Combustion Theory and Modelling

A study of the rate-controlled constrained-equilibrium dimension reduction method and its different implementations

Varun Hiremath [a] & Stephen B. Pope [a]

[a] Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, NY, 14853, USA
Version of record first published: 29 Jan 2013.

PLEASE SCROLL DOWN FOR ARTICLE

Taylor & Francis
Taylor & Francis Group

A study of the rate-controlled constrained-equilibrium dimension reduction method and its different implementations

Varun Hiremath* and Stephen B. Pope

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

The Rate-Controlled Constrained-Equilibrium (RCCE) method is a thermodynamic based dimension reduction method which enables representation of chemistry involving n_s species in terms of fewer n_r constraints. Here we focus on the application of the RCCE method to Lagrangian particle probability density function based computations. In these computations, at every reaction fractional step, given the initial particle composition (represented using RCCE), we need to compute the *reaction mapping*, i.e. the particle composition at the end of the time step. In this work we study three different implementations of RCCE for computing this reaction mapping, and compare their relative accuracy and efficiency. These implementations include: (1) RCCE/TIFS (Trajectory In Full Space): this involves solving a system of n_s rate-equations for all the species in the full composition space to obtain the reaction mapping. The other two implementations obtain the reaction mapping by solving a reduced system of n_r rate-equations obtained by *projecting* the n_s rate-equations for species evaluated in the full space onto the constrained subspace. These implementations include (2) RCCE: this is the classical implementation of RCCE which uses a direct projection of the rate-equations for species onto the constrained subspace; and (3) RCCE/RAMP (Reaction-mixing Attracting Manifold Projector): this is a new implementation introduced here which uses an alternative projector obtained using the RAMP approach. We test these three implementations of RCCE for methane/air premixed combustion in the partially-stirred reactor with chemistry represented using the n_s=31 species GRI-Mech 1.2 mechanism with n_r=13 to 19 constraints. We show that: (a) the classical RCCE implementation involves an inaccurate projector which yields large errors (over 50%) in the reaction mapping; (b) both RCCE/RAMP and RCCE/TIFS approaches yield significantly lower errors (less than 2%); and (c) overall the RCCE/TIFS approach is the most accurate, efficient (by orders of magnitude) and robust implementation.

Keywords: RCCE; dimension reduction; invariant manifold; RAMP; projections

1. Introduction

Recent advances on the experimental and theoretical fronts in the study of real fuel chemistry have led to more accurate chemical mechanisms of real fuels involving hundreds to thousands of species and thousands of reactions [1]. A major challenge in the numerical study of turbulent combustion problems is the accurate and efficient use of this detailed chemistry information in computations.

---

*Corresponding author. Email: vh63@cornell.edu

In a reacting flow computation, the species composition evolves by three processes: advection, diffusion and chemical reaction. Here we focus on the general class of solution methods in which a splitting scheme is used to account for these processes in separate fractional steps. In particular we focus on turbulent combustion modelling using Probability Density Function (PDF) methods [2], in which the chemical composition in the computational domain is represented by a large number of particles. The particle composition evolves due to mixing and chemical reaction, which are treated in two separate fractional steps. The main advantage of using PDF methods is that the chemical source term in the species evolution equation is represented exactly, which enables the use of detailed chemistry in combustion calculations. PDF methods are typically used in conjunction with Reynolds-averaged Navier–Stokes (RANS) or Large-Eddy Simulation (LES) based approaches to perform turbulent combustion simulations [3–6].

Turbulent combustion simulation using PDF methods with detailed chemistry (without any simplification) entails solving (at each reaction fractional step for each particle) a coupled set of Ordinary Differential Equations (ODEs) for the chemical species composition. These systems of ODEs are generally *stiff* owing to the wide range of chemical time-scales present in the system, and thus computing the solution is expensive. Owing to this high cost involved in representing chemistry in turbulent combustion simulations, incorporating detailed chemistry involving thousands of species is computationally prohibitive.

The current challenges of representing chemistry in turbulent combustion simulations and the various approaches used are highlighted in [7]. In the past two decades, numerous methods have been developed to tackle the high cost involved in representing combustion chemistry. These methods can be broadly classified into the following three categories.

(1) Mechanism Reduction: this includes methods designed to generate smaller skeletal mechanisms from the detailed mechanism by systematically removing unimportant species. Two prominent methods in this category are the Directed Relations Graph (DRG) [8]; and the DRG with error propagation (DRGEP) [9].

(2) Dimension Reduction: this includes methods used to represent chemistry using fewer 'represented' variables based on the detailed chemistry. Methods in this category include the Quasi Steady-State Assumption (QSSA) [10, 11]; Rate-Controlled Constrained-Equilibrium (RCCE) [12, 13]; Computational Singular Perturbation (CSP) [14]; Intrinsic Low-Dimensional Manifolds (ILDM) [15]; Trajectory-Generated Low-Dimensional Manifolds (TGLDM) [16]; and Invariant Constrained Equilibrium-Edge Pre-Image Curve (ICE-PIC) [17].

(3) Tabulation: this includes storage-and-retrieval based methods, such as In Situ Adaptive Tabulation (ISAT) [18, 19]; Piecewise Reusable Implementation of Solution Mapping (PRISM) [20]; and Artificial Neural Networks (ANNs) [21].

The aforementioned methods have been successfully applied in various combustion chemistry calculations, and they have enabled the use of detailed chemistry information in computations with acceptable levels of accuracy and efficiency.

In our research, we have focused on developing combined methodologies [22–25], which enables us to extract the best out of the aforementioned three categories, thereby further reducing the cost of chemistry computations. In particular, our recent efforts have been focused on developing a combined reduction–tabulation strategy [25, 26], which involves dimension reduction of chemistry using the RCCE method followed by tabulation using ISAT. This combined ISAT/RCCE approach can be used with both detailed and skeletal mechanisms. We have also developed an automated Greedy Algorithm with Local

Improvement (GALI) [25] to select 'good' represented species for performing dimension reduction with the RCCE method.

We have extensively tested this combined ISAT/RCCE/GALI methodology for methane and ethylene chemistry with chemical mechanisms involving 20 to 100 species in a Partially-Stirred Reactor (PaSR) [25, 27] and also in full-scale Large-Eddy Simulation (LES)/Probability Density Function (PDF) computations of Sandia Flame D [26]. The main conclusions drawn from these tests are that the ISAT/RCCE/GALI approach: (a) yields the same level of accuracy as other reduced (using QSSA) and skeletal mechanisms with relatively fewer represented species; and (b) results in speed-up by a factor of 2 to 15 relative to using ISAT alone [6, 25].

In the RCCE dimension reduction method, the chemistry involving n_s-species is represented in terms of fewer n_r-constraints. The reduced representation of chemistry using the RCCE dimension reduction method is denoted by an n_r-vector $\mathbf{r}$. In PDF based simulations of reacting flows using the RCCE method, given the initial reduced composition of a particle at the beginning of a reaction time step, $\mathbf{r}(0)$, the task is to compute the *reaction mapping*, $\mathbf{r}(t)$, at the end of the reaction time step t. There are different ways of implementing the RCCE dimension reduction method to obtain this reaction mapping, and the present paper studies the relative merits of these implementations.

Our implementation of RCCE [25] is different from the classical RCCE approach first introduced in [12] and further developed and tested in [13, 28–31]. The Close-Parallel Inertial Manifold (CPIM) method [32] describes yet another way of implementing RCCE. The main focus of this paper is to compare these different implementations of RCCE for their relative accuracy and efficiency. In particular we look at the following three implementations of RCCE:

(1) RCCE/TIFS (Trajectory In Full Space): this is the implementation used in our previous work [25–27] (described in Section 3.1);
(2) RCCE: this is the classical implementation introduced in [12] and further developed in [13, 28–31] (described in Section 3.2.1);
(3) RCCE/RAMP (Reaction-mixing Attracting Manifold Projector): this is a new implementation (based on the CPIM [32] method) proposed here (described in Section 5).

To compute the reaction mapping, the RCCE/TIFS implementation solves a system of n_s ODEs in the full composition space for all the species. In contrast, the RCCE and RCCE/RAMP implementations solve a reduced system of n_r ODEs for the constraints by *projecting* the full system of ODEs onto the constrained subspace. In this work we show that, for small reaction time steps, all the three aforementioned implementations yield similar levels of error. However, as the reaction time step increases, the RCCE/RAMP and RCCE/TIFS implementations yield orders-of-magnitude smaller errors than the RCCE implementation. We show that the projector used in the RCCE implementation is inaccurate, which results in large errors in the reaction mapping at large reaction time steps. We show that the RAMP approach provides a more accurate projector and significantly reduces the error. We also show that the RCCE/TIFS implementation is the most accurate, efficient and robust among the above three implementations.

The outline of the remainder of the paper is as follows: in Section 2 we describe the general framework and notation used for representing chemistry in our work; in Section 3 we give an overview of the RCCE dimension reduction method and describe the RCCE/TIFS and the classical RCCE implementations; in Section 4 we describe the projection issues involved in the implementation of RCCE; in Section 4.2 we describe the Close-Parallel

Inertial Manifold (CPIM) method [32] and then in Section 5 we extend it to the new Reaction-mixing Attracting Manifold Projector (RAMP) method; in Section 6 we describe the partially-stirred reactor used for testing the different implementations of RCCE; and in Section 7 we describe results to compare the relative accuracy and efficiency of the three implementations. Finally, in Section 8 we state our conclusions.

2. Chemistry representation

We consider a gaseous phase reacting flow consisting of n_s chemical species composed of n_e elements. The thermochemical state of the fluid (at a given position and time) is fully characterised by the pressure p, enthalpy h and an n_s-vector of species specific moles **z**. (It is often convenient to view **z** as a point in the n_s-dimensional composition space.)

For definiteness, we focus on the combined Large-Eddy Simulation (LES)/Probability Density Function (PDF) simulations of turbulent reacting flows, which has been our recent focus of attention [5, 6, 26]. In these simulations the thermochemical composition of the fluid is represented by a large number of particles in the computational domain. The particle chemical composition evolves due to *mixing* and *reaction*, which are treated in separate *fractional steps* [33].

Furthermore, here, in order to focus solely on the reaction fractional step and study different implementations of the RCCE method, we use the computationally cheaper representative test case of a Partially-Stirred Reactor (PaSR), which will be described in more detail in Section 6. The methodology described here, however, is applicable to other time-stepping based reacting flow simulations as well.

To simplify the exposition, here we consider an adiabatic and isobaric reaction fractional step, i.e. the enthalpy h and pressure p of a particle remain constant over the reaction fractional step. In addition, to simplify things further, we consider an isobaric flow so that pressure p is the same for all the particles. Hence the thermochemical state is fully characterised by **z** and h. (Note that these assumptions are made only to simplify the exposition and can be easily relaxed if needed. In general pressure p can vary from particle to particle, and changes in particle enthalpy h can be incorporated in other fractional steps in the flow solver.)

As a consequence of the aforementioned assumptions, in the reaction fractional step, the chemical composition of each of the particles in the computational domain evolves (at constant h and p) by the following set of rate-equations:

$$\frac{d\mathbf{z}}{dt} = \mathbf{S}(\mathbf{z}), \tag{1}$$

where the n_s-vector **S** denotes the *chemical source term* obtained from the chemical mechanism used for representing the chemistry. More precisely, the chemical source term is given as

$$\mathbf{S}(\mathbf{z}) \equiv \mathbf{S}(\mathbf{z}, T(h, \mathbf{z}), p). \tag{2}$$

Since the temperature T is known in terms of h and **z**, and the pressure p is assumed to be constant, we henceforth use the more concise notation $\mathbf{S}(\mathbf{z})$, where dependence on T and p is implicitly assumed.

Given the initial particle composition $\mathbf{z}(0)$ at $t = 0$, and a reaction fractional time step t, we refer to the composition at the end of the reaction fractional step $\mathbf{z}(t)$ as the *reaction*

mapping. The reaction mapping obtained by directly integrating the system of ODEs given by Equation (1) is referred to as the *exact solution* (since the ODE integration errors are relatively small). We use DDASAC [34] for performing ODE integration.

3. Rate-controlled constrained-equilibrium

Here we give a brief overview of the Rate-Controlled Constrained-Equilibrium (RCCE) dimension-reduction method and introduce the notation used in our implementation. A fuller description can be found in [25].

In the RCCE method, the chemistry is represented by a reduced number of n_r (with typically $n_r \ll n_s$) represented scalars or constraints. This *reduced representation* of chemistry is denoted by an n_r-vector $\mathbf{r}$.

In our implementation of RCCE, to represent the chemistry using a reduced representation we specify a set of n_{rs} represented species selected from the full set of n_s species present in the chemical mechanism. The reduced representation of chemistry is given as $\mathbf{r} = \{\mathbf{z}^r, \mathbf{z}^{u,e}\}$, where the n_{rs}-vector $\mathbf{z}^r$ denotes the species specific moles of the represented species and the n_e-vector $\mathbf{z}^{u,e}$ denotes the specific moles of elements in the unrepresented species (for element conservation). Thus the chemistry is represented in a reduced dimension of size $n_r = n_{rs} + n_e$ instead of the full dimension n_s. At any time t, the reduced representation $\mathbf{r}$ is related to the full representation $\mathbf{z}$ as

$$\mathbf{r}(t) = \mathbf{B}^T \mathbf{z}(t), \tag{3}$$

where $\mathbf{B}$ is a fixed $n_s \times n_r$ matrix determined by the choice of the represented species.

