A-Priori Tuning of Modified Magnussen Combustion Model

A. T. Norris,
NASA Langley Research Center, Hampton, Virginia, 23681

Abstract

In the application of CFD to turbulent reacting flows, one of the main limitations to predictive accuracy is the chemistry model. Using a full or skeletal kinetics model may provide good predictive ability, however, at considerable computational cost. Adding the ability to account for the interaction between turbulence and chemistry improves the overall fidelity of a simulation but adds to this cost. An alternative is the use of simple models, such as the Magnussen model, which has negligible computational overhead, but lacks general predictive ability except for cases that can be tuned to the flow being solved.

In this paper, a technique will be described that allows the tuning of the Magnussen model for an arbitrary fuel and flow geometry without the need to have experimental data for that particular case. The tuning is based on comparing the results of the Magnussen model and full finite-rate chemistry when applied to perfectly and partially stirred reactor simulations. In addition, a modification to the Magnussen model is proposed that allows the upper kinetic limit for the reaction rate to be set, giving better physical agreement with full kinetic mechanisms. This procedure allows a simple reacting model to be used in a predictive manner, and affords significant savings in computational costs for simulations.

Introduction

One of the challenges of performing simulations of reacting flows is the ability to model the chemical source term. The direct approach of modeling a full set of kinetic equations is extremely expensive computationally. For example, the GRI-Mech-3.0 mechanism for Methane combustion [1] consists of 53 species and 350 rate equations, which would be virtually intractable for anything but the simplest of flows. Because of the complexity of these full mechanisms, different techniques have been proposed to simplify them and thus, reduce the computational cost. Skeletal mechanisms, reduced mechanism and ILDM methods are examples of existing techniques that reduce the work required to evaluate the chemical source term.

However, it is also known that due to the nonlinear nature of the governing equations and chemical source terms for turbulent reacting flows, the source terms due to chemical reactions are not being calculated exactly. Specifically, using the mean species concentrations and mean temperature to calculate the mean species source term is not accurate, i.e.,

$$< S_j(c_i, T) > \neq S_j(< c_i >, < T >), \tag{1}$$

where S_j is the source term of the j-th species, c_i is the molar concentration of the ith species, T is the temperature and $<>$ denotes the average value of the quantity inside the angled brackets. Indeed it has been shown that the differences in the mean reaction rate associated with assuming equality in Eqn.(1) can be several orders of magnitude different from the exact solution [2].

To address this issue, several models and numerical approaches have been developed in order to resolve the modeling inaccuracy. Models include the Direct Quadrature Moment Method (DQMOM) [3], the Assumed PDF [4], the Transported PDF [5] and the Linear Eddy Model [6]. All these approaches attempt to model the effect that turbulence has on chemical reaction with varying degrees of fidelity. Unfortunately, as the modeling becomes more sophisticated, there is a corresponding increase in the computational cost, especially when coupled with the expense of chemical kinetics schemes.

One way to address the expense of kinetics modeling is to prescribe a simple method of accounting for the heat release of chemistry, and modify that to address the interaction between the turbulence and the chemistry source term. Such an approach was proposed by Magnussen in 1976, [7] and due to the simplicity of the model, the approach has proved very popular. However, the simplicity of the model does result in it lacking several desirable features, namely that it needs to be tuned to the flow that it is modeling and it lacks realizability in the modeled reaction rate.

Magnussen Model

Let us consider a chemical reaction defined by the mass fraction of fuel Y_f, oxidizer Y_o and products Y_p. The rate of fuel consumption for the Magnussen model is given by:

$$\frac{dY_f}{dt} = min \begin{cases} A\, Y_f/\tau_t \\ A\, Y_o/\tau_t r_s \\ AB\, Y_p/\tau_t(1+r_s) \end{cases} \tag{2}$$

where A and B are constants, r_s is the stoichiometric ratio of fuel to oxidizer and τ_t is the turbulent time scale. The two terms, $A\, Y_f/\tau_t$ and $A\, Y_o/\tau_t r_s$ give the rate of change of the fuel to be proportional to the availability of fuel or oxidizer, respectively. The constant A controls the magnitude of the rate, while the turbulent time scale in the equation models the effect of turbulence on reaction rates. As the turbulence becomes more intense, mixing of species will become more vigorous, and so a faster reaction can be sustained. The third term, $AB\, Y_p/\tau_t(1+r_s)$ sets the rate of fuel consumption to be proportional to the concentration of products in the gas mixture. This term is intended to account for the effect of ignition delay and is often neglected in non-premixed reaction calculations. [8]. The final reaction rate for the fuel is taken as the minimum of the three rate expressions and the reaction rates of the other species, Y_o and Y_p are simply proportional to the fuel reaction rate.

$$\begin{aligned} \frac{dY_o}{dt} &= r_s \frac{dY_f}{dt} \\ \frac{dY_p}{dt} &= -(1 + r_s) \frac{dY_f}{dt} \end{aligned} \tag{3}$$

An illustration of the Magnussen model is shown in Fig.(1) where the reaction rate of the product species is plotted against the mass fraction of the products for a stoichiometric mixture of fuel and oxidizer. The first two parts of Eqn.(2) describe the part of the curve to the right of the peak value, while the third part of the equation gives the curve to the left of the peak. The value of Y_p at which the peak reaction rate occurs is determined by the value of B, where a value of 1 would give a peak value at $Y_p = 0.5$. The effect on the reaction rate due to changing the value of the turbulent time scale is also shown, with the reaction rate increasing inversely proportional to the increase in τ_t. It should be noted that the shape of the reaction rate in Fig.(1) crudely matches what one would expect if the rate of a major product species was plotted against the concentration of that species.

The Magnussen model does have a couple of issues that limit its effectivness. First, it is not predictive. The model needs to be tuned to individual geometries, conditions and fuels. Hence, to perform a prediction of a reacting flow, one needs to know the answer beforehand to allow the model to be tuned to the specific case. Second, the model accounts for the interaction of turbulence and chemistry by simply making the reaction rate inversely proportional to the turbulence time scale. While for a certain range of time scales this is reasonable, chemical reactions do have an upper limit for the speed of reaction. Once that limit is reached, no matter how quickly the flow is mixed, the chemistry will not react any faster. However, the Magnussen model does not have such a limit. As the turbulent time scale goes to zero, the resulting reaction rate obtained from the Magnussen model approaches infinity.

