An automated target species selection method for dynamic adaptive chemistry simulations

Nicholas J. Curtis^a, Kyle E. Niemeyer^b, Chih-Jen Sung^{a,*}

^a*Department of Mechanical Engineering, University of Connecticut, Storrs, CT, 06269, USA*

^b*School of Mechanical, Industrial, and Manufacturing Engineering, Oregon State University, Corvallis, OR 97331, USA*

Abstract

The relative importance index (RII) method for determining appropriate target species for dynamic adaptive chemistry (DAC) simulations using the directed relation graph with error propagation (DRGEP) method is developed. The adequacy and effectiveness of this RII method is validated for two fuels: n-heptane and isopentanol, representatives of a ground transportation fuel component and bio-alcohol, respectively.

The conventional method of DRGEP target species selection involves picking an unchanging (static) set of target species based on the combustion processes of interest; however, these static target species may not remain important throughout the entire combustion simulation, adversely affecting the accuracy of the method. In particular, this behavior may significantly reduce the accuracy of the DRGEP-based DAC approach in complex multidimensional simulations where the encountered combustion conditions cannot be known a priori with high certainty. Moreover, testing multiple sets of static target species to ensure the accuracy of the method is generally computationally prohibitive. Instead, the RII method determines appropriate DRGEP target species solely from the local thermo-chemical state of the simulation, ensuring that accuracy will be maintained. Further, the RII method reduces the expertise required of users to select DRGEP target species sets appropriate to the combustion phenomena under consideration.

Constant volume autoignition simulations run over a wide range of initial conditions using detailed reaction mechanisms for n-heptane and isopentanol show that the RII method is able to maintain accuracy even when traditional static target species sets fail, and are even more accurate than expert-selected target species sets. Additionally, the accuracy and efficiency of the RII method are compared to those of static target species sets in single-cell engine simulations under homogeneous charge compression ignition conditions. For simulations using more stringent DRGEP thresholds, the RII method performs similarly to that of the static target species sets. With a larger DRGEP threshold, the RII method is significantly more accurate than the static target species sets without imposing significant computational overhead.

Furthermore, the applicability of the RII method to a DRG-based DAC scheme is discussed.

*Corresponding author

Email address: cjsung@engr.uconn.edu (Chih-Jen Sung)

Keywords: Mechanism reduction, Dynamic adaptive chemistry, Directed relation graph with error propagation, Target species selection

1. Introduction

The use of detailed reaction mechanisms is essential for high-fidelity predictions of important combustion phenomena such as pollutant emissions [1] and local flame extinction [2]—a cause of lean blowout [3]—as well as in the design of next-generation combustion devices such as homogeneous charge compression ignition (HCCI) engines [4]. However, the large size and high chemical stiffness of transportation-relevant fuel mechanisms prohibit their use in realistic simulations. For multidimensional reacting flow simulations, the chemistry time integration can take up to 75 % to 99 % of the total simulation time [5–8]. In order to utilize large chemical reaction mechanisms for transportation-relevant fuels in realistic simulations, accurate mechanism reduction and chemical stiffness removal strategies must be employed.

Detailed reaction mechanisms are constructed to be valid over a wide range of thermo-chemical states, and therefore tend to contain many species and reactions that are not important in all combustion regimes. A skeletal mechanism is created by removing species and reactions considered unimportant for the thermo-chemical state space under consideration. A comprehensive skeletal mechanism is constructed for a broad thermo-chemical state space, while a local skeletal mechanism is created for a specific, limited range of thermo-chemical states.

Several systematic techniques to generate skeletal mechanisms by removal of unimportant species and reactions from a detailed mechanism have been developed; the directed relation graph (DRG) method [9–11] and the directed relation graph with error propagation (DRGEP) [12] are two skeletal reduction methods popularly used due to their efficiency and reliability. Other commonly used skeletal reduction methods include sensitivity analysis [13–15], principal component analysis [16], level of importance analysis [17–19], and methods based on computational singular perturbation [20–22] modified for skeletal reduction [23, 24]. To generate even more compact skeletal mechanisms, the DRG and DRGEP methods are often combined with sensitivity analysis, as in the DRG-aided sensitivity analysis [11] and DRGEP-aided sensitivity analysis [25, 26] methods.

In order to further optimize comprehensive skeletal mechanisms, dimension reduction methods are often employed. For example, chemical timescale analysis is utilized to exploit the tendency of reaction mechanisms to be attracted to lower dimensional manifolds, constraining the dimensionality of the reaction mechanism and thus simplifying integration. Methods such as the quasi-steady state [27] and partial equilibrium approximations [28, 29] assume a species or reaction quickly reaches a steady state after initial transience, and thus can be solved for algebraically. More systematic dimension reduction methods include the computational singular perturbation [20–22] and the intrinsic low-dimensional manifold [30] methods, which analyze the Jacobian matrix to decouple the fast and slow chemical reaction modes to reduce chemical stiffness.

Finally, to accelerate chemical integration, tabulation methods that store and reuse previous solution information to cheaply construct approximate solutions of chemical integrations are often used. Common examples include in situ adaptive tabulation [31] (ISAT), piecewise reusable implementation of solution mapping [32], and artificial neural networks [33].

Most skeletal reduction approaches create a single, comprehensive skeletal mechanism for use over a prescribed range of conditions expected to be encountered. However, this approach is inherently at a disadvantage in multidimensional simulations because the same level of detail must be applied to the entire domain. Inside the flame zone, a highly detailed skeletal mechanism will likely be necessary to maintain accuracy, but in computational cells where combustion is mostly completed or weakly reacting a much smaller skeletal mechanism may be sufficient.

Recently, Liang et al. [34, 35] proposed a dynamic adaptive chemistry (DAC) scheme to exploit this observation, where the DRGEP skeletal reduction method was applied to a thermo-chemical state under consideration to generate a smaller, locally accurate skeletal mechanism. The local skeletal mechanism was then integrated for a single simulation time-step and discarded; this process was repeated at the next time step. The reduced expense resulting from integrating the smaller, locally accurate mechanism outweighed the overhead of the reduction method, leading to time savings overall. For this reason, reduction methods that scale linearly with the problem size (e.g., DRG, DRGEP, element flux analysis [36]) are typically used in DAC schemes.

Liang et al. [34] utilized this method to achieve a 30-fold speedup with high accuracy for single cell HCCI simulations of n-heptane. Later, they demonstrated the applicability of the method for n-heptane/isooctane/toluene blends in HCCI and homogeneous autoignition simulations [35]. By pairing an element flux analysis method with the DAC scheme, He et al. [36] achieved a 25-fold speedup in a simulation of n-pentane in a pairwise mixed stirred reactor; however, the overhead of the flux-based reduction method consumed nearly 20 % of the total simulation time. Yang et al. [37] paired the DAC scheme with the DRG method in turbulent methane flame simulations. They found that the DRG-based DAC method accurately reproduced the combustion process of a partially stirred reactor with significant levels of non-equilibrium chemistry. Furthermore, for simulations with longer flow time scales (e.g., 10^{-4} s or longer) it was more efficient to generate a larger skeletal mechanism to be used for the whole flow time step, rather than performing multiple reduction/integrations steps within a single flow time step due to the overhead of the mechanism reduction [37].

Tosatto et al. [38] formulated a DRG-based DAC scheme that additionally considered transport fluxes, achieving speedup factors of 5 and 10 for a steady JP-8 flame and a time-dependent ethylene flame, respectively. Gou et al. [39] paired a simplified version of the path flux analysis method [40] with the DAC scheme to develop a method of error control for DAC simulations. Data tabulated from simple zero-dimensional simulations were combined with a reaction progress variable—the mass fraction of oxygen—to automatically select appropriate reduction thresholds during the simulation. This method reached speedup factors of 5–100 with high accuracy. However, the use of pre-tabulated data and choice of reaction progress variable may lack generality in turbulent reacting systems where the combustion conditions are unknown a priori and mixing plays a much stronger role. Therefore, more investigation is needed in this

direction.