It is often convenient to view the n_s-dimensional composition space to be composed of the n_r-dimensional represented subspace (spanned by the columns of the constraint matrix $\mathbf{B}$ – size $n_s \times n_r$) and its orthogonal complement the n_u-dimensional *unrepresented* subspace (with $n_u = n_s - n_r$), such that together the represented and unrepresented subspaces span the entire composition space. This helps visualise the n_s-dimensional composition space in a 2D sketch (for example see Figure 1) indicated by the represented (denoted by **r**) and unrepresented (denoted by **u**) subspaces.

In general the reduced representation $\mathbf{r}$ can be any linear or nonlinear function of the full representation $\mathbf{z}$ as described in [31, 35]. However, our choice of the reduced representation for RCCE as described above makes the user interface very simple – the user only needs to specify a set of represented species and the rest is taken care of by the implementation. (This simple interface for RCCE has also been incorporated in the commercial CFD package ANSYS Fluent [36, 37].) In addition we have also developed an automated Greedy Algorithm with Local Improvement (GALI) [25], which can be used to select 'good' represented species for the RCCE method. Selection of good represented species or constraints is crucial for the overall accuracy of the RCCE method [27, 31]. In the remainder of the text, we use the simple reduced representation given by Equation (3) to describe different implementations of the RCCE method.

When using the RCCE method to represent chemistry in PDF based computations, the particles carry only the reduced representation of chemistry given by $\mathbf{r}$. In the reaction fractional step, given the initial reduced representation $\mathbf{r}(0)$, the task is to compute the *reduced reaction mapping* $\mathbf{r}(t)$ at the end of the reaction fractional time step t. In the following subsections we discuss the different implementations of RCCE used to compute the reduced reaction mapping.

Figure 1. Sketch of the composition space (indicated by represented **r** and unrepresented **u** sub-spaces) illustrating the reaction mapping computation using the RCCE/TIFS implementation. Given the initial reduced composition at $t = 0$ denoted by $\mathbf{r}(0)$, the reaction mapping $\mathbf{r}(t)$ is obtained in three steps (1) computing the constrained-equilibrium composition at $\mathbf{r}(0)$ denoted by $\mathbf{z}^{CE}(0)$; followed by (2) integrating the trajectory in full space (TIFS) to obtain $\mathbf{z}(t)$; followed by (3) reduction $\mathbf{r}(t) = \mathbf{B}^{T}\mathbf{z}(t)$.

3.1. *RCCE implementation using trajectory in full space*

Here we briefly describe our implementation of RCCE. A more detailed description can be found in [25, 27].

In our implementation, the reduced reaction mapping is computed by following the three steps (which are illustrated in the Figure 1 sketch) listed below:

(1) *species reconstruction*: given the reduced representation $\mathbf{r}(0)$, the constrained-equilibrium composition (at constant enthalpy h) is computed using CEQ [38] and is denoted as

$$\mathbf{z}^{CE}(0) \equiv \mathbf{z}^{CE}(\mathbf{r}(0), h), \tag{4}$$

where $\mathbf{z}^{CE}$ is a point on the Constrained-Equilibrium Manifold (CEM);
(2) *trajectory in full space*: starting from $\mathbf{z}^{CE}(0)$, the reaction trajectory (given by Equation 1) is integrated in the full space to obtain the reaction mapping $\mathbf{z}(t)$;
(3) *reduction*: from $\mathbf{z}(t)$, the *reduced reaction mapping* is obtained as $\mathbf{r}(t) = \mathbf{B}^{T}\mathbf{z}(t)$.

We henceforth refer to our implementation of RCCE as RCCE/TIFS – Trajectory In Full Space.

It is important to note here that the CEM is not an *invariant manifold*, i.e. a reaction trajectory originating from a point on the CEM does not necessarily remain on the manifold. Thus the reaction mapping $\mathbf{z}(t)$ obtained in step (2) of the above RCCE/TIFS implementation need not be on the CEM (as depicted in the Figure 1 sketch). However, it should be appreciated that the primary objective of the RCCE method being considered here is to obtain an accurate reduced reaction mapping $\mathbf{r}(t)$, and it does not matter if the reaction mapping in the full composition space, $\mathbf{z}(t)$, is not on the CEM. In fact, there exists an infinite number of compositions in the full composition space that yield the same reduced reaction mapping given by $\mathbf{r}(t) = \mathbf{B}^{T}\mathbf{z}(t)$. Here, using the TIFS approach, we seek to obtain a reaction mapping $\mathbf{z}(t)$ that yields an accurate reduced reaction mapping $\mathbf{r}(t)$.

3.2. *Classical RCCE implementation*

3.2.1. *Rate-equations for the constraints*

Here we describe the classical implementation of RCCE first introduced in [12], which involves solving a reduced set of n_r rate-equations for the constraints to compute the reduced reaction mapping $\mathbf{r}(t)$.

Given a composition $\mathbf{z}(t)$ at an instant of time, the reduced composition is given using Equation (3) as $\mathbf{r}(t) = \mathbf{B}^T\mathbf{z}(t)$. From this relation we get the rate-of-change of $\mathbf{r}(t)$ as

$$\frac{d\mathbf{r}}{dt} = \mathbf{B}^T\frac{d\mathbf{z}}{dt} = \mathbf{B}^T\mathbf{S}(\mathbf{z}). \tag{5}$$

Denoting the right-hand-side source vector by $\dot{\mathbf{r}}^e(\mathbf{z}) \equiv \mathbf{B}^T\mathbf{S}(\mathbf{z})$, we get the exact rate-equations for $\mathbf{r}$ as

$$\frac{d\mathbf{r}}{dt} = \dot{\mathbf{r}}^e(\mathbf{z}). \tag{6}$$

The classical RCCE implementation seeks to solve directly for the constraints $\mathbf{r}$ using a set of rate-equations based on Equation (6). To solve the rate-equations (6) explicitly for the constraints $\mathbf{r}$, we need a closure for the exact source vector on the right-hand-side $\dot{\mathbf{r}}^e(\mathbf{z})$ such that

$$\frac{d\mathbf{r}}{dt} = \dot{\mathbf{r}}(\mathbf{r}), \tag{7}$$

where we denote the right-hand-side approximated source vector by $\dot{\mathbf{r}}$ and henceforth refer to it simply as the *source vector*.

In the classical RCCE method [13], this closure is provided by assuming that $\mathbf{z}$ is always on the CEM – [13, equations 5.11 to 5.17] – i.e. $\mathbf{z} \equiv \mathbf{z}^{CE}(\mathbf{r})$, which yields

$$\dot{\mathbf{r}} \equiv \dot{\mathbf{r}}^{CE} \equiv \mathbf{B}^T\mathbf{S}(\mathbf{z}^{CE}(\mathbf{r})), \tag{8}$$

and gives an explicit set of rate-equations for the constraints $\mathbf{r}$ as

$$\frac{d\mathbf{r}}{dt} = \mathbf{B}^T\mathbf{S}(\mathbf{z}^{CE}(\mathbf{r})). \tag{9}$$

Now given $\mathbf{r}(0)$, the reduced reaction mapping $\mathbf{r}(t)$ can be obtained by directly integrating the reduced set of rate-equations given by Equation (9). We henceforth refer to this method as simply RCCE.

The closure provided by Equation (8) in the classical RCCE implementation [13] appears to be a simple and straightforward result. However, there is a logical flaw in the result given by Equation (8), which is based on inconsistent premises, namely

(1) $\mathbf{z}$ remains on the CEM, i.e. $\mathbf{z} \equiv \mathbf{z}^{CE}(\mathbf{r})$; and
(2) $d\mathbf{z}/dt = \mathbf{S}(\mathbf{z}^{CE}(\mathbf{r}))$.

The composition $\mathbf{z}$ remains on the CEM only if the chemical source term $\mathbf{S}(\mathbf{z}^{CE}(\mathbf{r}))$ is entirely in the CEM, i.e. the CEM is an *invariant manifold*. However, since the CEM is not an invariant manifold, the reaction trajectory starting from a point on the CEM does not

Figure 2.   Sketch of the composition space (indicated by represented **r** and unrepresented **u** sub-spaces) illustrating the projections involved in the RCCE method. Given a reduced composition denoted by **r**, the classical RCCE implementation computes the source vector $\dot{\mathbf{r}}$ by computing the chemical source term **S** at the constrained-equilibrium composition $\mathbf{z}^{CE}$ (on the CEM) and then projecting it back to the represented subspace giving $\dot{\mathbf{r}} = \mathbf{B}^T\mathbf{S}$. This implementation does not take into account the non-invariance of the CEM manifold. Alternatively, one could consider a projector denoted by **P**, which first projects the source vector **S** onto the tangent plane of the CEM (denoted by **T**) to account for the non-invariance, before projecting it back to the represented subspace to yield $\dot{\mathbf{r}} = \mathbf{B}^T\mathbf{P}\mathbf{S}$.

necessarily remain on the manifold, which makes the above RCCE assumptions inherently flawed.

This hidden flawed assumption implied by Equation (8) was first exposed in [32]. The notation used in [32] is different from our notation, and so for consistency we explain again this hidden assumption using our own notation. In the RCCE method, by solving a reduced system of rate-equations for the constraints, the full composition **z** is assumed to remain on the CEM. Thus, to obtain the source vector $\dot{\mathbf{r}}$, the chemical source term **S** is being implicitly *projected* onto the CEM. In the closure provided by Equation (8), the chemical source term **S** is implicitly being projected in the unrepresented subspace, i.e. the orthogonal complement of the represented subspace spanned by the columns of **B**. This hidden projection is illustrated in the Figure 2 sketch. It is not obvious if this is an accurate projection, and in fact it is shown in [32] that a more accurate projection is obtained by computing the reaction source vector on a Close-Parallel Inertial Manifold (CPIM). We present results in later sections which will confirm that the source vector approximation provided by the RCCE method is not accurate.

3.2.2.   *Rate-equations for the constraint potentials*

Numerical integration of the rate-equations for the constraints given by Equation (9) requires the computation of the constrained-equilibrium composition, $\mathbf{z}^{CE}(\mathbf{r})$, at each sub-step of the integration (as described in more detail in Section 7.4). This makes the numerical integration of Equation (9) expensive.

To reduce the computational cost, an alternative implementation of RCCE is described in [13, section 5.3] using the *rate-equations for the Lagrange multipliers*. This alternative implementation (also referred to as the *rate-equations for the constraint potentials*) transforms the rate-equations for the constraints (given by Equation 9) into rate-equations for the constraint potentials, thereby solving directly for the constraint potentials on the CEM and avoiding the need for computing the constrained-equilibrium composition at each sub-step

of the integration. This implementation has been further developed, implemented and tested in [28, 29, 31].

It is important to note here that the above two implementations of RCCE are *mathematically equivalent* and both the implementations make use of the closure provided by Equation (8) – see [13, equations 4.11, 5.19 and 5.20]. Numerically, however, the solution obtained by the two implementations may differ due to ODE integration errors. As the numerical integration sub-step time size approaches zero (i.e. for very small ODE integration error tolerance), the two implementations should yield the same solution and thus are consistent with each other. (This has been discussed in a previous work: refer to [23, Figure 1 and the accompanying discussion].) As mentioned later in Section 7.4, in this work we use DDASAC for ODE integration with a relatively small error tolerance of 10^{-8}. Hence we expect both the implementations to yield similar solutions (within the ODE integration error tolerance).

To assess the accuracy of the classical RCCE implementation, either of the aforementioned implementations can be used. In this work, we use the former implementation of RCCE, i.e. the implementation using the rate-equations for the constraints, which will henceforth be simply referred to as the RCCE implementation.

Unlike the RCCE/TIFS implementation, both the aforementioned implementations of RCCE (implicitly) attempt to follow the CEM in the full composition space accurately to obtain the reaction mapping $\mathbf{r}(t)$. However, as will be shown in later sections, the projection (Equation 8) used in the above implementations of RCCE yields an inaccurate reaction mapping.

In the next section we describe the CPIM method and then in the following section we describe the new Reaction-mixing Attracting Manifold Projector (RAMP), as an extension of the CPIM method, which provides a more accurate projection for implementing the RCCE method.

4.   Accurate projection for the RCCE method

It is clear that the exact rate-equations for the constraints are given by Equation (6); however, we need a closure for the right-hand-side source vector denoted by $\dot{\mathbf{r}}$.

In the previous section we showed that the closure provided by the RCCE method is inaccurate because the CEM is not an invariant manifold. One way to account for the non-invariance of the CEM is to replace the chemical source term $\mathbf{S}$ in Equation (8) by a projection of $\mathbf{S}$ onto the tangent plane of the CEM. Then, consistently, the composition remains on the CEM.

To this end, let us consider a general $n_s \times n_s$ projection matrix $\mathbf{P}$ such that the source vector $\dot{\mathbf{r}}$ is obtained (as shown in Figure 2) by

$$\dot{\mathbf{r}} = \mathbf{B}^{\mathrm{T}}\mathbf{PS}. \tag{10}$$

There are various choices available for the projection $\mathbf{P}$:

(1) project in the unrepresented subspace, as in RCCE;
(2) project in the CEM normal subspace;
(3) project in the 'fast' subspace, as given by ILDM [15] or CSP [14];

(4) project using a thermodynamic projector [39];

(5) project using a close-parallel manifold, as given by CPIM [32].

Projection (1) implied by RCCE is not accurate as will be shown in later sections. Projection (2) in the CEM normal subspace depends on the scaling of represented variables and thus contains arbitrariness. Different projections can be obtained if species composition is represented using mass fractions instead of specific moles. Projections (3) and (4) provide more accurate projections; however, their implementation is expensive and quite involved. Projection (5) provides a simple correction for the non-invariance of the CEM and an accurate projection, but has issues involving unrealizability and negative entropy production [32].