To address these issues, it is proposed that the Magnussen model be tuned against simple canonical reaction simulations, where the use of a full chemical mechanism is tractable, and extrapolate those results to a real flow situation. The simple canonical cases proposed are the Perfectly Stirred Reactor (PSR) and also the Partially Stirred Reactor (PaSR). In addition, the Magnussen model will be modified to address the issue of having no limit to the reaction rate. The process involved in tuning the model and limiting the rate will be addressed in the remainder of the paper. To illustrate the process, the chemical kinetics of JP7 and air will be considered and used to tune the Magnussen model. The full JP7 chemical mechanism is represented by a reduced mechanism of 16 species and 45 rate equations. [9]. This mechanism is a good representative of the sort of complex hydrocarbon kinetics model used in scramjet applications.

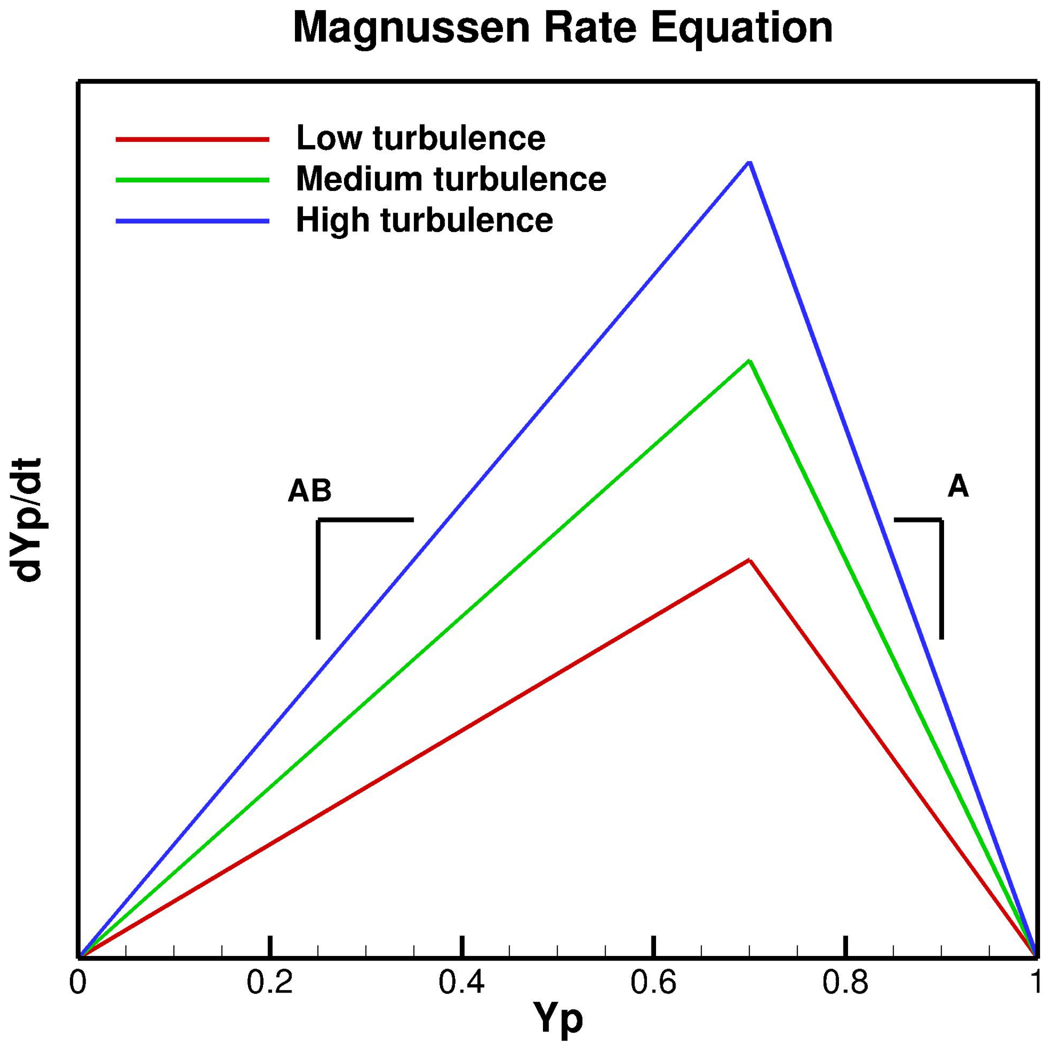

Figure 1: The rate of product creation dY_p/dt plotted against the product mass fraction Y_p showing the effect of varying turbulence time scale on the reaction rate.

PSR and PaSR

In this section, the concepts of the Perfectly Stirred Reactor (PSR) and the Partially-Stirred Reactor (PaSR) are defined.

A chemical reactor is a simple, zero-dimensional model that is very useful for the investigation of chemical reactions. It consists of a volume, V, at pressure, P, containing a mass m^r. Into this volume, fuel and oxidizer are supplied at a fixed rate, $\dot{m}^{in}$. Inside, the fuel and oxidizer mix instantly with the contents of the reactor and react. Finally, there is a fixed rate of mass leaving the reactor, $\dot{m}^{out}$. If one were to consider a set of N species with

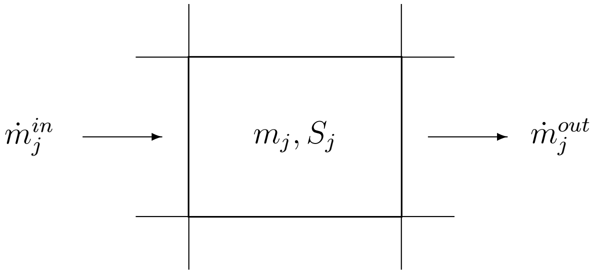

Figure 2: Schematic of a perfectly stirred reactor.

composition $\phi_i, i = 1, N$, a PSR can be described as:

$$\begin{aligned} \frac{dm^r}{dt} &= \dot{m}^{in} - \dot{m}^{out} \\ \frac{d\phi_i}{dt} &= S_i + \phi_i^{in} - \phi_i^{out} \end{aligned} \tag{4}$$

where S_i is the reaction source term for the i-th species and the superscripts *in* and *out* refer to the species entering and exiting the reactor volume. For a steady state solution, $d\phi_i/dt$ is set to zero and:

$$\begin{aligned} \dot{m}^{in} &= \dot{m}^{out} \\ S_i &= \phi_i^{in} - \phi_i^{out}. \end{aligned} \tag{5}$$

A sketch of the PSR is shown in Fig.(2).

