. The PPAC approach relies on an *a priori* partitioning of the composition space into a user-specified number of regions, over which suitable reduced chemical representations and chemical models are identified automatically using the Directed Relation Graph with Error Propagation method. A computational particle in the LES/PDF simulation evolves according to, and carries only the variables present in the reduced representation corresponding to the composition space region it belongs to. This region is identified using a low-dimensional binary tree search algorithm, thereby keeping the run-time overhead associated with the adaptive approach to a minimum. The methodology was tested in a non-premixed propane/air PaSR, showing good performances both in terms of CPU cost and memory requirements. In this work, the PPAC method is first applied to a different and larger kinetic mechanism for dodecane oxidation to provide a broader validation of the technique. An extension to include dimension reduction based on Rate-Constrained Controlled Equilibrium (RCCE) is then investigated to complement the DRGEP chemistry reduction performed in the pre-processing stage, and significantly lower the number of variables required to be carried by the PDF particles.

**Keywords:** Adaptive chemistry, DRGEP, RCCE, LES/PDF, Partially Stirred Reactor

## 1 Introduction

With a deeper understanding of chemical kinetics for hydrocarbon fuel combustion and improved mechanistic considerations, some of the latest published detailed chemical kinetic schemes are approaching  $10^4$  species and usually five times as many reactions [1]. However, integrating these advances in chemical kinetics with Computational Fluid Dynamics (CFD) tools to fully realize their potential in terms of improved understanding and optimization of practical combustion devices is numerically prohibitive. Typically, different modeling approaches for reactive turbulent flows

allow different upper limits of the number of variables that can be used to describe the chemistry. We specifically focus here on LES/particle PDF methods, which are shown to be computationally viable up to a few dozen variables [2].

Considering that thermodynamic conditions can vary widely in space and time for a given reactive flow configuration, and in any small range of temperature and compositions, only a few species are chemically active, one way forward is to employ the concept of *adaptive chemistry*, in which the chemical representation and kinetic equations are adapted to the local chemical activity and thermodynamic state. Hence, the chemistry description for different conditions may require many fewer species and reactions. Several adaptive strategies have been proposed recently (*e.g.* [3–15]). While those studies showed good results in terms of CPU cost reduction, very few of them directly address the issue of memory and storage requirements, a major limitation in LES/PDF simulations.

Recently, we proposed and assessed an adaptive chemistry methodology tailored to LES/PDF simulations, as described by Liang *et al.* [16]. The two primary defining characteristics of the methodology are that (1) the composition space is partitioned, and reduced representations and models defined in a pre-processing step, and (2) that the Directed Relation Graph with Error Propagation (DRGEP) is used to generate the reduced kinetic models. The methodology is therefore called Pre-Partitioned Adaptive Chemistry using Direction Relation Graphs with Error Propagation, or PPAC-DRGEP.

To provide context to the work described in this paper, a brief description of the method is provided in Sec. 2, focusing on the aspects most relevant to the subsequent sections. For additional details, the reader is referred to the original paper by Liang *et al.* [16] (preprint accessible online.) This is followed by the description of the two main contributions in this paper: 1) the application of PPAC to a larger chemical kinetic mechanism describing the oxidation of dodecane, to confirm that the methodology is indeed general (Sec. 3); and 2) the extension of the methodology to include dimension reduction based on Rate-Constrained Controlled Equilibrium (RCCE) [19], to complement the DRGEP chemistry reduction performed in the pre-processing stage. The potential of RCCE to further decrease the number of variables required to be carried by the PDF particles in PPAC is investigated (Sec. 4).

Note that all test cases described here are conducted in a Partially-Stirred Reactor (PaSR), which is (by construction) similar to a single cell in a LES/PDF simulation: it should be emphasized that the same methodology is directly applicable to an LES/PDF simulation.

