Model Simplification of Chemical Kinetic Systems Under Uncertainty

by

Thomas Michael Kyte Coles

BSc. Applied Mathematics, Imperial College London (2008)
M.S. Aeronautics and Astronautics, Stanford University (2009)

Submitted to the School of Engineering
in partial fulfillment of the requirements for the degree of

Master of Science in Computation for Design and Optimization

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2011

Author .....
School of Engineering
May 16, 2011

Certified by .....
Youssef M. Marzouk
Boeing Assistant Professor of Aeronautics and Astronautics
Thesis Supervisor

Accepted by .....
Nicolas Hadjiconstantinou
Associate Professor of Mechanical Engineering
Director, Computation for Design and Optimization Program

Model Simplification of Chemical Kinetic Systems Under Uncertainty

by

Thomas Michael Kyte Coles

Submitted to the School of Engineering
on May 16, 2011, in partial fulfillment of the
requirements for the degree of
Master of Science in Computation for Design and Optimization

Abstract

This thesis investigates the impact of uncertainty on the reduction and simplification of chemical kinetics mechanisms. Chemical kinetics simulations of complex fuels are very computationally expensive, especially when combined with transport, and so reduction or simplification must be used to make them more tractable. Existing approaches have been in an entirely deterministic setting, even though reaction rate parameters are generally highly uncertain. In this work, potential objectives under uncertainty are defined and then a number of studies are made in the hope of informing the development of a new uncertainty-aware simplification scheme. Modifications to an existing deterministic algorithm are made as a first step towards an appropriate new scheme.

Thesis Supervisor: Youssef M. Marzouk
Title: Boeing Assistant Professor of Aeronautics and Astronautics

Acknowledgments

This work was supported by funding from the Department of Energy, Office of Advanced Scientific Computing Research (ASCR) (DE-SC0003564).

I would like to thank Cosmin Safta (of Sandia National Labs) for the use of his *tchem* code, Habib Najm (also of Sandia) for useful discussions, and Xun Huan for the use of his polynomial chaos and quadrature code. Finally, I would like to thank my adviser, Prof. Youssef Marzouk, for his time, patience, guidance, and encouragement.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Chemical kinetics simulations are essential for studying reacting flow problems. Some key applications include modeling chemical production processes, atmospheric chemistry, and combustion engines. However, computational fluid dynamics (CFD) codes are very computationally expensive even without accounting for chemical kinetics, so modeling reacting flow in detail is an extremely challenging problem. This has become especially relevant recently, as there is now considerable interest in developing more efficient fuels and especially renewable biofuels; relevant chemical kinetics simulations contain many hundreds of species and thousands of reactions because the fuels consist of long chain hydrocarbons. Some particularly large examples can be seen in [4] (n-hexadecane with 225 species and 1841 reversible reactions), [3] (n-heptane with 561 species and 2538 reversible reactions), [66] (n-hexadecane with 2115 species and 8157 reversible reactions), [14] (methyl decanoate with 3036 species and 8555 reversible reactions), and [15] (combined methyl decanoate, methyl-9-decenoate, and n-heptane with 3299 species and 10806 reversible reactions).

In an effort to increase the tractability of these problems, there has been significant work in developing model simplification and reduction methods for chemical reaction mechanisms, but all existing algorithms assume that reaction rate parameters are known precisely. However, as will be explained in section 1.1.1, there is in fact considerable uncertainty in these parameters and the aim of this work is to investigate the impact of this uncertainty on model simplification. Note that the distinction

between model reduction and simplification will be explained in section 1.2.

This work focuses entirely on chemical kinetics in a zero-dimensional setting. These problems have generally been computationally tractable in the past when considered alone on a single run, but some of the large recent examples mentioned at the beginning of this chapter are very expensive. It used to only be practical to generate relatively small mechanisms for simpler problems because a lot of manual processing was required, but the number of large mechanisms under consideration is likely to grow in future with the recent development of automatic mechanism generation software; a brief overview of these methods and a number of references can be found in [28]. Furthermore, reduction and simplification techniques developed in zero dimensions also play an essential role in models containing transport, as the cost of the chemical kinetics simulations is now multiplied by a factor of the order of the number of grid points; this will be discussed in more detail in section 1.2.1. Optimization or sampling problems involving chemical kinetics are also often intractable without reduction or simplification because of the large numbers of runs required. Finally, it is difficult to develop physical insight into a very large reaction mechanism model, but zero-dimensional simplification can be valuable in giving insight into the various reaction pathways.

1.1 Reaction mechanisms and chemical kinetics

This section gives a brief overview of the theory of chemical reactions and the way that they are modeled mathematically; further details can be found in section 1.3.1 and [20]. Chemical processes are often mentioned in terms of one overall reaction, such as the following reaction of hydrogen with oxygen to give water and an enthalpy change ΔH:

$$2\mathrm{H}_2 + \mathrm{O}_2 \rightarrow 2\mathrm{H}_2\mathrm{O} + \Delta H \tag{1.1}$$

However, such a reaction does not suffice for a chemical kinetics simulation because it is lacking in details needed to accurately calculate the time evolution of temperature or species concentrations. Instead, a reaction mechanism must be used; this consists of a larger number of elementary reactions involving a variety of intermediate species. A complete hydrogen-oxygen mechanism provided in [68] consists of 9 species and 19 reversible reactions. A diagram demonstrating the species involved in each reaction can be seen in figure 1-1 and the first 3 reactions are shown explicitly below as examples:

$$\begin{array}{ccccccccc} \mathrm{H} & + & \mathrm{O_2} & \leftrightharpoons & \mathrm{O} & + & \mathrm{OH} & + & \Delta H_1 \\ \mathrm{O} & + & \mathrm{H_2} & \leftrightharpoons & \mathrm{H} & + & \mathrm{OH} & + & \Delta H_2 \\ \mathrm{H_2} & + & \mathrm{OH} & \leftrightharpoons & \mathrm{H_2O} & + & \mathrm{H} & + & \Delta H_3 \\ \vdots & & \vdots & & \vdots & & \vdots & & \vdots \end{array}$$

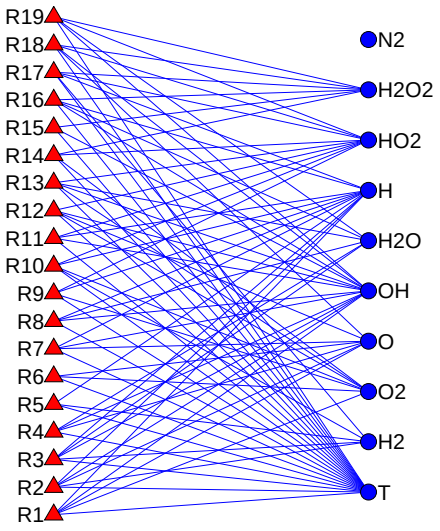

Figure 1-1: The relationship between species and reactions in the hydrogen-oxygen reaction mechanism of [68]. T represents temperature.

An ODE for the reaction mechanism can be expressed in terms of a state vector $\mathbf{z}$, which contains temperature T as its last entry and species molar concentrations C_j as its other entries. The ODE depends upon the rates of the above reactions and their stoichiometric vectors, which indicate the product and reactant species; the

product of a reaction's rate and stoichiometric vector gives the rate of change of each species due to that reaction. As an example, the stoichiometric vector $\mathbf{s}_1$ for the forward version of the first elementary reaction given above is (where all entries are zero except for those marked):

$$\Big[\ \cdots \ , \ \underbrace{-1}_{\text{H}} \ , \ \cdots \ , \ \underbrace{-1}_{\text{O}_2} \ , \ \cdots \ , \ \underbrace{1}_{\text{O}} \ , \ \cdots \ , \ \underbrace{1}_{\text{OH}} \ , \ \cdots \ , \ \frac{\Delta H_1}{\rho c_p} \ \Big]^{\mathrm{T}}$$

where ρ is density, c_p is specific heat capacity at constant pressure, and the factor ρc_p converts between enthalpy change and temperature change. When the reaction rates are denoted by F^i, the ODE for the reaction mechanism is written as:

$$\frac{d\mathbf{z}}{dt} = \mathbf{g}(\mathbf{z}) = \sum_{i=1}^{N_R} \mathbf{s}_i(\mathbf{z}) F^i(\mathbf{z}) \tag{1.2}$$

where N_R is the number of unidirectional reactions such that each reversible reaction is counted as two separate reactions. Note that both the reaction rates and stoichiometric vectors are denoted as depending on state, but the only state dependence in the stoichiometric vectors is in the ρc_p conversion factor. The right hand side $\mathbf{g}$ is sometimes known as the *chemical source term*, particularly when other terms are present in a CFD context, as described in section 1.2.1.

1.1.1 Reaction rate parameters

When the number of species if N_S, the reaction rates F^i are calculated as follows:

$$F^i(\mathbf{z}) = k_i(T) \prod_{j=1}^{N_S} C_j^{\nu_{ij}} \tag{1.3}$$

where ν_{ij} is the stoichiometric coefficient of species j as a reactant (but not as a product), T is stored in the last entry of the state vector $\mathbf{z}$, and the molar concentrations C_j are the remaining entries of $\mathbf{z}$. The parameter $k_i(T)$ is called the reaction constant and may also be dependent on pressure, but this will not be discussed further

here because all simulations in this work used constant pressure. The temperature dependence of the rate constant for a reverse reaction is usually determined from the rate constant for the forward reaction by using the equilibrium constant - details are given in [20]. The temperature dependence for the forward reaction is expressed in terms of the three parameters α, β, and E_A (activation energy) and can be modeled by the Arrhenius equation:

$$k_i = \alpha_i \, T^{\beta_i} \, e^{-\frac{Ea_i}{RT}} \tag{1.4}$$

where R is the universal gas constant. These parameters are generally highly uncertain and this work studies the impact of that uncertainty. The uncertainty arises because of the difficulty in measuring these parameters individually; although it is possible to isolate some small groups of reactions and occasionally even individual reactions, it is still necessary to fit multiple parameters from experimental data. Furthermore, there is uncertainty in measurements from sensor noise and there is often only a relatively small number of data points available due to the expense of conducting experiments. Finally, the Arrhenius equation is only an approximation and this can further complicate the analysis in regions where it is not valid.

Experimental techniques for determining these parameters are described in [20]; different techniques are suitable under different conditions, as summarized below:

Reactor Type	Pressure	Temperature (K)	Mixture Limits
Static/batch	atmospheric	< 1000	None
Stirred reactor	atmospheric-high	800-1400	Flammable
Plug-flow reactor	atmospheric-high	800-1400	None
Shock tube	atmospheric-high	> 1300	None
Flame	atmospheric-low	800-2500	Flammable

Further information on the use of shock tubes and guidelines for incorporating experimental data into reaction mechanisms can also be found in [5]. It should be noted that estimates for the rate parameters can sometimes also be calculated using quan-

tum chemistry. A database of rate constants and associated uncertainty factors (as defined in section 1.1.2) can be found in [2]. The Arrhenius parameters are generally modeled as independent random variables and independent uncertainty factors are given; this is unrealistic, but correlation information is not generally well known. However, there has been some recent work in this area, as discussed in [35].

The uncertainty in reaction rates can have a significant impact on chemical kinetics simulations; the example in figure 1-2 demonstrates the significant difference in ignition delay when changing the activation energy of only a single reaction in the hydrogen-oxygen mechanism by 5%, although it should be noted that this is a particularly sensitive reaction (as discussed in chapter 5). Uncertainty of many orders of magnitude may occur in the ignition delay for more complicated mechanisms. It is important to note at this point that rate parameter uncertainty only creates uncertainty in outputs related to rate, such as ignition delay. There is no uncertainty in the final equilibrium mass fractions and temperature, as these depend on thermodynamics and are independent of rate; these values can vary when there is uncertainty in species enthalpies of formation, but that is not considered in this work.

1.1.2 Uncertainty factors

The range of uncertainty in a parameter is reported in the chemical kinetics literature as an uncertainty factor; this is a convenient specification because it is only a single number. All parameters are assumed to have independent log normal distributions and (where u is the uncertainty factor for a particular parameter q):

$$\mathrm{P}\left(q < \frac{\mathrm{median}(q)}{u}\right) = 0.025 \qquad \mathrm{P}\left(q > u\,\mathrm{median}(q)\right) = 0.025 \tag{1.5}$$

This information is sufficient to determine the parameters of the lognormal distribution.

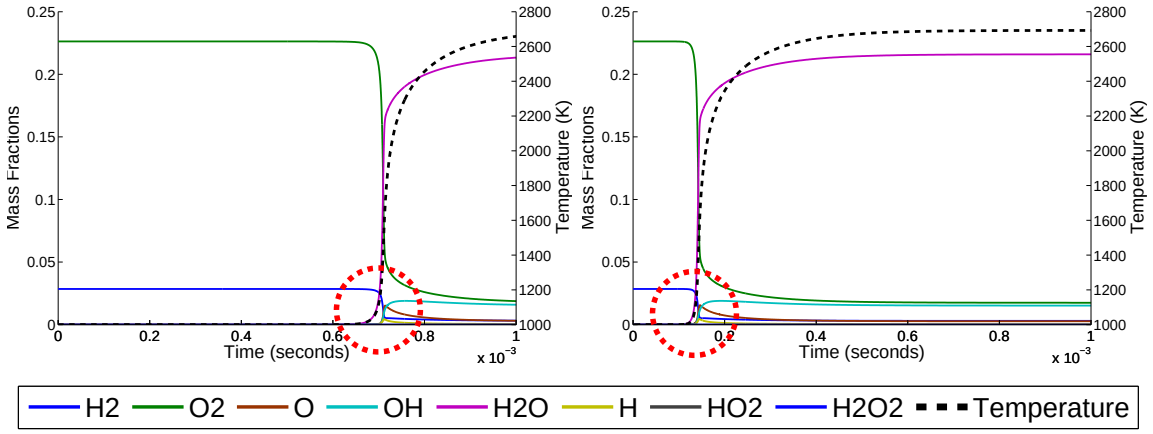

Figure 1-2: Reaction profiles for a hydrogen-oxygen system demonstrating the time evolution of species mass fractions and temperature, starting from a stoichiometric hydrogen-air mixture. Nitrogen was present, but is not shown because its mass fraction remains constant throughout. The plots are for two different values of the activation energy in the reaction $H + O_2 = O + OH$; these values differed by less than 5% from the nominal, but nevertheless caused a significant change in ignition time - this is highlighted with red circles.

1.1.3 Timescales and stiffness

Chemical kinetics systems are always characterized by a wide range of time scales in their dynamics, as demonstrated in figure 1-3. This causes stiffness and means that the ODE is expensive to evaluate, as an explicit solver cannot be used. The different timescales arise because of different reaction rates - some are fast and some are slow.

Figure 1-4 demonstrates why reduction is possible by showing solutions to a stiff 3-species ODE for a variety of initial conditions. This ODE was given in [62] and is designed to replicate the stiffness properties of a chemical system in 3D space to aid visualization. It is clear in the figure that the system rapidly relaxes to a 2D plane and then a line, but it is stiff because fast timescales are acting to constrain it to these lower dimensional manifolds; an explicit integrator would need to inefficiently capture these fast timescales, even though the solution evolves on a far slower timescale.

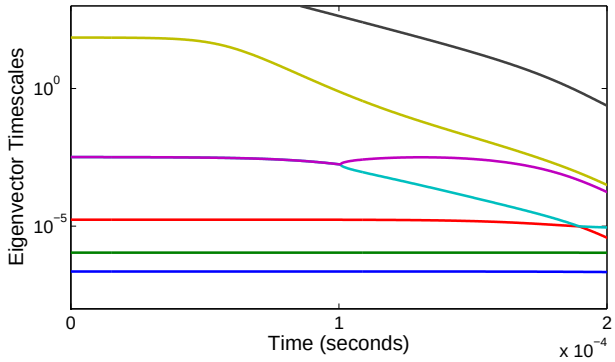

Figure 1-3: Timescales associated with the eigenvectors of the Jacobian of a linearized chemical source term in the pre-ignition region for a hydrogen-oxygen reaction with an initial stoichiometric mixture at 1000 K. Notice the wide range of values and the significant spacing between them.

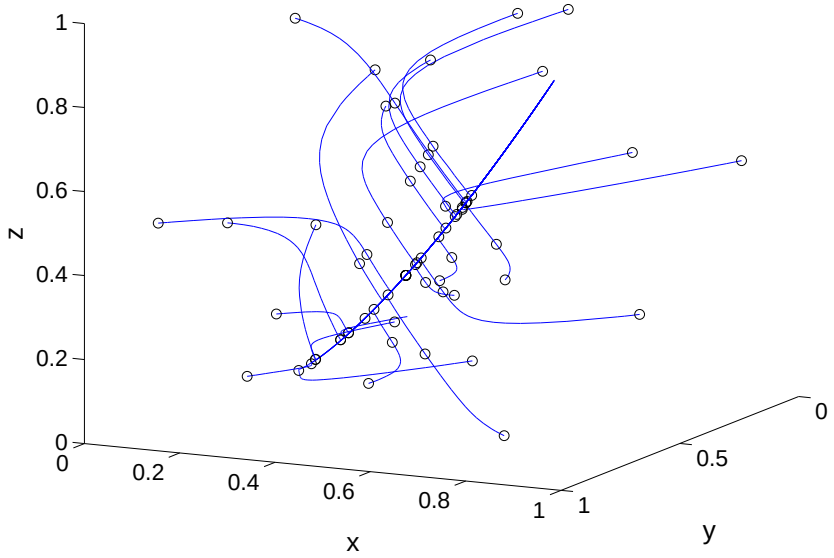

Figure 1-4: A stiff 3D ODE is solved at a variety of initial conditions. All trajectories converge initially to a 2D plane and then to a line; the symbol ∘ is used to denote initial conditions and convergence to lower dimensional manifolds. Movement to the line is very fast compared to movement along it; at this scale, there is no noticeable movement along the line in a timestep equal to the time required to move to the line from any of the initial conditions.

1.2 Model simplification and reduction

Strategies for reducing computational expense can generally be classified as either simplification or reduction:

- Reduction methods aim to reduce stiffness or otherwise improve computational efficiency. These attempt to capture movement along the slow manifold (as seen in figure 1-4) without the performance penalty associated with capturing the fast modes that constrain the system to the manifold. Some schemes also emphasize local dimension reduction in addition to stiffness removal - this is mentioned again briefly in section 1.2.1.

- Simplification methods aim to completely remove species and reactions from the original mechanism to form a *skeletal mechanism*. It is possible for some stiffness to be removed by simplification if particularly stiff reactions are deemed to be unimportant, but the emphasis here is on dimension reduction and lower computational expense at each timestep due to evaluating fewer reaction rates, as the evaluation of the exponential function for each reaction rate is the most expensive part of a chemical kinetics simulation. Note that the expense of a simulation tends to scale linearly with the number of reactions and quadratically with the number of species [61].

Simplification and reduction should not be used with low dimensional reaction mechanisms, as there is unlikely to be any justification for removing species or reactions without affecting the accuracy of the results. For this reason, no attempt has been made in this work to reduce the hydrogen-oxygen mechanism presented earlier in this chapter. However, the recently developed very high dimensional methods with thousands of species are ideal candidates for simplification and reduction.

1.2.1 In a CFD context

This work does not consider the effects of transport, but the ability to use a simplified or reduced model with a CFD code is the main source of motivation, as described in

this section. A CFD code adds transport to the model, so the governing ODE for the system is now expressed in terms of both position $\mathbf{x}$ and associated state $\mathbf{z_x}$:

$$\frac{\partial \mathbf{z_x}}{\partial t} = \mathbf{g}(\mathbf{z_x}) + \boldsymbol{\chi}(\mathbf{x}, \mathbf{z_x}) + \mathbf{D}(\mathbf{x}, \mathbf{z_x}) \tag{1.6}$$

where $\mathbf{g}$ is the chemical source term, $\boldsymbol{\chi}$ is the convection term, and $\mathbf{D}$ is the diffusion term. The timescales associated with diffusion and convection are generally in the range of 10^{-3} to 10^{-5} seconds. When comparing these values with figure 1-3, it is clear that this ODE is still stiff and that the stiffness is due to the chemical source term $\mathbf{g}$.

Operator splitting techniques can be used to reduce the impact of the stiffness on the transport terms; these would involve evaluating the chemical source term at smaller timesteps than the transport terms - this avoids the expense of evaluating the transport terms more often than necessary. However, the simulation as a whole remains stiff because the expense associated with evaluating the chemical source term at small timesteps has not been reduced; model reduction and simplification techniques are therefore applicable in this context. In fact, model reduction is now far more important than in the zero dimensional case because each node in the CFD code has an associated state and these states must be updated at every timestep by solving individual ODEs. When considering that there are many thousands of nodes, it is clear that the computational expense has risen by orders of magnitude compared to the zero dimensional case, even without considering the additional expense associated with evaluating the transport terms.

An overview of the interaction between a CFD code and a reduced chemical model is given in figure 1-5. Tabulation is used in an attempt to reduce the number of individual chemical kinetics ODEs that need to be solved at every timestep; if states at some of the nodes at the current timestep are similar to states that have been encountered earlier in the simulation, then the state at the next timestep can simply be cheaply loaded from the table rather than being expensively recomputed. Tabulation methods can also be used to reduce the expense of some of the additional processing

required by model reduction schemes, but that is not shown here and is described in chapter 3. A popular tabulation scheme used with chemical kinetics simulations is ISAT [41].

Tabulation is only practical for fairly low dimensional problems, so this provides further motivation for model reduction - the CFD code in the figure is shown to operate on a set of reduced variables. This also gives a significant reduction in memory requirements because each node requires only a much smaller state vector to be stored. However, reducing the dimension in the CFD code in this way requires additional care when implementing model reduction schemes, as most are designed to give reasonable projections to low dimensional manifolds for the chemical kinetics problem without considering the impact of these projections on the transport problem. Furthermore, a model reconstruction scheme is required so that the chemical source term can be evaluated when only the low dimensional parametrization is known; this will be discussed in more detail in chapter 2.

Model reduction schemes typically reduce a model to a manifold locally, so the dimension and correct parametrization (i.e. set of reduced variables) will change and therefore the set of reduced variables tracked by the CFD code must also be updated over time. Note that the use of an alternative model reduction scheme designed for stiffness removal rather than dimension reduction would also be beneficial in reducing computational expense to a lesser extent, but this would not reduce memory requirements or assist with tabulation.