For a given volume V and mass flow in and out, the residence time scale for the reactor can be calculated:

$$\tau_r = \frac{m^r}{\dot{m}^{in}} = \frac{m^r}{\dot{m}^{out}}. \tag{6}$$

This time scale can be thought of as the amount of time the species stay in the reactor before leaving or alternately as the ratio of the mass of fluid in the reactor to the mass of fluid entering or exiting the reactor.

The behavior of a PSR is consistent for all chemical reactions of interest to the combustion community. The mass contained in the reactor volume is given by ρV where ρ is the density of the fluid. For the case of a large residence time, ρV is large compared to $\dot{m}^{in}$ and so the fluid in the reactor has a long time to react before it has to leave. In this case, the reactor has a composition close to fully reacted. For the case where τ_r is small, the fluid has very little time to react before it must exit the reactor, and so the composition in the reactor is essentially unreacted.

In Fig.(3), a PSR simulation of a stoichiometric JP7-air reaction is shown for a pressure of 5 atmospheres and inflow temperatures of fuel and air of 750K. It can be seen that as expected, for large residence times the temperature of the reactor is that of a fully reacted, or equilibrium mixture. As the residence time becomes smaller, the temperature drops slightly before a sudden drop off to fully unreacted. This is typical of most hydrocarbon-air reactions The residence time when the mixture ceases to react is referred to as the blow-out limit.

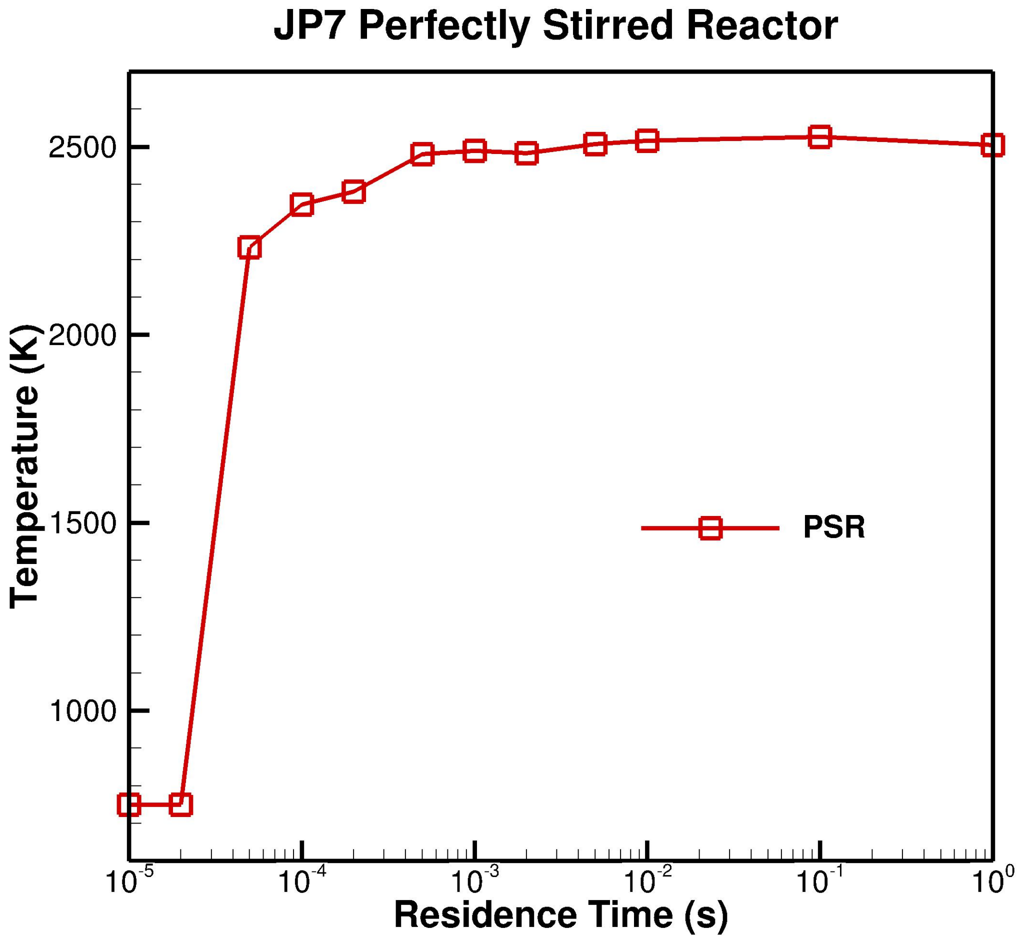

Figure 3: The temperature of the mixture in a perfectly-stirred reactor as a function of the residence time of the mixture.

In the description of the PSR, it should be noted that the mixing of the inflow species with those in the reactor is instantaneous, hence the ”Perfectly Stirred” expression used in the name. But in real life, nothing is ever perfectly mixed. There is always a level of mixing, or a rate of mixing that is present. To reflect this, another form of reactor is considered: The Partially-Stirred Reactor (PaSR). The PaSR is very similar to the PSR, except when the species enter the reactor, they mix with the existing species at a specified rate. This is quantified by the mixing time scale, τ_m and is modeled as being proportional to the turbulent time scale τ_t. Because the mixing in the reactor is not complete, the composition in the reactor is no longer homogenous. Instead it is a distribution of species at various stages of mixing and reaction. In essence, we have a simple model for the turbulence-chemistry interaction phenomena. The mixing is the turbulence and the chemistry is the reaction.