The projections given by ILDM, CSP and CPIM are based solely on thermochemistry. This can become problematic when reactions are not fast (e.g. at low temperatures) leading to unrealizability and negative entropy production. The Reaction-mixing Attracting Manifold Projector (RAMP) method introduced below helps address some of these issues.

Here we first present a mathematical formulation for the CPIM method (similar to that provided in [32], however using our own notation), and then introduce the RAMP method.

4.1. *Subspaces and projections*

In the following sections we work with subspaces and projections and so before proceeding further we describe here the notation used for denoting subspaces.

The columns of the constraint matrix $\mathbf{B}$ are not necessarily orthogonal, and it is more convenient to work in terms of orthonormal basis vectors for computing projections. Using the QR or SVD factorisation of $\mathbf{B}$ we obtain: (i) a set of n_r orthonormal basis vectors (for the represented subspace) denoted by matrix $\mathbf{R}$ of size $n_s \times n_r$ such that $\mathrm{span}(\mathbf{R}) = \mathrm{span}(\mathbf{B})$; and (ii) a set of n_u orthonormal basis vectors (for the orthogonal complement unrepresented subspace) denoted by matrix $\mathbf{U}$ of size $n_s \times n_u$ such that $\mathrm{span}(\mathbf{U}) = \mathrm{span}(\mathbf{B})^{\perp}$.

Another important subspace is formed by the tangent vectors of the CEM. Consider an $n_s \times n_r$ matrix $\mathbf{T}^{\mathrm{CEM}}$ whose columns are the tangent vectors of the CEM such that

$$\mathbf{T}^{\mathrm{CEM}} = \frac{\partial \mathbf{z}^{\mathrm{CE}}}{\partial \mathbf{r}}. \tag{11}$$

The column vectors of $\mathbf{T}^{\mathrm{CEM}}$ span the tangent subspace of the CEM; however, the column vectors need not be orthogonal. Using the QR or SVD factorisation of $\mathbf{T}^{\mathrm{CEM}}$ we obtain: (i) a set of n_r orthonormal basis vectors (for the CEM tangent subspace) denoted by matrix $\mathbf{T}$ of size $n_s \times n_r$ such that $\mathrm{span}(\mathbf{T}) = \mathrm{span}(\mathbf{T}^{\mathrm{CEM}})$; and (ii) a set of n_u orthonormal basis vectors (for the orthogonal complement CEM normal subspace) denoted by matrix $\mathbf{N}$ of size $n_s \times n_u$ such that $\mathrm{span}(\mathbf{N}) = \mathrm{span}(\mathbf{T}^{\mathrm{CEM}})^{\perp}$.

We now have two sets of subspaces (or basis vectors) which together span the full n_s-dimensional composition space

(1) the represented–unrepresented subspaces spanned by $\{\mathbf{R}, \mathbf{U}\}$, respectively;

(2) the CEM tangent-normal subspaces spanned by $\{\mathbf{T}, \mathbf{N}\}$, respectively.

It should be noted that the represented–unrepresented subspaces are fixed once the matrix $\mathbf{B}$ is known, whereas the CEM tangent-normal subspaces are local to the CEM and vary as we move along the CEM.

Figure 3. Sketch of the composition space (indicated by represented **r** and unrepresented **u** subspaces) illustrating the projections involved in the CPIM method. Given a reduced composition **r**, the RCCE implementation computes the source vector $\dot{\mathbf{r}}^{\mathrm{CE}}$ by projecting the chemical source term $\mathbf{S}^{\mathrm{CE}}$ computed at the constrained-equilibrium composition $\mathbf{z}^{\mathrm{CE}}$ onto the represented space yielding $\dot{\mathbf{r}}^{\mathrm{CE}} = \mathbf{B}^{\mathrm{T}}\mathbf{S}^{\mathrm{CE}}$. Alternatively, in the CPIM method, a manifold close-and-parallel to the CEM is considered to evaluate the chemical source term $\mathbf{S}^{\mathrm{CP}}$ (which lies entirely in the CEM and CPIM tangent space denoted by **T**). In the CPIM method, the source vector is given by $\dot{\mathbf{r}}^{\mathrm{CP}} = \mathbf{B}^{\mathrm{T}}\mathbf{S}^{\mathrm{CP}}$, which can be rewritten in terms of a projector **P** such that $\dot{\mathbf{r}}^{\mathrm{CP}} = \mathbf{B}^{\mathrm{T}}\mathbf{P}\mathbf{S}^{\mathrm{CE}}$.

Any n_{s}-vector **x** in the full space can be decomposed in the following two ways:

$$\mathbf{x} = \mathbf{R}\left(\mathbf{R}^{\mathrm{T}}\mathbf{x}\right) + \mathbf{U}\left(\mathbf{U}^{\mathrm{T}}\mathbf{x}\right), \tag{12}$$

or

$$\mathbf{x} = \mathbf{T}\left(\mathbf{T}^{\mathrm{T}}\mathbf{x}\right) + \mathbf{N}\left(\mathbf{N}^{\mathrm{T}}\mathbf{x}\right). \tag{13}$$

4.2. *Close-parallel inertial manifold*

In the CPIM method, it is hypothesised that there is a manifold close-and-parallel to the CEM which is invariant with respect to

$$\frac{\mathrm{d}\mathbf{z}^{\mathrm{CP}}(t)}{\mathrm{d}t} = \mathbf{S}\left(\mathbf{z}^{\mathrm{CP}}(t)\right), \tag{14}$$

where $\mathbf{z}^{\mathrm{CP}}$ is a point on the manifold. This means that any reaction trajectory originating from a point on the manifold, by hypothesis, remains on the manifold.

The steps involved in the CPIM method are illustrated in Figure 3. For a given constraint **r**, the constrained-equilibrium composition on the CEM is denoted by $\mathbf{z}^{\mathrm{CE}} \equiv \mathbf{z}^{\mathrm{CE}}(\mathbf{r})$. The chemical source term at $\mathbf{z}^{\mathrm{CE}}$ is denoted by $\mathbf{S}^{\mathrm{CE}} \equiv \mathbf{S}\left(\mathbf{z}^{\mathrm{CE}}\right)$.

Now we consider a point $\mathbf{z}^{\mathrm{CP}}$ on the CPIM such that

$$\mathbf{B}^{\mathrm{T}}\mathbf{z}^{\mathrm{CP}} = \mathbf{B}^{\mathrm{T}}\mathbf{z}^{\mathrm{CE}} = \mathbf{r}, \tag{15}$$

and we denote

$$\delta\mathbf{z}^{\mathrm{CP}} \equiv \mathbf{z}^{\mathrm{CP}} - \mathbf{z}^{\mathrm{CE}}. \tag{16}$$

Since $\mathbf{B}^{\mathrm{T}}\delta\mathbf{z}^{\mathrm{CP}} = 0$, this means that $\delta\mathbf{z}^{\mathrm{CP}}$ lies entirely in the unrepresented subspace, and so we can express

$$\delta\mathbf{z}^{\mathrm{CP}} = \mathbf{U}\delta\mathbf{u}^{\mathrm{CP}}, \tag{17}$$

where $\delta\mathbf{u}^{\mathrm{CP}}$ is an n_{u}-vector in the unrepresented subspace spanned by $\mathbf{U}$.

By the close-and-parallel assumption, $\mathbf{z}^{\mathrm{CP}}$ is close to $\mathbf{z}^{\mathrm{CE}}$, and so we can express $\mathbf{S}^{\mathrm{CP}} \equiv \mathbf{S}\left(\mathbf{z}^{\mathrm{CP}}\right)$ using a linear approximation about $\mathbf{S}^{\mathrm{CE}}$ as follows:

$$\mathbf{S}^{\mathrm{CP}} = \mathbf{S}^{\mathrm{CE}} + \mathbf{J}\delta\mathbf{z}^{\mathrm{CP}} = \mathbf{S}^{\mathrm{CE}} + \mathbf{J}\mathbf{U}\delta\mathbf{u}^{\mathrm{CP}}, \tag{18}$$

where $\mathbf{J}$ is the $n_{\mathrm{s}} \times n_{\mathrm{s}}$ Jacobian evaluated at $\mathbf{z}^{\mathrm{CE}}$

$$\mathbf{J} \equiv \frac{\partial\mathbf{S}\left(\mathbf{z}^{\mathrm{CE}}\right)}{\partial\mathbf{z}}. \tag{19}$$

Since the CPIM is invariant and parallel to CEM, at $\mathbf{z}^{\mathrm{CP}}$ the chemical source term $\mathbf{S}^{\mathrm{CP}} \equiv \mathbf{S}(\mathbf{z}^{\mathrm{CP}})$ must be in the CEM tangent subspace. This enforces that $\mathbf{N}^{\mathrm{T}}\mathbf{S}^{\mathrm{CP}} = \mathbf{0}$, which gives

$$\mathbf{N}^{\mathrm{T}}\mathbf{S}^{\mathrm{CE}} + \mathbf{N}^{\mathrm{T}}\mathbf{J}\mathbf{U}\delta\mathbf{u}^{\mathrm{CP}} = \mathbf{0}, \tag{20}$$

and solving for $\delta\mathbf{u}^{\mathrm{CP}}$ we obtain

$$\delta\mathbf{u}^{\mathrm{CP}} = \mathbf{L}^{-1}\mathbf{N}^{\mathrm{T}}\mathbf{S}^{\mathrm{CE}}, \tag{21}$$

where we denote

$$\mathbf{L} \equiv -\mathbf{N}^{\mathrm{T}}\mathbf{J}\mathbf{U}. \tag{22}$$

Substituting $\delta\mathbf{u}^{\mathrm{CP}}$ in Equation (17) we obtain

$$\mathbf{z}^{\mathrm{CP}} = \mathbf{z}^{\mathrm{CE}} + \mathbf{U}\mathbf{L}^{-1}\mathbf{N}^{\mathrm{T}}\mathbf{S}^{\mathrm{CE}}, \tag{23}$$

and using Equation (18) we get

$$\mathbf{S}^{\mathrm{CP}} = (\mathbf{I} + \mathbf{J}\mathbf{U}\mathbf{L}^{-1}\mathbf{N}^{\mathrm{T}})\mathbf{S}^{\mathrm{CE}}. \tag{24}$$

We now use $\mathbf{S}^{\mathrm{CP}}$ to compute the source vector $\dot{\mathbf{r}}$ in the rate-equations for the constraints equation (7) and denote it as

$$\dot{\mathbf{r}} \equiv \dot{\mathbf{r}}^{\mathrm{CP}} \equiv \mathbf{B}^{\mathrm{T}}\mathbf{S}^{\mathrm{CP}}. \tag{25}$$

Thus the rate-equations for the constraints equation (7) using the CPIM approach are given as

$$\frac{\mathrm{d}\mathbf{r}}{\mathrm{d}t} = \mathbf{B}^{\mathrm{T}}\mathbf{S}^{\mathrm{CP}}, \tag{26}$$

which can be re-written as

$$\frac{d\mathbf{r}}{dt} = \mathbf{B}^{\mathrm{T}}\mathbf{PS}\left(\mathbf{z}^{\mathrm{CE}}(\mathbf{r})\right), \tag{27}$$

where

$$\mathbf{P} \equiv \mathbf{I} + \mathbf{JUL}^{-1}\mathbf{N}^{\mathrm{T}}, \tag{28}$$

provides a more accurate projection of the chemical source term evaluated on the CEM onto the tangent subspace.

The main issues involved in the CPIM method as highlighted in [32] are:

(1) if the matrix $\mathbf{L}$ is ill-conditioned (especially at low temperatures) then the composition $\mathbf{z}^{\mathrm{CP}}$ can be unrealizable;
(2) the linear approximation of $\mathbf{S}^{\mathrm{CP}}$ (for ill-conditioned matrix $\mathbf{L}$) may lead to negative entropy production.

These issues were handled in the CPIM method [32] by considering a linear combination of $\mathbf{S}^{\mathrm{CP}}$ and $\mathbf{S}^{\mathrm{CE}}$ to ensure positive entropy production and realizability.

Here we present the Reaction-mixing Attracting Manifold Projector (RAMP) [7] as an extension to the CPIM approach, which provides a much simpler way of handling the unrealizability and negative entropy production issues.

5. Reaction-mixing attracting manifold projector

In the RAMP approach, similar to CPIM, a hypothetical manifold close-and-parallel to the CEM is considered; however, now this manifold is assumed to be invariant with respect to the following evolution equation:

$$\frac{d\mathbf{z}^{\mathrm{CP}}(t)}{dt} = \mathbf{S}\left(\mathbf{z}^{\mathrm{CP}}(t)\right) - \omega\left(\mathbf{z}^{\mathrm{CP}}(t) - \mathbf{z}^{\mathrm{CE}}(t)\right), \tag{29}$$

where $\mathbf{z}^{\mathrm{CP}}(t)$ is a point on the CPIM; $\mathbf{z}^{\mathrm{CE}}(t) \equiv \mathbf{z}^{\mathrm{CE}}\left(\mathbf{r} = \mathbf{B}^{\mathrm{T}}\mathbf{z}^{\mathrm{CP}}(t)\right)$ is a point on the CEM such that $\mathbf{B}^{\mathrm{T}}\mathbf{z}^{\mathrm{CP}} = \mathbf{B}^{\mathrm{T}}\mathbf{z}^{\mathrm{CE}}$; and ω is a specified mixing (relaxation) rate.

The inclusion of the additional mixing term in the evolution equation (29) is inspired by the general class of reaction–diffusion manifolds described in [7], and in particular the REDIM method [40]. The inclusion of the additional mixing term helps address the realizability issues encountered in the CPIM approach, where the evolution equation (14) contains only the reaction term.