1.2.2 Reasons for focusing on simplification

This work focuses entirely on simplification, although reduction methods are reviewed to provide context and interesting background information. Simplification and reduction are complementary steps; simplification is carried out first and then reduction can be applied later. Simplification simply involves removing species and reactions from the mechanism, so the simplified mechanism can then be used in a standard chemical kinetics code without changes. In contrast, reduction involves changing the code in an attempt to remove stiffness and possibly reduce the dimension of the prob-

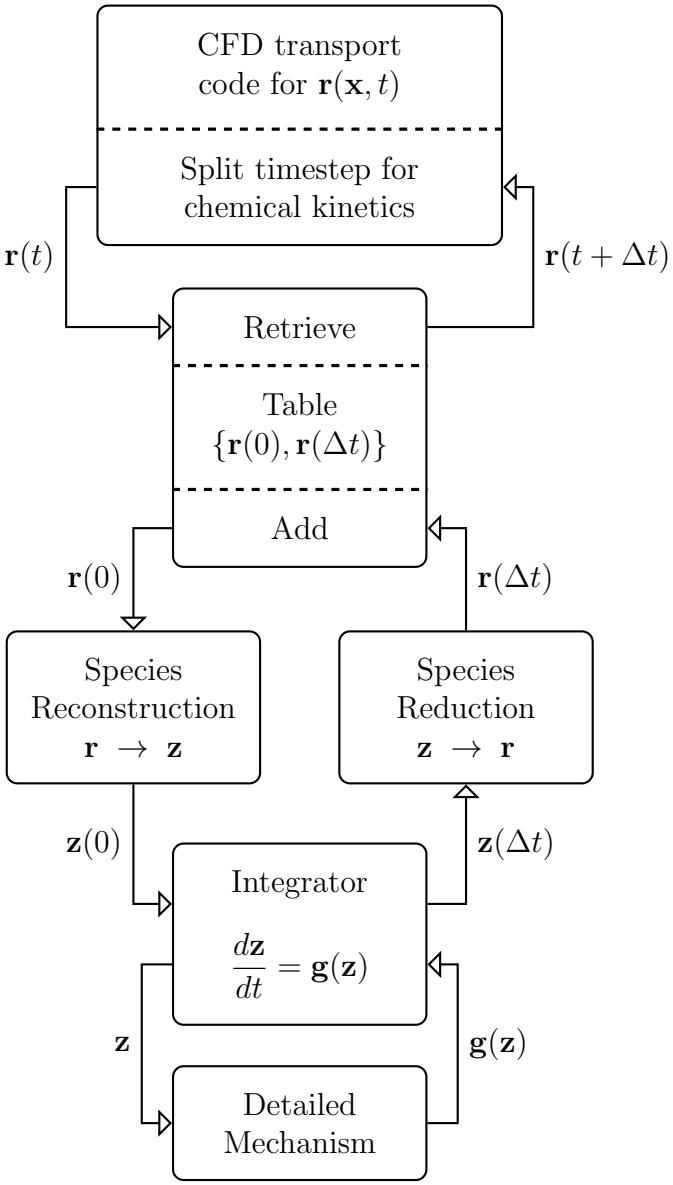

Figure 1-5: An example approach to incorporating detailed chemical kinetics into a CFD code, as described in [44]. Notice that tabulation is used to reduce the expense of evaluating the chemical source term if similar states are encountered at different times or positions. This becomes practical when a model reduction scheme is used to reduce the dimension of the problem by representing the state in terms of a low dimensional parametrization $\mathbf{r}$ instead of the full state $\mathbf{z}$. An alternative to the detailed mechanism calculation shown here is to integrate $\mathbf{r}$ directly - the chemical source term can be tabulated in terms of $\mathbf{r}$ and retrieved without reconstruction of $\mathbf{z}$.

lem. A code that implements reduction can also be used with a simplified mechanism without needing to start from the original full mechanism.

Not only can a simplified mechanism be used more easily by an end-user, but it is also certain to be cheaper - removing species and reactions always reduces computational expense, whereas reduction schemes must typically incorporate significant additional processing and stiffness removal may not compensate for this.

Finally, the impact of simplification on transport does not need to be separately considered, as the equations remain the same. However, reduction actually changes the equations and so care must be taken to ensure that changes to the chemical source term are compatible with the diffusion and convection terms; some model reduction methods do explicitly consider this, but others do not.

1.3 Computational tools

This work was carried out using code written in C++, but a variety of existing tools were also used:

- Chemical kinetics calculations used the library *tchem* [48], which was selected because it can calculate analytical Jacobian matrices for the chemical source term; these improve accuracy in the CSP method described in chapter 3.
- Integration was carried out with the library *CVODE* [16], which is a variable-order variable-step BDF solver.
- Random samples were generated with the *GNU scientific library* [9].
- Output was stored in *HDF5* files [54]; this is a convenient hierarchical data storage format.

1.3.1 Governing equations in more detail

An overview of the equations governing chemical kinetics was presented in section 1.1. The equations were expressed in terms of concentrations, as that allows for more

compact notation and therefore a simpler introductory explanation. This section contains more detailed equations in the form actually implemented in the code used for this work; all terms are now expressed in terms of mass fractions, as this form is more useful in numerical simulations when ensuring mass conservation. The following notation is used in this section:

- $\bar{W}$ is the molecular weight of the mixture.
- W_j is the molecular weights of species j.
- $\dot{\omega}_j$ is the molar production rate of species j.
- h_j is the molar enthalpy of species j.
- Y_j is the mass fraction of species j.
- C_j is the molar concentration of species j.
- $\mathbf{s}_i$ is the stoichiometric vector of reaction i.
- ρ is the mixture density.
- R is the universal gas constant.
- T is the temperature.
- N_R is the number of unidirectional reactions.
- N_S is the number of species.
- c_p is the specific heat capacity of the mixture at constant pressure.

The ODE for Y_j is as follows (and is equivalent to equation (1.2) for concentration units):

$$\rho \frac{dY_j}{dt} = \dot{\omega}_j W_j \tag{1.7}$$

where:

$$\dot{\omega}_j = \sum_{i=1}^{N_R} (\mathbf{s}_i)_j F^i \tag{1.8}$$

F^i can be evaluated using (1.3) after calculating the species molar concentrations:

$$C_j = \rho \frac{Y_j}{W_j} \tag{1.9}$$

The density also needs to be calculated; it can be found from the ideal gas law because the pressure is known and constant:

$$\rho = \frac{p\bar{W}}{RT} \tag{1.10}$$

where:

$$\bar{W} = \frac{1}{\sum_{j=1}^{N_S} \frac{Y_j}{W_j}} \tag{1.11}$$

An ODE for temperature is still needed; this was previously expressed only in terms of the enthalpy change for each reaction. It is now stated with an explicit enthalpy change calculation derived from the difference in enthalpies of products and reactants:

$$\rho c_p \frac{dT}{dt} = -\sum_{j=1}^{N_S} \dot{\omega}_j h_j \tag{1.12}$$

1.4 The aim of this work

There have been a number of prior studies into the effect of uncertainty on chemical reaction mechanisms, including [40] and [43]. However, these were aimed simply at assessing the impact of uncertainty on the results of chemical kinetic simulations; there was no analysis of the impact on model reduction or simplification schemes. This work is aimed at studying uncertainty in the simplification context, with a particular

emphasis on factors that need to be considered in the development of an uncertainty-aware simplification scheme. Results and discussion are presented in chapter 5 and chapter 6 describes an initial step towards incorporating uncertain information into a specific existing simplification algorithm.

1.4.1 Objectives under uncertainty

In this work, an uncertainty-aware simplification scheme is considered to have one of two objectives:

- Further simplification than in the deterministic setting by allowing for more error in outputs at nominal values; this is now justifiable because the full model output is itself uncertain and so it is not necessary to demand that the reduced model output should match it exactly. An appropriate error criterion is shown below:

$$\frac{|u^F - u^R|}{\sigma^F} \tag{1.13}$$

  where u^F is a full model output at nominal values, u^R is the corresponding output from the reduced model, and σ^F is the standard deviation of the full model output. Figure 1-6 demonstrates why this error criterion is a reasonable choice.

- Ensuring that the simplified model is capable of reproducing not only the nominal values of the full model outputs, but also their distributions under uncertainty. At the least, the reduced model output distribution should fall within the bounds of the full model output distribution. An appropriate error measure is the K-L divergence D_{KL}, which is a non-symmetric measure of the difference between two pdfs:

$$D_{KL}(\, P \,|\, Q \,) = \int_{-\infty}^{\infty} p(u) \log \frac{p(u)}{q(u)}\, du \tag{1.14}$$

  where p is the full model pdf, P is the full model cdf, q is the reduced model

pdf, and Q is the reduced model cdf. Note that the integrand becomes infinite whenever the support of the reduced model is not contained within the support of the full model; use of this measure therefore at least ensures that the bounds are reasonable. Figure 1-7 demonstrates the use of this error criterion.

These objectives need not be entirely mutually exclusive; for example, focusing on the first is likely to result in a smaller mechanism than when attempting to reproduce the full output pdf, but it would still be desirable to ensure that the pdf lies within reasonable bounds.

Although the first objective may appear to offer less accuracy than the second, it is still useful. Without simplification, CFD problems with large mechanisms may be completely intractable. It is hoped that improved simplification schemes can make them more tractable, eventually even in the context of uncertainty quantification where large numbers of samples will be required. However, only the first objective is currently relevant because it is not yet feasible to obtain more than a small number of samples.

1.5 The GRI-Mech 3.0 mechanism

The mechanism selected for study in this work was GRI-Mech 3.0 [52], which models methane-air combustion. It consists of 53 species and 325 reactions, most of which are reversible. This mechanism was selected because its size was considered large enough to allow for meaningful reduction without being large enough to be too computationally expensive to evaluate repeatedly in Monte Carlo simulations. Furthermore, it is used widely in the chemical kinetics literature and is often used in 'benchmark' tests of new schemes, so a study into the impact of uncertainty on the simplification of this mechanism is likely to be of broad interest.

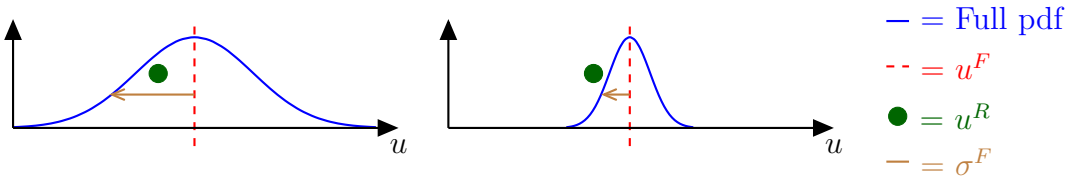

Figure 1-6: The simple error criterion (as explained in section 1.4.1) gives a measure of how well the reduced model prediction at nominal values fits within the output pdf of the full model. The situation on the left would be considered acceptable and this is reflected in the error criterion with scaling by a large σ^F to give a small error. The situation on the right would not be acceptable and this is reflected in the error criterion with scaling by a small σ^F to give a large error. Note that these pictures are simply cartoons to illustrate the above point and are not consistently normalized. Furthermore, outputs of interest would likely not have Gaussian pdfs and the full model nominal value prediction of any output need not coincide with the mean.

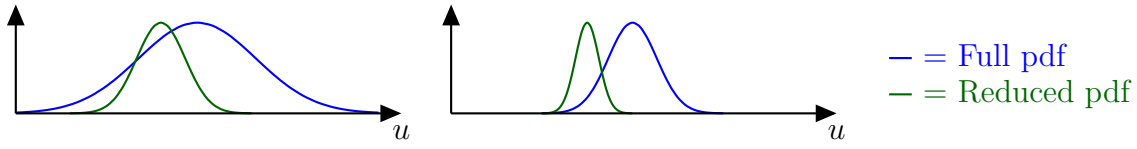

Figure 1-7: The K-L divergence measures the fit of two pdfs and strongly penalizes a reduced model output pdf that falls outside the bounds of the full model output pdf. Ideally, the reduced pdf would match the full pdf, but the situation on the left is still likely to be considered acceptable in many cases; the K-L divergence would be relatively low in this case. However, the situation on the right is always unacceptable and the K-L divergence would reflect this. Note that these pictures are simply cartoons to illustrate the above point and are not consistently normalized. Furthermore, outputs of interest would likely not have Gaussian pdfs like these.

Chapter 2

Literature Review

This chapter provides an overview of some of the more prominent model simplification and reduction strategies. Note that, although most methods incorporate some measure of local error, there are no global guarantees and all methods are heuristics. An overview of the field can be found in [39].

2.1 Overview by Lu and Law

A more recent overview of model reduction and simplification can be found in [28], which also describes a means of effectively using a combination of simplification schemes, primarily CSP and variants of DRG. CSP is discussed in detail in chapter 3, but there is no separate discussion of DRG because the simplification algorithm is identical to the Valorani CSP algorithm of section 3.3.2 other than in the use of a different metric. The DRG metric is inferior because, unlike the CSP metric, it does not consider system dynamics. However, it is cheaper than CSP and is therefore used to augment CSP simplification in [28]. Further details on DRG can be found in [26].

2.2 QSSA

This is the Quasi Steady-State Approximation method. It is the oldest of the methods presented here, having first been proposed in the early 1900s. Along with the Partial

Equilibrium method presented below, this is sometimes called one of the classical methods, but it is still used even in some of the most recent literature.

After an initial transient period, certain species are assumed to be in instantaneous equilibrium and this allows an algebraic constraint to be imposed. However, care must be taken when implementing this approximation because it does not actually imply that the time derivative of the concentration of a QSS species can be set to zero. Instead, it just provides an algebraic equation for the concentration of the QSS species in terms of the other species. When using the notation of (1.2), the QSSA equation for species i is:

$$g_i(\mathbf{z}) = 0 \tag{2.1}$$

Note that equation (2.1) is in fact a replacement for one of the terms in (1.2) and should not be substituted into (1.2). As noted above, this is not the same as setting the derivative to zero, as z_i can still change. A clear explanation of the difference is given in [22] and a graphical demonstration can be found in [8]. One of the main weaknesses of the method is in selection of species; this traditionally relies heavily on investigator experience. However, analysis, guidelines, and an inefficient species selection algorithm are presented in [58]. These authors also mention using QSSA to remove stiffness, but they conclude that it is generally not very effective unless too many species are removed for the model to be accurate. As the method targets species rather than reactions, this is not surprising.

The other major weakness of the method is that a solution is not guaranteed to exist. Using the notation of section 1.3.1, equation (2.1) can be written as:

$$0 = \dot{w}_i \tag{2.2}$$

or more explicitly as:

$$0 = \sum_{i=1}^{N_R} (\mathbf{s}_i)_j F^i = \sum_{i=1}^{N_R} (\mathbf{s}_i)_j \, k_i \prod_{j=1}^{N_S} C_j^{\nu_{ij}} \tag{2.3}$$

At this point, a non-linear solver must be used; Newton's method is recommended in the literature, but it not guaranteed to converge. A solution to this problem has been proposed in the form of LQSSA, which is briefly discussed in the following section.

2.2.1 LQSSA

LQSSA was introduced to improve QSSA by guaranteeing both the existence and uniqueness of a solution. It goes further than QSSA by making the assumption that the concentrations of QSS species are small. This means that any F^i involving more than one QSS reactant is small enough to be neglected. The remaining terms are then linear. If some of the QSS species concentrations are in fact not small, then they can either be removed from the list of QSS species (thereby resulting in a lower level of reduction) or else a hybrid linear/non-linear approach can be used.

Further details can be found in [27]. Note that the authors also describe a method for efficiently solving the linear equations created in the LQSSA method; this is known as QSSG and involves arranging the equations into a block upper triangular structure that can be solved in quadratic time.

2.3 Partial Equilibrium

The partial equilibrium approximation is considered as one of the classical methods along with QSSA. It is intended primarily for use in stiffness removal. It was often used incorrectly in very early papers by setting reaction rates directly equal to zero, but correct usage is now understood and will be briefly described below; more details are given in [22] and [23].

Rather than assuming that the net creation and destruction rates for a given species cancel out (as with QSSA), partial equilibrium assumes that the forward and reverse rates of a given reaction approximately cancel out; this gives an algebraic constraint on the concentrations of the species involved in the reaction. Unlike with QSSA, the algebraic constraint cannot be substituted for one of the species in the rate equations; this is because the criterion for partial equilibrium selection is only that

the net rate of the reaction should be small *when compared to the individual forward and reverse rates of the reaction itself and not of other reactions.* The reaction is considered exhausted at this point, but its rate of progress has not reached zero and is instead a small approximately constant value that may still be comparable to that of many of the remaining reactions. Rather than using the algebraic equation directly or attempting to calculate the constant, the reaction's rate equation should be differentiated to obtain an equation that can be safely used.

Substitutions among the original reaction rate equations can be used to remove the rate term for the exhausted reaction. The number of equations is then reduced by one, but this new differentiated partial equilibrium equation can be used to form a complete system without any of the stiff rate terms from the exhausted reaction.

2.4 ILDM - Intrinsic Low Dimensional Manifold

This scheme is designed to take advantage of time scale separation and the main principle is just a special case of CSP, which will be described in detail in section 3. Despite the same underlying theory, the implementation of the methods and subsequent modifications are distinct. The motivation for ILDM was to allow table look-up of the chemical source terms during a reacting flow simulation, as demonstrated in figure 1-5, whereas the original motivation for CSP primarily concerned stiffness removal. The original ILDM paper was [32] and additional references are mentioned below where appropriate.

The method involves finding the eigenvectors of the linearized Jacobian matrix of $\mathbf{g}$ in equation (1.2). A dimension M is selected such that the M most negative eigenvectors are deemed to be exhausted modes that are now conserved; a full explanation of this idea can be found in chapter 3 during the discussion of CSP, but a theoretical framework at that level of detail was not originally provided by the authors of ILDM.

When conserved directions have been found, the system is confined to a low dimensional manifold and can be parametrized by a reduced set of species. Although the original paper mentioned that the reduced set of variables can be arbitrarily chosen,

this later proved not to be the case; the variables must be chosen to ensure that the manifold is not multi-valued at any point in parameter space. An automatic method for determining appropriate parameters in terms of linear combinations of species was described in [30] and further improvements to manifold identification were made in [34].

As ILDM was designed for table look-up, the original authors provided implementation details in [31]. An adaptive table storage scheme was developed to avoid the unnecessary storage overhead associated with constant grid spacing throughout parameter space when the chemical source term may not change significantly in some areas. The table can be used to store eigenvectors in addition to the chemical source term to avoid some of the additional overhead that would otherwise be associated with ILDM.

2.5 The Method of Fraser and Roussel

Fraser and Roussel attempted to identify slow manifolds by analyzing trajectories [8, 38, 47]. Their method was designed to develop an algebraic expression for the manifold using symbolic functional iteration to provide better performance than table look-up. The authors treat only simple 2D and 3D examples, but the method can be used in higher dimensions. However, it can only identify a 1D manifold, no matter what the dimension of the system; this was later extended to higher dimensions by work referenced in section 2.5.2.

Although this method therefore appears to be very limited, it is nevertheless widely cited because it has a desirable feature in the form of the identification of an inertial manifold; this is a slow manifold that corresponds to a real trajectory and is globally attracting for the dynamics of the system. Crucially, motion along the manifold is determined by the full model without approximation, whereas a non-inertial manifold must use projection with a result highly dependent on the choice of projection operator; in the context of the ILDM discussion above, this would be dependent on the accuracy of the exhausted eigenvector approximation.

Despite this attractive feature, the Fraser and Roussel method is global in the sense that the construction of the entire manifold or at least a large part of it is generally required to identify a point in parameter space corresponding to a point in full composition space; this is because large numbers of trajectories must be calculated. This is a major disadvantage of the method.

2.5.1 A brief description of the method

The first step in the method is to substitute known constraints (including element number) into the chemical source term to reduce the number of variables. A species concentration C_j is selected as a parameter denoted by s and then ODEs for the remaining species concentrations can be found in phase space:

$$\frac{dC_i}{ds} = \frac{\dot{C}_i}{\dot{s}} \tag{2.4}$$

These equations are re-arranged in the following functional form:

$$C_i = f(C_i', C_{k\neq i,j}, s) \qquad \text{where} \qquad C_i' = \frac{dC_i}{ds} \tag{2.5}$$

Note that elementary reactions involve rate terms with algebraic rate orders that are generally no higher than 3 (most often only 1 or 2), so the equations can *always* be rearranged algebraically in this way. The method requires an initial estimate to be chosen for each C_i' (usually from a partial equilibrium or QSSA solution) and it then proceeds by iteration:

$${}^{n+1}C_i = f({}^{n}C_i', {}^{n}C_{k\neq i,j}, s) \tag{2.6}$$

An algebraic representation of the manifold is usually found with a small number of iterations on n. However, as with ILDM, the method breaks down if the manifold is multi-valued in s at any combination of C_i; finding a suitable species for s is not straightforward and the authors discuss an alternative approach involving linear combinations of species for s.

2.5.2 Improvements by Skodje and Davis

Significant improvements to the Fraser and Roussel method were made in [7, 51] to allow for the identification of higher dimensional manifolds. However, these authors found that it was necessary to use a numerical iteration scheme rather than algebraic because the algebraic expressions became too large for more complicated problems. It was noticed that the Fraser and Roussel manifold is generally more accurate than the ILDM manifold, by as much as 10% in highly curved regions; this demonstrates the advantage of using an inertial manifold. However, the ILDM manifold is a good starting point to use for Fraser and Roussel in order to guarantee convergence in general.

The above references also identified other related methods for finding inertial manifolds, but they have reached only a low level of development and do not appear to have been widely used. However, they provided some inspiration for later developments for ILDM, as referenced in that section. Some further details can be seen in [6].

2.6 Lumping

Rather than removing species altogether, lumping amalgamates them. The number of reactions may also be reduced because some are effectively duplicated when the species are amalgamated. This approach was first proposed in [65] for linear systems only, but has since been extended to cover non-linear systems too; an overview and analysis is provided in [56]. Note that the lumps themselves can also be either linear or non-linear combinations of species.

A major disadvantage of lumping is the loss of information concerning individual reactions and species. This has been found by [28] to only be acceptable in special cases, particularly for large hydrocarbons that have many isomers; these isomers typically have sufficiently similar properties for lumping to be used.

The above references deal with discrete lumping, which is the most useful form. However, continuum lumping methods also exist; these involve integrating a set of

species defined by a continuous parameter. They are not relevant in most situations, but they are intended to provide some tractability in ill-defined problems. This approach is covered in more detail in [39].

2.7 RCCE

This is rate-controlled constrained equilibrium. It was originally called the rate-controlled partial equilibrium method [19], but the name was later changed to avoid confusion with the partial equilibrium method mentioned above. RCCE identifies the rate controlling reactions because these instantaneously constrain the system. Note that any equilibrium condition is constrained; even when a reaction appears to have finished and the remaining species are in equilibrium, this is only equilibrium constrained by the slow time scale of nuclear reactions. All constraints here are linear and possibilities include: element number (when not dealing with nuclear reactions), bonded atom pairs, number of molecules, translational or vibrational energy (as energy transfer between molecules within a degree of freedom may be on much faster timescales compared to those in different degrees of freedom), and others given in [18]. Note that suitable constraints must be chosen by the user, as there is no automatic method. They can be identified from suitable combinations of reactions with slow rates.