Contino et al. [41, 42] proposed the tabulation of dynamic adaptive chemistry (TDAC) method that combines the strengths of tabulation methods (e.g., ISAT) and the DAC scheme. In the base ISAT method, when integrating the reaction mechanism at a thermo-chemical state, a database is first queried to determine if a similar state (and corresponding state after integration) are stored. If an appropriate state is found, an approximation to the integration of the queried state is cheaply constructed from the stored data. Otherwise, the reaction mechanism is directly integrated at the queried state and is used along with the resulting state to update the database. In the TDAC method, the direct integration of the reaction mechanism is accelerated using a DAC scheme resulting in further computational savings. A speedup factor of ~ 500 was reported for premixed combustion cases, and ~ 9 for non-premixed cases with high fidelity predictions of species concentrations and pressure traces. Ren et al. [43, 44] applied a similar scheme to computationally intensive partially stirred reactor simulations, showing up to ~ 1000 speedup factor for premixed cases and a 30 % improvement in computational efficiency over the base ISAT method for a non-premixed case.

Apart from the reduction thresholds explored by Gou et al. [39], the other major factor controlling the performance and accuracy of DRG/DRGEP-based DAC methods is the selection of target species (i.e., search-initiating species). These species are selected for their expected importance to the combustion processes under consideration, and control the reduction of the detailed mechanism. The skeletal mechanisms generated by DRG/DRGEP-based methods include only species whose removal would introduce large error into the production or consumption of these target species. Traditionally, these species are selected before the simulation and treated as target species throughout; however, as discussed by Shi et al. [8], this methodology may overestimate the importance of these target species in certain combustion regimes. For example, while the fuel molecule is almost always used as a target species, in high-temperature, post-ignition combustion almost all the large hydrocarbons have broken down into small molecules and the fuel molecule no longer plays an important role. To address this problem, Shi et al. [8] proposed an extended DAC scheme (EDAC) that switches between a small number of target species sets based on the local thermo-chemical state, resulting in an additional 8–10 % time savings in a three-dimensional direct-injection engine study using skeletal mechanisms for n-heptane. As the EDAC method is limited to a handful of target species sets, it may be inaccurate if unanticipated combustion conditions are encountered. Furthermore, it is not clear whether the EDAC method will easily extend to all problem types and reaction mechanisms, as much of the methodology for switching between the target species sets is empirically derived from zero-dimensional studies or taken from user experience.

This work will describe an automated target species selection method for dynamic adaptive chemistry simulations using only the local thermo-chemical state. First, the DAC scheme and DRGEP method will be outlined in more detail in Section 2, before examining current target species selection methods for the DRGEP method and further demonstrating the need for dynamic target selection. The relative importance index (RII) method of DRGEP target species selection will then be developed in Section 3. In Section 4, the RII method will be validated and its performance will be compared to the current method of static target species selection. Conclusions and suggestions for

future directions will be presented in Section 5.

2. Dynamic Adaptive Chemistry

2.1. Directed Relation Graph with Error Propagation

The DRGEP method as proposed by Pepiot-Desjardins and Pitsch [12] determines important species to be kept in the resulting skeletal mechanism based on their coupling to a list of target species (i.e., search-initiating species) that are expected to be key to the combustion processes under consideration. Each species in the detailed mechanism is represented as a vertex on a graph, and edges between species are defined using a direct interaction coefficient (DIC) that quantifies species coupling based on species production and consumption rates:

$$r_{AB} = \frac{\left| \sum_{i=1}^{N_R} \nu_{A,i} \omega_i \delta_{B,i} \right|}{\max(P_A, C_A)} , \quad (1)$$

where P_A , C_A , and $\delta_{B,i}$ are defined as:

$$P_A = \sum_{i=1}^{N_R} \max(0, \nu_{A,i} \omega_i) , \quad (2)$$

$$C_A = \sum_{i=1}^{N_R} \max(0, -\nu_{A,i} \omega_i) , \quad (3)$$

$$\delta_{B,i} = \begin{cases} 1 & \text{if reaction } i \text{ involves species } B , \\ 0 & \text{otherwise} , \end{cases} \quad (4)$$

$\nu_{A,i}$ is the net stoichiometric coefficient of species A in reaction i , ω_i is the net reaction rate of reaction i , and N_R is the number of reactions in the mechanism. The DIC measures the importance of species B to species A , and represents the error in the overall production of A that would result if B was removed from the mechanism. This definition of the DIC is modified from that of the original DRG method [9] in order to address shortcomings in situations with long chemical paths involving fast chemical modes [12].

The DRGEP method also considers the propagation of error caused by removal of a species along reaction pathways; a path dependent interaction coefficient from a target species T_j to a species B along a pathway p is defined as:

$$r_{T_j B, p} = \prod_{i=1}^{\text{length}(p)-1} r_{S_i S_{i+1}} , \quad (5)$$

where the i th edge of path p connects species S_i and S_{i+1} . The interaction coefficient between species B and target T_j is then defined as

$$R_{B, T_j} = \max_{\text{all paths } p} (r_{T_j B, p}) . \quad (6)$$

Finally, the overall interaction coefficient (OIC) is:

$$R_B = \max_{T_j \in \{Targets\}} R_{B,T_j} . \quad (7)$$

The species B is then considered active in the resultant skeletal mechanism if and only if

$$R_B \geq \varepsilon_{\text{DRGEP}} , \quad (8)$$

where $\varepsilon_{\text{DRGEP}}$ is a specified DRGEP threshold value (e.g., 10^{-4}).

Using this definition of the OIC, the error induced by removal of a species B must propagate along the graph pathways to reach the target species set; a species further away on the graph from the target set is more likely to be removed for this reason. In contrast, the DRG method uses a threshold ε_{DRG} to determine *whether* an edge exists between two species (the edge exists only if the DRG DIC is greater than ε_{DRG}), and a species B is kept in the skeletal mechanism if there exists a path from the target set to B . This binary truncation of the DIC eliminates valuable information on species couplings, and the DRG method will generally produce larger skeletal mechanisms than the DRGEP method as a result [12, 25].

Although the error propagation step of the DRGEP method makes a greater reduction extent possible, the tendency to remove species further away from the targets makes the accuracy of the DRGEP method highly dependent on the proper selection of target species. In Section 2.5 it is demonstrated that large errors can rapidly accumulate when improper target species are selected. In contrast, when using the DRG method any species with edges connecting to a target species will be included in the skeletal mechanism, regardless of their proximity to the target species, and the accuracy of the resulting skeletal mechanism will be less dependent on target species selection.

The direct interaction coefficients are calculated using the linear-time calculation approach (looping over reactions) proposed by Lu and Law [45]. Following the work of Niemeyer and Sung [46], Dijkstra’s algorithm (implemented with a binary heap) is used to calculate the OICs. Additionally, following ideas from the RBFS algorithm presented by Liang et al. [34], edges smaller than the DRGEP threshold are not expanded during the graph search, and the search exits when the maximum value on the heap is less than the DRGEP threshold.