For the RAMP approach, the invariance condition for the evolution equation (29) is given as

$$\mathbf{N}^{\mathrm{T}}\left[\frac{d\mathbf{z}^{\mathrm{CP}}(t)}{dt}\right] = \mathbf{0}. \tag{30}$$

Using the same notation as for the CPIM, we can express

$$\frac{d\mathbf{z}^{\mathrm{CP}}(t)}{dt} = \mathbf{S}^{\mathrm{CP}} - \omega\delta\mathbf{z}^{\mathrm{CP}}. \tag{31}$$

Following the same steps as for the CPIM method we can express $\mathbf{S}^{CP}$ by a linear approximation given by Equation (18), and $\delta\mathbf{z}^{CP} = \mathbf{U}\delta\mathbf{u}^{CP}$, which gives

$$\frac{d\mathbf{z}^{CP}(t)}{dt} = \mathbf{S}^{CE} + (\mathbf{J} - \omega\mathbf{I})\,\mathbf{U}\delta\mathbf{u}^{CP}. \tag{32}$$

Using the invariance condition we obtain

$$\mathbf{N}^T\mathbf{S}^{CE} + \mathbf{N}^T\,(\mathbf{J} - \omega\mathbf{I})\,\mathbf{U}\delta\mathbf{u}^{CP} = \mathbf{0}, \tag{33}$$

which gives

$$\delta\mathbf{u}^{CP} = \mathbf{L}^{-1}\mathbf{N}^T\mathbf{S}^{CE}, \tag{34}$$

where

$$\mathbf{L} \equiv \mathbf{L}(\omega) \equiv -\mathbf{N}^T\,(\mathbf{J} - \omega\mathbf{I})\,\mathbf{U}. \tag{35}$$

So the only difference from the CPIM method is the additional $\omega\mathbf{I}$ term in the definition of $\mathbf{L}$. This term makes $\mathbf{L}$ better conditioned, because by assumption if the close-parallel manifold is an *attracting* manifold (as is implicitly assumed), then the eigenvalues of the Jacobian $\mathbf{J}$ have negative real parts, which are further decreased by ω due to the introduction of the relaxation term in the evolution equation.

The new projector with the RAMP method (which is a function of the relaxation parameter ω) is given as

$$\mathbf{P}(\omega) \equiv \mathbf{I} + \mathbf{JUL}^{-1}\mathbf{N}^T, \tag{36}$$

with $\mathbf{L}$ given by Equation (35).

The source vector approximation given by the RAMP method is denoted as

$$\dot{\mathbf{r}} \equiv \dot{\mathbf{r}}^{CP}(\omega) \equiv \mathbf{B}^T\mathbf{P}(\omega)\mathbf{S}\left(\mathbf{z}^{CE}(\mathbf{r})\right). \tag{37}$$

The rate-equations for the constraints (7) using the RAMP approach are given as

$$\frac{d\mathbf{r}}{dt} = \mathbf{B}^T\mathbf{P}(\omega)\mathbf{S}(\mathbf{z}^{CE}(\mathbf{r})). \tag{38}$$

It is important to highlight here the following properties of the RAMP method for the two limiting values of the relaxation rate parameter ω.

(1) $\boldsymbol{\omega = 0}$. From the definition of the evolution equation (29), it is obvious that $\omega = 0$ corresponds to the CPIM approach, with the projector $\mathbf{P}(\omega = 0) = \mathbf{P}$ provided by the CPIM method.
(2) $\boldsymbol{\omega \to \infty}$. From the definition of $\mathbf{L}$ (Equation 35), we see that $\mathbf{L}^{-1}(\omega \to \infty) = \mathbf{0}$, i.e. the inverse of $\mathbf{L}$ is singular. This yields, from Equation (36), that $\mathbf{P}(\omega \to \infty) = \mathbf{I}$, and the RAMP method corresponds to the classical RCCE implementation, since $\dot{\mathbf{r}}^{CP}(\omega \to \infty) = \mathbf{B}^T\mathbf{P}(\omega \to \infty)\mathbf{S}^{CE} = \mathbf{B}^T\mathbf{S}^{CE} = \dot{\mathbf{r}}^{CE}$.

Figure 4. Sketch of the composition space (indicated by represented **r** and unrepresented **u** subspaces) illustrating reaction mapping computation using the three implementations of the RCCE method: RCCE/TIFS, RCCE and RCCE/RAMP. Given the initial reduced composition **r**(0), (1) the RCCE/TIFS implementation computes the reaction mapping by following the trajectory in full space starting from $\mathbf{z}^{CE}(0)$ to obtain $\mathbf{z}(t)$ followed by reduction to yield the reaction mapping $\mathbf{r}(t)$ [RCCE/TIFS] = $\mathbf{B}^T\mathbf{z}(t)$; (2) the RCCE implementation solves a reduced system of ODEs with source vector $\dot{\mathbf{r}} = \mathbf{B}^T\mathbf{S}$ to obtain $\mathbf{r}(t)$ [RCCE]; and (3) the RCCE/RAMP implementation solves the reduced system of ODEs (as in RCCE) using an alternative source vector $\dot{\mathbf{r}} = \mathbf{B}^T\mathbf{PS}$ to obtain $\mathbf{r}(t)$ [RCCE/RAMP].

In short this shows that the value of ω controls the closeness of the CPIM to the CEM in the RAMP approach. For $\omega = 0$, we obtain the CPIM method, and for $\omega \to \infty$, the CPIM collapses with the CEM and we get back the classical RCCE implementation. We will see in the results presented in the following sections that, by choosing an appropriate value of ω in the RAMP approach, we can address the unrealizability issues of the CPIM method and also account for the non-invariance of the CEM-based RCCE method.

6.   Partially-stirred reactor

In this section we describe the partially-stirred reactor (PaSR) test case used to study the three implementations of RCCE (illustrated in Figure 4) described in this work:

(1) RCCE/TIFS: the implementation used in our previous works (described in Section 3.1);
(2) RCCE: the classical RCCE implementation (described in Section 3.2.1); and
(3) RCCE/RAMP: the new implementation based on CPIM (described in Section 5).

To compare these different implementations of the RCCE method, we consider methane/air premixed combustion in a partially-stirred reactor (PaSR). A detailed description of the PaSR is provided in [18]: here we briefly describe the details pertinent to the current study.

The PaSR can be used to study PDF particle implementation applied to a statistically homogeneous flow. In the PaSR, the thermochemical composition is represented by a fixed number of particles. The particle composition evolves due to *mixing* and *reaction* in fractional steps. The mixing is implemented using a pairwise mixing model [18]. Our implementation of PaSR allows the particle chemistry to be represented in any one of the following ways:

Figure 5. Sketch depicting a test composition $\mathbf{z}(0)$ and its reaction mapping $\mathbf{z}(t)$ in the full composition space (indicated by the represented $\mathbf{r}$ and unrepresented $\mathbf{u}$ subspaces). The reaction makes the test composition move closer to a low-dimensional attracting manifold. In the represented subspace, $\dot{\mathbf{r}}^{e} = \mathbf{B}^{T}\mathbf{S}$ denotes the exact source vector and $\dot{\mathbf{r}}$ denotes a source vector approximation obtained using one of the RCCE implementations. In addition, $\mathbf{r}^{e}(t) = \mathbf{B}^{T}\mathbf{z}(t)$ denotes the exact reaction mapping and $\mathbf{r}(t)$ denotes the reaction mapping obtained using one of the RCCE implementations.

(1) detailed chemistry with n_{s} species, with reaction mapping computed using ODE integration (also referred to as *direct evaluation*);
(2) detailed chemistry with n_{s} species, with reaction mapping computed using ISAT;
(3) reduced chemistry with n_{r} represented variables using RCCE, with reaction mapping computed using the combined ISAT/RCCE approach [25].

For methane/air combustion in the PaSR, we use the same operating conditions as in [25]. The PaSR involves two inflowing streams: (1) a stoichiometric premixed stream of methane/air mixture at 600 K; and (2) a pilot stream of equilibrium products of composition of stream 1. The streams flow into the PaSR with a mass flow rate ratio of 0.95:0.05. Initially all the particles are set to the pilot stream composition. The pressure is atmospheric throughout. Other important parameters include: the number of particles, $N_{p} = 100$; the PaSR residence time, $\tau_{res} = 20$ ms;[1] the PaSR mixing time scale, $\tau_{mix} = 1$ ms; and the PaSR pairing time scale, $\tau_{pair} = 1$ ms. The chemistry is represented using the GRI-Mech 1.2 mechanism involving $n_{s} = 31$ species composed of $n_{e} = 4$ elements.

We perform a PaSR simulation with a reaction time step, $\Delta t = 0.033$ ms, and *reaction mapping* computed using ISAT (with error tolerance, $\epsilon_{tol} = 10^{-5}$) with the detailed mechanism. During the simulation, we save the compositions of the first N particles that result in an *add* in the ISAT table (and so are distinct) denoted by $\mathbf{z}^{(n)}(0)$ and their reaction mappings $\mathbf{z}^{(n)}(\Delta t)$ for $n = 1$ to N. We use these test compositions to study different implementations of the RCCE method. In this work, in all the tests, we use $N = 2500$ test compositions. Henceforth, we use t to denote a general reaction time step, and Δt to denote the exact time step $\Delta t = 0.033$ ms used in the PaSR test to compute the reaction mappings.

It is important to note here that, in the PaSR, *reaction* causes the particle compositions to move towards a low-dimensional attracting manifold, and *mixing* causes the particle compositions to be pulled away from this manifold (as illustrated in Figure 5). Hence the test compositions at the beginning of the reaction fractional step $\mathbf{z}^{(n)}(0)$ are expected to be away from the attracting manifold, and the compositions at the end of the step $\mathbf{z}^{(n)}(\Delta t)$ are

Table 1.   Sets of represented species obtained using GALI (with the 31-species GRI-Mech 1.2 mechanism) for dimension reduction of methane/air premixed combustion with RCCE for $n_{rs} = 9$ to 15 (obtained from [25, table 3]).

n_{rs}	Represented species
9	CH_4, CO_2, H_2, O_2, H, OH, O, CH_2O, CH_3
10	CH_4, CO_2, H_2, O_2, H, OH, O, CH_2O, C_2H_6, C_2H_4
11	CH_4, CO_2, H_2, O_2, H, OH, O, H_2O, CH_3, HO_2, CO
12	CH_4, CO_2, H_2, O_2, H, OH, O, CH_2O, CH_3, HO_2, CO, H_2O
13	CH_4, CO_2, H_2, O_2, H, OH, O, CH_3OH, CH_3, HO_2, CO, H_2O, CH_2CO
14	CH_4, CO_2, H_2, O_2, H, OH, O, CH_2O, CH_3, HO_2, CO, H_2O, CH_2CO, C_2H_5
15	CH_4, CO_2, H_2, O_2, H, OH, O, CH_2O, CH_3, HO_2, CO, H_2O, CH_2CO, C_2H_5, CH_2

expected to be closer to a low-dimensional attracting manifold. Unless otherwise specified explicitly, in the following tests to study different implementations of RCCE we use the test compositions from the beginning of the reaction fractional step $\mathbf{z}^{(n)}(0)$, which typically will be encountered in real PDF computations. We use the test compositions $\mathbf{z}^{(n)}(\Delta t)$ only in a few cases to study the validity of the CPIM approximation.

6.1. *PaSR tests to study RCCE implementations*

We use the test compositions saved from the PaSR run in the full dimension to study different implementations of the RCCE method. We perform RCCE tests over a range of values of n_{rs} from 9 to 15, corresponding to $n_r = n_{rs} + n_e$ in the range 13 to 19 (which yield less than 3% *reduction–tabulation error* [25]) with represented species selected using GALI. We use the same represented species as those listed in [25, table 3], which are obtained using GALI for the same test case – methane/air premixed combustion in PaSR with chemistry represented using the 31-species GRI-Mech 1.2 mechanism – as used in the current study. The relevant sets of represented species for $n_{rs} = 9$ to 15 used in this study are listed again in Table 1.

Now for a specified set of represented species for performing dimension reduction with RCCE, we form the constraint matrix $\mathbf{B}$. At each selected particle composition $\mathbf{z} \equiv \mathbf{z}^{(n)}(0)$ (and in some cases $\mathbf{z} \equiv \mathbf{z}^{(n)}(\Delta t)$) for $n = 1$ to $N = 2500$, we then compute

- the chemical source term, $\mathbf{S} \equiv \mathbf{S(z)}$;
- the exact source vector, $\dot{\mathbf{r}}^e = \mathbf{B}^T\mathbf{S}$ (as illustrated in Figure 5);
- the reduced composition, $\mathbf{r} = \mathbf{B}^T\mathbf{z}$;
- the constrained-equilibrium composition, $\mathbf{z}^{CE} \equiv \mathbf{z}^{CE}(\mathbf{r})$;
- the chemical source term, $\mathbf{S}^{CE} \equiv \mathbf{S(z^{CE})}$;
- the orthogonal projections of $\mathbf{S}$ and $\mathbf{S}^{CE}$ onto the CEM tangent plane, denoted by $\mathbf{S}_t$ and $\mathbf{S}_t^{CE}$, respectively; and the angles between S and S_t, and S^{CE} and S_t^{CE}, denoted by $< (S, S_t)$ and $< (S^{CE}, S_t^{CE}$, respectively, to assess the non-invariance of the CEM manifold;
- the source vector given by RCCE, $\dot{\mathbf{r}} \equiv \dot{\mathbf{r}}^{CE} = \mathbf{B}^T\mathbf{S}^{CE}$;
- the source vector given by RAMP, $\dot{\mathbf{r}} \equiv \dot{\mathbf{r}}^{CP}(\omega) = \mathbf{B}^T\mathbf{P}(\omega)\mathbf{S}^{CE}$, for different values of the relaxation rate ω;
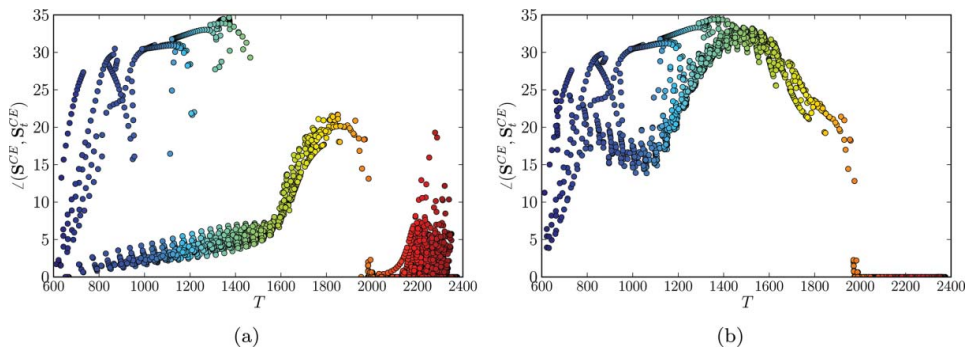
Figure 6.   Scatter plots of angle (in degrees) between the chemical source term $\mathbf{S}^{\mathrm{CE}}$ (evaluated on the CEM) and its orthogonal projection onto the CEM $\mathbf{S}_t^{CE}$ versus temperature T computed using the test compositions at $t = 0$ (a) and at $t = \Delta t$ (b).