To model the PaSR, we employ the Probability Density Function (PDF) method of modeling the reactor. The N species in the reactor are represented by a joint probability distribution of species, $\mathcal{F}$, which evolves due to molecular mixing, M, chemical reaction, S and the inflow and outflow species distributions $\dot{\mathcal{F}}^{in}$ and $\dot{\mathcal{F}}^{out}$, respectively. The evolution of the distribution of all the scalar quantities in the reactor, $\mathcal{F}$ as a function of time is given by

$$\frac{d\mathcal{F}(\psi)}{dt} = \sum_{i=1}^{N} \frac{d}{d\psi_i} \left([M_i(\psi) + S_i(\psi)]\mathcal{F}(\psi) \right) + \dot{\mathcal{F}}^{in}(\psi) - \dot{\mathcal{F}}^{out}(\psi) \tag{7}$$

where the ψ_i represents the scalar space of species i and the number of species is given by N. In a similar way as with the PSR, if the change in the distribution of species, $\mathcal{F}$ with time is zero the equation corresponds to a steady-state solution:

$$\sum_{i=1}^{N_s} \frac{d}{d\psi_i} \left([M_i(\psi) + S_i(\psi)]\mathcal{F}(\psi) \right) = \dot{\mathcal{F}}^{in}(\psi) - \dot{\mathcal{F}}^{out}(\psi) \tag{8}$$

And just like the PSR, a residence time can be defined in the same way as shown in Eqn.(6).

To model the PaSR process, a stochastic method is employed that takes Eqn.(7) and drives the left-hand term to zero. The distribution of species in the reactor is represented by an ensemble of notional fluid particles, each with a composition of species and fluid properties. At regular time intervals, δt, new particles are added to the ensemble, representing the inflow of species, taken from the distribution $\mathcal{F}^{in}$. At the same time, an identical number of particles are removed from the ensemble to account for the mass flow exiting the reactor. Reaction is accomplished by reacting the composition of each notional fluid particle, and mixing is performed via a stochastic mixing model. This process is continued until statistical steady state is reached for the mean species compositions.

There are several different stochastic mixing models available, such as the IEM model, [10], the Curl model [11], the modified Curl model [4] and the EMST model [12]. Theoretical considerations concerning the relative merits of the different mixing models are discussed in Subramaniam and Pope (1998) [12]. In this study, the Modified Curls Model was employed [4]. This is a particle interaction model based on the original Curl model [11]. In this model, a pair of particles, ϕ^p and ϕ^q are randomly selected from the ensemble, and over a time period δt their compositions are changed by:

$$\begin{aligned} \phi^p(t+\delta t) &= \phi^p(t) + \frac{1}{2} r(\phi^q(t) - \phi^p(t)) \\ \phi^q(t+\delta t) &= \phi^q(t) + \frac{1}{2} r(\phi^p(t) - \phi^q(t)) \end{aligned} \tag{9}$$

where $\phi^p(t)$ represents the composition of particle p at time t and r is a random number with a uniform distribution in the interval (0,1). In this case, however, the time step δt is

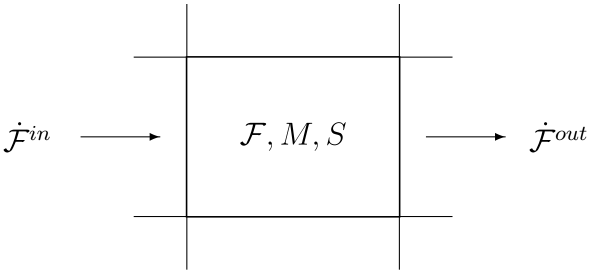

Figure 4: Schematic of a partially-stirred reactor.

a function of the turbulent time scale, τ_t and the number of particles in the ensemble, N_p, and is given by:

$$\delta t = \frac{\tau_t}{3N_p} \tag{10}$$

Thus, for a given time step, Eqn.(10) will give the number of particles that need to be mixed.

In the operation of the PSR and the PaSR models, the particles are initialized with fully mixed and reacted particles. The particles that are to be input into the reactor each time step are given a suitable composition corresponding to either fuel or air in stoichiometric proportion.

The overall time step used is based on the number of particles that are selected to be input, however, the mixing and reaction steps are often run at a smaller time step for several subiterations to ensure certain numerical constraints, such as the time step required for the mixing model. For each run, the code keeps iterating until the mean scalar quantities in the reactor have converged to a statistically constant value. To reduce the number of particles required for a certain level of statistical scatter, the mean scalar quantities are evaluated using a time-averaging technique [2] to reduce the amount of scatter in the results. Full details of this model are given in [13], however, the relevant part for this discussion is that the turbulent mixing is controlled by an inverse time scale proportional to ϵ/k. A diagram of the process is shown in Fig.(4).

Like the PSR, the behavior of the PaSR is a function of the residence time. A large residence time gives a fully reacted composition in the reactor while a short residence time results in blow-out. However, there is also the mixing time scale to consider. For a given residence time, the mixing time will alter the composition. A small mixing time (fast mixing) will give results similar to the PSR, but a slow mixing time will push the composition toward blowout. Fig.(5) shows the same simulation conditions of the PSR performed with a PaSR model.