## 2 PPAC Methodology Overview

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  $T$ , defined by the composition vector of size  $(n_s + 1)$ :  $\Phi \equiv \{\mathbf{Y}, T\}$ . In the PaSR simulation, the compositions of the large number of particles evolve in time in small time steps,  $\Delta t$ , in three fractional steps accounting for: inflow/outflow/pairing; mixing; and reaction. Similarly, in an LES/PDF simulation, particles are added and removed as appropriate at boundaries, and their composition changes in mixing and reaction fractional steps.

The PPAC-DRGEP combines an off-line, **pre-processing stage** (during which the composition space is partitioned and a set of reduced kinetic models is created), and an online, dynamic procedure, the **adaptive simulation**, in which the appropriate reduced representation and model to use for each particle at time  $t$  is identified.

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

1. *Database creation*: Using the PaSR and the detailed chemical mechanism (involving  $n_s$  species and  $n_r$  reactions), generate a database  $\mathcal{D}$  consisting of a large number  $n_{\mathcal{D}}$  of full compositions, representative of the compositions that occur during the PaSR or LES/PDF simulation of interest.
2. *Partitioning*: Partition a suitably-defined low-dimensional space (referred to as the *classifying space*),  $\mathbb{C}$ , into a specified number  $N_R$  of regions, the  $J$ th region being denoted by  $R_J$ , by iteratively introducing cutting hyperplanes in  $\mathbb{C}$  to isolate groups of sample compositions, *i.e.* subsets of  $\mathcal{D}$ , with similar short-term kinetics. The partition structure is stored as a binary tree.
3. *Reduced modeling*: For each region  $R_J$ , identify a reduced representation and kinetic model  $M_J$  involving a reduced set of  $n_s^J$  retained species and  $n_r^J$  reactions, with the expectation that  $n_s^J$  and  $n_r^J$  are significantly smaller than  $n_s$  and  $n_r$ , respectively. The  $n_s^J$ -vector of the mass fraction of the retained species is denoted by  $\mathbf{y}_J$ . In the adaptive simulation, the reduced composition of the  $n$ th particle is denoted by  $\phi^{(n)}$ , or reduced composition *based on model*  $M_J$  and is defined as

$$\phi_J^{(n)} = \{\mathbf{y}_J^{(n)}, T^{(n)}\}, \quad (1)$$

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 or LES/PDF can be computed using reduced representations and kinetic models dynamically chosen among this set. The steps involved in the **adaptive algorithm** are described in Fig. 1 for  $N_R = 4$ .

### 3 Application to non-premixed dodecane oxidation

The first objective of this paper is to demonstrate that the strategy outlined above, successfully tested for a 100-species propane mechanism, can be used to handle different, larger kinetic models with similar performances. This is done here in a non-premixed PaSR burning an overall stoichiometric mixture of  $n$ -dodecane and air. The parameters used in the study are provided in Sec. 3.1, followed by the description of the error measures used to assess the performances of the adaptive strategy (Sec. 3.2). Section 3.3 compares results obtained in this work for dodecane oxidation to those obtained in Liang *et al.* [16] for propane.

#### 3.1 PaSR configuration and parameters

We consider the detailed mechanism for  $n$ -dodecane developed by Narayanaswamy *et al* [17], which consists of  $n_s = 255$  species involved in  $n_r = 2289$  reactions. The PaSR has three inflow streams: pure fuel, oxidizer (air), and a pilot stream defined as the burnt products of a stoichiometric dodecane/air mixture in chemical equilibrium. The pressure is 1 bar throughout. The fuel and




Figure 1: 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 mixed according to the specified mixing model to obtain the composition after mixing  $\Phi^{(n),m}$  (2). Because of the mixing process, the representation after mixing,  $\Phi^{(n),m}$ , has moved to region  $R_4$ , which is determined from classification in the partition (3).  $\Phi^{(n),m}$  is reduced to the appropriate representation,  $\phi_{J=4}^{(n),m}$  (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  $R_4$ ).

air streams are at a temperature of 300K, while the pilot stream is at the adiabatic flame temperature. 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. All relevant parameters are included in Table 1.

Table 1: Specification of the parameters for the non-premixed propane/air piloted PaSR test case.