This method requires the assumption that changes in constraints are sufficiently slow for a system to evolve through a series of quasi-static states remaining close to the static equilibrium determined by the instantaneous values of the constraints. In other words, the relaxation time of molecular collision processes is short compared to the characteristic time for a change in constraints. This assumption is justified as being consistent with assumptions commonly made when deal with thermodynamics in general.

The method works by replacing the ODE of the form shown in equation (1.2) by a much smaller set of differential equations describing the time evolution of the constraints, which are effectively constants as far as slower reactions are concerned; these

new equations are found by differentiating the linear algebraic forms of the constraint equations. The differential equations are specified in terms of the reaction rates of the rate controlling reactions, so these must be known. However, the other reactions are not used, which is a major advantage because the accuracy of their reaction rate parameters is then irrelevant; only the controlling reaction rate parameters need be known accurately.

In some cases, only the constraint may be important (e.g. sometimes the constraint can be the concentration of a particular species of interest), so the constraint equations can just be integrated. In others, reconstruction of all species may be desired; this will be always be required at least partially in order to calculate all relevant reaction rates. If the number of known constraints is equal to the number of species, then the species concentrations can be found immediately by solving a linear system under the assumption that the constraints are instantaneously constant. However, only a small number of constraints are generally known and this is why the slowly changing assumption mentioned above is required; it is assumed that all concentrations have relaxed to instantaneous equilibrium and so can be calculated by maximizing the entropy for an adiabatic system (or minimizing the Gibbs free energy for an isothermal system) under the given constraint. Lagrange multipliers are used to enforce the constraint in the optimization problem.

Some of the controlling equations are not differential and are instead always constant; these include element number in the absence of transport or nuclear reactions. This does not affect the formulation of the problem, but it reduces the expense of integration. Finally, note that shifting equilibrium is a special case of RCCE in which constraints on element numbers are fixed and thermodynamic variables are allowed to change slowly while acting as instantaneous constraints.

2.8 ICE-PIC

The motivation for this method is to tabulate full states at specific reduced parameter values on an inertial manifold. Unlike other methods, there is a strong focus on species

reconstruction rather than the initial parametrization; this is because it is intended to fulfill the reconstruction role demonstrated in figure 1-5. ICE-PIC is a further development of pre-image curves and TGLDM, so these are presented first and the additional features of ICE-PIC development are summarized in section 2.8.3.

2.8.1 Pre-image curves (PICs)

The main idea of this approach involves considering points in the full composition space at earlier times that could have led to the current reduced set of species. It is unique in that it constructs this manifold locally, whereas other methods involving inertial manifolds are impractical in high dimensions due to being global, as described in the Fraser and Roussel section. Full details are given in [44] and a brief outline is presented here. A linear parametrization of the following form is used to convert from the full state vector $\mathbf{z}$ to a reduced state vector $\mathbf{r}$:

$$\mathbf{r} = \mathbf{B}^T\mathbf{z} \tag{2.7}$$

Note that $\mathbf{B}$ is a fat matrix and is generally composed of unit vector columns so that $\mathbf{r}$ is a reduced set of species rather than a combination of species. At a particular time, the method aims to reconstruct $\mathbf{z}$ from $\mathbf{r}$, noting that there are many possible combinations of $\mathbf{z}$ that would result in a given $\mathbf{r}$. A starting point is selected using the RCCE idea of maximizing the entropy under the constraints of constant element number and $\mathbf{r} = \mathbf{B}^T\mathbf{z}$. However, rather than selecting this as a solution, a pre-image curve is constructed - this is a curve through values of $\mathbf{z}$ at earlier times that would eventually result in a consistent $\mathbf{z}$ at the current time when integrated forward.

Values of $\mathbf{z}$ along the curve are parametrized by s, which is the arc length starting from 0 at the RCCE solution. Note that many possible curves could be chosen, but the method always selects the curve with minimum curvature so that it does not not stray far from the most likely solution in terms of entropy; the concentrations along this curve eventually tend to an asymptote as s increases towards its value at the boundary of the feasible region. The value at the asymptote is selected as the point

on the curve from which the current state originated. This is known as the pre-image. Construction of the curve requires integration of points along it using the full system in order to identify points that are pre-images and this is clearly computationally expensive, but values can be tabulated so that the expense is not repeated.

2.8.2 TGLDM

This method creates trajectory-generated low-dimensional manifolds. It is presented in [42] in the context of 2D manifolds, but an approach for extension to higher dimensions is also discussed.

The polar coordinates r and θ are chosen as the parametrization of the 2D low dimensional manifold. The initial conditions are determined by θ (known as the generator) and then r is the normalized arc length, so all trajectories follow a radial line from the circumference of a circle of radius 1 to equilibrium at the origin. The set of realizable compositions is bounded by non-negativity constraints, constraints for the conservation of element number, and additional constraints that define the manifold in composition space. If required, an additional constraint can be imposed to ensure that the boundary passes through the stoichiometric mixture.

Trajectories are generated at intervals along the boundary and then values are tabulated at various points along the trajectories for use in later calculations. There are three limitations: the first is that the set of major species (corresponding to the parametrization) must be chosen manually. The second is that there is no guarantee that the trajectories actually lie on the manifold; we might expect them to lie on the manifold when sufficiently close to equilibrium, but we are just starting from arbitrary points in composition space (subject to the constraints of the choice of major species) and so cannot expect the initial values to necessarily lie anywhere near the manifold. The final limitation is that, as with the Fraser and Roussel method, this method is global and a large portion of the manifold must be constructed to determine a particular composition.

2.8.3 Summary

ICE-PIC is described in [45] and combines the advantage of the simple TGLDM parametrization with the local nature of the pre-image curve (PIC) manifold. The method starts with a constrained equilibrium (CE) manifold, which is a manifold created under the RCCE approximation; the boundaries of this manifold are found and all trajectories emanating from these boundaries form an ICE (Invariant Constrained Equilibrium) manifold. The ICE manifold is then parametrized in the TGLDM manner. At first, it would seem that this is a global method and therefore impractical because a large number of trajectories would need to be generated along the ICE boundary in order to obtain the appropriate value of the TGLDM generator parameter (which corresponds to θ in the 2D example explained above). However, the PIC method can be used to get around this.

The intersection of the pre-image region with the CE manifold is found; this is a curve that can be followed until the boundary of the CE region is reached (instead of following the minimum curvature curve in the original PIC method). A trajectory from this point is then in the ICE manifold and the ICE-PIC solution is found as the value of the full composition on this trajectory when it becomes consistent with the reduced composition. The TGLDM parameters are the starting point on the ICE manifold and the distance moved along the trajectory until the correct reduced composition is reached, so the method is therefore local rather than global - it was possible to obtain the composition without needing to construct and explore a large region of the manifold.

Chapter 3

Computational Singular Perturbation

The Computational Singular Perturbation method (CSP) was originally developed in the context of model reduction [22, 23], but has since also been used for simplification [61, 60]; both uses will be described in this chapter. CSP explicitly exploits the time scale separation in chemical kinetics problems by identifying fast and slow modes in the reaction mechanism.

3.1 CSP Fundamentals

Mathematically, the CSP method aims to split the chemical source term of the ODE into fast and slow components; using the notation introduced in chapter 1, this is:

$$\frac{d\mathbf{z}}{dt} = \mathbf{g}(\mathbf{z}) = \mathbf{g}^{\text{fast}}(\mathbf{z}) + \mathbf{g}^{\text{slow}}(\mathbf{z}) \tag{3.1}$$

The new fast and slow chemical source terms are found by expressing the original chemical source term in terms of N modes, where N is the length of the state vector $\mathbf{z}$; this is 1 + number of species because temperature is also included and is treated as a species throughout this chapter. These modes will be indexed in order of their associated timescales such that $i = 1$ corresponds to the fastest mode. A number

M will be identified such that the first M modes are designated fast in the sense of being exhausted; the remaining modes are designated slow. If it is assumed that nuclear reactions are not also being considered, then the element number N_e is always conserved in chemical kinetics and therefore the slowest N_e modes are conserved and have infinite timescales. In general, the number of conserved modes (not necessarily only due to element number) will be denoted by N_c.

The modes will be denoted by $\mathbf{a}_i$ and each $\mathbf{a}_i$ will have an associated mode amplitude $f^i(\mathbf{z})$; this is found by projecting $\mathbf{g}$ onto $\mathbf{a}_i$. If the dual basis vectors for the $\mathbf{a}_i$ are denoted by $\mathbf{b}^i$, the $f^i(\mathbf{z})$ are given explicitly as:

$$f^i(\mathbf{z}) = \mathbf{b}^i \cdot \mathbf{g} \qquad\qquad \mathbf{b}^i \cdot \mathbf{a}_j = \delta^i_j \tag{3.2}$$

The chemical source terms can now be written explicitly with this new notation:

$$\mathbf{g}^{\text{fast}}(\mathbf{z}) = \sum_{i=1}^{M} \mathbf{a}_i(\mathbf{z}) f^i(\mathbf{z}) \qquad\qquad \mathbf{g}^{\text{slow}}(\mathbf{z}) = \sum_{i=M+1}^{N} \mathbf{a}_i(\mathbf{z}) f^i(\mathbf{z}) \tag{3.3}$$

Although this representation is mathematically meaningful, it is not as clearly physically meaningful as the more familiar physical representation in terms of stoichiometric vectors and reaction rates, as shown in chapter 1.

Now denote by $\mathbf{A}$ the matrix with the $\mathbf{a}_i$ as columns and denote by $\mathbf{B}$ the matrix with the $\mathbf{b}^i$ as rows. The choice of $\mathbf{a}_i$ and $\mathbf{b}^i$ will be seen by calculating the time derivative of the mode amplitudes from equation (3.2) and expressing it in terms of $\mathbf{A}$, $\mathbf{B}$, and the Jacobian $\mathbf{J}$ of $\mathbf{g}$:

$$\frac{d\mathbf{f}}{dt} = \mathbf{\Lambda f} \qquad\qquad \mathbf{\Lambda} = \left[\frac{d\mathbf{B}}{dt} + \mathbf{BJ}\right]\mathbf{A} \qquad\qquad \mathbf{J} = \frac{\partial \mathbf{g}}{\partial \mathbf{z}} \tag{3.4}$$

If $\mathbf{\Lambda}$ did not contain the time derivative of $\mathbf{B}$, it could be diagonalized by choosing the $\mathbf{a}_i$ to be the eigenvectors of $\mathbf{J}$. There would then be complete mode separation and each $\mathbf{a}_i$ would have a clearly identifiable timescale in the form of the reciprocal of the absolute value of the real part of the associated eigenvalue λ_i because the f_i

would evolve according to:

$$f_i(t) = f_i(0)\, e^{\lambda_i t} \tag{3.5}$$

where the real parts of the λ_i are assumed to be negative - see section 3.1.1 for a brief discussion of positive eigenvalues. The CSP vectors must be real and the treatment of complex eigenvectors is described in section 3.1.3.

The presence of the time derivative in (3.4) means that that the eigenvectors of $\mathbf{J}$ do not actually diagonalize $\mathbf{\Lambda}$, but they are nevertheless still used for the $\mathbf{a}_i$; they no longer provide complete mode separation, but the $\mathbf{a}_i$ are used only in fast and slow groups, so mode mixing is acceptable among the fast vectors and separately among the slow vectors. This means that non-zero terms can be present in the Jacobian in the upper left diagonal $M \times M$ fast block and lower right diagonal $(N-M)\times(N-M)$ slow block.

Significant mode mixing between the sets of fast and slow vectors is unacceptable because it would not allow for the separation required by the CSP representation; this would appear in the form of non-zero off-diagonal blocks. The authors of CSP found that such off-diagonal terms are generally negligible and hence do not significantly affect the CSP mode separation. The full CSP method also contains a refinement scheme to further depress these off-diagonal blocks, but that involves considerable computational expense and has been found not to be necessary to achieve good accuracy when simulating hydrogen and methane systems [25]. It was therefore not used in this work. Refinements were first introduced in [23], a full analysis was conducted in [69, 70], and implementation details can be found in [62].

It only remains to describe M, the number of fast modes. Fast modes are those that locally make no significant contribution to the solution trajectory in one timestep, as they have been exhausted and now act only to constrain the solution to a lower dimensional manifold rather than to move it along that manifold. The M calculation is therefore based on a user-defined tolerance of a small acceptable error step if the fast modes were to be ignored during a timestep of the length of the shortest timescale

among the remaining slow modes. The modes are removed in increasing order of timescale until this error tolerance is reached and this is mathematically described by (where τ_i denotes the timescale associated with $\mathbf{a}_i$):

$$M = \max\left\{ m \;\middle|\; \left| \tau_{m+1} \sum_{i=1}^{m} f^i A_{ji} \right| \leq \epsilon_{\text{abs}} + \epsilon_{\text{rel}} \left| z_j \right| \quad \forall \quad 1 \leq j \leq N \right\} \tag{3.6}$$

3.1.1 Positive eigenvalues in CSP

Unlike large negative eigenvalues, large positive eigenvalues do not represent fast exhaustive modes; instead, they represent explosive modes that dominate the dynamics of the system. As a result, fast modes must always be kept and not discarded by CSP. The simplest approach is to arrange the $\mathbf{a}_i$ in order of eigenvalue from smallest to largest; note that this is in terms of actual value rather than absolute value, so the $\mathbf{a}_i$ with most negative eigenvalues will appear first as the fastest modes. This ensures that the $\mathbf{a}_i$ with positive eigenvalues are not classified as fast modes in the exhausted sense.

3.1.2 Numerically conserved modes

Numerical routines for finding eigenvectors (such as LAPACK's `dgeev`) are generally highly inaccurate for small eigenvalues, especially when much larger eigenvalues are also present. This problem is particularly severe in chemical kinetics, as it is not uncommon to find eigenvalues of order 10^{15} or larger while also having conserved mode eigenvalues of zero and other modes with eigenvalues small enough to be indistinguishable from zero to within the accuracy of the routine. These additional modes will now be described as numerically conserved and will be included in N_c, although the number of modes in this category is state-dependent and so now N_c is no longer constant.

3.1.3 An efficient and accurate calculation of B

The $\mathbf{b}^i$ were defined above as a dual basis for the $\mathbf{a}_i$; this suggests that $\mathbf{B}$ can be found by inverting $\mathbf{A}$. However, this is clearly inefficient and is also inaccurate.

The eigenvectors for eigenvalue zero in the output from a numerical routine are not well-conditioned, especially when the null space is large. Even if the vectors were exact to within floating point precision, they would still likely not be well-conditioned because the solver is only designed to find enough vectors to span the space without any consideration of whether or not some of the chosen vectors are close to being linearly dependent. The substantial error for zero eigenvalues amplifies this problem, resulting in an $\mathbf{A}$ matrix with a very large condition number so that it cannot be accurately inverted to find $\mathbf{B}$.

To work around this, it was proposed in [25] to replace the eigenvectors for eigenvalues of zero with singular vectors for singular values of zero, as found using an SVD. In exact arithmetic, these vectors should span the same space. However, the SVD vectors should have far better numerical conditioning because they are orthogonal. Although it would allow for the accurate inversion of $\mathbf{A}$ to form $\mathbf{B}$, this procedure is even more expensive.

An accurate and more efficient alternative proposed in this work is to use the left eigenvectors of $\mathbf{J}$ to avoid inversion. If all eigenvalues are distinct and appropriate normalizations are used, the left eigenvectors provide a dual basis for the right eigenvectors; this is proved below for right eigenvectors $\mathbf{v}_i$, left eigenvectors $\mathbf{w}_i^{\mathrm{T}}$, and eigenvalues λ_i:

$$\begin{aligned} \mathbf{w}_i^{\mathrm{T}}\mathbf{J} = \lambda_i \mathbf{w}_i^{\mathrm{T}} \quad &\Rightarrow \quad \mathbf{w}_i^{\mathrm{T}}\mathbf{J}\mathbf{v}_j = \lambda_i \mathbf{w}_i^{\mathrm{T}}\mathbf{v}_j \\ \mathbf{J}\mathbf{v}_j = \lambda_j \mathbf{v}_j \quad &\Rightarrow \quad \mathbf{w}_i^{\mathrm{T}}\mathbf{J}\mathbf{v}_j = \lambda_j \mathbf{w}_i^{\mathrm{T}}\mathbf{v}_j \\ &\qquad \mathbf{w}_i^{\mathrm{T}}\mathbf{v}_j = 0 \quad \text{for} \quad \lambda_i \neq \lambda_j \end{aligned} \tag{3.7}$$

Complex eigenvalues

This proof holds for both real and complex eigenvalues. However, further work is required in the complex case because CSP vectors must be real. Consider the complex pair of eigenvalues λ_α and λ_β. Let these eigenvalues be $\mu_r \pm i\mu_i$ and associated eigenvectors be $\mathbf{u} \pm i\mathbf{v}$ where μ_r, μ_i, $\mathbf{u}$, and $\mathbf{v}$ are all real. It is clear that span{$\mathbf{u}$, $\mathbf{v}$} is an invariant subspace evolving on the timescale of $1/\mu_r$, so the $\mathbf{a}_{\alpha,\beta}$ can now be chosen as $\mathbf{u}$ and $\mathbf{v}$.

It does not immediately follow that $\mathbf{b}^{\alpha,\beta}$ can be set equal to the left eigenvectors (denoted by $\mathbf{w}^{\mathrm{T}} \pm i\mathbf{s}^{\mathrm{T}}$) without further work - they span the correct space and hence would satisfy the requirements of a dual basis with respect to all different eigenvalues, but numerical eigenvalue routines do not choose the complex normalizing constant on these vectors to ensure that:

$$\mathbf{w}^{\mathrm{T}}\mathbf{v} = 0 \qquad\qquad \mathbf{s}^{\mathrm{T}}\mathbf{u} = 0$$

To work around this, $\mathbf{b}^\alpha$ is chosen as whichever of $\mathbf{w}^{\mathrm{T}}$ and $\mathbf{s}^{\mathrm{T}}$ has the largest component in the direction of $\mathbf{a}_\alpha$. Then $\mathbf{b}^\beta$ is given by (assuming that $\mathbf{b}^\alpha = \mathbf{w}^{\mathrm{T}}$ was chosen):

$$\mathbf{b}^\beta = \mathbf{s}^{\mathrm{T}} - \left(\mathbf{s}^{\mathrm{T}}\mathbf{a}_\alpha\right)\mathbf{a}_\alpha^{\mathrm{T}} \tag{3.8}$$

Finally, appropriate normalizations are applied to $\mathbf{b}^\alpha$ and $\mathbf{b}^\beta$.

Repeated eigenvalues and defective matrices

If eigenvalues were to be repeated in a non-defective matrix, the same orthogonalization idea could be applied to the eigenvectors as for complex eigenvalues. In a defective matrix, generalized eigenvectors could be used instead of eigenvectors because they still span an invariant subspace, but the same orthogonalization idea would still need to be used. However, as described in section 3.1.4, repeated eigenvalues are not expected to occur in practice and therefore defective matrices also do not occur.

The exception to this rule is for the eigenvalue zero, which will be repeated N_c times. It is not necessary to calculate or use CSP vectors for a conserved mode, but it must be shown that these vectors satisfy the appropriate orthogonality relationship with the CSP vectors for the non-conserved modes to ensure that the approach creates an accurate $\mathbf{B}$ from $\mathbf{A}$. The conserved subspace is spanned by appropriate eigenvectors and possibly generalized eigenvectors if the matrix is defective in the eigenvalue zero. It was shown in equation (3.7) that eigenvectors satisfy the orthogonality relationship and it will now be demonstrated that the same result holds for generalized eigenvectors. If we denote a left eigenvector with eigenvalue zero by $\mathbf{v}_1$, a generalized left eigenvector by $\mathbf{v}_2$, and a non-conserved left eigenvector by $\mathbf{b}^i$ with non-zero eigenvalue λ_i, then $\mathbf{b}^i\mathbf{v}_1 = 0$ by (3.7) and the result follows:

$$\begin{aligned} \mathbf{b}^i\mathbf{J} = \lambda_i\mathbf{b}^i \quad &\Rightarrow \quad \mathbf{b}^i\mathbf{J}\mathbf{v}_2 = \lambda_i\mathbf{b}^i\mathbf{v}_2 \\ \mathbf{J}\mathbf{v}_2 = \mathbf{v}_1 \quad &\Rightarrow \quad \mathbf{b}^i\mathbf{J}\mathbf{v}_2 = \mathbf{b}^i\mathbf{v}_1 = 0 \\ &\qquad\quad \mathbf{b}^i\mathbf{v}_2 = 0 \end{aligned} \tag{3.9}$$

Similarly, now denote a right eigenvector with eigenvalue zero by $\mathbf{w}_1^{\mathrm{T}}$, a generalized right eigenvector by $\mathbf{w}_2^{\mathrm{T}}$, and a non-conserved right eigenvector by $\mathbf{a}_i$ with non-zero eigenvalue λ_i, then $\mathbf{w}_1^{\mathrm{T}}\mathbf{a}_i = 0$ by (3.7) and the result follows:

$$\begin{aligned} \mathbf{J}\mathbf{a}_i = \lambda_i\mathbf{a}_i \quad &\Rightarrow \quad \mathbf{w}_2^{\mathrm{T}}\mathbf{J}\mathbf{a}_i = \lambda_i\mathbf{w}_2^{\mathrm{T}}\mathbf{a}_i \\ \mathbf{w}_2^{\mathrm{T}}\mathbf{J} = \mathbf{w}_1^{\mathrm{T}} \quad &\Rightarrow \quad \mathbf{w}_2^{\mathrm{T}}\mathbf{J}\mathbf{a}_i = \mathbf{w}_1^{\mathrm{T}}\mathbf{a}_i = 0 \\ &\qquad\quad \mathbf{w}_2^{\mathrm{T}}\mathbf{a}_i = 0 \end{aligned} \tag{3.10}$$

So the orthogonality relationship also holds for generalized eigenvectors; it therefore also holds for any vector in the conserved subspace, which consists of all linear combinations of eigenvectors and generalized eigenvectors with eigenvalue zero.

3.1.4 Typical eigenvalue types in a chemical kinetics problem

The time evolution of the eigenvalues of the Jacobian of the hydrogen-oxygen combustion mechanism from [68] is shown in figure 3-1. This plot contains some of the typical features associated with combustion mechanisms in general, such as:

- The noise associated with zero eigenvalues is apparent; inexact floating point arithmetic means that two of the three zero eigenvalues are not identically zero.
- One eigenvalue sometimes becomes small enough to be almost indistinguishable from the zero eigenvalues; this can then be temporarily treated as a numerically conserved mode.
- Complex pairs of eigenvalues exist, so the code must be able to deal with these.
- Eigenvalues sometimes cross; these are the only times at which they are repeated. It is extremely unlikely that a timestep would fall precisely at the moment that two eigenvalues cross, so the code need not be able to deal with these and therefore need not be able to deal with a defective matrix. If this were ever to occur during a simulation, the most straightforward solution would be to choose a slightly different timestep.