It is instructive to show the PaSR results for a series of fixed ratios between the residence time of the reactor and the turbulence time scale. This ratio is a form of Dahmkolar number

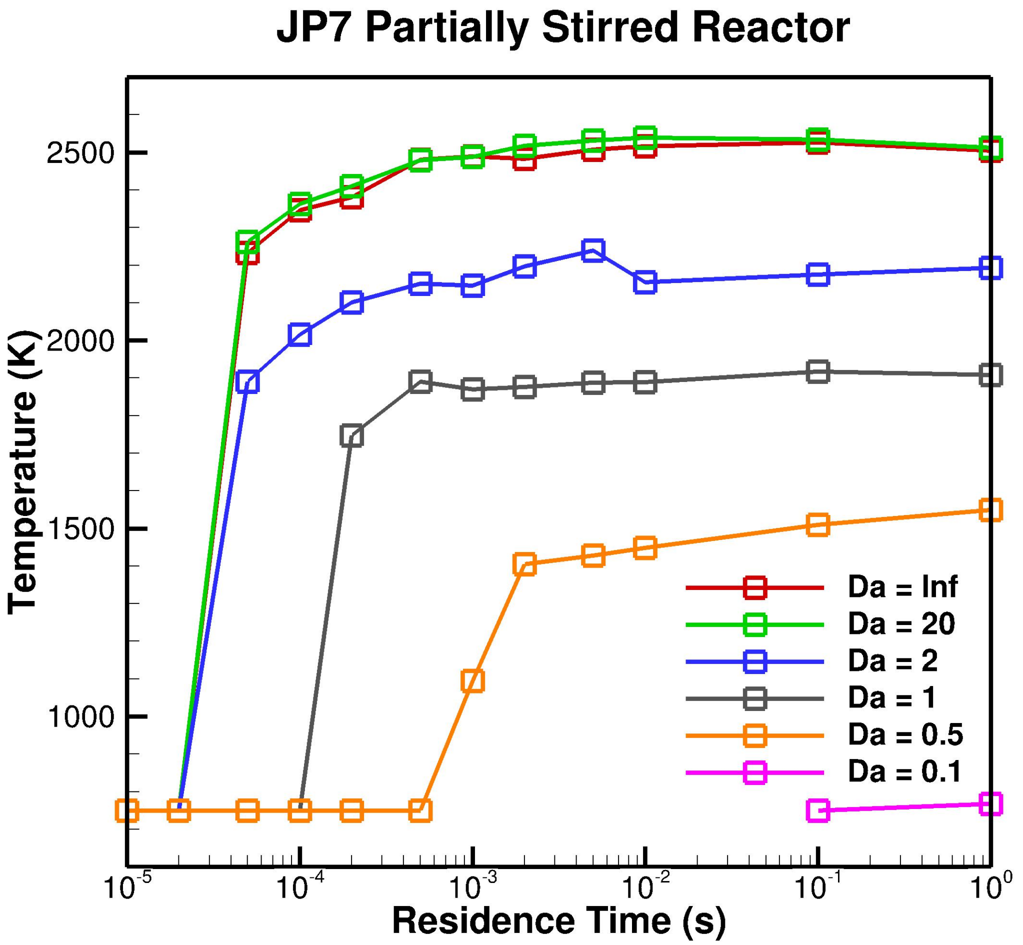

Figure 5: The temperature of the reactor mixture plotted as a function of residence time. The effect of varying the ratio of residence time to turbulence time is shown.

and is defined as:

$$Da = \tau_r/\tau_m \tag{11}$$

Several interesting results are shown in the PaSR simulation. First, the maximum temperature of the reaction is a function of Da. As Da gets large, it approaches the PSR result. But lower values of Da give lower temperatures until the point occurs where no reaction is sustained. This is probably one of the better illustrations of the importance of accounting for the interaction of turbulence and chemistry. If a Da was obtained for all the cells in a simulation, and the number was always above 20, a laminar chemistry approximation would be reasonable. But if a significant percentage of the cells were at lower numbers, the results would be suspect.

The other interesting observation is that all the different Da curves are similar in shape and all drop to zero by a similar residence time. This is due to the kinetic models having a maximum reaction rate, or thought of another way, there is a reaction time scale that reactions cannot drop below.

While the PSR and PaSR are useful and interesting tools in the study of reactions, they do have a more immediate connection to the issues of combustion modeling. Consider the description of the PSR: species enter a volume, mix perfectly, react and exit. This is exactly the same description as a single computational cell in a steady-state, finite-volume CFD simulation. In other words, the PSR is an exact model of what goes on in an individual CFD cell. Likewise, the PaSR is an exact description of what occurs in a single cell of an Eulerian scalar PDF simulation coupled with a finite-volume flow solver. Thus, it is reasonable to suggest that if a simple kinetics model used in a reactor gives similar results to that of a full kinetics scheme applied in a reactor, then the simple model will also give similar results in a full CFD simulation. Using this premise, the Magnussen model, if tuned to a finite rate mechanism in a reactor model, will give equivalent results in a full CFD simulation.

Tuning

The first question to be addressed is which reactor model should be used to tune the Magnussen model. Because the Magnussen model has a mixing model built into it, a simple comparison to the PSR cannot be made. And running the Magnussen model in the PaSR model would not work as the mixing model in the reactor would conflict with the built in mixing term in the Magnussen model. Thus, the best option is to compare the results of the Magnussen model used in the PSR with the finite rate results from the PaSR. In both cases, the simulations will be run with identical values of Da and the results compared.

The first part of the tuning process was to select a set of appropriate inflow fuel and air properties in addition to a representative pressure for the reacting region. Using the same values as the full mechanism in the previous section, fuel and air in stoichiometric proportions are chosen with a temperature of 750K and a pressure of 5 atmospheres. These conditions may be considered as typical of those in a scramjet combustor. Note that for high-speed flows, the air temperature is taken as a typical temperature in the region where combustion takes place, not necceserily the inlet air temperature. Using these initial conditions, PSR and PaSR simulations were performed using the finite rate mechanism and the results were shown earlier in Fig.(3) and (5).

The next stage of the tuning process is to create the thermal properties for the Magnussen synthetic species, fuel, oxidizer and product. The fuel and the oxidizer are simply the properties of the respective species, while the product is a mixture of H2O and CO2 in equal molar proportions. The mixture-averaged value for the specific heat of the product is obtained by

$$Cp^{mix} = \sum_{k=1}^{Ns} Cp_k X_k \tag{12}$$

where X_k is the molar fraction of species k.