Reactor characteristics		
Parameters	Name	Values
Pressure	$p$	1 bar
Number of particles	$n_p$	100
Pairing time	$\tau_{\text{pair}}$	1 ms
Mixing time	$\tau_{\text{mix}}$	1 ms
Residence time	$\tau_{\text{res}}$	10 ms
Time step	$\Delta t$	0.1 ms

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. After this initial run, we collect sample compositions for another ten residence times. The composition database  $\mathcal{D}$  consists of 10,000 compositions, obtained by randomly sampling from the  $10^5$  compositions collected.

### 3.2 Assessment of performances

It is important to define clearly the different types of errors and performance measures used to assess the accuracy of the adaptive runs.

- The *error tolerance*, or cut-off error  $\varepsilon_c$ , is used during *Reduced Modeling* in the pre-processing stage to determine for each region  $J$  a reduced model  $M_J$  that can describe the short-term time evolution of any composition belonging to that region with an error tolerance  $\varepsilon_c$ . This error tolerance is user-specified.
- The *incurred error* for a quantity  $X$ ,  $\varepsilon_X$  is the error observed during adaptive runs when comparing the resulting particle trajectories to those obtained using the detailed mechanism. This error has two sources: the reduced kinetic modeling, directly controlled by  $\varepsilon_c$ ; and the conversion error, which arises when the representation of a particle's composition changes from one reduced representation to another. The incurred error is measured for any quantity  $X$  (temperature or species mass fractions) as:

$$\varepsilon_X = \frac{\sum_{k=1}^{n_t} \sum_{n=1}^{n_p} |X_k^{(n),A} - X_k^{(n),D}|}{\sum_{k=1}^{n_t} \sum_{n=1}^{n_p} |X_k^{(n),D}|}. \quad (2)$$

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

- An additional measure relevant to PDF methods aims at assessing the computer memory cost of the adaptive simulations. It is defined as the relative number of species,  $\bar{n}$ , 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,  $n_s$ :

$$\bar{n} \equiv \frac{1}{n_t n_p} \sum_{k=1}^{n_t} \sum_{n=1}^{n_p} n_{s,k}^{(n),A}, \quad (3)$$

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

- Finally, we define the *relative time*  $t_{\text{rel}}$  of an adaptive simulation as its CPU time divided by that of the simulation with detailed chemistry.

Adaptive simulations are run for a wide range of error tolerances  $\varepsilon_c$  (therefore leading to a wide range of incurred errors  $\varepsilon_X$ ), and for two values of  $N_R$ , namely  $N_R = 1$  and  $N_R = 30$ . Note that the case  $N_R = 1$  is degenerate in that the adaptive methodology is in fact not adaptive: we refer to this as the *non-adaptive* case. Particle evolutions are recorded over a total time  $\tau$  equal to 10 residence times, corresponding to  $n_t = 1,000$  time steps. 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.

### 3.3 Results

Figure 2 shows how the incurred error measures for temperature,  $\varepsilon_T$ , and a major product, in this case  $\text{CO}_2$ ,  $\varepsilon_{\text{CO}_2}$  compare to one another. The correlation between the two being very strong,  $\varepsilon_T$  is used throughout the section as the representative incurred error  $\varepsilon$ .




Figure 2: Correlation between incurred error in temperature,  $\varepsilon_T$ , and incurred error in  $\text{CO}_2$ ,  $\varepsilon_{\text{CO}_2}$ , showing that  $\varepsilon_T$  is an appropriate representative measure of the errors introduced during the adaptive simulation.

**Computer time** Figure 3(a) shows the relative time  $t_{\text{rel}}$  as a function of the incurred temperature error for the different values of  $N_R$ . It may be observed that, as expected, the simulations using reduced models require less CPU time than those using the detailed mechanism (i.e.,  $t_{\text{rel}}$  is less than unity), and the adaptive case ( $N_R = 30$ ) require, typically, just two thirds of 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 20% of that for the detailed mechanism, and about 36% of that for the non-adaptive case.