3.2 CSP for reduction

CSP was not used in a reduction setting in this work, but reduction is described here briefly anyway for background information. CSP can used in this context in one of two ways:

- Integrate the chemical source term using only the slow modes; the fast term in (3.1) can be dropped. This should remove the stiffness associated with the fast modes without significant loss of accuracy.
- If it is assumed that there will locally be no movement in the direction of the fast modes, they can be used to construct conservation equations for the species.

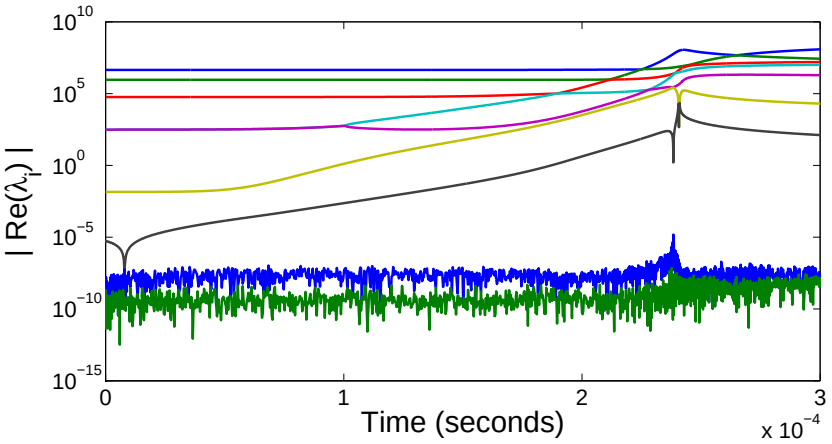

Figure 3-1: Time evolution of eigenvalues over the pre-ignition and ignition regions during the simulation of the hydrogen-oxygen combustion mechanism from [68] with initial conditions of 1000 K and stoichiometric mixture ratio. There are 10 variables including temperature, but only 9 lines are shown because one was identically zero. The three elements are H, N, and O. Note that absolute values are used here so that the results can be displayed on a logarithmic scale. The two lines that are initially merged represent a complex conjugate pair of eigenvalues and not a repeated eigenvalue.

> The ODE of (3.1) can then be rewritten as a DAE with reduced dimension, but it should be noted that DAE solvers are in general not as advanced and well developed as ODE solvers [1]. This type of reduction is generally associated with the G-Scheme [64], which is a recent method based on the same principles as CSP.

In both cases, some error will still be introduced. This is because the $\mathbf{a}_i$ vectors are only locally valid and depend upon the state. This means that the vectors may no longer adequately represent the fast and slow directions by the end of a timestep, but a correction can then be applied to bring the solution closer to the correct position on the manifold; more details are provided in [22, 23].

Although both cases above appear to offer computational savings, they are actually not beneficial in practice because of the expense of the additional CSP processing. However, these techniques can be combined with a tabulation scheme to avoid having to recompute the vectors at every timestep [25, 33, 37]; this has been found to give a decrease in computational expense compared to running the full model with no

reduction without significant loss of accuracy compared to standard CSP.

Both cases can offer good accuracy compared to the full model when used with an appropriate error threshold for choosing the fast modes, but it should be noted that they are heuristic methods with no error guarantees.

3.2.1 CSP radicals

The species removed during dimension reduction are known as CSP radicals; to ensure that it is valid to remove these species by substitution in the equations, they must be selected as the species with concentrations that change significantly in fast directions. Note that these CSP radicals are also often chemical radicals, but this need not always be the case. Mathematically, these species are identified as the M largest radical pointers, which are defined to be the diagonal elements of:

$$\sum_{i=1}^{M} \mathbf{a}_i \mathbf{b}^i \tag{3.11}$$

In other model reduction schemes, such as QSSA and Partial Equilibrium, it is often unclear as to which species should be removed. The correct choice even in these schemes is to use the CSP radical. This is explained in detail in [22] and a more thorough mathematical definition can be found in [63].

3.3 CSP for simplification

3.3.1 Definitions

CSP can be used to provide diagnostic information from the full model; this information can then be used to develop simplification schemes. Two types of diagnostic information are mentioned frequently in the literature - participation indices and importance indices:

- Participation indices measure the relevance of a reaction k to a CSP vector i. These are not used in this work, but are presented here for completeness (where

N_R is the number of unidirectional reactions such that each reversible reaction is counted as two separate reactions):

$$P_k^i = \frac{\mathbf{b}^i \cdot \mathbf{s}_k F^k}{\sum_{j=1}^{N_R} |\mathbf{b}^i \cdot \mathbf{s}_j F^j|} \tag{3.12}$$

- Importance indices measure the relevance of a reaction k to a species i in either a fast or slow sense by projecting the stoichiometric vector of k onto the fast or slow CSP vectors. Mathematically, these are given by:

$$(I_k^i)_{\text{slow}} = \left| \frac{\sum_{s=M+1}^{N-N_c} [\mathbf{a}_s]_i (\mathbf{b}^s \cdot \mathbf{s}_k) F^k}{\sum_{j=1}^{N_R} \left| \sum_{s=M+1}^{N-N_c} [\mathbf{a}_s]_i (\mathbf{b}^s \cdot \mathbf{s}_j) F^j \right|} \right| \tag{3.13}$$

$$(I_k^i)_{\text{fast}} = \left| \frac{\sum_{s=1}^{M} [\mathbf{a}_s]_i (\mathbf{b}^s \cdot \mathbf{s}_k) F^k}{\sum_{j=1}^{N_R} \left| \sum_{s=1}^{M} [\mathbf{a}_s]_i (\mathbf{b}^s \cdot \mathbf{s}_j) F^j \right|} \right| \tag{3.14}$$

A further normalization should be carried out to give importance indices between 0 and 1 because this allows for more meaningful comparison of importance indices under different conditions [60]. Without this normalization, the range of values (and even orders of magnitude) varies widely, especially at different values of M. The I_k^i defined above are rescaled for each i by dividing by $\max\{I_m^i \,|\, 1 \leq m \leq N_R\}$.

Note that fast importance is only considered to be a valid concept for radical species, which were defined above in section 3.2.1. Both importance measures are only considered to be meaningful for non-trace species, where trace species are those with only a negligible concentration [61].

3.3.2 The Valorani algorithm

This algorithm makes use of CSP importance indices to identify species and reactions that should be retained in a reaction mechanism [61]. A set of target species of interest must be declared in advance; this is denoted by S_0 and will generally contain some of the major species and temperature. The algorithm works by applying a threshold η

to importance indices, so this must also be decided in advance. The algorithm itself is as follows (where R denotes reaction sets, S denotes species sets, and the suffix rad denotes only the current CSP radicals in the set):

S_{global} is empty

R_{global} is empty

for all initial conditions and times of interest indexed by i **do**

$S_i^0 = S_0$

R_i^0 is empty

$j = 0$

repeat

$R_i^{j+1} = \{k \,|\, (I_k^p)_{\text{slow}} > \eta \text{ for any } p \in S_i^j\} \bigcup \{k \,|\, (I_k^p)_{\text{fast}} > \eta \text{ for any } p \in (S_i^j)_{\text{rad}}\}$

$S_i^{j+1} = \{$ species involved in reactions in R_i^{j+1} $\}$

$j = j + 1$

until $S_i^j = S_i^{j-1}$

$S_{\text{global}} = S_{\text{global}} \bigcup S_i^j$

$R_{\text{global}} = R_{\text{global}} \bigcup R_i^j$

end for

Reduced sets of species and reactions have now been found. Note that an optional additional step is to recover all remaining reactions that involve only the species in the final reduced set; computational expense generally scales only linearly with number of reactions and quadratically with number of species, so full reaction recovery is relatively cheap and is always performed in this work whenever the Valorani algorithm is used.

3.3.3 Performance of the Valorani algorithm

The algorithm was used in [60, 61] to gain insight into mechanism structure by varying the threshold parameter η - it was found that reactions involved in the same path of the mechanism tended to be removed at similar threshold values. Furthermore, it

was found that the algorithm could be used to generate simplified mechanisms for GRI-Mech 3.0 that would speed up computation by up to 8 times before errors in the reaction profile became unacceptable.

This algorithm is however only a heuristic and good performance is not always guaranteed. To demonstrate this, the ignition delay was selected as an error measure and then the Valorani algorithm was compared to an exhaustive combinatorial search over all possible reduced mechanisms of every given size that still contained S_0. Such a search was highly computationally expensive and was applied to a modified GRI-Mech 3.0 mechanism; the original mechanism was too large for a combinatorial search to be tractable, so nitrogen chemistry was removed to reduce it to only 34 species. Although essential for tracking pollutants, the nitrogen chemistry does not have a significant impact on the combustion of methane. Even with the reduced size, it was still not possible to run a combinatorial search for some intermediate mechanism sizes in a reasonable time.

Figure 3-2 demonstrates that it is possible to find better mechanisms than with the Valorani algorithm at any given size. However, the Valorani algorithm ran in only minutes, compared to days for the exhaustive combinatorial search. Furthermore, the ignition delay is only one error measure and does not track significant differences in reaction profile or even ensure that equilibrium conditions do not change significantly, although it was also found that differences in equilibrium were not apparent except for some of the smallest mechanisms that could only be generated with the exhaustive search and not with the Valorani algorithm. From this it would seem reasonable to conclude that significant improvements to the Valorani algorithm are possible for the larger mechanism sizes, but performance of the Valorani algorithm is already likely to be sufficient here anyway due to the small relative error in ignition delay.

Finally, it is unlikely that any heuristic algorithm would be able to achieve the small errors shown in the exhaustive search for the largest mechanism sizes; figure 3-3 demonstrates that this small error is confined to a very small number of reduced mechanisms, whereas the error in the majority is in the higher range found by the Valorani algorithm.

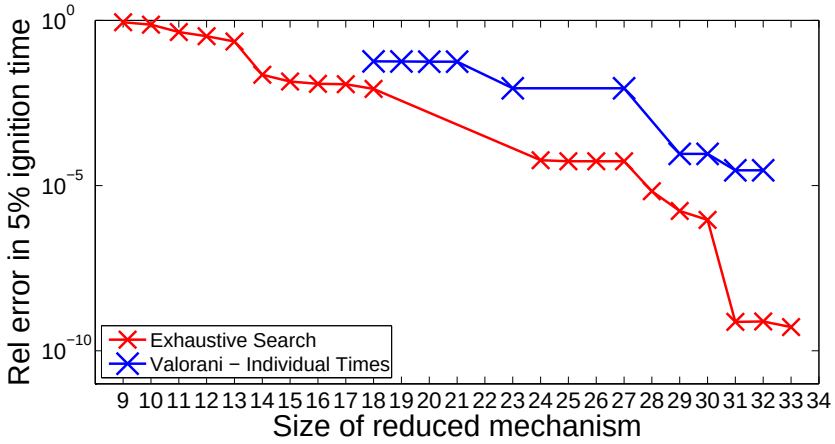

Figure 3-2: Error in ignition delay (defined as a temperature rise to 5% of peak) for simplified mechanisms generated with the Valorani algorithm compared to those generated by an exhaustive combinatorial search for each mechanism size. Full mechanism is GRI-Mech 3.0 with nitrogen chemistry removed; this contains 34 species. Initial temperature was 1000 K and error is the average over five initial states with equal amounts of methane and hydrogen and uniformly spaced mixture ratios from 0.6 to 2.

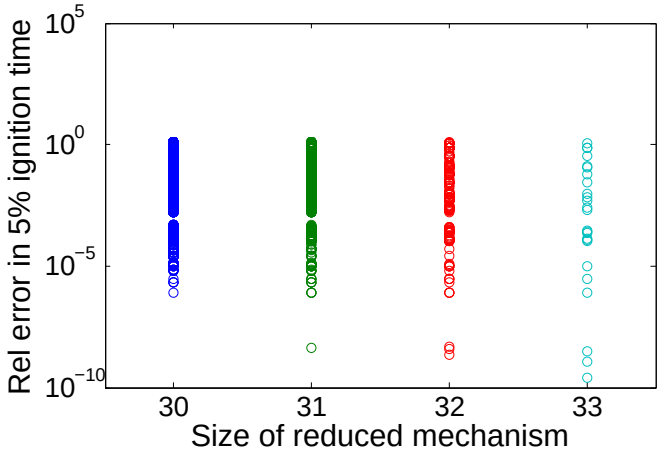

Figure 3-3: Error in ignition delay (defined as a temperature rise to 5% of peak) for every possible reduced mechanism of a few sizes under the same conditions as figure 3-2. Those not shown are clustered together in the same way as size 30, but a small number of mechanisms significantly outperform the others for the three largest sizes; this is reflected in the superior performance of these three specific sizes in figure 3-2.

Chapter 4

Uncertainty Quantification Tools

A variety of uncertainty quantification tools can be used to examine the impact of uncertainty in reaction rate parameters on outputs of interest. The output used for examples in this section is ignition time, but other outputs are mentioned in section 5.2 and chapter 7. Note that the sensitivity results and polynomial chaos expansion were in terms of log ignition time to ensure non-negativity.

4.1 Monte Carlo Simulation

Monte Carlo simulations can be used to quantify how uncertainty in the rate parameters propagates through the chemical kinetics model. They simply involve evaluating the model at a large number of samples until convergence is achieved in the output pdfs of interest. However, convergence is slow and so large numbers of expensive function evaluations are typically required - convergence in expected value is only proportional to $1/\sqrt{n}$, where n is the number of samples [24]. In this work, 25000 samples was found to be a reasonable number to use for the mechanisms and outputs under consideration. As an example, figure 4-1 demonstrates the convergence of pdfs of ignition delay for GRIMech 3.0 under specific initial conditions.

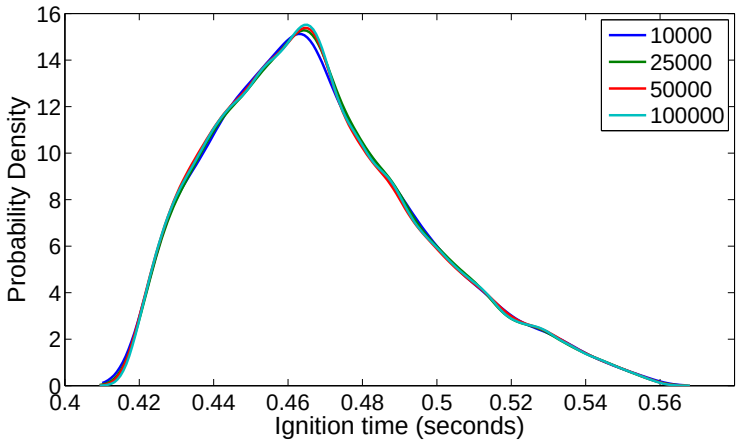

Figure 4-1: Convergence of Monte Carlo sampling of ignition delay output for GRIMech 3.0 with uncertainty in the hydrogen-oxygen reactions at initial temperature of 1000 K in a stoichiometric methane-air mixture. The distribution of the input random variables is as described later in chapter 5. The numbers in the legend are different numbers of samples.

Random samples were generated using the GNU Scientific Library (GSL), which implements a number of different algorithms for generating random samples from a choice of distributions [9]. In terms of guaranteed lack of correlation between samples, the most reliable random number generator available in GSL is RANLUX; details on this algorithm can be found in [17, 29].

4.2 Polynomial Chaos

A Polynomial Chaos (PC) representation is a surrogate for a full model that depends on uncertain parameters. The full model is expensive to evaluate, but a polynomial chaos model is simply a polynomial function of random variables and so it is cheap to evaluate. The fundamentals of PC are described in this section, but more extensive descriptions of PC and other related techniques can be found in [24, 21, 67, 36].

Denote by $\boldsymbol{\xi}$ the vector containing the H *independent* random variables of interest; these are the uncertain rate parameters in this work. Denote the outputs of interest by $\boldsymbol{u}(\boldsymbol{\xi}, \boldsymbol{\theta})$, where $\boldsymbol{\theta}$ contains the deterministic parameters that include initial conditions

and rates of non-random reactions. The order P polynomial chaos representation of $\boldsymbol{u}(\boldsymbol{\xi}, \boldsymbol{\theta})$ is in terms of multi-dimensional basis functions $\Psi_{\boldsymbol{i}}(\boldsymbol{\xi})$ and coefficients $\boldsymbol{u_i}(\boldsymbol{\theta})$:

$$\boldsymbol{u}(\boldsymbol{\xi}, \boldsymbol{\theta}) = \sum_{||\boldsymbol{i}||_1 \leq P} \boldsymbol{u_i}(\boldsymbol{\theta})\, \Psi_{\boldsymbol{i}}(\boldsymbol{\xi}) \tag{4.1}$$

The multi-dimensional basis functions $\Psi_{\boldsymbol{i}}(\boldsymbol{\xi})$ are products of single dimensional basis functions $\phi_q(\xi_j)$ (where q denotes order) and are determined by the entries of the multi-index $\boldsymbol{i}$, as shown:

$$\Psi_{\boldsymbol{i}}(\boldsymbol{\xi}) = \prod_{j=1}^{H} \phi_{i_j}(\xi_j) \tag{4.2}$$

The $\phi_q(\xi_j)$ are generally chosen as orthogonal polynomials to reduce computational expense because it is then straightforward to find the $\boldsymbol{u_i}(\boldsymbol{\theta})$ from:

$$\mathbb{E}\left[\boldsymbol{u}(\boldsymbol{\xi}, \boldsymbol{\theta})\Psi_{\boldsymbol{i}}(\boldsymbol{\xi})\right] = \boldsymbol{u_i}(\boldsymbol{\theta})\, \mathbb{E}\left[\Psi_{\boldsymbol{i}}(\boldsymbol{\xi})\right] \tag{4.3}$$

The expectation on the right can be evaluated analytically, but the expectation on the left must be evaluated numerically with either an intrusive or non-intrusive approach. An intrusive approach involves reformulating the governing ODE to directly calculate PC coefficients instead of the original output; this is generally the most accurate and efficient approach. However, it requires the code to be completely rewritten and, more significantly, it is not straightforward when the output of interest is not an explicit function of the state vector governed by the ODE. For example, a method for computing a PC expansion of the ignition time is not clear.

An alternative non-intrusive approach requires computing the expectation integral through quadrature. This is more expensive, but the main ODE solver code does not need to be changed and it is straightforward to evaluate the log ignition time PC expansion. This approach was used in this work and a brief overview of techniques is given in section 4.2.2.

Having decided on using a non-intrusive approach, it only remains to select an

appropriate polynomial order. An order 4 polynomial was found to be sufficient in most cases and this is demonstrated in figure 4-2. Note that pdfs were verified for many representative sets of reactions and initial conditions under consideration to give confidence that an appropriate order was selected.

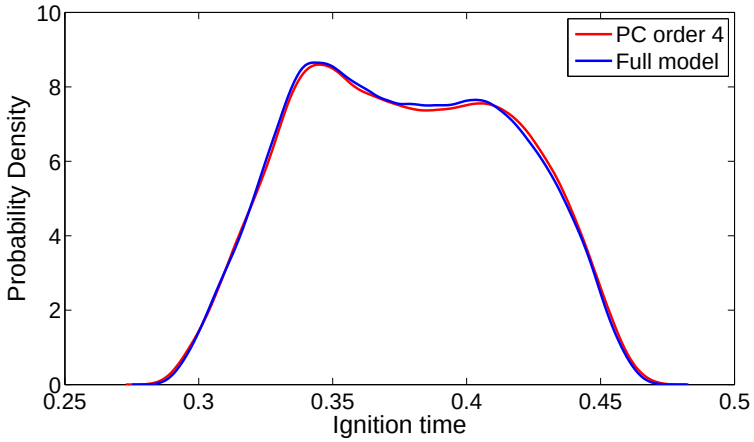

Figure 4-2: A pdf constructed from a polynomial chaos expansion of order 4 matches the full model pdf for GRIMech 3.0 when uncertainty is introduced into the CH4-CH3 reactions for an initial temperature of 1000 K with a stoichiometric methane-air mixture. The distribution of the input random variables is as described later in chapter 5.

4.2.1 Augmenting the PC basis

Polynomial chaos expansions are only capable of approximating smooth functions. This means that they cannot be used to represent importance indices directly because of the discontinuous changes that occur when M changes, as seen in equations (3.13) and (3.14). This issue is discussed again in chapter 5.

Forming PC expansions of the CSP vectors is also problematic because there are discontinuities in eigenvectors when eigenvalues cross. If the locations of these discontinuities are known, then the PC basis polynomials can be augmented with functions designed to fit them. However, this involves considerable additional computational expense because these new functions are not orthogonal to the polynomials. Furthermore, it has been found that these functions must capture the location of the discontinuity precisely and do not improve accuracy if there is even a small error in

the location. They are therefore generally not useful in practice because the output must always be known exactly in order to generate a PC representation, but clearly the PC representation is not needed when the output is already known. More information on stochastic eigenvalue problems in relation to polynomial chaos can be found in [12, 13].

There has also been recent work on attempting to exploit the structure of the Jacobians in chemical kinetics problems to create polynomial chaos approximations to the CSP vectors without the problems associated with discontinuities. This will not be discussed further here, but more information can be found in [49].

4.2.2 Quadrature

It is useful to provide a brief overview of quadrature-based techniques for computing expectation integrals for polynomial chaos in the non-intrusive case. The most obvious approach is to evaluate multi-dimensional integrals simply by taking tensor products of a 1D quadrature rule, but this is very inefficient and rapidly becomes computationally intractable as the number of dimensions increases.

Sparse quadrature schemes are a more effective alternative; the boundaries of the space are filled with quadrature points according to the 1D rule, while only a relatively small number of points are used in the interior by an extrapolation of the quadrature rule to high dimensions. The sets of quadrature points are generally built in hierarchies, so the level of grid refinement can be increased until the integrals computed at two successive refinement levels yield values that are within a user-specified error tolerance of each other. This procedure is demonstrated in figure 4-3, where examples of sparse grids at different refinement levels are shown. More details can be found in [10], which also contains references for other non-quadrature approaches, such as Monte Carlo integration.