One other item needs to be considered. Because a full mechanism will contain minor species, the fully burnt temperature of the finite rate mechanism will be generally a couple of hundred degrees lower than the fully burnt Magnussen temperature. To account for this, the thermal properties of the product species are adjusted so the resulting fully burnt peak temperature is the same as the finite rate temperature. The thermal properties are given by the 7-constant polynomials of Gordon and McBride [14],

$$\begin{aligned}
\frac{Cp_k}{RT} &= a1_k + a2_k T + a3_k T^2 + a4_k T^3 + a5_k T^4 && (13)\\
\frac{H_k}{RT} &= a1_k + \frac{a2_k T}{2} + \frac{a3_k T^2}{3} + \frac{a4_k T^3}{4} + \frac{a5_k T^4}{5} + \frac{a6_k}{T} && (14)\\
\frac{S_k}{RT} &= a1_k log_e T + a2_k T + \frac{a3_k T^2}{2} + \frac{a4_k T^3}{3} + \frac{a5_k T^4}{4} + a7_k && (15)
\end{aligned}$$

where Cp is the constant-pressure specific heat, H is the enthalpy and S is the entropy of species k. The universal gas constant is represented by R and T is the temperature.

The equilibrium temperature of the two mechanisms can be compared and the ratio between them can be given as X,

$$T_{eq}^{mag} = X T_{eq}^{full} \qquad (16)$$

then the constants a1, a2, etc. are adjusted as follows:

$$\begin{aligned}
a1'_k &= a1_k\\
a2'_k &= a2_k X\\
a3'_k &= a3_k X^2\\
a4'_k &= a4_k X^3\\
a5'_k &= a5_k X^4\\
a6'_k &= a6_k / X\\
a7'_k &= a7_k + a1_k log_e(X) && (17)
\end{aligned}$$

This process is somewhat crude and is a subject of further work. Also, note that this method of adjusting the thermal properties of the product species does not account for the change in mean molecular weight due to the presence of minor species in the full kinetic mechanism. Again this is an area that could be improved upon with further work.

Using the scaled thermal properties for the product species, the results of the reactor simulation are shown in Fig.(6) using Magnussen constants of $A = 0.5$ and $B = 0.0$ that were obtained from an un-tuned simulation. Comparing the results to those of the finite rate simulation in Fig.(5), there are several obvious differences. First, the Magnussen model curves show no drop off, or blow out. This is an artifact of the model having a reaction rate dependant linearly on the inverse turbulent time scale. Thus, as observed earlier, there is no

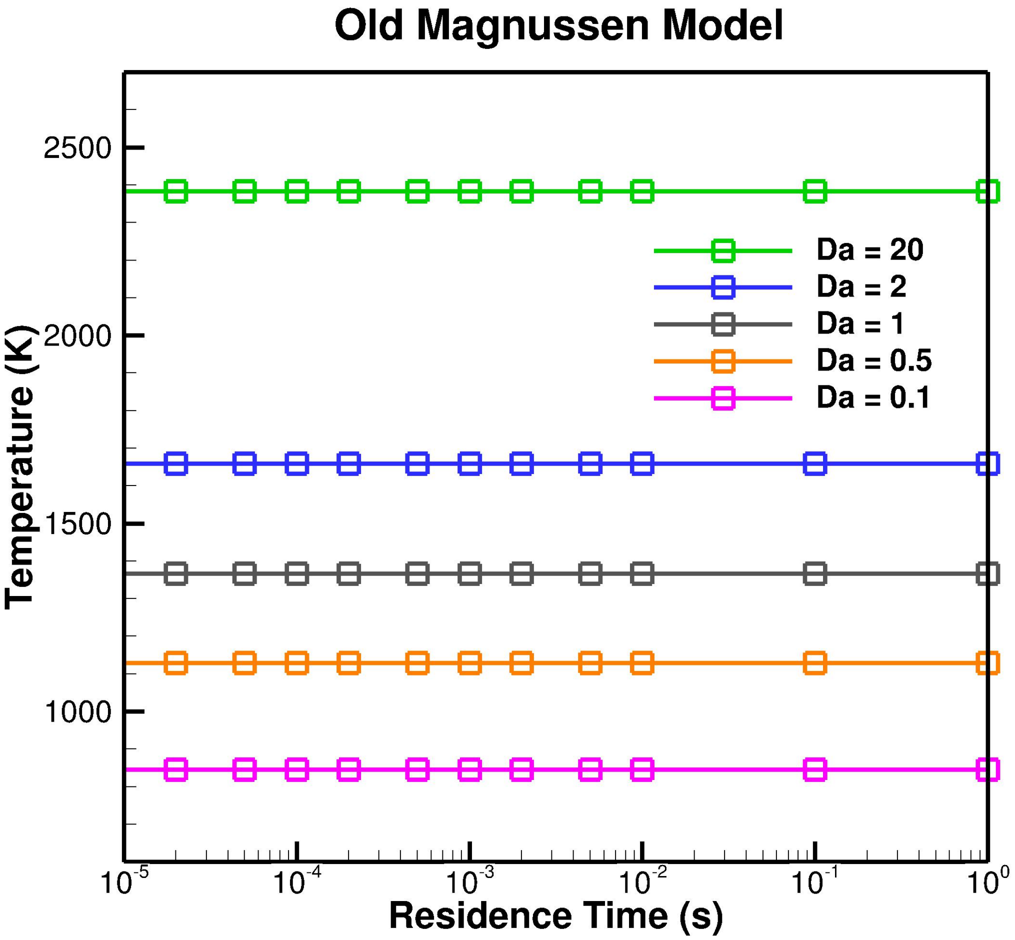

Figure 6: The temperature of a reactor mixture plotted against residence time for the Magnussen model using values tuned against an existing flow solution.

upper limit to the reaction rate. The second difference is that, while the high levels of Da match for high residence times, the lower values do not match.

The issue of mismatched values for the different Da curves can be resolved by adjusting the Magnussen model constants. The effect of the A constant is to adjust the temperature spacing between the different Da numbers, while the B constant determines what value of Da will not support burning. The A constant was adjusted first to set the spacing between the higher values of Da and then the B value was adjusted to ensure the lower values were in the correct location. This resulted in values of $A = 1.7$ and $B = 1.2$. The reactor result for the adjusted A and B constants is shown in Fig.(7). It is clear that the results agree quite well with the results of the finite rate PaSR simulation (See Fig.(5)) for the fully burnt parts of the plot.