**Computer memory** Storage requirements of the PaSR simulations, as measured by the relative number of species  $\bar{n}$ , are shown in Fig. 3(b) for both adaptive ( $N_R = 30$ ) and non-adaptive ( $N_R = 1$ ) cases. This shows, for any specified incurred temperature error in the adaptive PaSR simulations, the degree of reduction that can be achieved by the method. Both curves are monotonically decreasing with increasing incurred error  $\varepsilon$ , 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). As above, considering  $\varepsilon_T = 3 \times 10^{-3}$  as a typical desired error level, 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 ( $N_R = 1$ ).

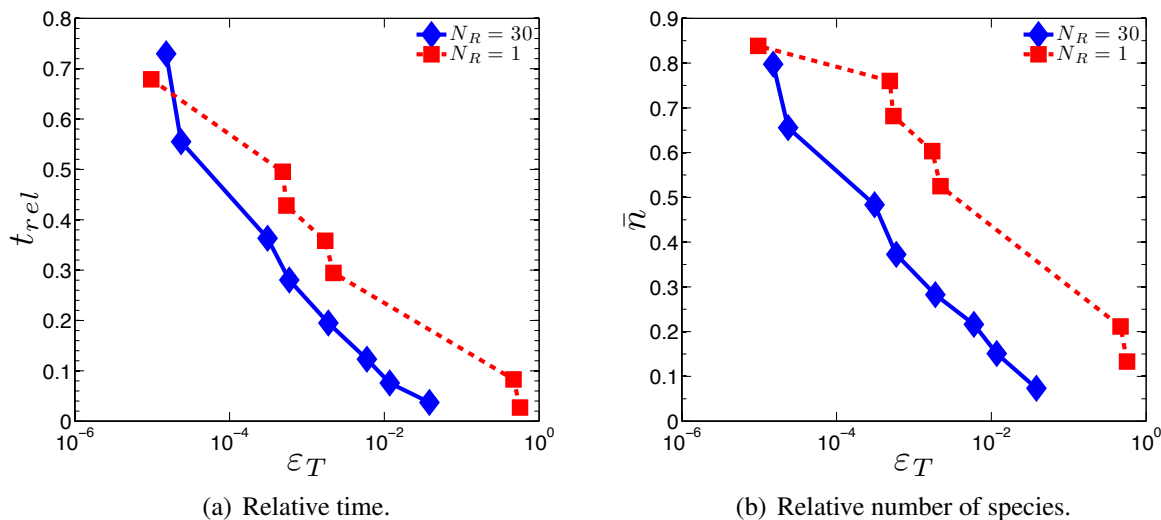


Figure 3: Relative time  $t_{rel}$  and number of species  $\bar{n}$  (adaptive to detailed) vs. the incurred error in temperature  $\varepsilon_T$ . Comparison between the non-adaptive case,  $N_R = 1$  (red squares) and adaptive case with  $N_R=30$  (blue diamonds).

**Comparison with previous results.** While the above results show clear benefits of the adaptive approach, even in a PaSR with pair-wise mixing (and it is expected that those 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), it is interesting to compare performances of the PPAC-DRGEP approach when applied to different mechanisms. This is done in Figure 4, where results obtained previously for a non-premixed propane/air PaSR (115 species, 1300 reactions, [16]) are plotted on top of the above dodecane results (255 species, 2289 reactions). The qualitative behavior is quite similar, with the non-adaptive case yielding both higher relative times and number of species. It can also be observed that performances in the propane case are significantly better for low values of the incurred error, which may indicate a non-optimal choice of parameters in DRGEP. However, these differences for a large part disappear (with better performances even observed for the dodecane case) in the range of incurred errors most likely of interest for LES/PDF, namely,  $10^{-4} < \varepsilon < 10^{-2}$ .