- the exact reaction mapping $\mathbf{r}^{\mathrm{e}}(t) = \mathbf{B}^{\mathrm{T}}\mathbf{z}(t)$ and reaction mappings $\mathbf{r}(t)$ using the three implementations of RCCE (as illustrated in Figures 4 and 5).

In the following section we look at various scatter plots to analyse these data, and in addition quantify and compare the errors involved in the three implementations of RCCE.

7.   Results

7.1.   *Non-invariance*

In this section we examine the 'degree' of non-invariance of the CEM, by looking at the angle between the chemical source term $\mathbf{S}^{\mathrm{CE}}$ and its orthogonal projection $\mathbf{S}_t^{\mathrm{CE}}$ on the CEM (denoted by $< (\mathrm{S}^{\mathrm{CE}}, \mathrm{S}_t^{\mathrm{CE}})$). Figure 6 shows scatter plots of $\angle(\mathbf{S}^{\mathrm{CE}}, \mathbf{S}_t^{\mathrm{CE}})$ versus temperature T computed using the test compositions saved at $t = 0$ and $t = \Delta t$. If the CEM were an invariant manifold, then we would have $\angle(\mathbf{S}^{\mathrm{CE}}, \mathbf{S}_t^{\mathrm{CE}}) = 0$. However, in these scatter plots we see that $\angle(\mathbf{S}^{\mathrm{CE}}, \mathbf{S}_t^{\mathrm{CE}})$ is as large as $35°$ at both $t = 0$ and $t = \Delta t$. This confirms that the CEM is not an invariant manifold, and in fact the reaction trajectories could be moving away from the manifold at large angles. This non-invariance introduces a large error in the RCCE implementation, in which $\mathbf{S}^{\mathrm{CE}}$ is directly projected onto the constrained subspace without accounting for this non-invariance.

We now examine the angle between the chemical source term $\mathbf{S}$ (computed at the test composition) and its orthogonal projection onto the CEM $\mathbf{S}_t$ (denoted by $\angle(\mathbf{S}, \mathbf{S}_t)$), which gives a measure of the orientation of the chemical source term $\mathbf{S}$ relative to the CEM. Figure 7 shows a scatter plot of $\angle(\mathbf{S}, \mathbf{S}_t)$ versus temperature T computed using the test compositions saved at $t = 0$ and $t = \Delta t$. Here we see that at $t = 0$, the angle $\angle(\mathbf{S}, \mathbf{S}_t)$ is relatively large (around $10°$) because the test compositions are pulled away from the attracting manifold due to mixing. However, at $t = \Delta t$, the angle $\angle(\mathbf{S}, \mathbf{S}_t)$ is very small (less than $2°$ for $T > 1000\,\mathrm{K}$), which shows that the reaction takes the compositions closer to an attracting manifold, and this attracting manifold is nearly parallel to the CEM (because the angle is measured relative to the CEM). This observation confirms the CPIM idea that there exists an invariant manifold close-and-parallel to the CEM. This is the reason why the CPIM and its extension RAMP are able to give a better approximation for the RCCE source vector (as quantified in the results included below).

Figure 7. Scatter plot of angle (in degrees) between the chemical source term **S** and its orthogonal projection onto the CEM $\mathbf{S}_t$ versus temperature T computed using the test compositions at $t = 0$ (a) and at $t = \Delta t$ (b).

7.2. *Realizability and entropy production*

Here we briefly examine the realizability and entropy production issues highlighted in the CPIM work [32] and mentioned in Section 4.2, which are also pertinent to the extended RAMP approach.

The RCCE/TIFS implementation computes the reaction mapping following the reaction trajectory in the full space by integrating the full system of ODEs (Equation 1) using the chemical source term $\mathbf{S}(\mathbf{z})$. Since there is no projection involved in this method, there are no realizability or negative entropy production issues in this implementation.

The RCCE implementation uses a reduced system of ODEs obtained by projecting the chemical source term $\mathbf{S}(\mathbf{z})$ directly onto the constrained subspace given as $\dot{\mathbf{r}} = \mathbf{B}^{\mathrm{T}}\mathbf{S}(\mathbf{z})$. During the computation of the reaction mapping, the chemical composition evolves through a series of constrained-equilibrium compositions on the CEM, and it is shown in [13] that this implementation ensures (mathematically) non-negative entropy production and realizability.

The CPIM and RAMP approaches use an alternative projection **P** which need not necessarily ensure non-negative entropy production as described in [32]. There are two main concerns:

(1) for an ill-conditioned matrix $\mathbf{L}(\omega)$, the composition $\mathbf{z}^{\mathrm{CP}}(\omega)$ may not be realizable; and
(2) the linear approximation $\mathbf{S}^{\mathrm{CP}}(\omega)$ (for ill-conditioned matrix $\mathbf{L}(\omega)$) may lead to negative entropy production.

However, since $\mathbf{z}^{\mathrm{CP}}(\omega)$ does not directly appear in the definition of the projector $\mathbf{P}(\omega)$ given by Equation (36), the realizability of $\mathbf{z}^{\mathrm{CP}}(\omega)$ is not a major concern. It is only the linear approximation to the chemical source term $\mathbf{S}^{\mathrm{CP}}(\omega)$ that directly influences the projector, $\mathbf{P}(\omega)$.

To analyse these issues we compute the following quantities using the saved test compositions:

(1) the minimum species composition in $\mathbf{z}^{\mathrm{CP}}(\omega)$ denoted by $\min(\mathbf{z}^{\mathrm{CP}})$; and
(2) the ratio of the entropy production given by the RAMP approach denoted by $\dot{s}^{\mathrm{CP}}$ to the actual entropy production at the test composition denoted by $\dot{s}$ (which are defined below).

The entropy production rate $\dot{s}$ at a test composition $\mathbf{z}$ is given as

$$\dot{s} = \boldsymbol{\eta}^{\mathrm{T}}(\mathbf{z})\mathbf{S}(\mathbf{z}), \tag{39}$$

where $\boldsymbol{\eta}$ is the entropy gradient vector (at constant enthalpy h and pressure p) given as

$$\boldsymbol{\eta} = \left.\frac{\partial s}{\partial \mathbf{z}}\right|_{h,p} = \mathbf{s} - \frac{\mathbf{h}}{T}, \tag{40}$$

where $\mathbf{s}$ and $\mathbf{h}$ are molar entropies and enthalpies, respectively.

The entropy production given by the RAMP approach denoted by $\dot{s}^{\mathrm{CP}}$ is given as

$$\dot{s}^{\mathrm{CP}} = \boldsymbol{\eta}^{\mathrm{T}}(\mathbf{z}^{\mathrm{CE}})\mathbf{P}\mathbf{S}(\mathbf{z}^{\mathrm{CE}}). \tag{41}$$

First we look at the realizability issue. We consider the test compositions saved at $t = 0$, and compute the minimum species composition in $\mathbf{z}^{\mathrm{CP}}$ denoted by $\min(\mathbf{z}^{\mathrm{CP}})$ at each test composition. For $\mathbf{z}^{\mathrm{CP}}$ to be realizable, we must have $\min(\mathbf{z}^{\mathrm{CP}}) \geq 0$. Figure 8 shows the value of $-\min(\mathbf{z}^{\mathrm{CP}})$ for a range of values of ω from 0 to $10^9\,\mathrm{s}^{-1}$. Each subplot shows only the test compositions for which $\mathbf{z}^{\mathrm{CP}}$ is unrealizable, i.e. $-\min(\mathbf{z}^{\mathrm{CP}}) > 0$. The title of each subplot indicates the percentage of test compositions (out of the overall 2500) for which we have an unrealizable $\mathbf{z}^{\mathrm{CP}}$, i.e. $-\min(\mathbf{z}^{\mathrm{CP}}) > 0$. In addition, in parentheses, we indicate the percentage of these unrealizable compositions for which we have the value of $-\min(\mathbf{z}^{\mathrm{CP}})$ above a reference value of 10^{-12} (indicative of round-off error). For this test case, we notice that, at small values of ω, over 60% of the compositions are unrealizable and some of them have relatively large negative compositions in the order -10^{-4}. However, as the value of ω increases, the percentage of unrealizable compositions decreases to less than 7% for values of $\omega > 10^5$, and less than 30% of these unrealizable compositions have $\min(\mathbf{z}^{\mathrm{CP}}) < -10^{-12}$. The maximum magnitude of the negative species composition decreases to a value of less than 10^{-6} for $\omega \geq 10^5\,\mathrm{s}^{-1}$ and a value less than 10^{-22} for $\omega \geq 10^8\,\mathrm{s}^{-1}$. As described earlier, larger values of ω pull the CPIM manifold closer to the CEM and improve the conditioning of the $\mathbf{L}(\omega)$ matrix, thereby making more compositions $\mathbf{z}^{\mathrm{CP}}(\omega)$ realizable. Nevertheless, as mentioned earlier, since $\mathbf{z}^{\mathrm{CP}}(\omega)$ does not directly appear in the definition of the projector $\mathbf{P}(\omega)$, realizability of $\mathbf{z}^{\mathrm{CP}}(\omega)$ is not a major concern.

Next we look at the ratio of entropy productions, $\dot{s}^{\mathrm{CP}}/\dot{s}$, as shown in Figure 9 for the same range of values of ω. We notice that, for the current test case, none of the test compositions yields a negative entropy production using the RAMP approach (even at $\omega = 0$, which corresponds to the CPIM method). However, we do notice that, at small values of ω and at low temperatures, the entropy production is significantly underpredicted (by about 80%) using the RAMP approach compared to the actual entropy production at the test composition. We also notice that, for values of $\omega > 10^7$, the entropy production (predicted by the RAMP method) is overpredicted by 4 to 8 times in the temperature range around 1000 K. At large values of ω, $\mathbf{z}^{\mathrm{CP}}$ approaches $\mathbf{z}^{\mathrm{CE}}$ and $\mathbf{S}^{\mathrm{CP}}$ approaches $\mathbf{S}^{\mathrm{CE}}$, which in the temperature range around 1000 K is misaligned with the CEM and $\mathbf{S}$ (as inferred from Figures 6 and 7), and this might be a possible reason for the overprediction of entropy in this temperature range. We analysed these quantities using test compositions saved at $t = \Delta t$ as well (results not included), and observed a similar behaviour.

In summary, compared to the CPIM method [32], the RAMP approach provides a good control over realizability and entropy production using the relaxation rate parameter ω.
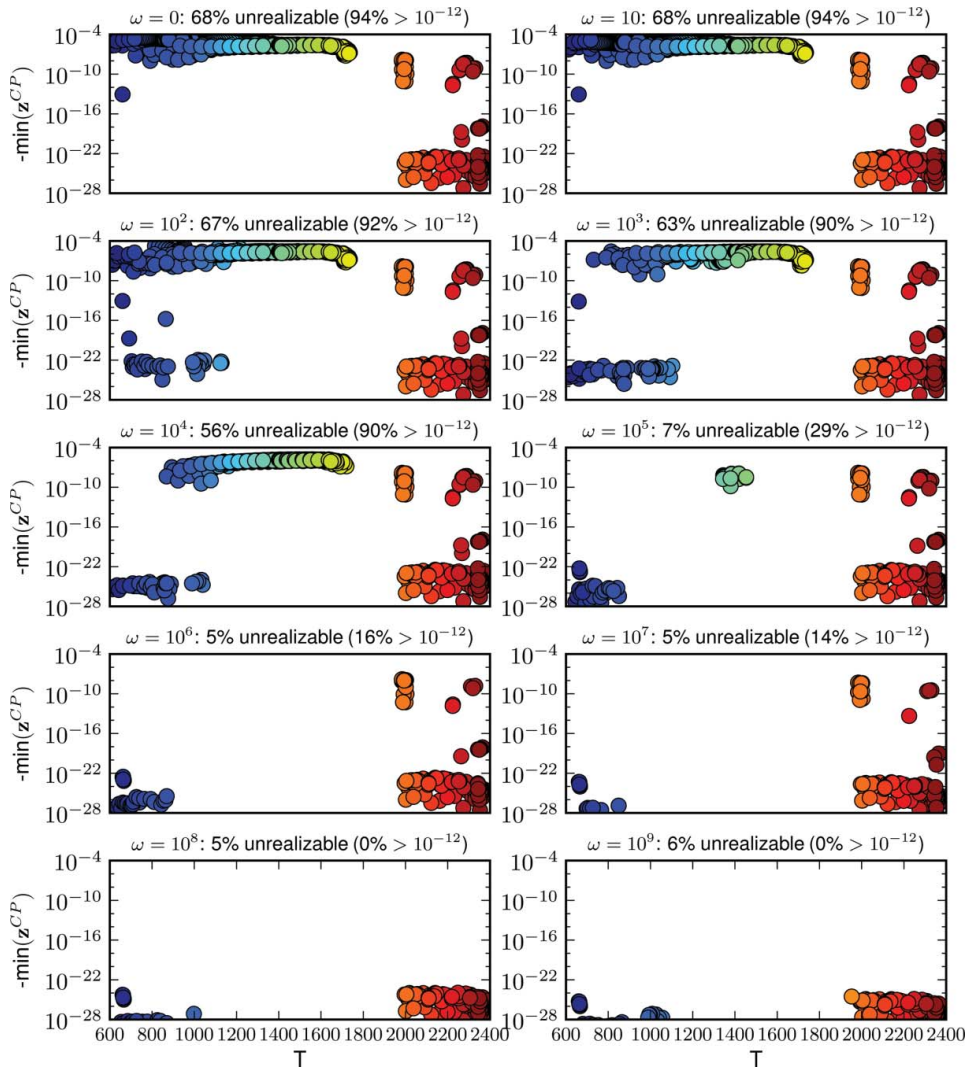
Figure 8. Scatter plots of the minimum species composition in $\mathbf{z}^{CP}$ (composition on the CPIM obtained using the RAMP approach) versus temperature T computed using the saved test compositions at $t = 0$ for a range of values of relaxation rate ω from 0 to $10^9 \, \mathrm{s}^{-1}$.