Finally, dimension adaptive sparse quadrature (DASQ) controls the level of grid refinement in each direction depending on the degree of non-linearity of the function in that dimension. This can clearly result in substantial computational time and memory savings in many cases, especially when very high dimensional functions are

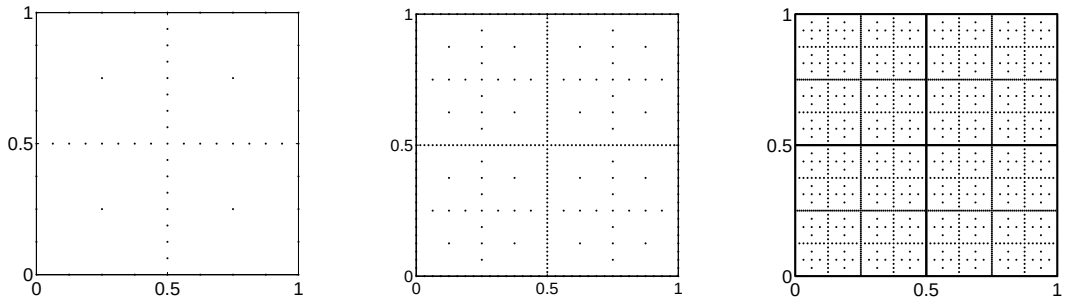

Figure 4-3: Clenshaw-Curtis sparse grids for a 2D input with increasing levels of refinement from left to right.

highly non-linearly in only a small number of dimensions. A DASQ scheme based on Clenshaw-Curtis quadrature was used in this work to calculate the PC integrals. A full description of the method and can be found in [11].

4.3 Global Sensitivity Analysis

Global sensitivity analysis determines the extent to which the input parameters are responsible for the output variance, both individually and in groups. First order and total effect sensitivities were used in this work and will be discussed here briefly; more details and descriptions of other sensitivity indices can be found in [53].

When a scalar output u depends on input random variables $\boldsymbol{\xi}$, the contribution of a particular ξ_i to the variance of u can be determined by starting with the following decompositions (where the notation $\boldsymbol{\xi}_{\sim i}$ indicates all elements of $\boldsymbol{\xi}$ except for ξ_i):

$$\mathrm{var}_{\boldsymbol{\xi}}(u) = \mathrm{var}_{\xi_i}\left(\mathbb{E}_{\boldsymbol{\xi}_{\sim i}}[u \,|\, \xi_i]\right) + \mathbb{E}_{\xi_i}\left[\mathrm{var}_{\boldsymbol{\xi}_{\sim i}}(u \,|\, \xi_i)\right] \tag{4.4}$$

$$\mathrm{var}_{\boldsymbol{\xi}}(u) = \mathrm{var}_{\boldsymbol{\xi}_{\sim i}}\left(\mathbb{E}_{\xi_i}[u \,|\, \boldsymbol{\xi}_{\sim i}]\right) + \mathbb{E}_{\boldsymbol{\xi}_{\sim i}}\left[\mathrm{var}_{\xi_i}(u \,|\, \boldsymbol{\xi}_{\sim i})\right] \tag{4.5}$$

The first term on the right of equation (4.4) is the variance explained by ξ_i acting alone. Similarly, the first term on the right of equation (4.5) is the variance explained by the elements of $\boldsymbol{\xi}_{\sim i}$ acting either alone or on in combination with each other, but not with ξ_i. The second term on the right of this equation therefore indicates the variance contribution from ξ_i both when acting alone and in combination with any

of the other random variables. After rescaling, these are used to define first order S_i and total effect S_{T_i} sensitivity indices for ξ_i:

$$S_i = \frac{\mathrm{var}_{\xi_i}\left(\mathbb{E}_{\boldsymbol{\xi}_{\sim i}}[u \,|\, \xi_i]\right)}{\mathrm{var}_{\boldsymbol{\xi}}(u)} \tag{4.6}$$

$$S_{T_i} = \frac{\mathbb{E}_{\boldsymbol{\xi}_{\sim i}}\left[\mathrm{var}_{\xi_i}(u \,|\, \boldsymbol{\xi}_{\sim i})\right]}{\mathrm{var}_{\boldsymbol{\xi}}(u)} = 1 - \frac{\mathrm{var}_{\boldsymbol{\xi}_{\sim i}}\left(\mathbb{E}_{\xi_i}[u \,|\, \boldsymbol{\xi}_{\sim i}]\right)}{\mathrm{var}_{\boldsymbol{\xi}}(u)} \tag{4.7}$$

These are useful for model simplification because they indicate the impact of removing reactions on the output pdf; this has implications for a successful model reduction strategy. Further discussion of the utility of sensitivity indices will be presented in chapters 5 and 7.

4.3.1 Combined effect sensitivities

First order sensitivities measure the effect of a random variable acting on the variance alone. Total effect sensitivities measure the effect of a random variable on the variance whether acting alone or in combination with any of the other variables. More detailed information can be found by calculating the combined effect of specific groups of random variables acting together; explicit forms for these combined effect sensitivities will not be given here, as they are not used individually in this work and details can be found in [53]. However, the overall significance of these combined effects can be identified simply by looking at the sum of the S_i or S_{T_i}; this can be seen by decomposing the variance with the ANOVA-HDMR (Analysis of Variance High Dimensional Model Representation), noting that n is the stochastic dimension and multiple subscript suffices identify the species involved in a particular combined effect sensitivity index:

$$1 = \sum_i S_i + \sum_i \sum_{j>i} S_{ij} + \sum_i \sum_{j>i} \sum_{k>j} S_{ijk} + \ldots + S_{1,2,\ldots,n} \tag{4.8}$$

The left hand side is 1 because sensitivity indices are all scaled by $\mathrm{var}_{\boldsymbol{\xi}}(u)$, as in the S_i and S_{T_i} definitions above. All sensitivity indices are positive and S_{Ti} is the sum

of all sensitivities that include i (either by itself or with others), so the following are both indicators of the size of the combined effects:

$$\left(\sum_i S_{T_i}\right) - 1 \qquad\qquad 1 - \sum_i S_i \tag{4.9}$$

4.3.2 Calculating Sensitivity Indices

The sensitivity indices given in equations (4.6) and (4.7) can be calculated easily with polynomial chaos expansions. An outline of the derivation for S_i is presented in this section with respect to a scalar output u (which may depend on additional deterministically chosen parameters that are not explicitly indicated here). Some preliminary results are required:

$$\mathbb{E}_{\boldsymbol{\xi}_{\sim i}}[u \,|\, \xi_i] = \sum_{||\mathbf{j}||_1 \le P} u_{\mathbf{j}} \, \mathbb{E}_{\boldsymbol{\xi}_{\sim i}}\left[\Psi_{\mathbf{j}}(\boldsymbol{\xi}) \,\middle|\, \xi_i\right] \tag{4.10}$$

$$\begin{aligned}
\text{where } \mathbb{E}_{\boldsymbol{\xi}_{\sim i}}\left[\Psi_{\mathbf{j}}(\boldsymbol{\xi}) \,\middle|\, \xi_i\right] &= \mathbb{E}_{\boldsymbol{\xi}_{\sim i}}\left[\prod_{m=1}^{H} \phi_{j_m}(\xi_m) \,\middle|\, \xi_i\right] \\
\text{(by independence)} \qquad &= \mathbb{E}_{\boldsymbol{\xi}_{\sim i}}\left[\phi_{j_i}(\xi_i) \,\middle|\, \xi_i\right] \prod_{m \ne i} \mathbb{E}_{\boldsymbol{\xi}_{\sim i}}\left[\phi_{j_m}(\xi_m)\right] \\
&= \phi_{j_i}(\xi_i) \prod_{m \ne i} \delta_{0 j_m}
\end{aligned} \tag{4.11}$$

The last step above assumed that the order 0 orthogonal polynomials are normalized to 1. The variance can now be computed by invoking independence again:

$$\mathrm{var}_{\xi_i}\left(\mathbb{E}_{\boldsymbol{\xi}_{\sim i}}[u \,|\, \xi_i]\right) = \mathbb{E}_{\xi_i}\left[\left(\mathbb{E}_{\boldsymbol{\xi}_{\sim i}}[u \,|\, \xi_i]\right)^2\right] - \mathbb{E}_{\xi_i}\left[\mathbb{E}_{\boldsymbol{\xi}_{\sim i}}[u \,|\, \xi_i]\right]^2 \tag{4.12}$$

$$\text{where } \mathbb{E}_{\xi_i}\left[\mathbb{E}_{\boldsymbol{\xi}_{\sim i}}[u \,|\, \xi_i]\right] = u_{\mathbf{0}} \tag{4.13}$$

$$\begin{aligned}
\text{and } \mathbb{E}_{\xi_i}\left[\left(\mathbb{E}_{\boldsymbol{\xi}_{\sim i}}[u \,|\, \xi_i]\right)^2\right] &= \mathbb{E}_{\xi_i}\left[\sum_{||\mathbf{j}||_1 \le P} \sum_{||\mathbf{k}||_1 \le P} u_{\mathbf{j}} \phi_{j_i}(\xi_i) \, u_{\mathbf{k}} \phi_{k_i}(\xi_i) \prod_{m \ne i} \delta_{0 j_m} \prod_{n \ne i} \delta_{0 k_n}\right] \\
&= \sum_{||\mathbf{j}||_1 \le P} u_{\mathbf{j}}^2 \, \mathbb{E}_{\xi_i}\left[\phi_{j_i}^2\right] \prod_{m \ne i} \delta_{0 j_m}
\end{aligned} \tag{4.14}$$

Note that $\text{var}_{\boldsymbol{\xi}}(u)$ can be computed with similar working, so S_i is:

$$S_i = \frac{\text{var}_{\xi_i}\left(\mathbb{E}_{\boldsymbol{\xi}_{\sim i}}[u \mid \xi_i]\right)}{\text{var}_{\boldsymbol{\xi}}(u)} = \frac{\sum\limits_{\substack{||\mathbf{j}||_1 \leq P \\ ||\mathbf{j}||_1 > 0}} u_{\mathbf{j}}^2 \, \mathbb{E}_{\xi_i}\left[\phi_{j_i}^2\right] \prod\limits_{m \neq i} \delta_{0 j_m}}{\sum\limits_{\substack{||\mathbf{j}||_1 \leq P \\ ||\mathbf{j}||_1 > 0}} u_{\mathbf{j}}^2 \prod\limits_{l=1}^{H} \mathbb{E}_{\xi_i}\left[\phi_{j_l}^2\right]} \tag{4.15}$$

Similarly, S_{T_i} is:

$$S_{T_i} = 1 - \frac{\text{var}_{\boldsymbol{\xi}_{\sim i}}\left(\mathbb{E}_{\xi_i}[u \mid \boldsymbol{\xi}_{\sim i}]\right)}{\text{var}_{\boldsymbol{\xi}}(u)} = \frac{\sum\limits_{\substack{||\mathbf{j}||_1 \leq P \\ ||\mathbf{j}||_1 > 0}} u_{\mathbf{j}}^2 \, \delta_{0 j_i} \prod\limits_{m \neq i} \mathbb{E}_{\xi_i}\left[\phi_{j_m}^2\right]}{\sum\limits_{\substack{||\mathbf{j}||_1 \leq P \\ ||\mathbf{j}||_1 > 0}} u_{\mathbf{j}}^2 \prod\limits_{l=1}^{H} \mathbb{E}_{\xi_i}\left[\phi_{j_l}^2\right]} \tag{4.16}$$

4.3.3 Local Sensitivity Analysis

Although there has been recent work involving global sensitivity analysis in chemical kinetics (for example, see [71]), the term *sensitivity analysis* has more often been used in theliterature to refer to local sensitivity analyses. These are not variance-based techniques and instead involve examining local derivatives at different states. This is less sophisticated because it only accounts for the sensitivity of an output to a parameter at specific points and not over the entire range. It also does not implicitly consider the range of uncertainty. Neither of these short-comings is present in the variance-based global sensitivity analysis presented above. More details on these local techniques and situations in which they have been used can be found in [59, 58, 57, 55, 43].

4.4 Simplification at Quadrature Points

By simplifying a mechanism separately at each quadrature point in the domain of uncertainty, the differences and similarities between all possible reduced mechanisms can be observed. In particular, the intersection and union of these mechanisms can

be found; the difference in their sizes gives an indication of the number of degrees of freedom available in the choice of a simplified mechanism of given size in this region. If that number is close to zero, then it is a strong indicator that a largely deterministic algorithm will be suitable without significant extra work to account for uncertainty. The frequency with which different species occur and the distribution of reduced mechanisms among the quadrature points can also provide insight into simplification under uncertainty. These ideas will be explored in more detail in chapters 5 and 7.

Note that simplification will be carried out at DASQ points to reduce computational overhead. The DASQ routine is designed to compute an integrand and chooses the points based on the degree of non-linearity of the integrand in each direction; the integrand selected for this investigation was the log ignition time.

Chapter 5

An Investigation of the Relationship between Sensitivity and Importance of Reactions

This section explores the relationship between sensitivity and importance of reactions. Existing deterministic simplification methods consider some measure of importance or participation, but information on variance-based sensitivity is needed to extend these ideas to an uncertain context. Although no new scheme is presented in this chapter, it is hoped that the ideas given here will be able to inform the development of such a scheme.

5.1 Case studies

As mentioned in chapter 1, the GRI-Mech 3.0 methane-air combustion mechanism is used for all case studies in this work. As it is not computationally feasible to meaningfully study uncertainty in all 325 reactions at once, smaller groups had to be selected. Figure 5-1 demonstrates the different paths through the mechanism from methane to its products; this aided the selection of appropriate groups of uncertain reactions:

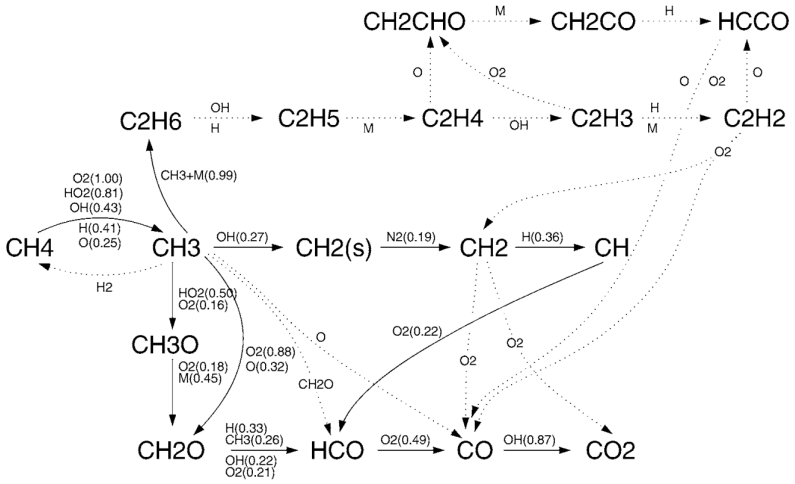

Figure 5-1: Pathways through the GRI-Mech 3.0 mechanism from methane to its products. Image courtesy of Habib N. Najm, Sandia National Laboratories.

- CH_4 to CH_3: these reactions clearly form a distinct group that affects every pathway in the mechanism.

- Hydrogen-oxygen reactions: these reactions are not explicitly shown in the figure because they do not involve products of methane, but it is clear that hydrogen, oxygen, and their radicals are involved at every step. Introducing uncertainty into the reactions that only involve these species should therefore have a significant effect on the entire mechanism.

Changes to activation energy generally have a more noticeable impact than changes to the other Arrhenius parameters, so only uncertainties in activation energy were considered. This had the additional benefit of further reducing the stochastic dimension because some of the reactions in the two groups have zero activation energy and therefore are not affected when uncertainty is introduced, as all uncertainties are relative to the original values. For reference, the uncertain reactions in these two groups are given in tables 5.1 and 5.2.

Uncertainty factors for reaction rate parameters were defined in section 1.1.2. Although uncertainty factors with log normal distributions are available, this study used uniform distributions for convenience and altered the standard uncertainty factor definition to give the bounds of the uniform distribution rather than 2.5% and

Number	Equation	E_A (cal/mol)
11	$O + CH4 \leftrightharpoons OH + CH_3$	8600
52	$H + CH_3(+M) \leftrightharpoons CH_4(+M)$	536
53	$H + CH_4 \leftrightharpoons CH_3 + H_2$	10840
98	$OH + CH_4 \leftrightharpoons CH_3 + H_2O$	3120
139	$CH_2 + CH_4 \leftrightharpoons 2CH_3$	8270
150	$CH_2(S) + CH_4 \leftrightharpoons 2CH_3$	-570
157	$CH_3 + H_2O_2 \leftrightharpoons HO_2 + CH_4$	5180
161	$CH_3 + CH_2O \leftrightharpoons HCO + CH4$	5860
162	$CH_3 + CH_3OH \leftrightharpoons CH_2OH + CH_4$	9940
163	$CH_3 + CH_3OH \leftrightharpoons CH_3O + CH_4$	9940
164	$CH_3 + C_2H_4 \leftrightharpoons C_2H_3 + CH_4$	9200
165	$CH_3 + C_2H_6 \leftrightharpoons C_2H_5 + CH_4$	10450
303	$CH_3 + CH_3CHO \rightarrow CH_3 + CH_4 + CO$	5920
317	$CH_3 + C_3H_8 \leftrightharpoons C_3H_7 + CH_4$	7154

Table 5.1: Reactions corresponding to uncertainty in methane chemistry.

Number	Equation	E_A (cal/mol)
3	$O + H_2 \leftrightharpoons H + OH$	6260
5	$O + H_2O_2 \leftrightharpoons OH + HO_2$	4000
38	$H + O_2 \leftrightharpoons O + OH$	17041
44	$H + HO_2 \leftrightharpoons O + H_2O$	671
45	$H + HO_2 \leftrightharpoons O_2 + H_2$	1068
46	$H + HO_2 \leftrightharpoons 2OH$	635
47	$H + H_2O_2 \leftrightharpoons HO_2 + H_2$	5200
48	$H + H_2O_2 \leftrightharpoons OH + H_2O$	3600
84	$OH + H_2 \leftrightharpoons H + H_2O$	3430
86	$2OH \leftrightharpoons O + H_2O$	-2110
87	$OH + HO_2 \leftrightharpoons O_2 + H_2O$	-500
88	$OH + H_2O_2 \leftrightharpoons HO_2 + H_2O$	427
89	$OH + H_2O_2 \leftrightharpoons HO_2 + H_2O$	29410
115	$2HO_2 \leftrightharpoons O_2 + H_2O_2$	-1630
116	$2HO_2 \leftrightharpoons O_2 + H_2O_2$	12000

Table 5.2: Reactions corresponding to uncertainty in hydrogen-oxygen chemistry.

97.5% intervals. All uncertainty factors were 1.25 unless mentioned otherwise. The sole exception to this is reaction 38 in hydrogen-oxygen uncertainty; the system is highly sensitive to this reaction and so an uncertainty factor of 1.05 was found to be more reasonable. Although these uncertainty factors are artificial for this particular mechanism, they do not affect the validity of the conclusions of this work, as the aim is to make observations that are hopefully broadly applicable; further tests can then be carried on other reaction mechanisms.

The initial conditions used in the studies in this section are mentioned separately in each case. These generally consist of an a temperature of 1000 K or 1200 K with an initial stoichiometric mixture, either of methane with air or methane and hydrogen with air. Pressure is constant and is always chosen to be 1 atmosphere.

5.2 Importance Indices

For this work, the outputs of interest are the importance indices and ignition time (or log ignition time for polynomial chaos expansions and sensitivity analysis, as mentioned in chapter 4). However, importance indices refer to values at specific timesteps and concentrations. As the reaction profiles are different at each rate sample, it is not valid to compare importance indices at the same timestep between samples. One possible approach is to identify an alternative progress variable, such as temperature or particular species concentration, but this does not necessarily capture an equivalent stage of the reaction for all outputs - the effect of a change in reaction rate may be more dramatic for some species than others. Furthermore, these progress variables are not guaranteed to be monotonic.

Time-averaged and maximum importance indices were considered instead. These are not as meaningful as individual importance indices at each step, but they can justifiably be compared between samples. Both are computed individually for each species-reaction pair over the scaled importance indices at every timestep for specific initial conditions. Maximum importance indices retain the useful scaling between 0 and 1, but they are not smooth and hence are not suitable for use with polynomial

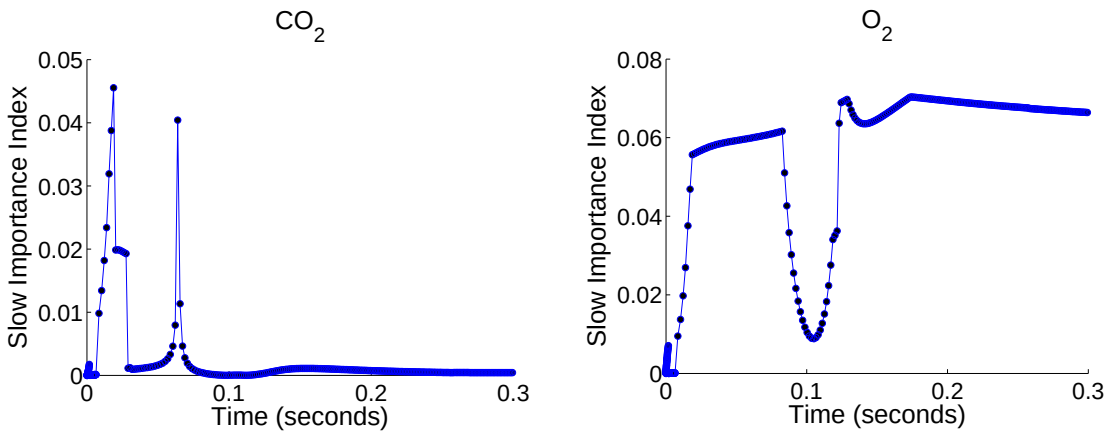

Figure 5-2: Time evolution of two example importance indices for the reaction $O + CH_4 \rightarrow OH + CH_3$ in GRI-Mech 3.0 with initial conditions of 1000 K and a stoichiometric methane-air mixture.

chaos, as described in section 4.2. As a result, time-averaged importance indices are used instead in this work. It was found that averaging over a complete reaction profile resulted in a loss of information, as the impact of reactions that are only briefly important during the ignition region becomes diluted if the reaction takes longer to reach equilibrium; examples of these brief periods of importance can be seen in the sample CO_2 importance index time evolution plot in figure 5-2, which is contrasted with a plot in which importance is spread over a larger region. To mitigate this effect, time-averaging was performed separately over the pre-ignition, ignition, and post-ignition regions.

The notation $(I_k^i)^{\text{avg}_r}$ will be used to denote time-averaged importance indices in region r, where r is one of the three regions just listed. Note that time-averaging of fast importance indices is only over timesteps at which the species i is a radical, as the fast importance index is not meaningful elsewhere - this was explained in chapter 3.