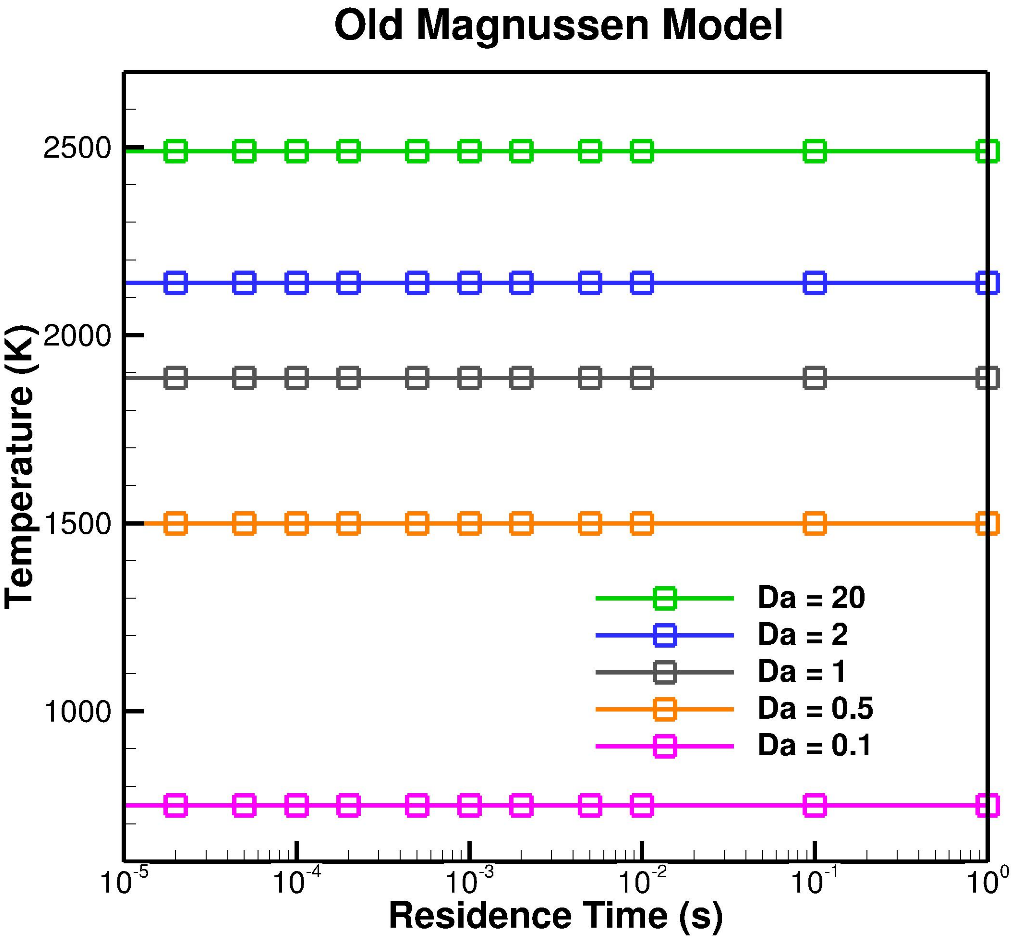

Figure 7: The temperature of a reactor mixture plotted against the residence time showing the effect of adjusting the A constant to fit the finite-rate results.

To address the issue of no blow-out, the Magnussen model is modified to include a limit to the mixing time scale. Essentially the value of τ_t is not permitted to drop below a certain value, τ_{min}. This value can be set via trial and error, but is generally similar in magnitude to the residence time where reaction ceases to occur. In this case, $\tau_t = 5.0e-05$. The results of the reactor simulation using both the tuned A and B constants as well as the turbulent time scale limit are shown in Fig.(8). The results show that the new constants give a good representation of the finite rate PaSR results, and the Magnussen model can be considered tuned for this set of conditions.