2.2. Dynamic Adaptive Chemistry Scheme Formulation

In the dynamic adaptive chemistry scheme, the starting mechanism is first reduced (e.g., using the DRGEP method), resulting in a skeletal mechanism with m active species (superscript a) and n inactive species (superscript i). For a fixed mass, known volume system, the chemical kinetics equations can be expressed as:

$$\begin{pmatrix} \dot{y}_1^a \\ \vdots \\ \dot{y}_m^a \\ \dot{T} \end{pmatrix} = \begin{pmatrix} f_1(T, p, y_1^a, \dots, y_m^a, y_1^i, \dots, y_n^i) \\ \vdots \\ f_m(T, p, y_1^a, \dots, y_m^a, y_1^i, \dots, y_n^i) \\ f_{m+1}(T, p, y_1^a, \dots, y_m^a, y_1^i, \dots, y_n^i) \end{pmatrix} , \quad (9)$$

where T represents the temperature, p the pressure, y_j^a the mass fraction of the j th of m active species, and y_k^i the mass fraction of the k th of n inactive species. A reaction

is considered active (and thus considered in the right-hand-side functions f_i) if and only if all participating species are active in the skeletal mechanism. It is noted that although inactive species do not participate directly in any active reaction, their removal can induce serious errors in third-body and pressure-dependent reactions. In order to minimize the size of the ODE system while accounting for these third body effects, the DAC scheme calculates the net species production rates of the active species, as well as the derivatives of any state variables (e.g., temperature, pressure) from the local thermo-chemical state—including inactive species—as seen in Eq. (9).

The skeletal mechanism is then integrated for one simulation time-step, holding the inactive species fixed, and the resulting composition is stored. As the simulation time-step for a typical computational fluid dynamics (CFD) simulation is small (e.g., 10^{-5} s) the local skeletal mechanism can be assumed to be valid for the whole time-step. In some situations the flow time-step may be significantly larger, invalidating this assumption. In such cases, multiple reduction steps within a single flow time-step are needed [37].

2.3. Dynamic Adaptive Chemistry Implementation

The DAC scheme used in this work follows that of Liang et al. [34], as seen in Eq. (9). All calculations are completed based on the open-source chemical kinetics software Cantera [47], modified to allow dynamic mechanism reduction as well as to enable dynamic adaptive chemistry integration. The simulation time-step is set to 5×10^{-6} s for all cases, and the integrator (including Jacobian) is reinitialized at every time-step in order to account for the changing problem size. This reinitialization adversely affects the performance of homogeneous simulations because the Jacobian must be recalculated at the beginning of each simulation time-step; however, it is realistic for operator-splitting/fractional-step schemes commonly used in reacting flow simulations where transport invalidates such saved information. Engine simulations are run using constant time steps. At each time step, the crank angle and piston velocity are calculated using the following engine parameters taken from Sjöberg et al. [48]: a displacement volume of 0.981 L, a ratio of connecting rod length to crank radius of 3.2, and a compression ratio of 14. For constant-volume autoignition runs, the ignition delay is determined as the time where temperature reached 400 K greater than the initial temperature. Simulation wall-clock times are reported as the average over 25 runs. All simulations were performed on a single core of a 12-core Xeon X5650 processor.

2.4. Conventional Target Species Selection Methods

In order to investigate the performance of static target species sets, version 2 of the LLNL detailed n-heptane mechanism with 561 species and 4564 reactions [49–51] was used; n-heptane is an important primary reference fuel (PRF) for gasoline that has been extensively studied in the literature. n-Heptane exhibits strong negative temperature coefficient (NTC) behavior, as well as two-stage ignition processes in the low temperature chemistry regime [52]. These factors require a high-fidelity DAC scheme, as large errors can accumulate rapidly otherwise. As such, the n-heptane mechanism is a suitable choice to demonstrate the problems with static target species selection.

The performance and accuracy of the DAC scheme are dependent on the reduction method utilized; for the DRGEP method, these are functions of the target species chosen and the DRGEP threshold ($\varepsilon_{\text{DRGEP}}$) used. Due to the error propagation step, the DRGEP method is more sensitive to the proper selection of target species than the original DRG method. Traditionally, target species have been chosen based on their expected importance to important combustion processes. Typical choices of target species include the fuel, oxygen, combustion products (e.g., CO_2), and certain key radicals and intermediates known to be good indicator species (e.g., H , OH , CO , HO_2) [35, 37, 38]. In addition, species such as NO_x and other pollutants may be added to the target species if high accuracy in emission predictions are required.

A target species set that is not modified through addition or removal of target species throughout a simulation will be termed “static” (unchanging) in this context. One common practice is to remove a species from the target set when its mass fraction falls below some small cutoff (e.g., 10^{-30} [34]). In this work, the static target set category includes cases where the only changes are removals of target species from the set based on mass fraction criteria.

Most attempts to determine appropriate target species sets have consisted of directly comparing the accuracy and performance of a small number of static target species sets. However, it is often difficult to choose a single static target species set that will be appropriate for all combustion processes. As demonstrated by Shi et al. [8], the choice of static target species can overestimate the importance of the targets for certain combustion stages. For example, during the post-ignition period most of the large hydrocarbons have been broken down into small molecules, and thus the fuel species should no longer be considered as a target species. The subsequent section will further detail the problems with static target species sets.

2.5. Performance and Range of Applicability of Static Target Species Sets

In a multidimensional combustion problem many different combustion modes will be encountered. For example, these include local ignition and extinction, as well as a variety of temperatures, pressures, fuel blends, and local equivalence ratios [2]. In order for a DAC scheme to produce high-fidelity results in such a simulation, it should remain accurate at all the possible combustion conditions that may be encountered. In their original DAC study, Liang et al. [34] compared the performance of several static target species sets in single cell, adiabatic, HCCI engine simulations of n-heptane ($\text{n-C}_7\text{H}_{16}$). It was found that the target species set $\{\text{n-C}_7\text{H}_{16}, \text{CO}, \text{HO}_2\}$ was capable of accurately reproducing both the pressure and major species traces with $\varepsilon_{\text{DRGEP}} = 10^{-4}$, and mass fraction cutoff of 10^{-30} .

Figure 1 compares the performance of the target species set $\{\text{n-C}_7\text{H}_{16}, \text{CO}, \text{HO}_2\}$ to that of $\{\text{n-C}_7\text{H}_{16}, \text{CO}, \text{HO}_2, \text{OH}\}$ for constant-volume autoignition simulations of n-heptane at 20 atm, in the temperature range of 860 K–1000 K, and equivalence ratios of $\phi = \{0.5, 1, 2\}$. As seen in Fig. 1a, adding OH to the target species set reduces the maximum percent error in ignition delay from 22 % to 12.5 %. Figure 1b compares the temperature traces of the two static target species sets for 900 K, 20 atm, and $\phi = 0.5$ with that of the full mechanism; the target set $\{\text{n-C}_7\text{H}_{16}, \text{CO}, \text{HO}_2\}$ greatly underpredicts the ignition delay for first stage ignition, and the resulting combustion process differs significantly from that of the full mechanism. Adding OH to the target species

set results in proper prediction of first-stage ignition and a reduction of the overall error in ignition delay prediction. As demonstrated here, adding OH to the target species set can be necessary to accurately predict low-temperature ignition; in a multidimensional simulation where local ignition events may be encountered, this could be vital to the accuracy of the simulation.

However, the use of the target set $\{\text{n-C}_7\text{H}_{16}, \text{CO}, \text{HO}_2, \text{OH}\}$ is not without cost. Figure 2 compares the performance of the two static target sets at HCCI conditions adapted from Liang et al. [34], using the engine parameters from Sjöberg et al. [48] described above. Both target sets reproduced the ignition delay predictions (in terms of crank angle degree) of the full detailed mechanism with high accuracy, resulting a maximum ignition delay error of 0.23° . However, examining the simulation wall times shows that the target set $\{\text{n-C}_7\text{H}_{16}, \text{CO}, \text{HO}_2, \text{OH}\}$ is 5–10 % slower than $\{\text{n-C}_7\text{H}_{16}, \text{CO}, \text{HO}_2\}$. The use of a static target species set with suboptimal execution speed may be necessary in order to ensure that a simulation remains accurate for low-temperature ignition.