7.3. *Accuracy*

In this section we look at three measures of error to compare the accuracy of the three implementations of RCCE. The three measures of error (which are described in more detail in the following sections) include:

(1) **source vector error**: this is a measure of error in the source vector $\dot{\mathbf{r}}$ used in the RCCE and RCCE/RAMP implementations relative to the exact source vector $\dot{\mathbf{r}}^e$;

(2) **reaction mapping error**: this is a measure of error in the reaction mapping $\mathbf{r}(t)$ obtained using the three implementations of RCCE, relative to the exact reaction mapping $\mathbf{r}^e(t)$ obtained without using the RCCE dimension reduction;

Figure 9.   Scatter plots of ratio of entropy production rates $\dot{s}^{\mathrm{CP}}$ (obtained using the RAMP approach) to $\dot{s}$ (at the test composition) versus temperature T computed using the saved test compositions at $t = 0$ for a range of values of relaxation rate ω from 0 to $10^9\,\mathrm{s}^{-1}$.

(3) **reduction–tabulation error**: this a measure of error in the reaction mapping obtained using the combined ISAT/RCCE methodology (as described in our previous work [6]) with the three implementations of RCCE.

7.3.1.   *Source vector error*

We define the error in the source vector approximation $\dot{\mathbf{r}}(t)$ relative to the exact source vector $\dot{\mathbf{r}}^{\mathrm{e}}(t)$ (as illustrated in Figure 5) by

$$\epsilon(\dot{\mathbf{r}}(t)) = \frac{[\dot{\mathbf{r}}(t) - \dot{\mathbf{r}}^{\mathrm{e}}(t)]_{\mathrm{rms}}}{[\dot{\mathbf{r}}^{\mathrm{e}}(t)]_{\mathrm{rms}}}, \tag{42}$$
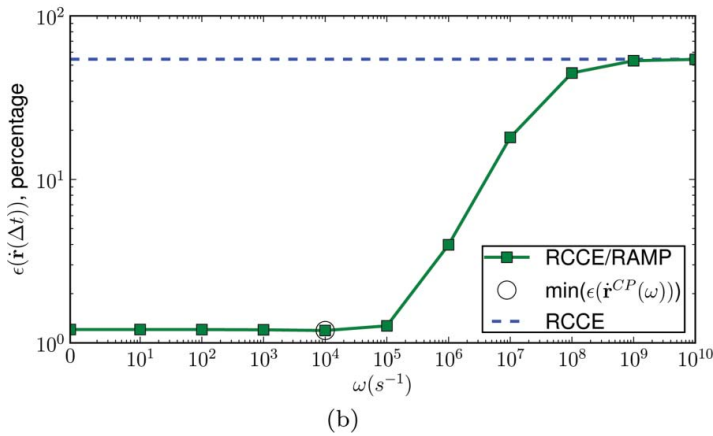
Figure 10. Source vector error in the RCCE and RCCE/RAMP (for various values of ω) implementations at $n_{rs} = 11$ computed using test compositions saved at $t = 0$ (a) and $t = \Delta t$ (b).

where the operator $[\mathbf{x}]_{rms}$ is defined as

$$[\mathbf{x}]_{rms} = \sqrt{\frac{1}{N} \sum_{n=1}^{N} \|\mathbf{x}^{(n)}\|^2},$$
(43)

with $\|\mathbf{x}\|$ denoting the vector 2-norm. We measure the error $\epsilon(\dot{\mathbf{r}}(t))$ at two discrete times, $t = 0$ and $t = \Delta t$, using the saved test compositions.

Figure 10 shows the source vector error $\epsilon(\dot{\mathbf{r}}(t))$ computed using the saved test compositions at $t = 0$ and $t = \Delta t$ for the RCCE and RCCE/RAMP implementations using $n_{rs} = 11$ represented species (listed in Table 1). For the RAMP implementation, the error is computed for a range of values of relaxation rate ω from 0 to 10^{10} s^{-1}.

At $t = 0$, we notice that both the RCCE and RCCE/RAMP implementations incur the same error (around 20%) for small values of ω. We see a slight reduction in the source vector error using the RAMP approach (to about 11%) near $\omega = 10^7$ s^{-1}. Presumably this value of ω provides the best approximation for the test compositions at $t = 0$ using the RAMP

approach. As the value of ω further increases, the error using the RAMP implementation again increases and reaches the same value as with RCCE at $\omega = 10^9\,\mathrm{s}^{-1}$.

At $t = 0$, the test compositions are away from the attracting manifold (due to mixing) and thus both RCCE and RCCE/RAMP yield a similar level of accuracy. However, at $t = \Delta t$, we notice that for small values of ω the source vector error using RCCE/RAMP is considerably smaller (around 2%) than the error incurred using RCCE (over 50%). The test compositions at $t = \Delta t$ are closer to an attracting manifold which is well approximated by the CPIM used in the RAMP approach. The RCCE method incurs large errors due to the inaccurate projection used to compute the source vector. The RAMP approach, however, yields a more accurate projection (using the CPIM), which in turn provides a more accurate source vector approximation as confirmed by these error measurements. We do notice that, as the value of ω increases beyond $10^4\,\mathrm{s}^{-1}$, the error incurred by the RCCE/RAMP approach starts increasing and approaches the same level as RCCE at $\omega = 10^9\,\mathrm{s}^{-1}$. This, as explained earlier, is because, for larger values of ω, the CPIM is pulled closer to the CEM, and the RAMP approach yields the same source vector as the RCCE approach. In this case the minimum error is attained at $\omega = 10^4\,\mathrm{s}^{-1}$. We computed the source vector error for a few other values of n_{rs} (not shown), and observed a similar behaviour. The RCCE/RAMP and RCCE approaches yield similar levels of error at $t = 0$; however, RCCE/RAMP yields significantly lower errors at $t = \Delta t$, and for this test the minimum error is always achieved around $\omega = 10^4\,\mathrm{s}^{-1}$. For this reason, in the following tests, we use a fixed value of $\omega = 10^4\,\mathrm{s}^{-1}$ in the RCCE/RAMP implementation.

We think the value of ω which provides the best source vector approximation should be indirectly related to the flow mixing time scale, τ_{mix}. The mixing time scale determines the extent to which particle compositions can move away from the attracting manifold in the composition space. Similarly, the ω value in the RAMP approach controls the position of the CPIM relative to the CEM. Hence, we think, a value of ω in the range $0.1/\tau_{\mathrm{mix}}$ to $10/\tau_{\mathrm{mix}}$ should provide good source vector approximation. In this work we have $\tau_{\mathrm{mix}} = 1\,\mathrm{ms}$, which yields the good range of values of ω to be from 100 to $10{,}000\,\mathrm{s}^{-1}$.

Figure 11 shows the source vector error $\epsilon(\dot{\mathbf{r}}(t))$ computed for a range of values of n_{rs} (using the represented species listed in Table 1) at $t = 0$ and $t = \Delta t$. For the RCCE/RAMP implementation, we use a fixed value of $\omega = 10^4\,\mathrm{s}^{-1}$. Here again we see that at $t = 0$ the errors using the two implementations are comparable; however, at $t = \Delta t$, RCCE/RAMP yields an order-of-magnitude smaller error than the RCCE.

7.3.2. *Reaction mapping error*

Here we compare the relative error in the reaction mapping $\mathbf{r}(t)$ obtained using different implementations of RCCE relative to the exact reaction mapping $\mathbf{r}^{\mathrm{e}}(t)$ (as illustrated in Figure 5). For this, at each test composition $\mathbf{z} \equiv \mathbf{z}^{(n)}(0)$ we define $\mathbf{r}(0) = \mathbf{B}^{\mathrm{T}}\mathbf{z}$ and then compute the reaction mapping using the following methods (as illustrated in Figures 4 and 5).

- **Direct evaluation**: the exact reaction mapping, $\mathbf{r}^{\mathrm{e}}(t) = \mathbf{B}^{\mathrm{T}}\mathbf{z}(t)$, where $\mathbf{z}(t)$ is obtained using ODE integration in the full space starting from $\mathbf{z}$;
- **RCCE/TIFS**: $\mathbf{r}(t) = \mathbf{B}^{\mathrm{T}}\mathbf{z}(t)$, where $\mathbf{z}(t)$ is computed using ODE integration in the full space starting from $\mathbf{z}^{\mathrm{CE}}(\mathbf{r}(0))$;
- **RCCE**: $\mathbf{r}(t)$ obtained by integrating the reduced system of ODEs (7) for the constraints starting from $\mathbf{r}(0)$ with the source vector $\dot{\mathbf{r}} \equiv \dot{\mathbf{r}}^{\mathrm{CE}}$; and
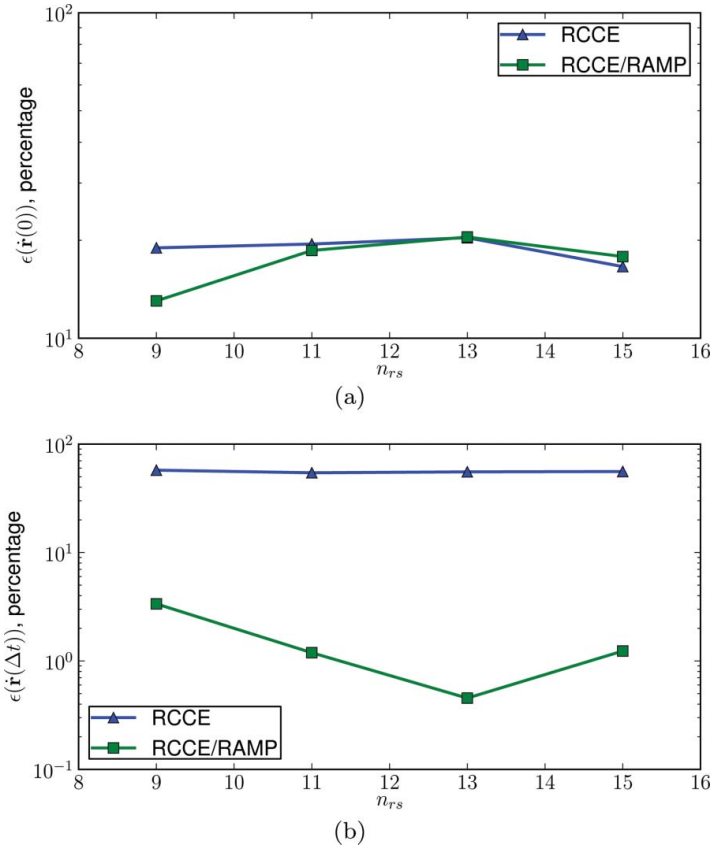
Figure 11.   Source vector error in the RCCE and RCCE/RAMP (for fixed $\omega = 10^4$) implementations at various values of n_{rs} computed using test compositions saved at $t = 0$ (a) and $t = \Delta t$ (b).

- **RCCE/RAMP**: $\mathbf{r}(t)$ obtained by integrating the reduced system of ODEs (7) for the constraints starting from $\mathbf{r}(0)$ with the source vector $\dot{\mathbf{r}} \equiv \dot{\mathbf{r}}^{CP}$.

We define the error in the reaction mapping $\mathbf{r}(t)$ as

$$\epsilon(\mathbf{r}(t)) = \frac{[\mathbf{r}(t) - \mathbf{r}^e(t)]_{\text{rms}}}{[\mathbf{r}^e(t) - \mathbf{r}(0)]_{\text{rms}}}, \tag{44}$$

where the rms error is computed as before using Equation (43).

Figure 12 shows the reaction mapping error $\epsilon(\mathbf{r}(t))$ for varying reaction time steps from $t = 10^{-12}$ to $t = 10^{-2}$ s using the three implementations of RCCE with $n_{rs} = 11$ and $n_{rs} = 15$ represented species (listed in Table 1). The RAMP implementation uses a fixed value of $\omega = 10^4 \, \text{s}^{-1}$. For small reaction time steps, we notice that all three implementations yield similar reaction mapping error.