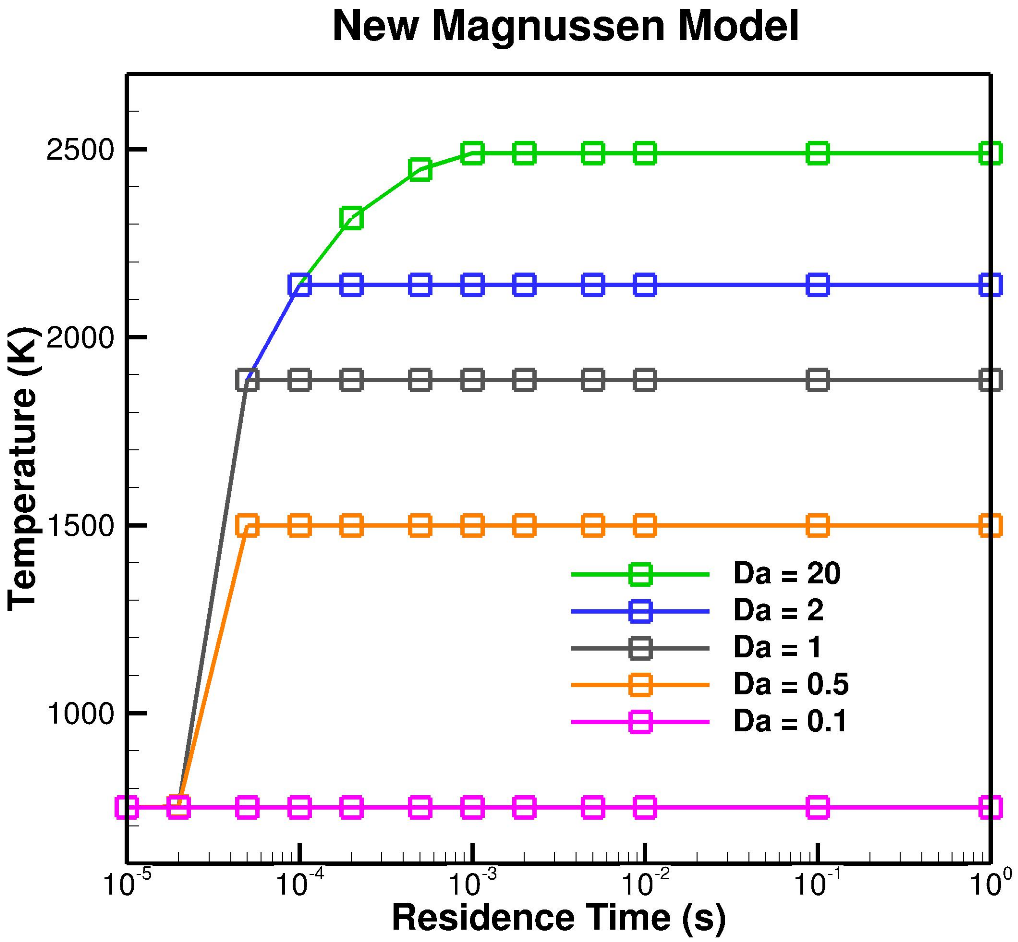

Figure 8: The temperature of a reactor mixture plotted against the residence time showing the effect of adding a turbulence time scale limit to the equation.

Summary and Future Work

In this paper, a method has been proposed that would allow a modified version of the Magnussen model to be employed in a more predictive manner than before. Using simple reacting flow simulations of perfectly and partially stirred reactors, the Magnussen model can be tuned to behave much like the full mechanism. However, this tuning process comes with several caveats:

1/ The tuning process described is developed for subsonic reacting flows. While careful choice of initial conditions will allow this to work with supersonic flows, the presence of shocks and expansions in the reacting flow region may result in poor performance. An extended model, with the Magnussen constants as a function of velocity is one option if this issue needs to be addressed.

2/ The effect of minor species in the full mechanism was compensated for by adjusting the product thermal properties. However, no attempt was made to adjust for the change in mean molecular weight due to dissociation. This may be important.

3/ The tuning of the A and B constants is done via a trial and error process. A more rigorous process should be looked into.

4/ While the methodology of tuning the Magnussen model has been described, no examples have been presented. This will be addressed in the future.

The goal of this paper was to develop a method of making the Magnussen model more predictive and the process described gives a logical approach to that. In addition, the method provides a useful tool to investigate a reacting flow to see if the assumption of laminar chemistry (no turbulence-chemistry interaction) is valid.

Bibliography

[1] Smith, G. P., Golden, D. M., Frenklach, M., Moriarty, N. W., Eiteneer, B., Goldenberg, M., Bowman, C. T., Hanson, R. K., Song, S., William C. Gardiner, J., Lissianski, V. V., and Qin, Z., `http://www.me.berkeley.edu/gri_mech/`, Accessed 2016-03-01.

[2] Norris, A. T. and Hsu, A. T., "Comparison of PDF and Moment Closure Methods in the Modeling of Turbulent Reacting Flows," *30th AIAA/ASME/SAE/ASEE Joint Propulsion Conference, Indianapolis, Indiana*, 1994, pp. AIAA–94–3356.

[3] Fox, R. O., *Computational Models for Turbulent Reacting Flows*, Cambridge University Press, Cambridge, 2003.

[4] Janicka, J., Kolbe, W., and Kollmann, W., "Closure of the Transport Equation for the Probability Density Function of Turbulent Scalar Fields," *J. Non-Equilib. Thermodyn.*, Vol. **4**, 1979, pp. 47–66.

[5] Pope, S. B., "PDF Methods for Turbulent Reactive Flows," *Prog. Energy Combust. Sci.*, Vol. **11**, 1985, pp. 119–192.

[6] Kerstein, A. R., "A linear eddy model of turbulent scalar transport and mixing," *Combust. Sci. Tech.*, Vol. **60**, 1988, pp. 391–421.

[7] Magnussen, B. F. and Hjertager, B. H., "On Mathematical Modeling of Turbulent Combustion With Special Emphasis on Soot Formation and Combustion," *Sixteenth Symposium (International) on Combustion*, 1976, pp. 719–729.

[8] Norris, A. T., Jentink, T. N., Eklund, D. R., Hagenmaier, M. A., and Cox-Stouffer, S. K., "CFD Analysis of the SED SJX61-6 Engine: Comparison to Experimental Data

(U)," *JANNAF 44th Combustion, 32nd Airbreathing Joint Meeting, Arlington, Virginia*, 2011, pp. JANNAF–672.

[9] Gokulakrishnan, P., Pal, S., Klassen, M., Hamer, R., Roby, R., Kozaka, O., and Menon, S., "Supersonic Combustion Simulation of Cavity-Stabilized Hydrocarbon Flames using Ethylene Reduced Kinetic Mechanism." *AIAA/ASME/SAE 42nd Joint Propulsion Conference, Sacramento, CA*, 2006, pp. AIAA–2006–5092.

[10] Dopazo, C. and O'Brian, E. E., "An Approach to the Autoignition of a Turbulent Mixture," *Acta Astronautica*, Vol. **1**, 1974, pp. 1239–1266.

[11] Curl, R. L., "Dispersed Phase Mixing: 1. Theory and Effects in Simple Reactors," *A. I. Ch. E. J.*, Vol. **9**, 1963, pp. 175–181.

[12] Subramanian, S. and Pope, S. B., "A Mixing Model for Turbulent Reactive Flows Based on Euclidean Minimum Spanning Trees," *Combustion and Flame.*, Vol. **115**(4), 1998, pp. 487–514.

[13] Chen, J. Y., "Stochastic Modeling of Partially Stirred Reactors," *Combustion Science and Technology.*, Vol. **122**(1-6), 1995, pp. 63–94.

[14] McBride, B. J., Gordon, S., and Reno, M. A., "Coefficients for Calculating Thermodynamic and Transport Properties of Individual Species," Tech. Rep. NASA Technical Memorandum 4513, NASA: Office of Management. Scientific and Technical Information Program, 1993.