#### 4 Combined DRGEP/RCCE methodology in PPAC

The PPAC procedure described above relies on the DRGEP chemistry reduction technique to simultaneously decrease computer time and memory requirement. DRGEP's main mode of action is to eliminate from the models  $M_J$  all species and reactions that do not significantly impact combustion chemistry, as measured by the prediction of user-specified targets. To further decrease CPU and memory costs, PPAC-DRGEP can be combined with additional dimension/reduction techniques to further decrease the number of variables to consider, and storage/retrieval algorithms to accelerate the expensive evaluation of the chemical source terms. Of special interest here is the combined rate-controlled constraint equilibrium/in situ adaptive tabulation (RCCE/ISAT) methodology previously developed by Hiremath *et al.* [2, 18], since it has been shown to be very efficient

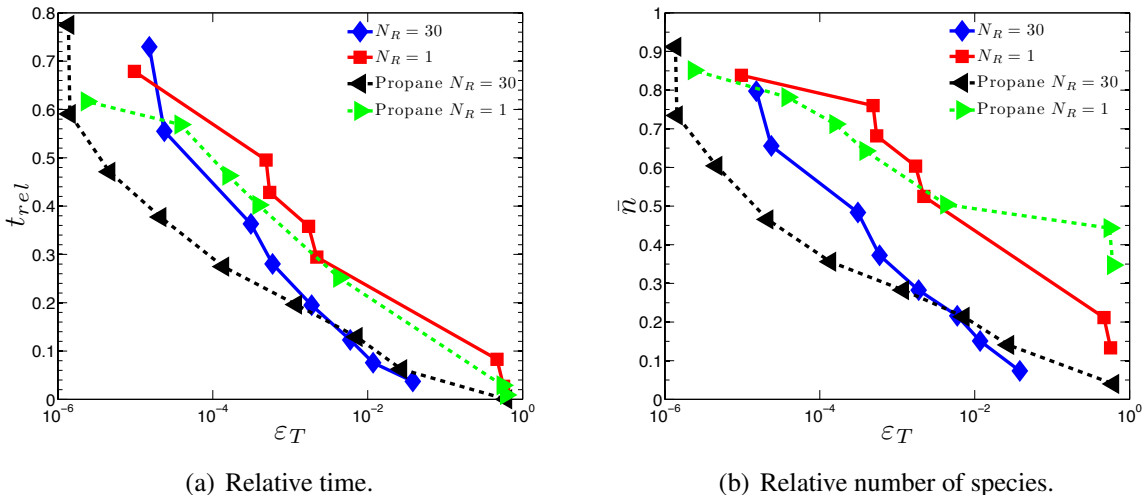


Figure 4: Relative time  $t_{rel}$  and number of species  $\bar{n}$  (adaptive to detailed) vs. the incurred error in temperature  $\varepsilon_T$ . Comparison between simulations using a 255-species dodecane mechanism (diamonds/squares) and a 115-species propane mechanism (triangles).

in the LES/PDF computational framework. As a first step toward full integration of PPAC-DRGEP and RCCE/ISAT, we focus on the potential of the RCCE approach to further decrease computer memory requirements. Integration with ISAT, and therefore demonstration of further CPU cost reduction, is left to future work.

#### 4.1 Methodology

In the conventional RCCE framework [2, 18], the user specifies a number  $n_{rep}$  of major species (or, more generally, “represented variables”, which are linear combinations of species), with the assumption that  $n_{rep}$  is much smaller than the number of species  $n_s$  of the detailed kinetic mechanism. Then, the RCCE reduction method identifies a low-dimensional manifold (of dimension  $n_{rep}$ ) in the  $n_s$ -dimensional species space; and, by assumption, the composition occurring in the reactive flow are confined to this low-dimensional manifold, which is parameterized by the major species.