From the definition of the reaction mapping error Equation (44), we find that as $t \to 0$ we get

$$\lim_{t \to 0} \mathbf{r}(t) = \mathbf{r}(0) + t\dot{\mathbf{r}}(0) \quad \text{and} \quad \lim_{t \to 0} \mathbf{r}^e(t) = \mathbf{r}(0) + t\dot{\mathbf{r}}^e(0), \tag{45}$$
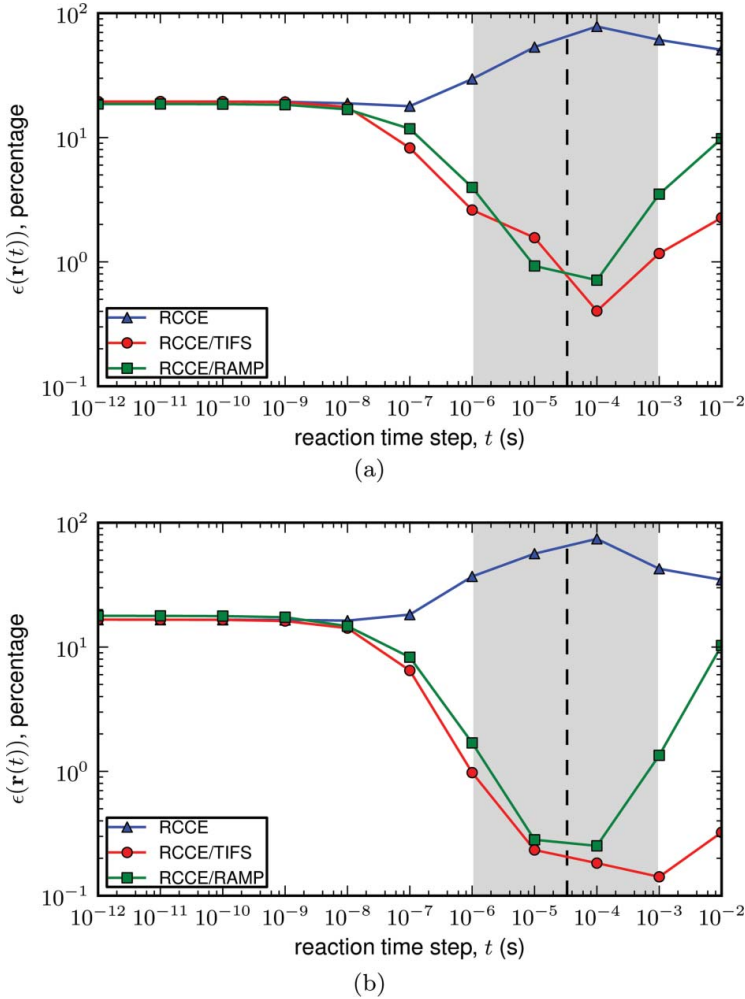
Figure 12. Reaction mapping error for varying reaction time steps t at $n_{rs} = 11$ (a) and $n_{rs} = 15$ (b) using the three implementations: RCCE/TIFS, RCCE and RCCE/RAMP. The grey region shows the typical range of values of t (from 1 μs to 1 ms) used in real LES/PDF computations. The dashed line indicates $t = \Delta t$.

which gives

$$\lim_{t \to 0} \epsilon(\mathbf{r}(t)) = \frac{[\dot{\mathbf{r}}(0) - \dot{\mathbf{r}}^e(0)]_{\text{rms}}}{[\dot{\mathbf{r}}^e(0)]_{\text{rms}}}, \tag{46}$$

and hence

$$\lim_{t \to 0} \epsilon(\mathbf{r}(t)) = \epsilon(\dot{\mathbf{r}}(0)). \tag{47}$$

So for small reaction time steps, the reaction mapping error $\epsilon(\mathbf{r}(t))$ for the RCCE and RCCE/RAMP implementations is the same as the corresponding source vector error $\epsilon(\dot{\mathbf{r}}(0))$ at the same value of n_{rs}, as seen in Figure 11. Also, the reaction mapping error for the

Figure 13. Reaction mapping error for large reaction time steps t at $n_{\rm rs} = 11$ and $n_{\rm rs} = 15$ using the RCCE/TIFS implementation.

RCCE/TIFS implementation is the same as the error for the RCCE implementation, since both use the chemical source term $\mathbf{S}^{\rm CE}$ computed at $\mathbf{z}^{\rm CE}(\mathbf{r}(0))$ for ODE integration.

At both $n_{\rm rs} = 11$ and $n_{\rm rs} = 15$, we notice that, as the reaction time step increases, the error in the RCCE/TIFS and RCCE/RAMP implementations starts decreasing. This is because in the RCCE/TIFS implementation we solve the full system of (unconstrained) $n_{\rm s}$ ODEs, which provides an accurate reaction mapping; and in the RCCE/RAMP implementation the RAMP approach provides an accurate approximation for the source vector at large time steps (as seen in Figures 10 and 11). However, we notice that the error in the RCCE implementation grows or remains constant for larger reaction time steps. This is because the RCCE approach does not provide a good approximation for the source vector (as seen in Figure 11), and so the error in the reaction mapping grows for large reaction time steps. These results show that both RCCE/TIFS and RCCE/RAMP approaches are accurate and yield similar levels of error: however, the RCCE approach yields significantly larger errors at large reaction time steps. Figure 12 highlights the typical range of values of reaction time step t from 1 μs to 1 ms that we may use in real LES/PDF computations. We notice that in this range the RCCE/TIFS and RCCE/RAMP implementations both yield less than 4 and 2% errors at $n_{\rm rs} = 11$ and $n_{\rm rs} = 15$, respectively. However, the RCCE implementation yields over 30% error.

We notice that the error in the RCCE/TIFS and RCCE/RAMP implementations increases for time steps beyond $t = 10^{-4}$ s. Ideally, the reaction mapping error should approach zero for large reaction time t because both $\mathbf{r}(t)$ and $\mathbf{r}^{\rm e}(t)$ in the definition of $\epsilon(\mathbf{r}(t))$ in Equation (44) should reach the equilibrium composition as $t \to \infty$. The ODE integration using DDASAC becomes expensive at large reaction time steps because the solution is computed by taking many smaller sub-steps (as explained later in Section 7.4). Thus we examine the reaction mapping error at very large reaction time steps using the RCCE/TIFS implementation alone and the results are presented in Figure 13. We find that, as expected, the reaction mapping error approaches zero at very large reaction time steps over 10^3 s. The relatively large variations in the reaction mapping error for time steps beyond $t = 10^{-4}$ and

Figure 14.   Combined reduction–tabulation error for a fixed reaction time step, $t = \Delta t$, at various values of n_{rs} using the three implementations of RCCE: RCCE/TIFS, RCCE and RCCE/RAMP with ISAT. The tabulation error (without reduction) is indicated by the solid line labelled ISAT.

the long time required for the error to reach zero can be attributed to different ignition delay times for different test compositions. Since the test compositions obtained from PaSR have a wide range of initial temperatures from 600 K to equilibrium temperature (around 2400 K), the ignition delay time for different test compositions can vary by orders of magnitude thereby requiring a very long time for all the test compositions to reach equilibrium.

In summary, these results show that both the RCCE/TIFS and RCCE/RAMP implementations yield significantly smaller error than the RCCE implementation.

7.3.3.   *Reduction–tabulation error*

In [25], we describe our combined dimension reduction and tabulation ISAT/RCCE methodology. In this combined methodology, the reaction mapping computation using RCCE is tabulated using the ISAT algorithm to save computational time in large-scale LES/PDF computations [26]. Here we measure the combined reduction–tabulation error incurred using this ISAT/RCCE methodology with RCCE implemented using the three methods described. We use the same definition of error as in our previous work [25], given as

$$\epsilon_{RT} = \frac{[\mathbf{z}^r(\Delta t) - \mathbf{z}^r_{RT}(\Delta t)]_{rms}}{[\mathbf{z}^r(\Delta t) - \mathbf{z}^r(0)]_{rms}}, \tag{48}$$

where $\mathbf{z}^r(0)$ and $\mathbf{z}^r(\Delta t)$ denote the composition of the represented species obtained from $\mathbf{z}(0)$ and $\mathbf{z}(\Delta t)$, respectively, and $\mathbf{z}^r_{RT}$ denotes the reaction mapping obtained using ISAT/RCCE. This error is measured for a fixed reaction time step, $t = \Delta t = 0.033$ ms, at different values of n_{rs}. The rms error is computed by considering $N = 10^5$ test compositions from a PaSR simulation (the same as in [25]).

Figure 14 shows the combined reduction–tabulation error using a fixed reaction time step $t = \Delta t$ for different values of n_{rs} for the three implementations of RCCE. (The

RAMP implementation uses a fixed value of $\omega = 10^4\,\text{s}^{-1}$.) A fixed ISAT error tolerance of $\epsilon_{\text{tol}} = 10^{-5}$ is used, which yields less than 1% tabulation error as shown in Figure 14. As we observed in the reaction mapping error results, here again we notice that the errors incurred by the RCCE/TIFS and RCCE/RAMP implementations are comparable and significantly smaller than the RCCE implementation. For $n_{\text{rs}} > 9$, the RCCE/TIFS and RCCE/RAMP implementations yield less than 2% error; however, the RCCE implementation yields over 50% error.

7.4. *Computational efficiency*

In this section we compare the computational performance of the three implementations of RCCE. Each of the three implementations of RCCE involves solving a system of ODEs to compute the reaction mapping. We use DDASAC [34] to perform ODE integration with an absolute error tolerance of 10^{-8}.

Given a system of ODEs

$$\frac{d\mathbf{x}}{dt} = \mathbf{f}(\mathbf{x}),\tag{49}$$

with initial condition $\mathbf{x}(0)$ and a time duration t, DDASAC returns the solution $\mathbf{x}(t)$ within a specified error tolerance. DDASAC uses a variable time step predictor–corrector algorithm – involving backward-difference formulas (BDFs) for the predictor and modified Newton iterations for the corrector – to compute the solution. To solve a system of ODEs of the form Equation (49) using DDASAC, we need to provide functions to compute the right-hand-side source vector $\mathbf{f}(\mathbf{x})$ and the Jacobian $\mathbf{J}(\mathbf{x})$ defined as

$$\mathbf{J}(\mathbf{x}) = \frac{\partial \mathbf{f}(\mathbf{x})}{\partial \mathbf{x}}.\tag{50}$$

For the Jacobian, we also have the option to use DDASAC's built-in finite difference approximation for the Jacobian.

For a given initial condition $\mathbf{x}(0)$ and time duration t, DDASAC takes multiple variable time steps to compute the solution, $\mathbf{x}(t)$. At each step the source vector $\mathbf{f}(\mathbf{x})$ is evaluated; however, the Jacobian is re-evaluated only when needed (based on an error estimate). The number of sub-steps and Jacobian evaluations depends on the stiffness of the ODE equations and the specified error tolerance.

Considering the three implementations of RCCE, in the RCCE/TIFS implementation we solve a system of n_{s} ODEs given as

$$\frac{d\mathbf{z}}{dt} = \mathbf{S}(\mathbf{z}),\tag{51}$$

and in the RCCE and RCCE/RAMP implementations we solve a reduced system of n_{r} ODEs given as

$$\frac{d\mathbf{r}}{dt} = \dot{\mathbf{r}}(\mathbf{r}),\tag{52}$$

where in the RCCE implementation $\dot{\mathbf{r}} = \dot{\mathbf{r}}^{\text{CE}}$ and in the RCCE/RAMP implementation $\dot{\mathbf{r}} = \dot{\mathbf{r}}^{\text{CP}}$.

Figure 15.   Average CPU time required to compute the reaction mapping and other quantities involved in the RCCE/TIFS, RCCE and RCCE/RAMP implementations at $n_{\mathrm{rs}} = 11$ with the reaction time step $t = \Delta t$. The quantities shown (from the bottom) include: the constrained-equilibrium composition, $\mathbf{z}^{\mathrm{CE}}$; the CEM tangent space, $\mathbf{T}$; the Jacobian, $\mathbf{J}$; the chemical source term, $\mathbf{S}$; the source vector in the RCCE implementation, $\dot{\mathbf{r}}^{\mathrm{CE}}$; the source vector in the RCCE/RAMP implementation, $\dot{\mathbf{r}}^{\mathrm{CP}}$; and finally, at the top, the reaction mapping $\mathbf{r}(t)$ using the three implementations.

For the RCCE/TIFS implementation, we provide a function generated using ADIFOR [41] to evaluate the Jacobian. However, due to the relatively complex steps involved in the evaluation of $\dot{\mathbf{r}}^{\mathrm{CE}}$ and $\dot{\mathbf{r}}^{\mathrm{CP}}$ in the RCCE and RCCE/RAMP implementations, respectively, using ADIFOR to generate functions for their Jacobian is not the best approach (as it results in the generation of many subfunctions). Hence we use DDASAC's built-in finite difference approximation for the Jacobian in the RCCE and RCCE/RAMP implementations. (The finite difference approximation provides a slightly less accurate Jacobian compared to ADIFOR.)

Figure 15 compares (on a log scale) the average cost of evaluating the source vector for the three implementations; and the overall cost of evaluating the reaction mapping using the three implementations. We see that the cost of evaluating the chemical source term $\mathbf{S}$ (source vector for the RCCE/TIFS implementation) is only around 30 μs. Compared to this the cost of evaluating the source vector for the RCCE implementation, $\dot{\mathbf{r}}^{\mathrm{CE}}$, is around 200 μs as it involves computing the constrained-equilibrium composition $\mathbf{z}^{\mathrm{CE}}$, which dominates the cost. The cost of evaluating the source vector for the RCCE/RAMP implementation, $\dot{\mathbf{r}}^{\mathrm{CP}}$, goes further up to about 1000 μs. This is because, in the RAMP approach, to compute the projector we need to compute the Jacobian $\mathbf{J}$ and the CEM tangent subspace $\mathbf{T}$, which are expensive to evaluate.