Furthermore, this is just one case where the static target species set $\{\text{n-C}_7\text{H}_{16}, \text{CO}, \text{HO}_2\}$ may fail. Consider for example that in diffusive systems, important highly mobile species such as the H radical may need to be added to the target species set. This is demonstrated in Fig. 3, where a stream of stoichiometric n-heptane/air mixture at 300 K, 1 atm flows into a constant-pressure reactor and is ignited by a pilot stream of the H radical at the same temperature. The static target species set $\{\text{n-C}_7\text{H}_{16}, \text{CO}, \text{HO}_2, \text{OH}\}$ is not sufficient to capture this ignition process; instead, the H radical must be added to the target species set in order to maintain accuracy. Although a special situation, this example demonstrates the inability of a static target species set to adapt to combustion conditions dominated by a species not included in the target set. In order to ensure accuracy is maintained in a multidimensional simulation, other species corresponding to major combustion or physical processes may need to be added to the target set. For static target species sets, this requires the user to determine which species may be needed, potentially adversely affecting performance if the proper selection is not made.

These issues with static target species selection motivate us to develop an automated target species selection method which can adapt to changing combustion conditions, ensuring that accuracy and execution speed are maintained. In due course, we will demonstrate that the developed method overcomes these problems and provides a promising alternative to static target species selection.

Finally, it is noted that although the accuracy of the DRG method is less sensitive to the selection of target species (as discussed in Section 2.1), DRG-based DAC schemes will also benefit from use of an automated target species selection method. For example, as discussed previously the inclusion of the fuel molecule in the target species set during post-ignition may cause the resulting local skeletal mechanisms to be larger than necessary. Further, for cases where combustion is controlled by a species not included in the target species set (e.g. analogous to the example in Fig. 3), large errors will still accumulate without manual determination of important species for the target species set. The application of the developed automated target species selection method to DRG-based DAC schemes will address these issues, and is a topic that merits future investigation.

2.6. Dynamic Target Species Selection

The above observations suggest that no single static target species set will be optimal, in line with the work of Shi et al. [8] who proposed the EDAC scheme as a solution to this problem. The EDAC scheme changes the target species set based on the local thermo-chemical state in order to better respond to changing combustion conditions and maintain computational efficiency. However, the EDAC scheme uses empirically derived cutoffs of two proposed progress equivalence ratios as well as the local temperature to switch between various target species sets. As a result, it may be difficult to extend to other combustion problems and reaction mechanisms. Additionally, the EDAC scheme utilizes only a small fixed number of target species sets; the accuracy of the EDAC scheme will suffer if combustion conditions are encountered where none of the target species sets are appropriate.

As such, the goal of this work is to develop an automated method of determining appropriate target species solely from the local thermo-chemical state with the following characteristics:

1. The method will not be limited to a small fixed number of target species sets in order to ensure that accuracy and efficiency are maintained even for unanticipated combustion conditions.
2. The method will be generally applicable, and easy to extend to new problem types and reaction mechanisms.
3. The method will reduce the required user expertise necessary to determine appropriate DRGEP target species.

3. The Relative Importance Index Method

3.1. Sample Ignition Studies

In order to develop this new method of dynamic target species selection, two constant-volume autoignition studies will be examined using version 2 of the LLNL detailed n-heptane mechanism [49–51]. The sample cases will be taken at initial conditions of 700 K and 1000 K, $\phi = 1$, and 2 atm in order to explore the effects of different dominant combustion processes (low-temperature and high-temperature combustion, respectively). The low-temperature case additionally exhibits a two-stage ignition process. This provides an opportunity to investigate the effect of initial temperature and changing chemical pathways on the developed method.

3.2. Graph Structure of the Directed Relation Graph with Error Propagation

The DRGEP adjacency matrix in this work is formatted in row-major format: the value at row i , column j is the DRGEP direct interaction coefficient (DIC) representing the dependence of species i on species j , i.e., r_{ij} . Figure 4 shows an example adjacency matrix generated for the GRI 3.0 mechanism [53]. A species A is considered a neighbor of species B if they both participate in at least one reaction together.

The goal of r_{AB} is to quantify the dependence of species A on B ; therefore, if r_{AB} is large, B is considered important to the production or consumption of A . We consider the sum of the DICs of the neighbors of species A in column A (light grey line in Fig. 4):

$$\text{Column Sum}(A) = \sum_{B_i \in \{\text{neighbors}(A)\}} r_{B_i A} . \quad (10)$$

A species with a large column sum will be important to the production or consumption of many other species, and therefore will be very active in the mechanism. However, since the column sum of a species depends on its number of neighbors, the column sum cannot be used directly to make a fair comparison between different species. Yet, a sense of its utility in tracking the activity of a species can be observed by comparing column sums of the same species in different ignition cases, as depicted in Fig. 5.

In the 700 K ignition case (Fig. 5a), much of the n-heptane is consumed during first-stage ignition; as a result, the column sum of n-heptane drops off significantly at this point. In addition, several large fuel fragments (e.g., i-C₃H₆CO, C₄H₇CO1-4, n-C₃H₁₁CO; see [51] for species dictionary) decompose almost entirely in CO-forming reactions; these pathways shut off near first-stage ignition and are responsible for the corresponding drop in the column sum of CO (note: the magnitude of this drop is masked somewhat by use of a log scale). In the low-temperature region before first-stage ignition, the column sums of O₂ and OH are particularly large as a result of the enhanced R-OH and R-O₂ chemistry. The column sum of HO₂ is lower in this region compared to its column sum after first-stage ignition, indicating it is less important to the mechanism at this lower temperature state. The column sum of CO₂ is small in the low-temperature region, as it is only formed through a handful of reactions. After first-stage ignition, the column sums of all species except n-heptane reach levels similar to that of the 1000 K case, indicative of the temperature effect on the changing strength of chemical pathways.

The absence of a multistage ignition event in the 1000 K ignition case (Fig. 5b) manifests in far more gradual changes in the column sums. The column sum of n-heptane steadily declines as it is consumed. As in the 700 K ignition case, the column sum of CO is initially bolstered due to the decomposition of large fuel fragments (e.g., a-C₃H₅CO, i-C₃H₇CO, C₄H₇CO1-4) forming CO; as these reactions slow down, the column sum of CO declines, becoming stable until climbing slowly leading up to ignition as CO becomes important for CO₂ creation. The column sum of OH stays roughly constant throughout ignition, while the column sum of HO₂ declines approaching ignition as one of its primary consumption pathways, via reactions with n-heptane, slows down. Finally, the column sum of CO₂ is fairly constant until increasing leading up to ignition as CO₂ formation increases.

3.3. Target Species Roles

Common choices of static target species tend to fall into one of two categories: important radicals such as OH and HO₂, and important reactants/products such as the fuel, O₂, CO₂, and CO. Figure 6 presents a sample DRGEP reduction of the n-heptane mechanism using either OH or n-heptane as the sole target species, with $\epsilon_{\text{DRGEP}} = 10^{-5}$.

By examining the neighbors of n-heptane on the reduction graph, it becomes clear that n-heptane only directly participates with a handful of species, but tends to be strongly dependent on each. Most species included in the resultant skeletal mechanism are kept due to their importance to the first ring of species strongly linked to n-heptane.

A species behaving in this manner will be called a “locally important” target species, due to its strong but localized link to a few species. In this mechanism, CO, CO₂, and n-heptane all behave as locally important target species.

The OH reduction graph exhibits an almost entirely different pattern: most of the species included are one step away from OH on the graph with fairly weak direct interaction coefficient links. As a result, relatively few species not directly neighboring OH are included by its choice as a target species. Further, the included species that are more than one step away on the graph from OH tend to be strongly linked to species that directly neighbor OH. This type of behavior will be termed a “globally important” target species, due to its direct but weak link to many species. It has been found that most radicals (e.g., OH, HO₂) behave as globally important target species; in addition, non-radical species involved in many reactions (e.g., O₂) tend to behave as globally important targets as well.