5.3 Global Sensitivity Analysis Results

Log ignition time global sensitivity analyses were run for the GRI-Mech 3.0 mechanism in a number of cases; importance index sensitivity has not been examined so far, but

will be the subject of future work. The two uncertainty cases given in the introduction to this chapter were tested at various initial conditions and a representative set of results is shown in figure 5-3. It was clear in every case that only a very small number of the uncertain reactions contributed to the variance because both the first order and total effect sensitivity indices of the remaining reactions were zero or at least close to zero. This is particularly noteworthy in the case of hydrogen-oxygen uncertainty because almost all of the variance is due to reaction 38, even though its uncertainty factor is smaller than those of the other reactions.

Furthermore, it was clear that the degree of interaction between reactions was negligible in five of the six examples shown here - the first order sensitivities all summed to slightly less than one and the total effect sensitivities all summed to slightly more than one. The combined effect was still low in the remaining example, as it was responsible for less than 10% of the output variance.

5.3.1 Comparison to importance

Importance here was measured at nominal values only. Even then, importance indices are always local in nature and apply to specific species-reaction pairs at specific states, so a measure of total importance for each reaction is not immediately clear. The measure used here for the slow importance of reaction k is (where r represents region - pre-ignition, ignition, or post-ignition):

$$F^k_{\text{slow}}\{\max_r[\,(I^i_k)^{\text{avg}_r}_{\text{slow}}\,]\} \tag{5.1}$$

where F^k_{slow} can be i for temperature only, max over i in S_0, or mean over i in S_0

S_0 now consists of temperature and the set of major species. This may not be an ideal measure of total importance in all situations, but it serves here as an indicator of whether or not a reaction is ever important. Furthermore, as will be seen shortly when results are presented, the average slow importance over S_0 is always high whenever the maximum is high or the temperature importance is high, which suggests that a reaction important to one species in S_0 is often important to most of the others

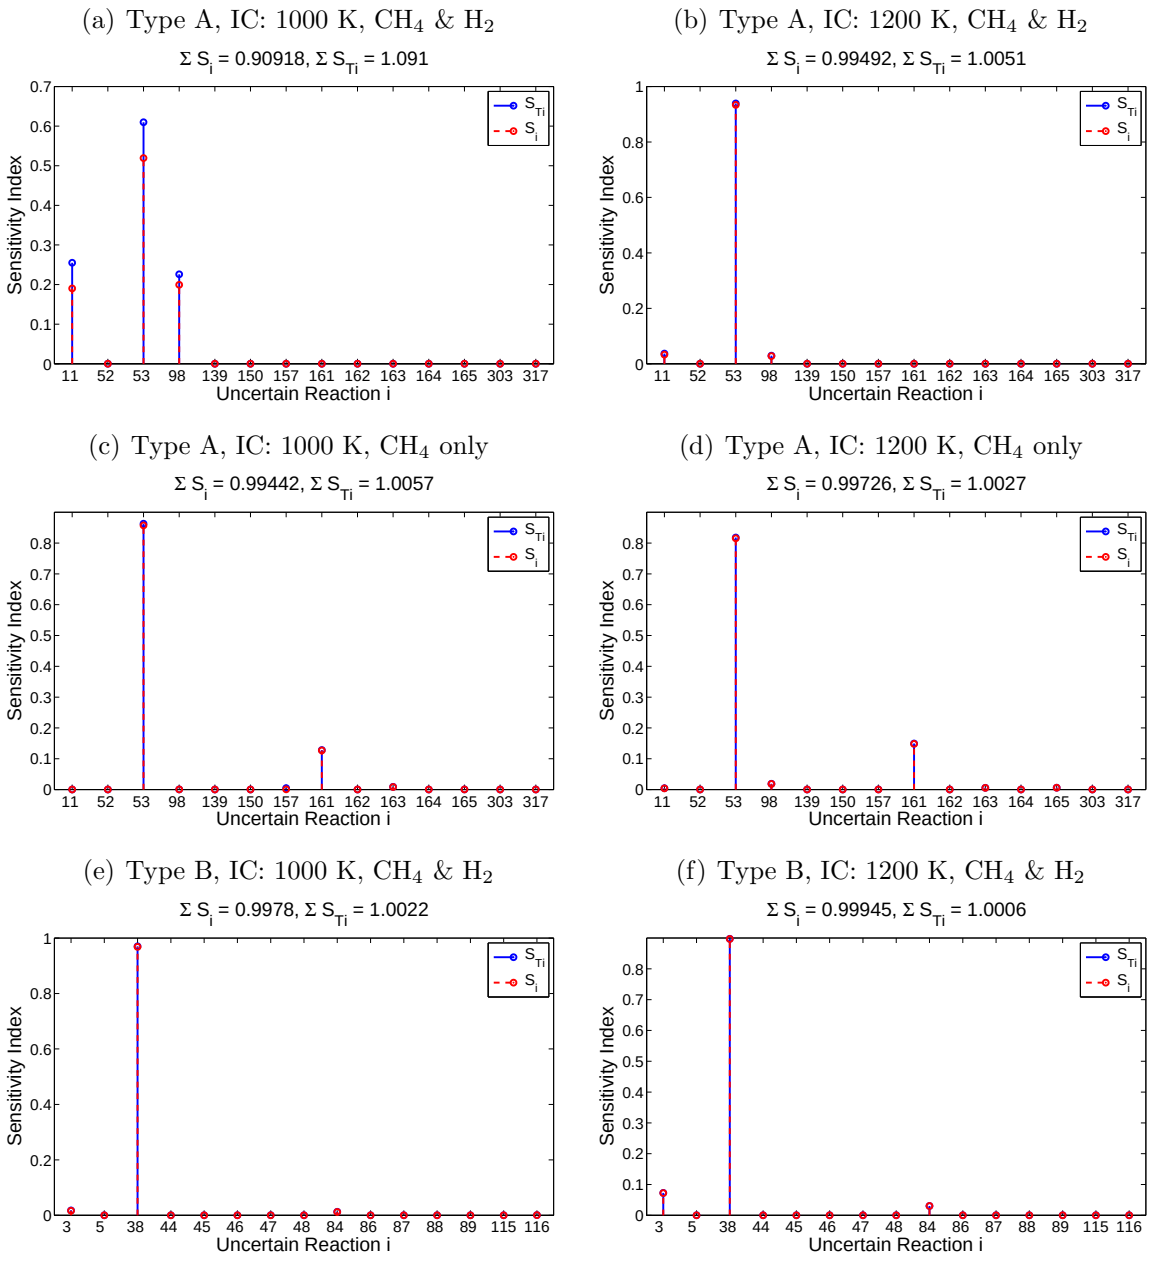

Figure 5-3: Sensitivity indices for log ignition time of GRI-Mech 3.0 at stoichiometric mixture ratios. Type A refers to uncertainty in methane chemistry and type B to uncertainty in hydrogen-oxygen chemistry. Initial conditions (IC) are indicated in individual figure subtitles. The differences of the summations of these sensitivity indices from 1 give an indication of the sensitivity of log ignition time to reactions when acting together rather than alone.

too; this consistency indicates that the importance measure is meaningful and does not just represent the impact on a small number of specific species that may not be relevant in all applications. This tendency for temperature and some major species to have similar importance is demonstrated by figure 5-4 for one specific GRI-Mech 3.0 reaction. The explanation for this effect is that the temperature and many of the major species tend to depend strongly on each other and so reactions that are important to one can generally be expected to also be important to the others.

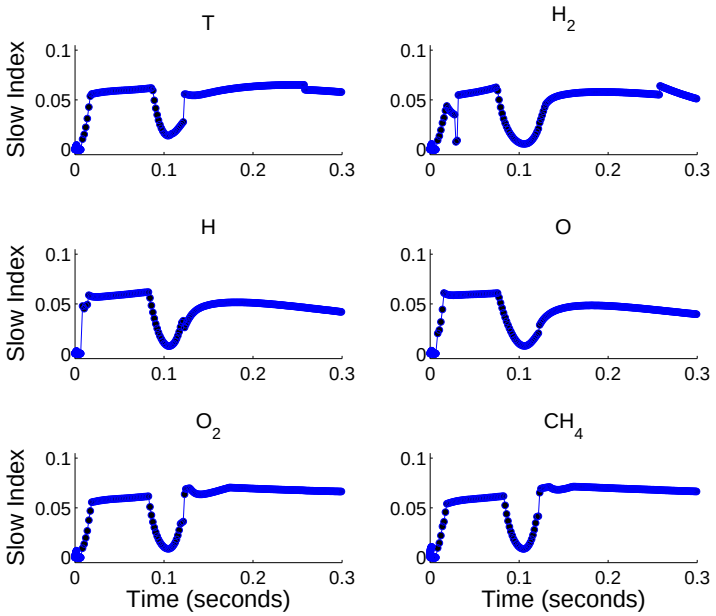

Figure 5-4: Time evolution of importance indices of temperature and some major species for the reaction $O+CH_4 \rightarrow OH+CH_3$ in GRI-Mech 3.0 with initial conditions of 1000K and a stoichiometric methane-air mixture. This is a demonstration of the similar importance of temperature and a number of major species in many reactions. The pre-ignition, ignition, and post-ignition regions can clearly be seen here after accounting for the initial transient, which is likely due to an absence of O in the initial conditions.

A measure of fast importance must be applied to a set of CSP radicals S_{rad} rather than to S_0, as discussed in chapter 3. The set of radicals changes at each timestep, but it was convenient to define S_{rad} to be only the set of species that were designated radicals most often; these were defined as the species that were designated radicals over a total time equivalent equal to at least 25% of the longest time in which any

particular radical was designated. The fast importance measure was then:

$$F^k_{\text{fast}}\{\max_r[\,(I^i_k)^{\text{avg}_r}_{\text{fast}}\,]\} \tag{5.2}$$

where F^k_{fast} can be max over i in S_{rad} or mean over i in S_{rad}

Importance results are presented in figures 5-5 and 5-6. The cases shown here correspond to those used for the sensitivity index study in figure 5-3. It is interesting to note from these plots that important reactions need not be sensitive, but all sensitive reactions had non-zero importance. Plots 5-5(c) and 5-5(d) further indicate that the most important reaction need not be the most sensitive. It is therefore clear that an uncertainty-aware simplification algorithm must consider the impact of sensitivity as a separate concept from importance.

5.3.2 Anisotropy

A brief study was carried out to check for a link between sensitivity and anisotropy, which was represented here by differences in the numbers of unique values of each parameter in the DASQ grid. Results are shown in figure 5-7 and it appears that, at least under this measure, there is no consistent link over the six examples between anisotropy and sensitivity. This is not contradictory, as the level of DASQ refinement in each direction is a measure of the degree of non-linearity - this concept is distinct from sensitivity, as a linear function can be very sensitive if it has a steep gradient.

5.3.3 Analysis

It was noted above that all sensitive reactions were also important; this should be expected in general because it would not be possible for a reaction to affect the ignition delay if it was of no importance to the temperature or major species. The explanation for the zero sensitivity result on many of the important reactions is that the system was largely constrained by other rate limiting reactions and hence moderate changes to the rate of these insensitive reactions had no significant effect - the

(a) Type A, IC: 1000 K, CH_4 & H_2

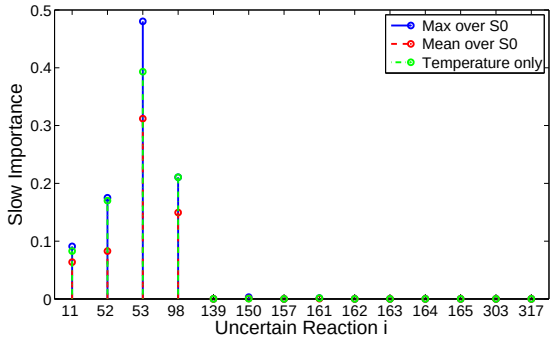

(b) Type A, IC: 1200 K, CH_4 & H_2

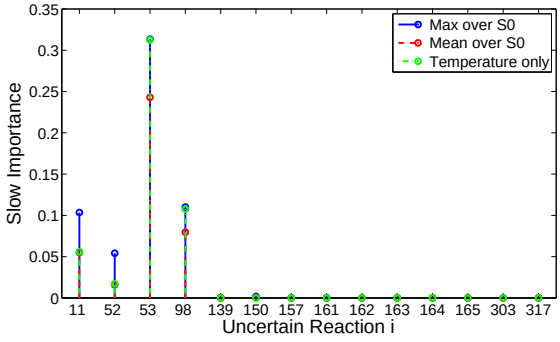

(c) Type A, IC: 1000 K, CH_4 only

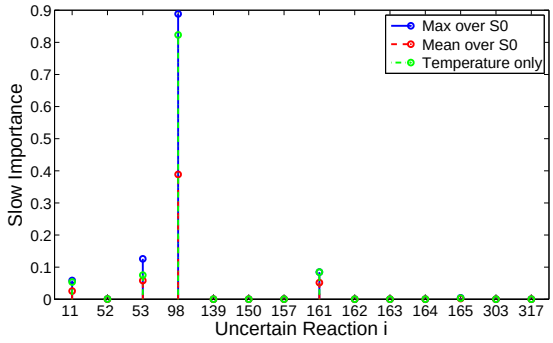

(d) Type A, IC: 1200 K, CH_4 only

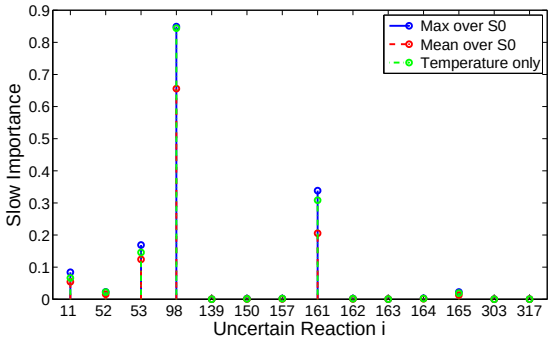

(e) Type B, IC: 1000 K, CH_4 & H_2

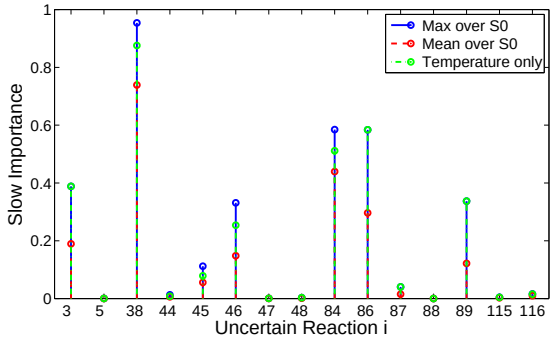

(f) Type B, IC: 1200 K, CH_4 & H_2

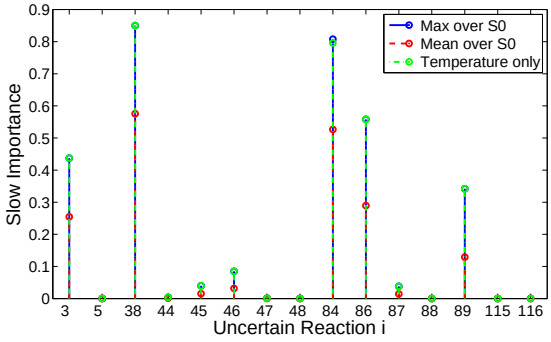

Figure 5-5: Measures of slow importance for conditions corresponding to the sensitivity analysis plots of figure 5-3.

(a) Type A, IC: 1000 K, CH_4 & H_2

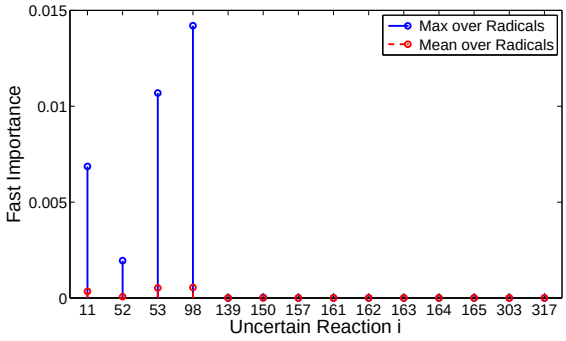

(b) Type A, IC: 1200 K, CH_4 & H_2

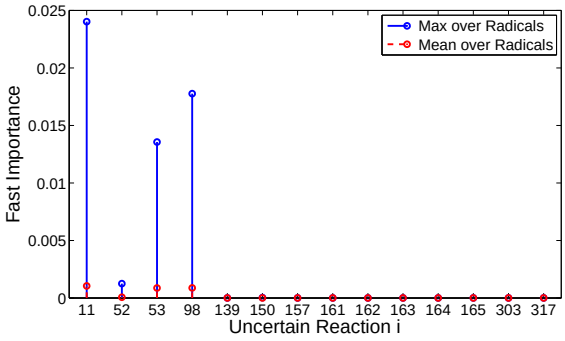

(c) Type A, IC: 1000 K, CH_4 only

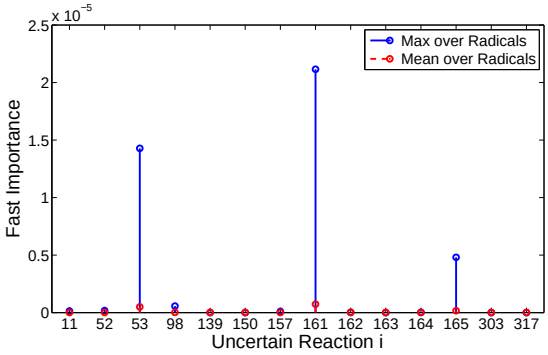

(d) Type A, IC: 1200 K, CH_4 only

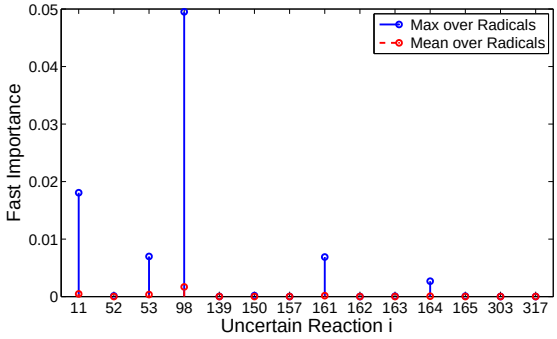

(e) Type B, IC: 1000 K, CH_4 & H_2

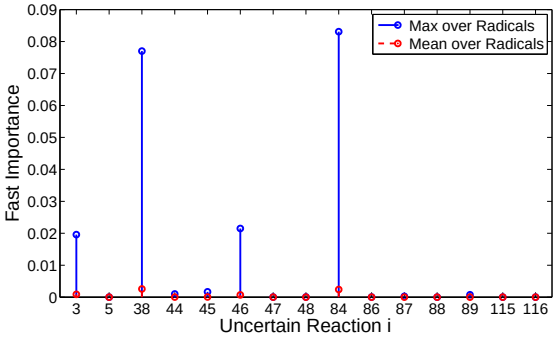

(f) Type B, IC: 1200 K, CH_4 & H_2

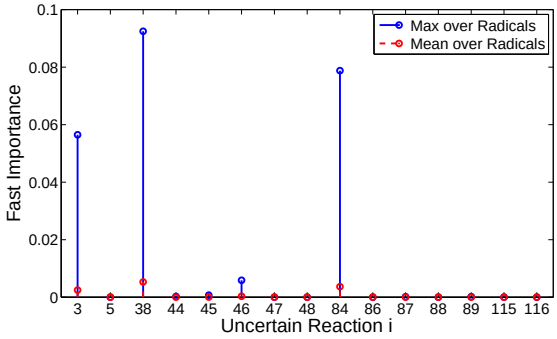

Figure 5-6: Measures of fast importance for conditions corresponding to the sensitivity analysis plots of figure 5-3.

(a) Type A, IC: 1000 K, CH_4 & H_2

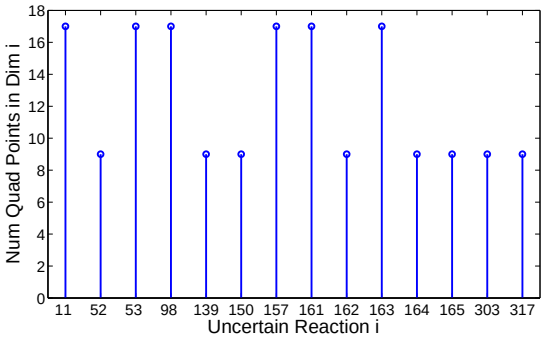

(b) Type A, IC: 1200 K, CH_4 & H_2

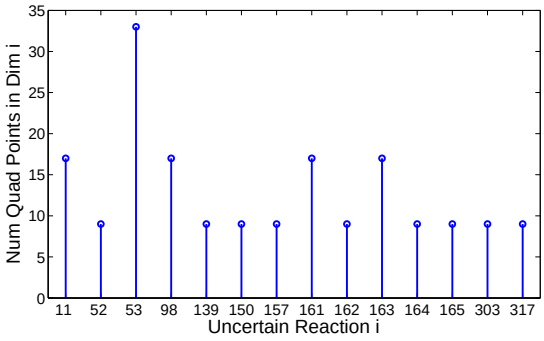

(c) Type A, IC: 1000 K, CH_4 only

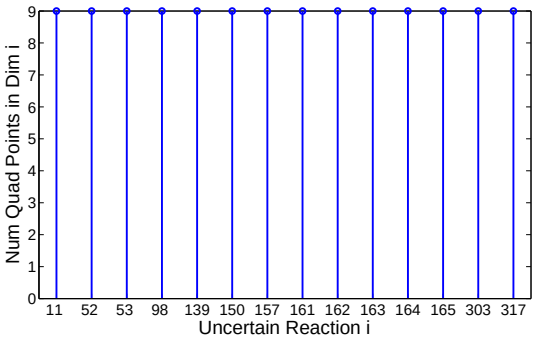

(d) Type A, IC: 1200 K, CH_4 only

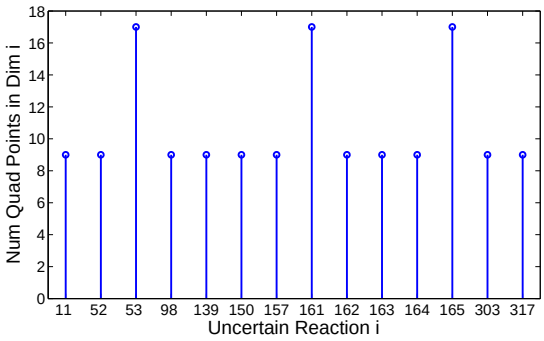

(e) Type B, IC: 1000 K, CH_4 & H_2

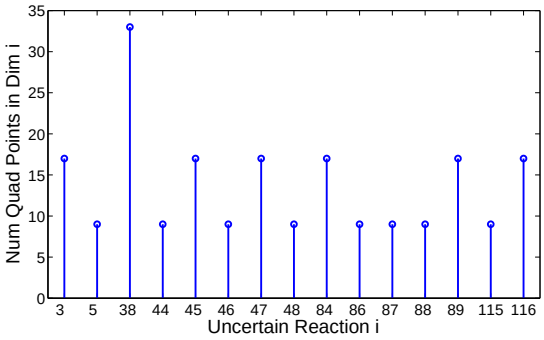

(f) Type B, IC: 1200 K, CH_4 & H_2

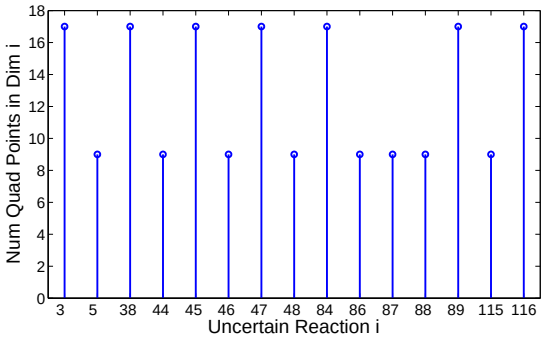

Figure 5-7: Measures of anisotropy for conditions corresponding to the sensitivity analysis plots of figure 5-3.