Here, rather than considering the detailed mechanism as starting point, the RCCE procedure is applied to the reduced models  $M_J$  identified by DRGEP in each region of the partition. Accordingly, the PPAC-DRGEP reduced representation  $\phi_J = \{y_J, T\}$  in region  $R_J$  is replaced by the RCCE representation  $\phi_J^R$  defined as:

$$\phi_J^R \equiv \{y_J^R, z^{u,e}, T\} \quad (4)$$

where  $y_J^R$  is the mass fraction of the  $n_{rep}^J$  represented species for region  $J$ , and  $z^{u,e}$  is the  $n_e$ -vector of the specific moles of the elements in the unrepresented species (atom conservation). Species reconstruction from  $\phi_J^R$  to  $\phi_J$  is performed by calculating the constrained-equilibrium composition for the given constraints of represented species using CEQ library [19], and the reconstruction of  $\phi_J$  to full representation is done using the usual PPAC conversion tools.

The represented species for each region  $J$ , forming a subset of the retained species in region



$J$ , is chosen as part of the pre-processing stage using a greedy algorithm. Initially, all retained species are identified as represented species. Species are then removed iteratively from the set of represented species based on how they impact the accuracy of the CEQ reconstruction step. The algorithm stops when the CEQ reconstruction applied to all composition in the pre-processing database yields errors larger than the 10% DRGEP error cut-off  $\varepsilon_c$ . In the adaptive simulation, particles carry only represented species and unrepresented elements instead of the DRGEP-derived  $\phi_J$  representation.

## 4.2 Results

The non-premixed propane/air PaSR used as main test case in previous work [16] is repeated here with an identical choice of parameters and the addition of the RCCE procedure. This results in two major effects: 1) on the relative number of species required, since only  $\phi_J^R$  needs to be carried by the particles, and 2) on the observed error levels, since the CEQ reconstruction step does introduce an additional source of error. The performances of the combined DRGEP/RCCE approach are compared to the original PPAC-DRGEP in Fig. 5 for both incurred temperature and  $\text{CO}_2$  errors. We observe a significant shift of the curves toward lower errors, for example decreasing the number of variables needed to achieve an incurred error of  $10^{-3}$  by close to a factor of two.

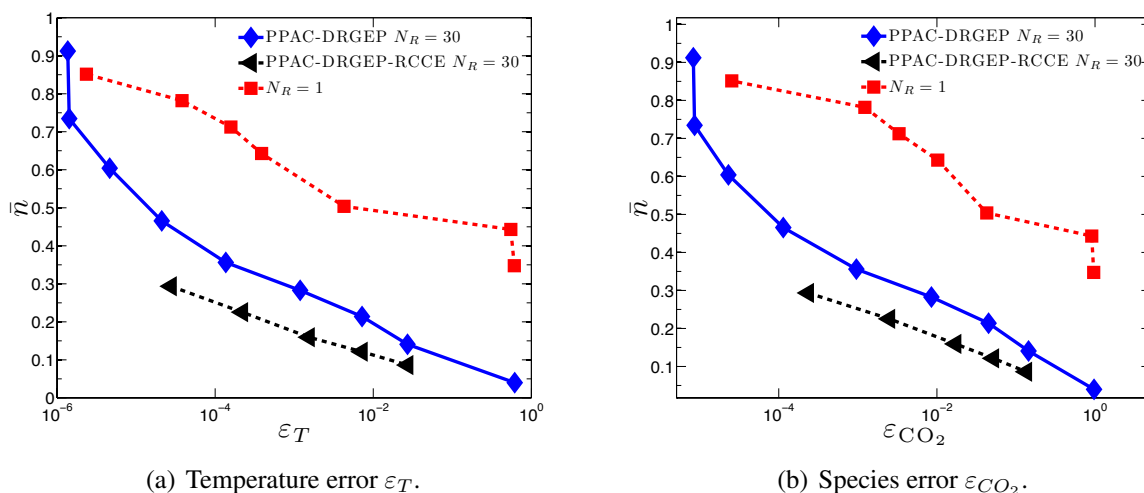


Figure 5: Relative number of species  $\bar{n}$  (computed based on the number of represented variables ( $n_{\text{rep}}^J + n_e$ ) carried by particles) as a function of the incurred temperature error (left) and  $\text{CO}_2$  error (right) in a non-premixed propane/air PaSR. Comparison between the non-adaptive case ( $N_R = 1$ , DRGEP only, red squares) and the adaptive cases ( $N_R = 30$ ) with RCCE (black triangles) and without RCCE (blue diamonds).