Whether a species behaves as a locally or globally important target is largely controlled by its number of neighbors—i.e., the number of species it directly interacts with through reactions—in the mechanism. Table 2 lists the number of neighbors and corresponding reactions for each of the discussed target species. In general, global target species will tend to have many neighbors in a reaction mechanism. By examining Eq. (10), it can be inferred that globally important target species will have larger column sums than locally important target species in general; as such, the column sum alone is insufficient to compare potential target species.

3.4. Column Sum Normalization Using Row Average

The magnitude of the column sum of a species will roughly depend on its number of neighbors in a mechanism. Therefore, a normalizing factor that tends to be smaller for species with many active neighbors will be necessary. In the n-heptane mechanism used, the OH radical participates in 1051 reactions with 519 different species, while n-heptane is involved in 146 reactions with 36 species as shown in Table 2. OH reacts with so many different species that most neighboring species will only be involved in a relatively small fraction of OH’s total production and consumption. However, n-heptane reacts with relatively fewer species, meaning that most of its neighboring species may control an appreciable fraction of its total consumption and production. From Eq. (1), it can be inferred that the average DRGEP coefficient from OH to its neighbors:

$$\overline{r_{\text{OH}}} = \frac{\sum_{B_i \in \{\text{neighbors}(\text{OH})\}} r_{\text{OH}, B_i}}{\sum_{B_i \in \{\text{neighbors}(\text{OH})\}} \delta_{\text{OH}, B_i}}, \quad (11)$$

where

$$\delta_{\text{OH}, B_i} = \begin{cases} 1 & \text{if } r_{\text{OH}, B_i} > 0 \text{ and} \\ 0 & \text{if } r_{\text{OH}, B_i} = 0, \end{cases} \quad (12)$$

will be much smaller than that of n-heptane, i.e.,:

$$\overline{r_{\text{OH}}} \ll \overline{r_{\text{n-C}_7\text{H}_{16}}}. \quad (13)$$

Consequently, the Row Average of a species A is defined as:

$$\text{Row Average}(A) = \frac{\sum_{B_i \in \{\text{neighbors}(A)\}} r_{AB_i}}{\sum_{B_i \in \{\text{neighbors}(A)\}} \delta_{AB_i}}. \quad (14)$$

Figure 7 shows the row averages for common target species, where it is seen that the relationship predicted by Eq. (13) holds; in both ignition cases, the row average of OH is almost an order of magnitude smaller than that of n-heptane. Further, the species expected to be locally important targets (e.g., CO, n-C₇H₁₆, CO₂), as listed in Table 2, have row averages around an order of magnitude larger than the globally important target species (e.g., O₂, OH, HO₂) for both ignition cases.

3.5. Definition of the Relative Importance Index

Using the row average as a normalizing factor, the relative importance index (RII) is defined as:

$$\text{RII}(A) = \text{Row Average}(A) \times \text{Column Sum}(A) \quad (15)$$

Figure 8 examines the RII values of commonly used target species for the sample ignition cases. In both cases, the RII of OH remains one of the largest throughout the simulation. In the 700 K ignition case (Fig. 8a), the RII of n-heptane declines as it begins to be consumed during first-stage ignition, while the RII of HO₂ rises as new pathways open leading up to the first ignition event. In the low-temperature chemistry region, the RII of OH and O₂ are significantly higher than in the second-stage induction process, again due to the enhanced R–OH and R–O₂ activities. The RII of CO drops during first-stage ignition due to the closing of creation pathways, e.g., via the destruction of CH₃CO and C₂H₅CO, as well as the destruction of large fuel fragments as discussed in Section 3.2, and increases slowly leading up to the final ignition event as CO becomes important for CO₂ formation. Finally, the RII of CO₂ remains relatively low throughout, increasing slightly during first-stage ignition and continuing to increase leading up to the final ignition event.

In the 1000 K case (Fig. 8b), the RII values of HO₂ and n-heptane match each other closely for the beginning of the induction period, as HO₂ consumption is strongly coupled to n-heptane reactions; closer to the point of ignition, this phenomenon stops as HO₂ begins to react with C₂H₄ and CO instead. Near the beginning of the simulation, OH consumption and production are strongly tied to a few reactions (e.g., formation of heptyl radicals via reactions with n-heptane), but as the system approaches the point of ignition, OH consumption and production becomes weakly tied to reactions with many species; as a result, the RII of OH drops slowly throughout the simulation. Similar to its column sum, the RII of CO declines as its creation via the decomposition of large fuel fragments slows down, and later increases slowly leading up to ignition. Finally, CO₂ has a small RII for most of the simulation, but slowly increases leading up to ignition, similar to the 700 K case.

As demonstrated above, the RII of a species balances how strongly a species depends on the majority of its neighbors with how strongly its neighbors depend on it. As such, the use of the RII method allows both globally important and locally important target species to be selected.

3.6. Target Selection Process

The RII of each species can be calculated from the DRGEP coefficients computed during the reduction process. Moreover, it is trivial to add the calculation of active neighbors, column sums, and row sums (later combined with the active neighbors to form the row average) to the pre-existing DRGEP calculation loop, meaning that this method can be executed with minimal overhead. The only additional work required compared to the use of a static target set is a single loop that visits all of the species in the mechanism in order to put this information together to calculate the RII values. In this loop, each species mass fraction is tested; if it is greater than a minimum threshold, the species is inserted into a priority queue, with priority equal to its RII value. When the queue size becomes larger than the number of RII targets to be selected, the species with the minimum RII value is popped off the queue. Depending on the priority queue implementation, the cost of each insertion/removal pair from the priority queue scales linearly or logarithmically with the maximum size of the queue (the number of RII target species to be selected). However, since the number of RII target species is a small fixed number, the overall cost can be considered constant time in asymptotic analysis, i.e., $\mathcal{O}(1)$. Therefore, the total additional overhead of this method compared to using a static target species set scales linearly with the number of species in the mechanism: $\mathcal{O}(N_S)$. As the asymptotic cost of the DRGEP method scales linearly with the number of reactions in the mechanism, $\mathcal{O}(N_R)$, and the number of reactions tends to scale as $N_R \approx 5 \times N_S$ [2], the overhead involved in calculating the RII values has a minimal impact on the overall speed of the mechanism reduction process.

Consider the hypothetical case where a species A participates in only two reactions, one that produces A and another that consumes A . Species A will naturally depend heavily on the few species it reacts with, and have a relatively large row average (compared to those shown in Fig. 7) as a result. If the neighboring species of A also participate in relatively few reactions, then the DRGEP coefficients $r_{B_i A}$ for $B_i \in \{\text{neighbors}(A)\}$ will tend to be closer to 1.0, as each neighbor B_i is likely to depend strongly on A . The resulting column sum of A will then be reasonably similar in magnitude to the number of neighbors of A . As discussed in Section 3.5, the RII responds to how active a species is in various parts of a simulation; however, for such a species A the RII may be large regardless. In this case, it is difficult to determine the importance of A from the RII value alone, so an additional criteria will be needed as a supplement. An example detailing this case is presented in the appendix.

If species A is only weakly reacting it will typically be present in small quantities, whereas when it is strongly reacting it will tend to be present in much larger quantities. Therefore, applying a mass fraction threshold in combination with the RII value can help to properly identify the importance of species. For example, by using a relatively large mass fraction threshold (e.g., 10^{-8}), such a species A will not be considered unless it is strongly reacting. Although a more rigorous filtering method may be more effective in eliminating these potentially biased RII values, this simple combined method proves effective in our studies.