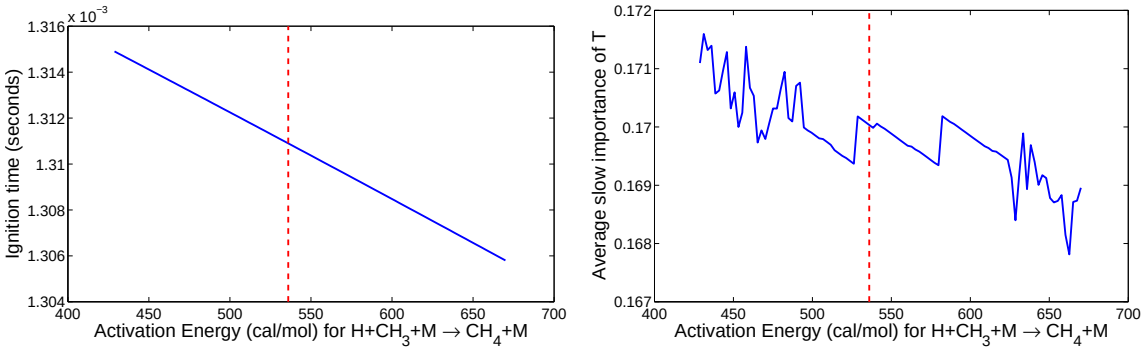

Figure 5-8: The insensitivity of temperature to the activation energy of reaction 52 at an initial temperature of 1000 K with both methane and hydrogen present. Note the narrow range of values on the y-axes compared to figure 5-9. The average slow importance here is the maximum of the time-average over the pre-ignition, ignition, and post-ignition regions. The dashed vertical lines indicates the nominal activation energy.

constraint was still active. This can be demonstrated by examining how temperature importance and ignition time change as reaction rate parameters are varied. The reactions selected for this demonstration were from uncertainty in methane chemistry at an initial temperature of 1000 K with both methane and hydrogen present. Plots are shown in figure 5-8 for reaction 52 (which is important, but not sensitive) and in figure 5-9 for reaction 98 (which is both important and sensitive). Comparing these two figures clearly shows the difference between a reaction that is constrained and a reaction that is not.

It is possible to remove a constraint by a sufficiently large change in reaction rate in one direction - the appropriate direction depends on whether the constraint is on the products or the reactants. However, the required change is clearly larger than the range of uncertainty if the reaction is insensitive. Similarly, a rate-controlling reaction is itself likely to become constrained if the rate is changed to a sufficient extent. This is demonstrated in figure 5-10, where a much larger range of activation energies is used to show that a plateau is eventually reached; the reaction is constrained in this region.

It was noted earlier that there is no significant combined effect sensitivity; this is potentially highly significant, as it means that a model simplification scheme designed

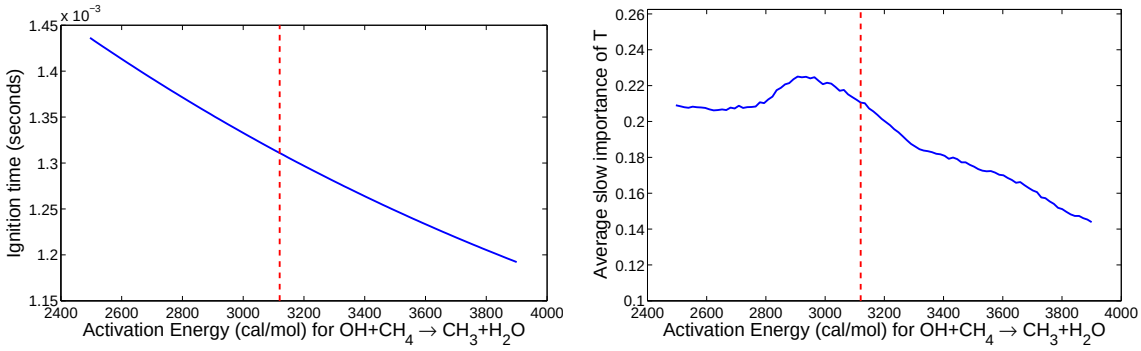

Figure 5-9: The sensitivity of temperature to the activation energy of reaction 98 at an initial temperature of 1000 K with both methane and hydrogen present. Note the wide range of values on the y-axes compared to figure 5-8. The average slow importance here is the maximum of the time-average over the pre-ignition, ignition, and post-ignition regions. The dashed vertical lines indicates the nominal activation energy.

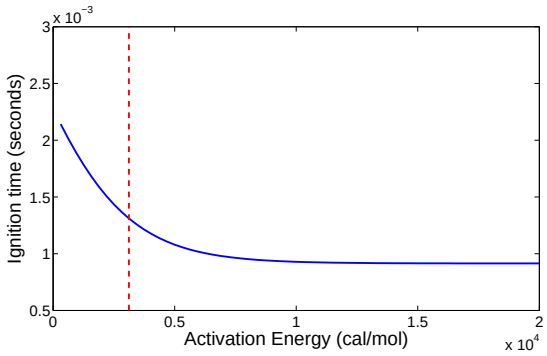

Figure 5-10: The sensitivity of temperature to the activation energy of reaction 98 at an initial temperature of 1000 K with both methane and hydrogen present in a stoichiometric mixture. A very large range of uncertainty has now been applied to demonstrate that the reaction eventually becomes constrained, even though it was previously seen to be rate-controlling in the region around the nominal value, which is indicated by the dashed vertical line.

to consider uncertainty need only treat uncertain reactions individually rather than in large high dimensional groups. This is a major advantage because characterizing probability distributions or forming polynomial chaos expansions is very cheap in a single dimension compared to higher dimensions due to the 'curse of dimensionality.' However, this lack of significant combined effect sensitivity remains to be confirmed in general with further experimentation over a wide variety of mechanisms.

5.4 Simplification at Quadrature Points

Results for simplifying the mechanism separately at each quadrature point are given in figure 5-11 for two initial conditions under two different ranges of uncertainty. It is unsurprisingly shown that there are more degrees of freedom at all threshold values when the uncertainty is increased; there should therefore be low confidence in the applicability of simplified models generated at nominal values when significant uncertainty is present.

The more interesting result is that a significant percentage of the simplified model species are present at all quadrature points; this can hopefully be used to reduce the computational expense of a model simplification algorithm that acts under uncertainty. Similarly, species not present in the union can definitely be excluded from a simplified model. Although some expense is required to identify the union and intersection, it is not significant. The standard uncertainty plots in figure 5-11 required only 37 function evaluations each. The higher uncertainty plots required fewer than 300 evaluations - this is negligible compared to the many thousands of evaluations required in Monte Carlo simulations and is also cheaper than calculating only a single coefficient in a polynomial chaos expansion, as only a low DASQ tolerance was required.

Note that convergence of the intersection and union plots was reasonable even with these low numbers of function evaluations and low DASQ tolerance; as an example, figure 5-12 is a convergence demonstration for one of the standard uncertainty plots. This figure was produced by gradually decreasing the error tolerance on the

(a) 1000 K, standard uncertainty

(b) 1000 K, high uncertainty

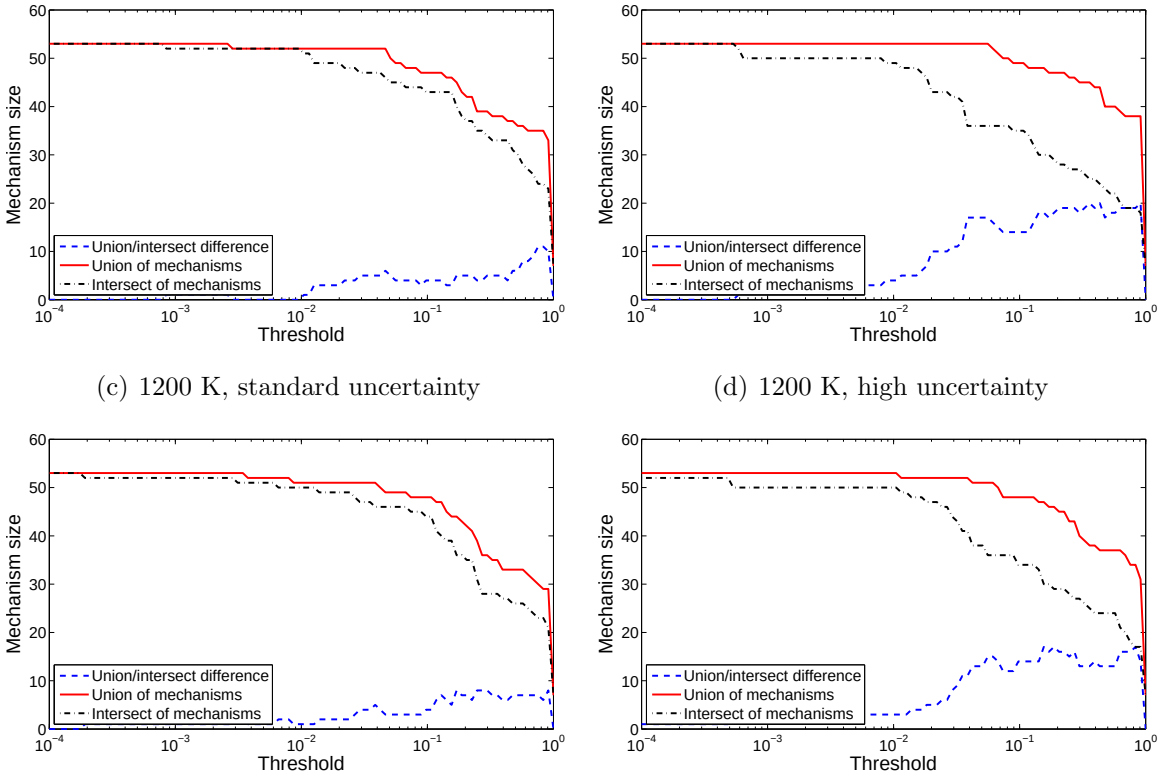

(c) 1200 K, standard uncertainty

(d) 1200 K, high uncertainty

Figure 5-11: Sizes of the union and intersection of simplified mechanisms for simplification at quadrature points with uncertainty in methane chemistry and an initial stoichiometric mixture containing both hydrogen and methane. High uncertainty refers to an uncertainty factor of 2 on all uncertain reactions instead of the standard 1.25.

(a) 37 Function Evaluations

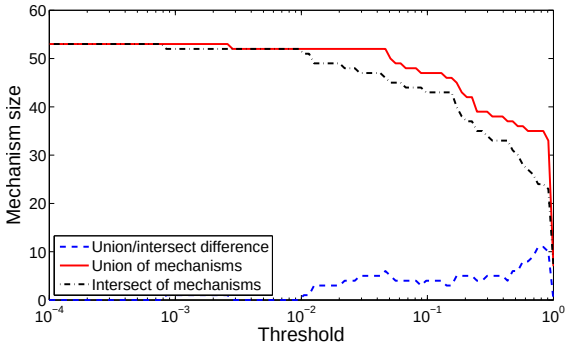

(b) 55 Function Evaluations

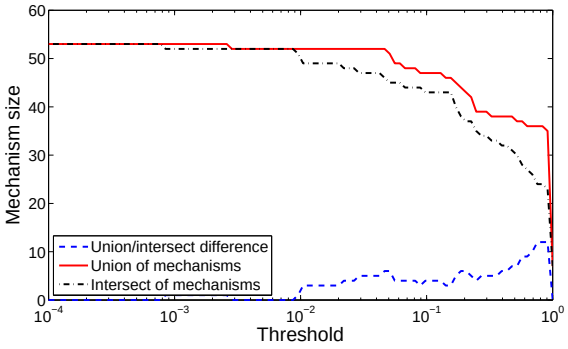

(c) 217 Function Evaluations

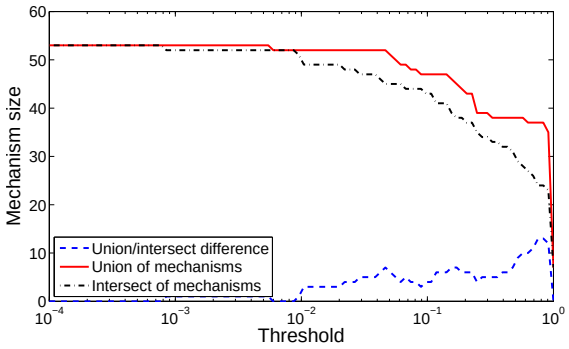

(d) 917 Function Evaluations

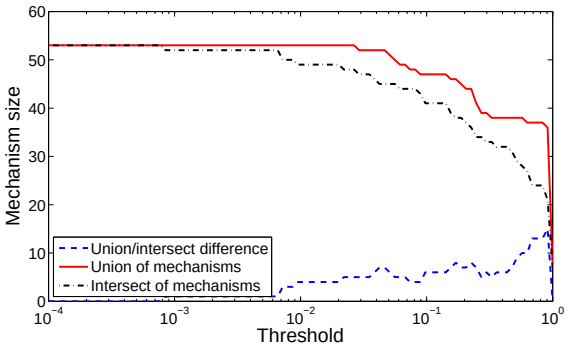

Figure 5-12: Quadrature point simplification for the case of figure 5-11(a) repeated at various more accurate quadrature levels; the original figure was for 37 function evaluations. There is no significant change when the level of accuracy increases, so the original figure was already close to the converged solution.

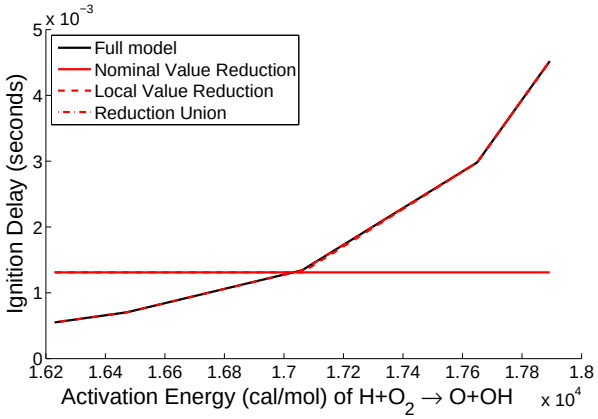

Figure 5-13: Simplification at a threshold of 0.05 for uncertainty only in reaction 38, where simplified mechanisms at the quadrature points were of sizes 47-48. The inadequacy of nominal value simplification in an uncertain context can clearly be seen, as the nominal value model fails to predict ignition delay when the activation energy changes. The mechanisms reduced individually at each local value are a much closer match and the union of these mechanisms matches the full model precisely.

integral computed by the DASQ routine in order to generate a larger number of grid points. There is no significant change when the number of function evaluations is increased well beyond 37, even though the error tolerance was decreased by 4 orders of magnitude in the most accurate plot.

5.4.1 The need for uncertainty-aware simplification

Reduction at quadrature points in only a single dimension can be used to see the difference in performance between simplification at a nominal value and simplification at each quadrature point; the mechanism generated at a nominal value fails to predict the ignition delay when the activation energy of a highly uncertain reaction changes even slightly. This demonstrates why an uncertainty-aware algorithm is desirable to increase confidence in the quality of the simplified model.

5.5 Summary

It has been demonstrated in this section that a simplification algorithm needs to be aware of uncertainty in order to provide reliable results. The stochastic dimension of

the problem is very high and this would generally indicate significant (and perhaps unreasonable) computational expense. However, it has been demonstrated here that an output of interest (ignition delay) is sensitive to only a small number of inputs. This result was also recently confirmed in [50]. As a result, the stochastic dimension can be reduced dramatically before applying an uncertainty-aware simplification algorithm.

The uncertain inputs appear to act individually rather than together in most cases. This suggests that each stochastic dimension can be considered individually in separate one dimensional problems at far lower computational expense than would be required for a single high dimensional problem. Even if this does not prove to be generally true for all other mechanisms, the expense of the sensitivity analysis is relatively low and so it will always be worth examining a mechanism to see if it can be treated cheaply in this way.

Simplification at quadrature points demonstrated that it is possible to very cheaply identify the number of degrees of freedom in a simplification problem. This result could potentially be used to reduce the cost of a new uncertainty-aware simplification algorithm.

The most significant conclusion is that any new uncertainty-aware simplification algorithm must involve a consideration sensitivity in the global variance-based sense. Relying only on a measure of importance is insufficient to determine the impact on the ignition delay when a reaction is removed, especially when looking to preserve the output pdf in the simplified model.

Chapter 6

Modifying the Valorani algorithm with CVaR

It has been shown that an uncertainty-aware simplification scheme must take account of sensitivity if it is to be reliable. The modification presented in this chapter is an early attempt at introducing uncertainty into the existing Valorani algorithm; it does not consider sensitivity and so is certainly not a final answer, but it proposes a solution to the issue of how to treat importance indices under uncertainty. It was mentioned in chapter 5 that time-averaging importance indices separately across the pre-ignition, ignition, and post-ignition regions allows for meaningful comparison between samples of different uncertain rate parameters, but a distribution of importance indices now exists instead of values at a single point and so a modification to the algorithm is needed.

6.1 CVaR

A threshold must be applied to the importance indices in the Valorani algorithm; a single value is therefore needed to characterize the distribution. Using the mean is not suitable because it would result in an unreasonably low value that fails to identify many of the occasions on which a given reaction is important; if a simplified mechanism is to be valid throughout the domain of uncertainty, it must include

reactions that are important with even fairly low (but still non-negligible) probability. As the majority of the distribution must be captured, a more appropriate choice would therefore be to use a confidence level γ (also known as the VaR - Value at Risk) with probability α, where $\alpha = 0.95$ would likely be the minimum suitable value:

$$\mathrm{P}\left((I_k^i)^{\mathrm{avg}_r} > \gamma \right) = \alpha \tag{6.1}$$

However, there are still some outstanding issues with this measure:

- Sensitivity to sampling error and random number generator quality; this is particularly significant for reactions with a VaR that is close to the threshold, as there is a sudden cut-off between being important and not important.
- Sensitivity to arbitrary threshold choice; there is no guidance on choosing the threshold in the Valorani algorithm and it is designed to be varied to produce mechanisms of different quality in order to reveal details about the structure of the original mechanism, as explained in chapter 3. If the VaR is even just slightly lower than this arbitrary threshold, then it would not be deemed important. This is related to the third bullet below.
- No measure of risk; in using the VaR as an overall importance measure, it is assumed that the remaining (1-α) fraction of importance indices is not sufficiently more important than the VaR to introduce significant error in any region of the uncertainty domain. However, this may not be the case and so a good measure should account for the amount by which the importance of a reaction exceeds the VaR, as it likely that it would be necessary to include a reaction with importance that ever exceeds the VaR by a large amount at any time, even if in only a small number of samples.

These issues can be addressed by using the CVaR (Conditional Value at Risk) instead [46]. This is defined as:

$$\mathrm{CVaR} = \mathbb{E}\left[(I_k^i)^{\mathrm{avg}_r} \,|\, (I_k^i)^{\mathrm{avg}_r} > \gamma \right] \tag{6.2}$$

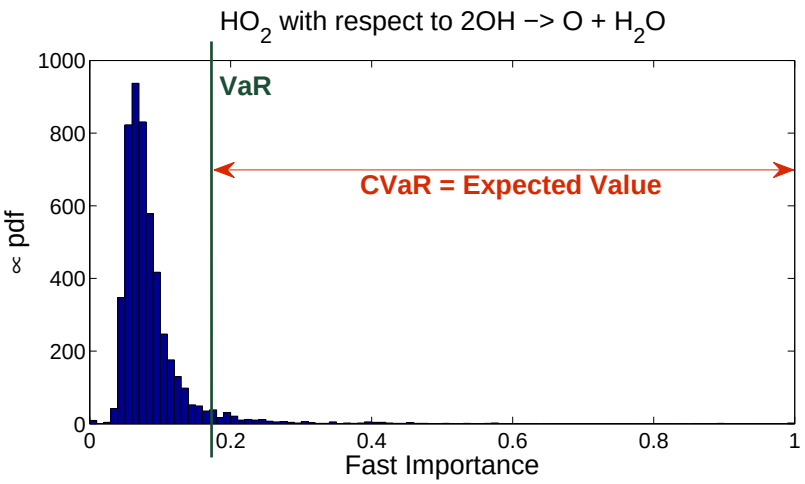

Figure 6-1: The difference between VaR and CVaR for a specific example at an initial temperature of 1000 K with an initial stoichiometric mixture of air with methane and hydrogen.

This expectation is now a measure of risk; this addresses both the second and third bullets listed above. Furthermore, the expectation operation provides additional smoothing and this helps to address the first bullet. Figure 6-1 uses a specific example to demonstrate the difference between the VaR and CVaR.

6.2 The updated algorithm

The Valorani algorithm of section 3.3.2 is repeated here with the CVaR modification, where the CVaRs γ^{slow} and γ^{fast} are calculated by Monte Carlo sampling over the chosen uncertain rate parameters:

S_{global} is empty

R_{global} is empty

for all initial conditions of interest indexed by i **do**

$S_i^0 = S_0$

R_i^0 is empty

$j = 0$

repeat

Compute γ^{slow} given by: $\mathrm{P}\left((I_k^i)_{\text{slow}}^{\text{avg}_r} > \gamma^{\text{slow}} \right) = \alpha$

Compute γ^{fast} given by: $\mathrm{P}\left((I_k^i)_{\text{fast}}^{\text{avg}_r} > \gamma^{\text{fast}} \right) = \alpha$

$C^s = \mathbb{E}\left[(I_k^i)_{\text{slow}}^{\text{avg}_r} \,|\, (I_k^i)_{\text{slow}}^{\text{avg}_r} > \gamma^{\text{slow}} \right]$

$C^f = \mathbb{E}\left[(I_k^i)_{\text{fast}}^{\text{avg}_r} \,|\, (I_k^i)_{\text{fast}}^{\text{avg}_r} > \gamma^{\text{fast}} \right]$

$R_i^{j+1} = \{k \,|\, C^s > \eta \text{ for any } p \in S_i^j\} \bigcup \{k \,|\, C^f > \eta \text{ for any } p \in S_i^j\}$

$S_i^{j+1} = \{$ species involved in reactions in R_i^{j+1} $\}$

$j = j + 1$

until $S_i^j = S_i^{j-1}$

$S_{\text{global}} = S_{\text{global}} \bigcup S_i^j$

$R_{\text{global}} = R_{\text{global}} \bigcup R_i^j$

end for

Note that it is no longer necessary to explicitly specify the set of radicals for the species involved in the fast importance index check because they have already been taken into account by the averaging process, as explained in chapter 5.