The cost of evaluating the reaction mapping is directly related to the cost of evaluating the source vector, and, for this reason, we see that the RCCE/TIFS implementation is cheapest (around 10^4 μs) and the RCCE/RAMP implementation is the costliest (around 10^6 μs, two orders of magnitude more than RCCE/TIFS). The RCCE implementation is around five times more expensive than the RCCE/TIFS implementation.

For the RCCE implementation, there exists a mathematically equivalent alternative implementation in terms of the constraint potentials as described in [13, 28, 29, 31].

This implementation reduces the cost of evaluating the source vector by avoiding the constrained-equilibrium calculations by transforming the ODEs in terms of the constraint potentials. This implementation, however, poses some numerical issues as reported in [31], because constraint potentials can attain very large values (the constraint potential value corresponding to a species with zero concentration must be infinite). Special transformation and pre-conditioning methods [31] are used to resolve these problems. This alternative implementation may be computationally less expensive; however, since this implementation is mathematically equivalent to our current implementation of RCCE, the accuracy is not improved by this alternative implementation.

In summary, these results show that the RCCE/TIFS implementation is the most efficient among the three implementations of RCCE.

7.5. *Robustness*

In addition, for a method to be accurate and efficient, it is also important for the method to be robust and fail-safe. In large-scale LES/PDF computations, we compute the reaction mapping in the order $\mathcal{O}(10^{12})$ times, and we want our ISAT/RCCE implementation to return an accurate, realizable reaction mapping every time without failing.

In our tests with the three implementations of RCCE, we find the RCCE/TIFS approach to be the most robust. We have tested this implementation using the partially-stirred reactor for a wide range of chemical mechanisms, represented species and testing conditions, and the method has never failed [27, 25]. It has also been recently tested for performing large-scale LES/PDF simulations of Sandia Flame D [26], and again the method worked without any issues.

However, in the current study we encountered some test cases (at certain values of n_{rs}, e.g. $n_{rs} = 10$, for which results are not included) where both the classical RCCE and RCCE/RAMP implementations failed to provide a realizable reaction mapping or failed to converge within DDASAC. Both these implementations solve a reduced system of ODEs for the constraints by projecting the chemical source term onto the constrained subspace. The projected source vector on some occasions is found to yield negative unrealizable constraint compositions during the DDASAC sub-steps. We were able to resolve some of these cases by using a smaller error tolerance in DDASAC (which in turn forces DDASAC to take smaller sub-steps); however, the overall implementation still failed for a few test points.

7.6. *Comparison with previous works*

Here we have shown that the classical RCCE implementation yields significant errors in the reaction mapping at large reaction time steps. However, many of the previous works using the classical RCCE implementation report good accuracy [28–31]. There are two key differences in the implementation of RCCE used in these previous works compared to our classical RCCE implementation used in this work. In the previous works:

(1) the rate-equations for the constraint potentials are solved to compute the reaction mapping; and
(2) general linear combinations of species compositions are used as constraints.

As pointed out earlier, the rate-equations for the constraint potentials are mathematically equivalent to the rate-equations for the constraints, and thus the two approaches should

yield the same reaction mapping for similar constraints. The accuracy of the RCCE method is very sensitive to the choice of the constraints [27, 29, 31], and thus one possible reason for achieving good accuracy in these previous works could be attributed to better selection of linear constraints. We understand that the selection of linear constraints may in some cases provide more accurate results with RCCE. Ideally, it would be very insightful to test the three implementations of RCCE described in this work for linear constraints. However, our current implementation of ISAT/RCCE [25] would require significant changes to incorporate linear constraints, and hence we are unable to perform these tests at this stage. Nevertheless, it should be noted that, in the previous CPIM work [32], it is shown that even for the constraint potentials based implementation using general linear constraints, the CPIM method yields better accuracy than the RCCE implementation.

Moreover, it should be noted that most of the previous works [28, 30, 31] rely more on qualitative comparison of temperature and species profiles against time to assess the accuracy of the RCCE implementation versus the detailed mechanism, which is a relatively weak test for determining the overall accuracy of the RCCE dimension reduction method. Furthermore, to the authors' knowledge, the sensitivity of the accuracy of the RCCE implementation to reaction time step has not been studied carefully in any of the previous works.

8. Conclusions

In this study, we looked at three different implementation of the RCCE dimension reduction method:

(1) RCCE/TIFS: involving the solution of the full system of n_s ODEs;
(2) RCCE: involving the solution of a reduced system of n_r ODEs for the constraints, with source vector evaluated by projecting the chemical source term evaluated on the CEM directly onto the represented subspace;
(3) RCCE/RAMP: involving an implementation similar to RCCE, but evaluating the source vector using a more accurate RAMP approach (based on the CPIM method).

From the results presented in this work we can draw the following conclusions:

- the RCCE/TIFS implementation is the most accurate, robust and efficient among the three implementations of RCCE;
- the RCCE implementation is based on an inaccurate projection of the chemical source term onto the represented subspace, which does not take into account the non-invariance of the CEM manifold and thus yields large errors;
- the RAMP approach provides a more accurate projection for the source vector evaluation (based on the CPIM method) which significantly reduces the error;
- for the methane/air test case considered in this work, both the RCCE/TIFS and RCCE/RAMP implementations yield less than 2% error in the reaction mapping compared to over 50% using the RCCE implementation;
- computationally, however, the RCCE/RAMP is an order of magnitude more expensive than the RCCE/TIFS and RCCE implementations due to the need for expensive Jacobian evaluations in the RAMP approach.

The RAMP approach, even though being expensive

- elucidates the inaccuracies in the projection employed in the classical RCCE implementation;

- demonstrates the accuracy of the CPIM approximation;
- provides an alternative framework for the implementation of RCCE (and possibly other related dimension reduction methods) based on the invariant manifold concepts.

Acknowledgments

Note

1. The residence time was misreported in [25] to be $\tau_{res} = 10$ ms.

References

[1] S. Sarathy, C. Westbrook, M. Mehl, W. Pitz, C. Togbe, P. Dagaut, H. Wang, M. Oehlschlaeger, U. Niemann, K. Seshadri, P. Veloo, C. Ji, F. Egolfopoulos, and T. Lu, *Comprehensive chemical kinetic modeling of the oxidation of 2-methylalkanes from C7 to C20*, Combust. Flame 158 (2011), pp. 2338–2357.

[2] S.B. Pope, *PDF methods for turbulent reactive flows*, Prog. Energy Combust. Sci. 11 (1985), pp. 119–192.

[3] M. Muradoglu, P. Jenny, S.B. Pope, and D.A. Caughey, *A consistent hybrid finite-volume/particle method for the PDF equations of turbulent reactive flows*, J. Comput. Phys. 154 (1999), pp. 342–371.

[4] V. Raman, R.O. Fox, and A.D. Harvey, *Hybrid finite-volume/transported PDF simulations of a partially premixed methane–air flame*, Combust. Flame 136 (2004), pp. 327–350.

[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 (2011), pp. 1319–1330.

[6] 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 (2012), pp. 3096–3109.

[7] S.B. Pope, *Small scales, many species and the manifold challenges of turbulent combustion*, Proc. Combust. Inst. 34 (2013), doi: http://dx.doi.org/10.1016/j.proci.2012.09.009

[8] T. Lu and C.K. Law, *A directed relation graph method for mechanism reduction*, Proc. Combust. Inst. 30 (2005), pp. 1333–1341.

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

[10] M. Bodenstein and S.C. Lind, *Geschwingigkeitder bildung des bromwasserstoffs aus seinen elementen*, Z. Phys. Chem. 57 (1906), p. 168.

[11] M.D. Smooke, *Reduced Kinetic Mechanisms and Asymptotic Approximations for Methane–Air Flames*, Springer-Verlag, Berlin, 1991.

[12] J.C. Keck and D. Gillespie, *Rate-controlled partial equilibrium method for treating reacting gas-mixtures*, Combust. Flame 17 (1971), pp. 237–241.

[13] J.C. Keck, *Rate-controlled constrained equilibrium theory of chemical reactions in complex systems*, Prog. Energy Combust. Sci. 16 (1990), pp. 125–154.

[14] S.H. Lam and D.A. Goussis, *The CSP method for simplifying kinetics*, Int. J. Chem. Kinetics 26 (1994), pp. 461–486.

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

[16] S.B. Pope and U. Mass, *Simplifying chemical kinetics: trajectory-generated low-dimensional manifolds*, Mech. & Aerosp. Engng Rep. FDA 93-11, Cornell University, Ithaca, NY, 1993.

[17] Z. Ren, S.B. Pope, A. Vladimirsky, and J.M. Guckenheimer, *The invariant constrained equilibrium edge pre-image curve method for the dimension reduction of chemical kinetics*, J. Chem. Phys. 124 (2006), Art. No. 114111. Available at: http://dx.doi.org/10.1063/1.2177243.

[18] S.B. Pope, *Computationally efficient implementation of combustion chemistry using in situ adaptive tabulation*, Combust. Theory Model. 1 (1997), pp. 41–63.

[19] L. Lu and S.B. Pope, *An improved algorithm for in situ adaptive tabulation*, J. Comput. Phys. 228 (2009), pp. 361–386.

[20] S.R. Tonse, N.W. Moriarity, N.J. Brown, and M. Frenklach, *PRISM: Piecewise Reusable Implementation of Solution Mapping. An economical strategy for chemical kinetics*, Israel J. Chem. 39 (1999), pp. 97–106.

[21] F.C. Christo, A.R. Masri, E.M. Nebot, and S.B. Pope, *An integrated PDF/neural network approach for simulating turbulent reacting systems*, Proc. Combust. Inst. 26 (1996), pp. 43–48.

[22] Q. Tang, J. Xu, and S.B. Pope, *PDF calculations of local extinction and NO production in piloted-jet turbulent methane/air flames*, Proc. Combust. Inst. 28 (2000), pp. 133–139.

[23] 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 (2002), pp. 1411–1417.

[24] Z. Ren, V. Hiremath, and S.B. Pope, *Dimension reduction and tabulation of combustion chemistry using ICE-PIC and ISAT*, in *6th US National Combustion Meeting*, 17–20 May 2009, Ann Arbor, MI, 2009.

[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 (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 (2013). doi: 10.1016/j.proci.2012.06.004

[27] V. Hiremath, Z. Ren, and S.B. Pope, *A greedy algorithm for species selection in dimension reduction of combustion chemistry*, Combust. Theory Model. 14 (2010), pp. 619–652.

[28] D. Hamiroune, P. Bishnu, M. Metghalchi, and J.C. Keck, *Rate-controlled constrained-equilibrium method using constraint potentials*, Combust. Theory Model. 2 (1998), pp. 81–94.

[29] W. Jones and S. Rigopoulos, *Rate-controlled constrained equilibrium: formulation and application to nonpremixed laminar flames*, Combust. Flame 142 (2005), pp. 223 – 234.

[30] M. Janbozorgi and H. Metghalchi, *Rate-controlled constrained-equilibrium theory applied to the expansion of combustion products in the power stroke of an internal combustion engine*, Int. J. Thermodyn. 12 (2009), pp. 44–50.

[31] M. Janbozorgi, S. Ugarte, H. Metghalchi, and J.C. Keck, *Combustion modeling of mono-carbon fuels using the rate-controlled constrained-equilibrium method*, Combust. Flame 156 (2009), pp. 1871–1885.

[32] Q. Tang and S.B. Pope, *A more accurate projection in the rate-controlled constrained-equilibrium method for dimension reduction of combustion chemistry*, Combust. Theory Model. 8 (2004), pp. 255–279.

[33] H. Wang, P.P. Popov, and S.B. Pope, *Weak second order splitting schemes for Lagrangian Monte Carlo particle methods for the composition PDF/FDF transport equations*, J. Comput. Phys. 229 (2010), pp. 1852–1878.

[34] M. Caracotsios and W.E. Stewart, *Sensitivity analysis of initial value problems with mixed ODEs and algebraic equations*, Comput. Chem. Engrg 9 (1985), pp. 359–365.

[35] G.P. Beretta, J.C. Keck, M. Janbozorgi, and H. Metghalchi, *The rate-controlled constrained-equilibrium approach to far-from-local-equilibrium thermodynamics*, Entropy 14 (2012), pp. 92–130.

[36] Z. Ren, G.M. Goldin, V. Hiremath, and S.B. Pope, *Reduced description of reactive flows with tabulation of chemistry*, Combust. Theory Model. 15 (2011), pp. 827–848.

[37] Z. Ren and G.M. Goldin, *An efficient time scale model with tabulation of chemical equilibrium*, Combust. Flame 158 (2011), pp. 1977–1979.

[38] S.B. Pope, *The computation of constrained and unconstrained equilibrium compositions of ideal gas mixtures using Gibbs function continuation*, Rep. FDA 03-02, Cornell University, Ithaca, NY, 2003.

[39] A.N. Gorban and I.V. Karlin, *Uniqueness of thermodynamic projector and kinetic basis of molecular individualism*, Physica A 336 (2004), pp. 391–432.

[40] V. Bykov and U. Maas, *Problem adapted reduced models based on reaction–diffusion manifolds (REDIMs)*, Proc. Combust. Inst. 32 (2009), pp. 561–568.

[41] ADIFOR 2.0, *Automatic Differentiation of Fortran*; software available at: http://www.mcs.anl.gov/research/projects/adifor/.