For the 700 K ignition case, Fig. 9 examines the target species selected by the RII method with three RII target species and a mass fraction cutoff of 10^{-8} . During first-stage ignition (Fig. 9a), n-heptane and O_2 are initially selected as target species

simply because they are the reactants. Around 0.005 s, two seven carbon ketones ($n\text{-C}_7\text{H}_{14}\text{O}_3$ 3-5 and $n\text{-C}_7\text{H}_{14}\text{O}_3$ 4-2) are selected as target species by the RII method; these ketones play a large role in OH creation in this region, via both their creation and destruction, and lay along major pathways in the low-temperature breakdown of *n*-heptane. At approximately 0.007 s, $\text{C}_2\text{H}_5\text{COCH}_2$ —a major product of the decomposition of these ketones—begins to be selected as a target; $\text{C}_2\text{H}_5\text{COCH}_2$ additionally lies along an important path for early CO_2 production, forming a bulk of the CH_2CO and C_2H_5 in the system. Once the concentration of OH builds up, it is continually selected until the final ignition event. Radicals such as OH and other species that participate in many reactions in the mechanism are often important even in small concentrations. Further, not selecting these species as targets can greatly increase error, e.g., as in the low-temperature ignition example shown in Fig. 1 where selecting OH as a target species significantly improved accuracy. For this reason, a relaxed mass fraction cutoff of 10^{-15} was used for any species that participated in more than 10 % of the reactions in the mechanism or neighbored more than 50 % of the species in the mechanism. Applying this relaxed mass fraction cutoff results in the selection of OH as a RII target species near the beginning of the simulation.

During second-stage ignition (Fig. 9b), the HCO created during the first stage begins to react, and it is selected as a target species until HCO levels are depleted. Around the same time (approximately 0.012 s), formaldehyde (CH_2O)—an important HCO precursor—begins to be selected as a target, and will continue to be selected for most of the second stage induction period. The selection of $\text{C}_7\text{H}_{14}\text{O}_2$ -5 as a target species is an example of the biased RII values discussed previously (and detailed in the appendix). Although involved in the breakdown of the leftover large fuel fragments from first-stage ignition, this chemistry is of relatively minor importance to the mechanism at this point and therefore $\text{C}_7\text{H}_{14}\text{O}_2$ -5 would likely not be considered as a target species given a more rigorous RII bias detection method. After the final ignition event the chemistry is dominated by the $\text{H}_2\text{--O}_2$ and CO oxidation reactions, and hence HO_2 , HCO, and H_2O_2 are chosen as target species.

In the 1000 K ignition case, shown in Fig. 10, *n*-heptane is again initially selected as a target species. At this higher temperature, destruction of *n*-heptane is initially dominated by reactions with HO_2 . In addition, HO_2 is the primary producer of H_2O_2 and an important secondary creation source for OH through reactions with CH_3 and C_2H_5 , resulting in the selection of HO_2 as a target species for much of the induction period. Initially, a large portion of CO and OH production occurs via CH_2O formation reactions, resulting in its brief selection as a target species; closer to ignition, CH_2O is selected again as it becomes involved in the two primary HCO production pathways. Shortly thereafter, CH_2O is replaced as a target species by HCO as it becomes the primary producer of CO and HO_2 , while becoming the primary consumer of OH. Near ignition, CH_2CHO begins to be selected as a target species; as the primary product of the decomposition of several large ketones (e.g., $n\text{-C}_7\text{H}_{14}\text{O}_3$ 1-3, $n\text{-C}_5\text{H}_{10}\text{O}_3$ 1-3), CH_2CHO reacts almost exclusively with oxygen, to form CH_2O , OH, and CO. In addition to being an important secondary path for the creation of CH_2O , it provides an important path to the CO_2 creation reactions. After ignition, the same targets were selected as in the 700 K case post ignition: HO_2 , HCO, and H_2O_2 , covering $\text{H}_2\text{--O}_2$ and CO oxidation reactions.

3.7. Adaptability to Combustion Conditions

In Section 2.5, it was found that a static target species set not containing the H radical could not capture the ignition of a stream of stoichiometric n-heptane/air flowing into a constant-pressure reactor by a pilot stream of the H radical (Fig. 3). In Fig. 11, the temperature traces of a three RII target species criterion are compared to those of the static target sets for the same problem. As discussed in Section 3.6, highly connected species were identified as species participating in 10 % or more of the total reactions, or neighboring 50 % of total species. For this mechanism, the species identified by these criteria were H, HO₂, O₂, and OH and a relaxed mass fraction cutoff of 10⁻¹⁵ was used for each. The RII target species criteria has no difficulty picking up the H radical induced ignition, demonstrating an ability to adapt to different combustion conditions. It is noted that when using a DRG-based DAC scheme in this example, the RII target species criteria can accurately predict ignition while a static target species set that does not include the H radical will rapidly accumulate large errors, similar to the DRGEP-based DAC scheme without considering the H radical. This further demonstrates that DRG-based DAC schemes will benefit from the use of an automated target selection method.

4. Validation and Performance

4.1. n-Heptane Constant Volume Autoignition Simulations

As seen in Section 2.5, a target species set that is accurate for single-cell HCCI engine simulations can often be quite inaccurate in predicting constant-volume ignition delay. Therefore, the accuracy of the RII method will first be assessed with constant-volume ignition delay predictions over a wide range of initial conditions in order to validate the method on a more stringent problem set. The target sets were tested at 5 atm and 20 atm, equivalence ratios of $\phi = \{0.5, 1, 2\}$, and initial temperatures ranging from 640 K to 1200 K.

Table 3 compares the accuracy of target sets with three, four, and five RII targets to those of the static target sets $\{\text{n-C}_7\text{H}_{16}, \text{CO}, \text{HO}_2\}$ and $\{\text{n-C}_7\text{H}_{16}, \text{CO}, \text{HO}_2, \text{OH}\}$. The mass fraction cutoffs were set to 10⁻³⁰ for the static targets sets and 10⁻⁸ for the RII target species criteria. As in Section 3.7, a relaxed mass fraction cutoff of 10⁻¹⁵ was used for the highly connected species H, HO₂, O₂, and OH.

It is seen from Table 3 that the RII target criteria are capable of accurately predicting the ignition delays for the entire range of initial conditions. Furthermore, the RII target species criteria are more accurate than the static target sets, having maximum errors of 10.6 % to 5.45 % for 3 to 5 RII targets compared to 12.5 % for $\{\text{n-C}_7\text{H}_{16}, \text{CO}, \text{HO}_2, \text{OH}\}$ and 22.5 % for $\{\text{n-C}_7\text{H}_{16}, \text{CO}, \text{HO}_2\}$. It is noted that the maximum error does not decrease monotonically with increasing numbers of RII target species; the five RII target species criterion has a slightly greater maximum error than the four RII target species criterion. However, the average error over all cases does decrease monotonically with larger RII target species criteria.

4.2. *n*-Heptane HCCI Engine Simulations

Next, the performance and accuracy of the RII method will be compared to that of the two static target species sets in single-cell engine simulations at HCCI conditions adapted from Liang et al. [34] (Table 1). Figure 12 compares the accuracy and wall-clock time of the two static target species sets with three, four, and five RII target species criteria, as used in Section 4.1. The simulation wall-clock times shown are normalized by the longest run for each starting condition.

Figure 12a compares the static target sets to the RII target species criteria, with a DRGEP threshold of $\varepsilon_{\text{DRGEP}} = 10^{-4}$. The RII target species criteria are more accurate for the low equivalence ratio case (NH1), but are less accurate for the larger equivalence ratio cases (NH2, NH3). Nevertheless, the maximum error for the static targets is 0.23° , while the maximum error for the RII target species criteria ranges from 0.35° to 0.23° for three to five RII targets, respectively. The speed of the RII simulations is roughly the same as the static target sets; although for the $\phi = 1.2$ case (NH3), the RII target species criteria are approximately 15–20 % slower than the static target sets.