## 5 Conclusions

The recently developed pre-partitioned adaptive chemistry PPAC approach tailored for LES/particle PDF simulations has been tested on a larger kinetic mechanism for dodecane oxidation containing more than 250 species and nearly 2300 reactions. Similar observations have been made compared to the propane/air configuration extensively studied in the original work [16], including excellent

correlations between various incurred errors, and significant time and memory reduction compared to both the detailed mechanism, and a non-adaptive reduced simulation. The PPAC procedure has also been extended to include dimension-reduction in the pre-processing chemistry reduction stage. Using the resulting combined DRGEP/RCCE approach, the number of variables required to maintain the incurred errors in the adaptive simulation below a given threshold has been reduced by a factor of two compared to the DRGEP procedure alone. To fully reap the benefits of the RCCE approach, PPAC will be combined with ISAT in future work, with an expected significant decrease in CPU time for the adaptive simulations.

## 6 Acknowledgements

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.

## References

- [1] O. Herbinet, W. J. Pitz, and C. K. Westbrook. *Combust. Flame*, 154 (2008) 507–528.
- [2] V. Hiremath, S. R. Lantz, H. Wang, and S. B. Pope. *Proc. Combust. Inst.*, 34 (2013) 205 – 215.
- [3] L. Liang, J. G. Stevens, and J. T. Farrell. *Proc. Combust. Inst.*, 32 (2009) 527–534.
- [4] F. Contino, H. Jeanmart, T. Lucchini, and G. D’Errico. *Proc. Combust. Inst.*, 33 (2011) 3057–3064.
- [5] L. Tosatto, B. A. V. Bennett, and M. D. Smooke. *Combust. Flame*, 158 (2011) 820–835.
- [6] T. Løvas, S. Navarro-Martinez, and S. Rigopoulos. *Proc. Combust. Inst.*, 33 (2011) 1339–1346.
- [7] B. J. Debusschere, Y. M. Marzouk, H. N. Najm, B. Rhoads, D. A. Goussis, and M. Valorani. *Combust. Theory Modelling*, 16 (2012) 173–198.
- [8] M. Singer and W. Green. *Appl. Numer. Math.*, 59 (2009) 272–279.
- [9] I. Banerjee and M. Ierapetritou. *Combust. Flame*, 144 (2006) 619–633.
- [10] H. Yang, Z. Ren, T. Lu, and G. M. Goldin. *Combust. Theory Modelling*, 17 (2013) 167–183.
- [11] Z. Ren, C. Xu, T. Lu, and M. A. Singer. *J. Comput. Phys.*, 263 (2014) 19–36.
- [12] A. Viggiano and V. Magi. *Appl. Energ.*, 113 (2014) 848–863.
- [13] X. Goo, Z. Chen, W. Sun, and Y. Ju. *Combust. Flame*, 160 (2013) 225 – 231.
- [14] Z. Ren, Y. Liu, T. Lu, L. Lu, O. O. Oluwole, and G. M. Goldin. *Combust. Flame*, 161 (2014) 127–137.
- [15] K. He, I. P. Androulakis, and M. G. Ierapetritou. *Energy Fuels*, 25 (2011) 3369–3376.
- [16] Y. Liang, S. B. Pope, and P. Pepiot. An Adaptive Methodology for the Efficient Implementation of Combustion Chemistry in Particle PDF Methods. (2015), submitted. Preprint accessible at <https://tcg.mae.cornell.edu/pubs.html>.
- [17] K. Narayanaswamy, P. Pepiot, and H. Pitsch. *Combust. Flame*, 161 (2014) 866 – 884.
- [18] V. Hiremath, Z. Ren, and S. B. Pope. *Combust. Flame*, 158 (2011) 2113–2127.
- [19] D Hamiroune, P Bishnu, M Metghalchi, and J C Keck. *Combust. Theory Modelling*, (1998) 81 – 94.