6.3 Results and conclusion

The new CVaR-based algorithm has been found to perform better than the original Valorani algorithm for pdf reproduction in many cases. This is demonstrated by figures 6-2 and 6-3, in which simplified mechanisms were generated both by the Valorani algorithm in a deterministic setting and by the CVaR-based algorithm over uncertain rate parameters; the output pdfs of these two mechanisms were then compared when

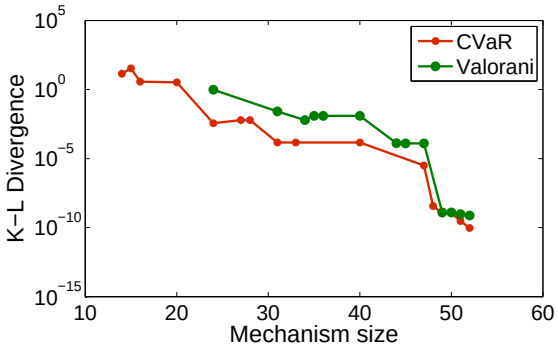

Figure 6-2: A K-L divergence comparison demonstrating that the CVaR-based algorithm outperforms the original Valorani algorithm. Lower K-L divergence corresponds to lower error. Note that the smallest mechanism sizes with large error in the CVaR algorithm could not be generated with the Valorani algorithm. Sampling error in calculating the K-L divergence was negligible and barely visible at this scale. Note that the values below 10^{-8} effectively reflect full convergence - this is demonstrated when comparing these values with figure 6-3. The conditions here were GRI-Mech 3.0 with methane uncertainty, initial temperature of 1000 K, and an initial stoichiometric mixture of hydrogen and methane with air.

sampling over the same rate parameters. It is shown both by the K-L divergence and the qualitative pdf comparisons that the CVaR-based algorithm selected better models at smaller mechanism sizes, as their output pdfs are a much better match with the full model than those of the original Valorani algorithm.

Despite the success of the CVaR algorithm in this example and others, there are cases in which its performance is not better than or even slightly worse than the Valorani algorithm. However, the author has not yet found any situations in which CVaR is substantially worse than Valorani in the same way that Valorani is substantially worse than CVaR in the above example. Further work is needed to incorporate a measure of sensitivity into the algorithm and to explore the use of a different measure of importance with the CVaR, as time-averaging is not ideal - the clear normalization between 0 and 1 no longer applies and so it is not immediately clear that it is meaningful to compare time-averaged importance indices from different reactions to the same threshold.

The computation of the expectation is expensive and requires a large number of samples, but it is hoped that polynomial chaos expansions of time-averaged importance indices could be used to increase performance if they are found to be sufficiently

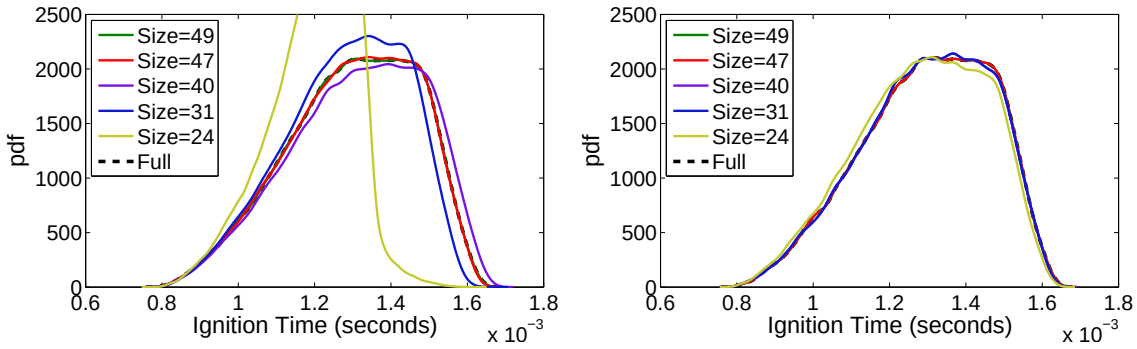

Figure 6-3: A pdf comparison demonstrating that the CVaR-based algorithm outperforms the original Valorani algorithm. The reduced model output pdfs are a much better match for the full model pdf at small mechanism sizes for CVaR than Valorani. This qualitative comparison corresponds to selected mechanism sizes from the quantitative comparison in figure 6-2 under the same conditions.

smooth - this was briefly discussed in chapter 4. Running time for the algorithm is now orders of magnitude longer than for the original algorithm in a deterministic setting. It also takes far longer than simply forming the union of deterministically reduced mechanisms at quadrature points in the uncertainty domain, as discussed in chapter 5. However, this union is likely to be needlessly large and, although such an approach works under uncertainty and is likely to result in accurate pdf reproduction, it does not take advantage of the uncertainty to allow for additional error. In contrast, the CVaR-based approach is designed to satisfy this latter goal by allowing for error when there is uncertainty by using a value of α lower than 1.

Chapter 7

Overall Conclusion and Future Work

It has been demonstrated that an uncertainty-aware simplification scheme must consider both the importance of reactions and the sensitivity of outputs to those reactions. However, an important additional step in this analysis will be to identify the sensitivities of importance indices; this could allow for a CVaR-like approach, but at each individual timestep without needing to use the unsatisfactory time-averaging procedure - the thresholding could now be in terms of the sensitivity without needing to worry about comparisons of importance indices between rate parameter samples.

Further work on simplification at quadrature points is likely to present a promising alternative path towards an improved algorithm. Rather than simply forming unions and intersections, simplified models at the quadrature points can be examined further to find the proportion of points at which certain species are included and a more detailed analysis would also consider the relevance of quadrature weights.

With both of the above approaches, a wider range of outputs can be tested to give more information about the quality of the simplified models that are produced. These may include the concentrations of trace species of interest, which would most likely be pollutants. Different species could be used in the set S_0 when examining importance indices; the radical HCO may be particularly appropriate here, as its concentration is known to be particularly sensitive to ignition time. L_2 error between the full and

simplified reaction trajectories could also be considered, perhaps with an offset to account for the difference in ignition times. Finally, equilibrium concentrations can be examined in more detail. The equilibrium values do not depend on reaction rate, as they are determined entirely by thermodynamic considerations. However, they are affected when species and reactions are removed from the model. Equilibrium values were examined for all simplified mechanisms considered in detail in this work and it was found that there was no significant deviation from the full model. Significant changes are nevertheless likely with the smallest simplified mechanisms.

It was found that ignition time was not sensitive to all reactions that were important. It will be interesting to see if the rate parameters of the remaining important reactions are sensitive inputs for some of these other outputs of interest. However, it is likely that none of the outputs will be sensitive to some important reactions because of the rate limiting constraints that were discussed in chapter 5.

The observation that combined effect sensitivities are negligible is likely to allow the computational expense of a new method to be reduced dramatically, although this will require further development of the above methods to determine an effective means of combining simplified mechanisms generated from a series of 1D analyses rather than creating the final simplified mechanism at once from a single multi-dimensional analysis.

More generally, further work is needed to formalize the trade-off between output pdf reproduction and deterministic error, as the above ideas seem better suited to instead creating computationally efficient methods that can create simplified mechanisms with the ability to produce output pdfs. Adjusting thresholds or the number of quadrature points in these ideas will hopefully allow for taking advantage of uncertainty to allow more error in a deterministic setting, but a method to quantify this trade-off has not yet been created.

Bibliography

[1] U. M. Ascher and L. R. Petzold. *Computer Methods for Ordinary Differential Equations and Differential-Algebraic Equations*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1st edition, 1998. 59

[2] D. L. Baulch. Evaluated Kinetic Data for Combustion Modeling: Supplement II. *Journal of Physical and Chemical Reference Data*, 34(3):757, 2005. 26

[3] H. J. Curran, P. Gaffuri, W. J. Pitz, and C. K. Westbrook. A comprehensive modeling study of n-heptane oxidation. *Combustion and Flame*, 114(1-2):149 – 177, 1998. 21

[4] P. Dagaut, S. Gaïl, and M. Sahasrabudhe. Rapeseed oil methyl ester oxidation over extended ranges of pressure, temperature, and equivalence ratio: Experimental and modeling kinetic study. *Proceedings of the Combustion Institute*, 31(2):2955 – 2961, 2007. 21

[5] D. F. Davidson and R. K. Hanson. Interpreting shock tube ignition data. *International Journal of Chemical Kinetics*, 36(9):510–523, September 2004. 25

[6] M. J. Davis. Low-Dimensional Manifolds in Reaction-Diffusion Equations. 2. Numerical Analysis and Method Development. *The Journal of Physical Chemistry A*, 110(16):5257–5272, 2006. 45

[7] M. J. Davis and R. T. Skodje. Geometric investigation of low-dimensional manifolds in systems approaching equilibrium. *The Journal of Chemical Physics*, 111(3):859–874, 1999. 45

[8] S. J. Fraser. The steady state and equilibrium approximations: A geometrical picture. *The Journal of Chemical Physics*, 88(8):4732, 1988. 40, 43

[9] M. Galassi, J. Davies, J. Theiler, B. Gough, G. Jungman, M. Booth, and F. Rossi. *GNU Scientific Library: Reference Manual*. Network Theory Ltd., February 2003. 33, 66

[10] T. Gerstner and M. Griebel. Numerical integration using sparse grids. *Numerical Algorithms*, 18:209–232, 1998. 69

[11] T. Gerstner and M. Griebel. Dimension-adaptive tensor-product quadrature. *Computing*, 71(1):65–87, 2003. 70

[12] R. Ghanem and D. Ghosh. Efficient characterization of the random eigenvalue problem in a polynomial chaos decomposition. *International Journal for Numerical Methods in Engineering*, 72(4):486–504, 2007. 69

[13] D. Ghosh and R. Ghanem. Stochastic convergence acceleration through basis enrichment of polynomial chaos expansions. *International Journal for Numerical Methods in Engineering*, 73(2):162–184, 2008. 69

[14] O. Herbinet, W. J. Pitz, and C. K. Westbrook. Detailed chemical kinetic oxidation mechanism for a biodiesel surrogate. *Combustion and Flame*, 154(3):507 – 528, 2008. 21

[15] O. Herbinet, W. J. Pitz, and C. K. Westbrook. Detailed chemical kinetic mechanism for the oxidation of biodiesel fuels blend surrogate. *Combustion and Flame*, 157(5):893 – 908, 2010. 21

[16] A. C. Hindmarsh and R. Serban. *User Documentation for CVODE v2.6.0*, 2009. 33

[17] F. James. RANLUX: A Fortran implementation of the high-quality pseudorandom number generator of Lüscher. *Computer Physics Communications*, 79(1):111–114, 1994. 66

[18] J. Keck. Rate-controlled constrained-equilibrium theory of chemical reactions in complex systems. *Progress in Energy and Combustion Science*, 16(2):125–154, 1990. 46

[19] J. Keck and D. Gillespie. Rate-controlled partial-equilibrium method for treating reacting gas mixtures. *Combustion and Flame*, 17(2):237–241, October 1971. 46

[20] Robert J. Kee, Michael E. Coltrin, and Peter Glarborg. *Chemically Reacting Flow: Theory and Practice*. John Wiley & Sons, Inc., 2005. 22, 25

[21] O. M. Knio and O. P. Le Maître. Uncertainty propagation in CFD using polynomial chaos decomposition. *Fluid Dynamics Research*, 38(9):616–640, September 2006. 66

[22] S. H. Lam. Using CSP to Understand Complex Chemical Kinetics. *Combustion Science and Technology*, 89(5):375–404, March 1993. 40, 41, 51, 59, 60

[23] S. H. Lam and D. A. Goussis. The CSP method for simplifying kinetics. *International Journal of Chemical Kinetics*, 26(4):461–486, April 1994. 41, 51, 53, 59

[24] O. P. Le Maître and O. M. Knio. *Spectral Methods for Uncertainty Quantification: With Applications to Computational Fluid Dynamics*. Springer, 2010. 65, 66

[25] J. C. Lee, H. N. Najm, S. Lefantzi, J. Ray, M. Frenklach, M. Valorani, and D. A. Goussis. A CSP and tabulation-based adaptive chemistry model. *Combustion Theory and Modelling*, 11(1):73–102, February 2007. 53, 55, 59

[26] T. Lu and C. K. Law. A directed relation graph method for mechanism reduction. *Proceedings of the Combustion Institute*, 30(1):1333–1341, January 2005. 39

[27] T. Lu and C. K Law. Systematic approach to obtain analytic solutions of quasi steady state species in reduced mechanisms. *The Journal of Physical Chemistry A*, 110(49):13202–8, December 2006. 41

[28] T. Lu and C. K. Law. Toward accommodating realistic fuel chemistry in large-scale computations. *Progress in Energy and Combustion Science*, 35(2):192 – 215, 2009. 22, 39, 45

[29] M. Lüscher. A portable high-quality random number generator for lattice field theory simulations. *Computer Physics Communications*, 79(1):100–110, 1994. 66

[30] U. Maas. Efficient calculation of Intrinsic Low-Dimensional Manifolds for the simplification of chemical kinetics. *Computing and Visualization in Science*, 81:69–81, 1998. 43

[31] U. Maas and S. B. Pope. Implementation of simplified chemical kinetics based on intrinsic low-dimensional manifolds. *Symposium (International) on Combustion*, 1992. 43

[32] U. Maas and S. B. Pope. Simplifying Chemical Kinetics: Intrinsic Low-Dimensional Manifolds in Composition Space. *Combustion and Flame*, (88):239–264, 1992. 42

[33] Y. M. Marzouk, B. Debusschere, H. N. Najm, D. A. Goussis, M. Valorani, and M. Frenklach. Time Integration of Reacting Flows with CSP Tabulation. In *Second International Workshop on Model Reduction in Reacting Flows*, Notre Dame, IN, 2009. 59

[34] J. Nafe and U. Maas. A general algorithm for improving ILDMs. *Combustion Theory and Modelling*, 6(4):697–709, December 2002. 43

[35] T. Nagy and T. Turanyi. Uncertainty of arrhenius parameters. *International Journal of Chemical Kinetics*, 43(7):359–378, 2011. 26

[36] H. N. Najm. Uncertainty Quantification and Polynomial Chaos Techniques in Computational Fluid Dynamics. *Annual Review of Fluid Mechanics*, 41(1):35–52, January 2009. 66

[37] H. N. Najm, J. C. Lee, M. Valorani, D. A. Goussis, and M. Frenklach. Adaptive chemical model reduction. *Journal of Physics: Conference Series*, 16:101–106, January 2005. 59

[38] A. H. Nguyen and S. J. Fraser. Geometrical picture of reaction in enzyme kinetics. *Journal of Chemical Physics*, (February):186–193, 1989. 43

[39] M. S. Okino and M. L. Mavrovouniotis. Simplification of Mathematical Models of Chemical Reaction Systems. *Chemical reviews*, 98(2):391–408, April 1998. 39, 46

[40] B. Phenix, J. Dinaro, M. Tatang, J. Tester, J. Howard, and G. McRae. Incorporation of parametric uncertainty into complex kinetic mechanisms: Application to hydrogen oxidation in supercritical water. *Combustion and Flame*, 112(1-2):132–146, January 1998. 35

[41] S. B. Pope. Computationally efficient implementation of combustion chemistry using *in situ* adaptive tabulation. *Combustion Theory and Modelling*, 1:41–63, January 1997. 31

[42] S. B. Pope and U. Maas. Simplifying Chemical Kinetics: Trajectory-Generated Low-Dimensional Manifolds. Technical Report FDA 93-11, Cornell University, 1993. 49

[43] M. T. Reagan, H. N. Najm, P. P. Pébay, O. M. Knio, and R. G. Ghanem. Quantifying uncertainty in chemical systems modeling. *International Journal of Chemical Kinetics*, 37(6):368–382, June 2005. 35, 73

[44] Z. Ren and S. B. Pope. Species reconstruction using pre-image curves. *Proceedings of the Combustion Institute*, 30(1):1293–1300, January 2005. 12, 32, 48

[45] Z. Ren, S. B. Pope, A. Vladimirsky, and J. M. Guckenheimer. The invariant constrained equilibrium edge preimage curve method for the dimension reduction of chemical kinetics. *The Journal of chemical physics*, 124(11):114111, March 2006. 50

[46] R. T. Rockafellar and S. Uryasev. Conditional value-at-risk for general loss distributions. *Journal of Banking*, 26(7):1443–1471, July 2002. 96

[47] M. R. Roussel and S. J. Fraser. Geometry of the steady-state approximation: Perturbation and accelerated convergence methods. *The Journal of Chemical Physics*, 93(2):1072, 1990. 43

[48] Safta, C. *tchem*. Being prepared for public release by Sandia National Laboratories; anticipated in late 2011. 33

[49] M. Salloum, H. N. Najm, and O. M. Knio. Analysis and Reduction of a Simplified Stochastic Chemical System via CSP. In *Third International Workshop on Model Reduction in Reacting Flows*, Corfu, Greece, 2011. 69

[50] R. T. Skodje, D. Zhou, and M. J. Davis. A Mechanism Improvement Strategy Using Global Sensitivity Analysis. In *Third International Workshop on Model Reduction in Reacting Flows*, Corfu, Greece, 2011. 93

[51] Rex T. Skodje and M. J. Davis. Geometrical Simplification of Complex Kinetic Systems. *The Journal of Physical Chemistry A*, 105(45):10356–10365, November 2001. 45

[52] Gregory P. Smith, David M. Golden, Michael Frenklach, Nigel W. Moriarty, Boris Eiteneer, Mikhail Goldenberg, C. Thomas Bowman, Ronald K. Hanson, Soonho Song, William C. Gardiner, Vitali V. Lissianski, and Zhiwei Qin. *GRI-Mech 3.0*. http://www.me.berkeley.edu/gri_mech. 37

[53] I. M. Sobol. Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. *Mathematics and Computers in Simulation*, 55:271–280, 2001. 70, 71

[54] The HDF Group. *Hierarchical data format version 5*. http://www.hdfgroup.org/hdf5, 2000-2010. 33

[55] A. S. Tomlin, M. Pilling, T. Turanyi, J. Merkin, and J. Brindley. Mechanism reduction for the oscillatory oxidation of hydrogen: Sensitivity and quasi-steady-state analyses. *Combustion and Flame*, 91(2):107–130, November 1992. 73

[56] J. Toth, G. Li, H. Rabitz, and A. S. Tomlin. The Effect of Lumping and Expanding on Kinetic Differential Equations. *SIAM Journal of Applied Mathematics*, 57(6):1531–1556, 1997. 45

[57] T. Turanyi, T. Berces, and S. Vajda. Reaction rate analysis of complex kinetic systems. *International Journal of Chemical Kinetics*, 21(2):83–99, February 1989. 73

[58] T. Turanyi, A. S. Tomlin, and M. J. Pilling. On the error of the quasi-steady-state approximation. *The Journal of Physical Chemistry*, 97(1):163–172, January 1993. 40, 73

[59] S. Vajda, P. Valko, and T. Turanyi. Principal component analysis of kinetic models. *International Journal of Chemical Kinetics*, 17(1):55–81, January 1985. 73

[60] M. Valorani, F. Creta, F. Donato, H. N. Najm, and D. A. Goussis. Skeletal mechanism generation and analysis for n-heptane with CSP. *Proceedings of the Combustion Institute*, 31(1):483–490, January 2007. 51, 61, 62

[61] M. Valorani, F. Creta, D. A. Goussis, J. C. Lee, and H. N. Najm. An automatic procedure for the simplification of chemical kinetic mechanisms based on CSP. *Combustion and Flame*, 146(1-2):29–51, July 2006. 29, 51, 61, 62

[62] M. Valorani, D. A. Goussis, F. Creta, and H. N. Najm. Higher order corrections in the approximation of low-dimensional manifolds and the construction of simplified problems with the CSP method. *Journal of Computational Physics*, 209(2):754–786, November 2005. 27, 53

[63] M. Valorani, H. N. Najm, and D. A. Goussis. CSP analysis of a transient flame-vortex interaction time scales and manifolds. *Combustion and Flame*, 134(1-2):35–53, July 2003. 60

[64] M. Valorani and S. Paolucci. The G-Scheme: A framework for multi-scale adaptive model reduction. *Journal of Computational Physics*, 228(13):4665–4701, July 2009. 59

[65] J. Wei and J. C. W. Kuo. Lumping Analysis in Monomolecular Reaction Systems. Analysis of the Exactly Lumpable System. *Industrial & Engineering Chemistry Fundamentals*, 8(1):114–123, February 1969. 45

[66] C. K. Westbrook, W. J. Pitz, O. Herbinet, H. J. Curran, and E. J. Silke. A comprehensive detailed chemical kinetic reaction mechanism for combustion of n-alkane hydrocarbons from n-octane to n-hexadecane. *Combustion and Flame*, 156(1):181 – 199, 2009. 21

[67] D. Xiu. Fast Numerical Methods for Stochastic Computations: A Review. *Communications in Computational Physics*, 5(2):242–272, 2009. 66

[68] R. A. Yetter, F. L. Dryer, and H. Rabitz. A Comprehensive Reaction Mechanism For Carbon Monoxide/Hydrogen/Oxygen Kinetics. *Combustion Science and Technology*, 79(1):97–128, September 1991. 11, 13, 23, 58, 59

[69] A. Zagaris, H. G. Kaper, and T. J. Kaper. Analysis of the Computational Singular Perturbation Reduction Method for Chemical Kinetics. *Journal of Nonlinear Science*, 14(1):59–91, January 2004. 53

[70] A. Zagaris, H. G. Kaper, and T. J. Kaper. Fast and Slow Dynamics for the Computational Singular Perturbation Method. *Multiscale Modeling & Simulation*, 2(4):613, 2004. 53

[71] T. Ziehn, K. J. Hughes, J. F. Griffiths, R. Porter, and A. S. Tomlin. A global sensitivity study of cyclohexane oxidation under low temperature fuel-rich conditions using HDMR methods. *Combustion Theory and Modelling*, 13:589–605, August 2009. 73