The inaccuracy of the RII target species criteria in cases NH2 and NH3 is the result of the relatively large mass fraction cutoff used (i.e., 10^{-8}). This large mass fraction cutoff delays the selection of CH_2CHO and $\text{CH}_2\text{CH}_2\text{CHO}$ (early CO_2 precursors) as target species, which in turn delays ignition. By lowering the mass fraction cutoff (e.g., to 10^{-10}) the accuracy of the RII target species criteria becomes equivalent to that of the static target species sets; however, the computational efficiency of the RII method is decreased due to issues with biased RII values (as described in Section 3.6), and the static target species sets perform up to 30 % faster for cases NH2 and NH3. Although use of the larger mass fraction cutoff induces higher error in these cases, in most situations it provides a good balance between accuracy and efficiency.

Figure 12b compares the various target sets with a DRGEP threshold of $\varepsilon_{\text{DRGEP}} = 10^{-3}$. At this larger DRGEP threshold, the RII target species criteria are considerably more accurate than the static targets for case NH1, with a maximum error of $\sim 0.5^\circ$ compared to $\sim 1.4^\circ$ for the static target sets (both with and without OH). For cases NH2 and NH3, the accuracy of RII target species criteria are comparable to the static target sets. The simulation wall-clock times show again that the speeds of RII and static target sets are roughly equivalent for this higher DRGEP threshold. A speedup factor of approximately $4.5\text{--}6\times$ over the detailed mechanism was achieved for *n*-heptane with both the static target species sets and RII target species criteria.

4.3. Isopentanol Mechanism

Isopentanol is a next-generation biofuel with the potential to greatly reduce NO_x and particulate emissions; as biofuels are produced from renewable sources, net CO_2 emissions can be reduced as well [54]. Less hygroscopic, corrosive, and miscible in water than ethanol [54, 55], isopentanol is more compatible with the existing fuel infrastructure. In addition, isopentanol has a volumetric energy density over 30 % higher than that of ethanol [56]. These factors make isopentanol a promising alternative for future combustion devices, either by itself or blended with gasoline.

Isopentanol, a C_5 alcohol with a methyl branch, exhibits largely different chemistry than a typical alkane fuel, including minimal NTC behavior and a single-stage ignition

process. As such, it is a good choice to demonstrate the general applicability of the RII method. A detailed mechanism from Sarathy et al. [54] with 360 species was used to study constant-volume autoignition problems and single-cell HCCI simulations.

4.4. Isopentanol Constant Volume Simulations

First, the accuracy of the RII method will be compared to that of static target species sets using constant-volume simulations with initial conditions of 5 atm and 20 atm, 700 K to 1100 K, and $\phi = \{0.38, 0.6, 1.0\}$. The static target sets again used a mass fraction cutoff of 10^{-30} , while the RII target species criteria used a cutoff of 10^{-8} . For this mechanism, the highly connected species (as determined by the criteria detailed in Section 3.6) were CH_3 , H , HO_2 , O_2 , and OH ; the mass fraction cutoff for these species was relaxed to 10^{-15} . For all cases, the DRGEP cutoff threshold $\varepsilon_{\text{DRGEP}}$ was set to 5×10^{-3} .

Table 4 compares the error statistics for the various target species sets. It is seen that including OH in the static target species set again increases the accuracy from a maximum error of 12.3 % to 9.51 %, while the RII target species criteria are even more accurate with a maximum error of 6.86 % to 5.31 % for the three to five RII target species criteria, respectively. The average percent error over all the cases is significantly lower for the RII target species criteria as well, ranging from 2.83 % to 2.00 %, compared to 4.43 % to 3.40 % for the static target sets.

4.5. Isopentanol HCCI Engine Simulations

Next, the RII method was tested with single-cell engine simulations at HCCI conditions listed in Table 5, adapted from those given by Yang et al. [55]. The static target sets tested were the same as those used in the constant-volume autoignition simulations.

To achieve a simplified exhaust gas recirculation (EGR) process, the engine cycle was first simulated at the conditions listed. At 122° ATDC (corresponding to exhaust valve opening), the state of the engine cell was extracted and assumed to be cooled and throttled to 367 K and 202 kPa (i.e., the exit temperature and pressure of the EGR cooling cycle in [55]), before being mixed adiabatically on a per mass basis with the inlet gas. In addition, the amount of isopentanol was adjusted to keep the charge mass equivalence ratio,

$$\phi_m = \frac{(F/C)}{(F/A)_{\text{stoich}}} , \quad (16)$$

constant following the procedure of Yang et al. [55]. (F/C) is the mass ratio of fuel to total charge mass (i.e. fresh air and recycled exhaust gas) and $(F/A)_{\text{stoich}}$ is the stoichiometric fuel/air mass ratio.

Figure 13 compares the errors in ignition delay and wall-clock times of the various target species sets; the latter were again normalized by the longest time in each case. At the smaller DRGEP threshold (10^{-3}), the static target species sets have a maximum error of 0.24° to 0.28° with and without OH, respectively, while the RII target species criteria have a maximum error of approximately 0.05° (Fig. 13a). The RII target species criteria performed around 5 % to 10 % slower at this DRGEP threshold.

For a higher DRGEP threshold (5×10^{-3}), the RII targets species criteria perform more accurately than the static target species sets for almost all conditions shown in

(Fig. 13b). The static target sets result in maximum errors of 1.28° to 1.56° with and without OH, respectively, while the RII target species criteria produce maximum errors of 0.25° to 0.19° for three to five RII targets, respectively. In terms of computational efficiency, the various target sets performed similarly at this DRGEP threshold, with the static target sets operating faster for some conditions and the RII target species criteria faster for others. A speedup factor of approximately $2-3\times$ over the detailed mechanism was achieved for isopentanol with both the static target species sets and RII target species criteria.

5. Conclusions and Future Work

The current methodology of DRGEP-based DAC simulations typically utilizes a static (unchanging) set of target species for the DRGEP method. However, as demonstrated in Section 2.5, while using a single static target species set (e.g., $\{n\text{-C}_7\text{H}_{16}, \text{CO}, \text{HO}_2\}$) may work well for simple simulations, this approach could suffer in accuracy for different problems and/or combustion conditions. In an attempt to improve accuracy of a static target species set in general, additional key species can be added (e.g., OH); however, this requires expert knowledge of the chemical mechanism at hand and may adversely affect performance. Further, for a complex multidimensional simulation the range of combustion conditions encountered cannot be known with certainty a priori. While multiple cases could be run to test different target species sets, the prohibitive computational expense of this operation makes it difficult to ensure the applicability of a selected target species set.

In this work, a novel method of automatically determining appropriate target species for DRGEP-based DAC simulations was developed and implemented that relies solely on the local thermo-chemical state: the relative importance index (RII) method. As shown in Section 3.6, the RII method selects target species based on their relevance to the local chemistry. For single-cell HCCI engine simulations of n-heptane and isopentanol, the RII method was demonstrated to match the accuracy of conventional static target species sets at small DRGEP thresholds, while maintaining higher accuracy at larger DRGEP thresholds. In all cases, RII performed similarly in terms of execution speed to the static targets sets. In addition, the same RII target species criteria were shown to be more accurate than static target sets in the more stringent cases of constant-volume autoignition studies for both the n-heptane and isopentanol mechanisms. These factors make the RII method a promising candidate to maintain accuracy and ensure efficient performance for complex multidimensional simulations. Additionally, automating the target selection process with RII greatly reduces the user knowledge required to select appropriate DRGEP target species, simplifying the use of a DRGEP-based DAC method for new problem types and reaction mechanisms.

Further improvement to the RII method is likely to be achieved through the following avenues. First, the RII bias filtering method discussed in Section 3.6 is largely empirical, and while sufficient in the context of this work, a more rigorous method may improve the accuracy of the RII method. Second, a natural extension for the RII method is the automatic selection of a DRGEP threshold derived from the local thermo-chemical state. In the DAC-based simulation efforts demonstrated thus far,

the DRGEP threshold value is typically a static value selected based on past experience, which provides only rudimentary a priori error control. As a result, the DRGEP threshold is often set to a conservative value in order to ensure accuracy; however, as shown by Shi et al. [8], the DRGEP threshold can be increased in certain cases—e.g., after combustion has completed—with minimal effect on the accuracy of the simulation while improving execution speed. Gou et al. [39] proposed an error controlled DAC scheme with some success; however, it is noted that their method relied on data tabulated from homogeneous ignition studies and therefore may not be applicable for general multidimensional simulations. Therefore, more investigation is needed in this direction. Finally, although the accuracy of the DRG method is less sensitive to the proper selection of target species, there are cases where DRG-based DAC schemes would benefit from the use of an automated target selection method, e.g. by removing the fuel molecule from the target species set after ignition or in situations where combustion is dominated by a single species not included in the static target species set. To improve the generality of the RII method, its applicability to DRG-based DAC schemes should be thoroughly investigated.

Acknowledgments

This work was supported by the Combustion Energy Frontier Research Center, an Energy Frontier Research Center funded by the US Department of Energy, Office of Science, Office of Basic Energy Sciences under award number DE-SC0001198, and the National Science Foundation Graduate STEM Fellows in K-12 Education program.

Appendix: Simple Example of RII Biasing

Consider the reaction pathways presented in Fig. 14, which gives an example of two steps in a long chain of large hydrocarbon breakdown reactions. Species A participates in only two reactions: formation through the decomposition of B_1 (forming A and B_4), and a reaction with B_3 forming B_2 . Note that the neighboring species of A in this simple example are B_1 – B_4 . It will be demonstrated that the RII of a species A in this situation can be large enough that the species may be selected as a target species even if A is only weakly reacting. Hence for a species with only a few neighbors, an additional criterion to determine the importance of a species is needed to supplement the RII.

If the rate of production of A (P_A) is greater than the rate of consumption of A (C_A), from Eq. (1) the DICs r_{AB_1} and r_{AB_4} must be unity as $P_A > C_A$. In this case, the DICs r_{AB_2} and r_{AB_3} will simply be equal to the ratio of consumption to production of A , $\frac{C_A}{P_A}$. Therefore:

$$\text{Row Average}(A) = \frac{2 + 2 \times \frac{C_A}{P_A}}{4}. \quad (17)$$

If on the other hand, A is being consumed faster than it is produced it follows that:

$$\text{Row Average}(A) = \frac{2 + 2 \times \frac{P_A}{C_A}}{4}. \quad (18)$$

From which it is seen the row average of A is bounded between 0.5 and 1.0. Even if the production and consumption of A is highly unbalanced, the lower bound on the row average of A is very large compared to those shown in Fig. 7.

Further, if B_1 – B_4 react in relatively few reactions, as will typically be the case for a long chain of large hydrocarbon breakdown reactions, the column sum of A will naturally tend to be large. This phenomenon is analogous to the discussion in Section 3.4; as each neighboring species B_i has few neighbors, A will likely be involved in a large portion of the production or consumption of B_i . Hence the DICs making up the column sum of A , r_{B_1A} , r_{B_2A} , r_{B_3A} , and r_{B_4A} , will tend to be closer to 1.0 in value. This implies that although the Column Sum(A) is bounded between 0.0 and 4.0, it may be closer to 4.0.

Therefore, the RII(A) can easily be larger than the RII's of commonly used target species presented in Fig. 8. Consider, if we simply take the row average and column sums to be the midpoints of their bounds (i.e., Row Average(A) = 0.75, and Column Sum(A) = 2.0), the RII(A) will be equal to 1.5, comparable to that of OH and greater than that of HO₂ in the sample ignition cases (Fig. 8). Finally it is noted that the above discussion applies whether A is strongly reacting or not; hence, in such a situation the RII(A) may not accurately reflect the true importance of A .

Such behavior is only possible for a species that participates in just a handful of reactions, with only a few neighbors. For a species with more neighbors, the RII will better reflect the relative activity of the species as seen in Section 3.5, because it becomes unlikely that the production and consumption rates of the species and its neighbors will exhibit the strongly coupled behavior demonstrated in this example. Further, since the species A in this example is only a single stage in a long breakdown chain, it is not

expected to exist in high concentrations for long periods. Instead it is more likely that species A will be present in large quantities only when it is strongly reacting; hence using a relatively large mass fraction cutoff (e.g., 10^{-8}) will tend to eliminate species A from consideration as a target species when it is only weakly reacting.

References

- [1] S. M. Correa, *Proc. Combust. Inst.* 27 (1998) 1793–1807.
- [2] T. Lu, C. K. Law, *Prog. Energy Combust. Sci.* 35 (2009) 192–215.
- [3] S. J. Shanbhogue, S. Husain, T. Lieuwen, *Prog. Energy Combust. Sci.* 35 (2009) 98–120.
- [4] M. Yao, Z. Zheng, H. Liu, *Prog. Energy Combust. Sci.* 35 (2009) 398–437.
- [5] S. R. Tonse, N. W. Moriarty, M. Frenklach, N. J. Brown, *Int. J. Chem. Kinet.* 35 (2003) 438–452.
- [6] L. Liang, S.-C. Kong, C. Jung, R. D. Reitz, *J. Eng. Gas Turb. Power* 129 (2007) 271–278.
- [7] Y. Shi, W. H. Green, H.-W. Wong, O. O. Oluwole, *Combust. Flame* 159 (2012) 2388–2397.
- [8] Y. Shi, L. Liang, H.-W. Ge, R. D. Reitz, *Combust. Theory Model.* 14 (2010) 69–89.
- [9] T. Lu, C. K. Law, *Proc. Combust. Inst.* 30 (2005) 1333–1341.
- [10] T. Lu, C. K. Law, *Combust. Flame* 144 (2006) 24–36.
- [11] T. Lu, C. K. Law, *Combust. Flame* 154 (2008) 153–163.
- [12] P. Pepiot-Desjardins, H. Pitsch, *Combust. Flame* 154 (2008) 67–81.
- [13] H. Rabitz, M. Kramer, D. Dacol, *Annu. Rev. Phys. Chem.* 34 (1983) 419–461.
- [14] T. Turányi, *New J. Chem.* 14 (1990) 795–803.
- [15] T. Turányi, *J. Math. Chem.* 5 (1990) 203–248.
- [16] S. Vajda, P. Valko, T. Turányi, *Int. J. Chem. Kinet.* 17 (1985) 55–81.
- [17] T. Løvås, D. Nilsson, F. Mauss, *Proc. Combust. Inst.* 28 (2000) 1809–1815.
- [18] T. Løvås, F. Mauss, C. Hasse, N. Peters, *Proc. Combust. Inst.* 29 (2002) 1403–1410.
- [19] T. Løvås, *Combust. Flame* 156 (2009) 1348–1358.
- [20] S. Lam, D. Coussis, *Proc. Combust. Inst.* 22 (1989) 931–941.
- [21] S. Lam, *Combust. Sci. Tech.* 89 (1993) 375–404.
- [22] S. Lam, D. Goussis, *Int. J. Chem. Kinet.* 26 (1994) 461–486.
- [23] M. Valorani, F. Creta, D. A. Goussis, J. C. Lee, H. N. Najm, *Combust. Flame* 146 (2006) 29–51.

- [24] M. Valorani, F. Creta, F. Donato, H. N. Najm, D. A. Goussis, *Proc. Combust. Inst.* 31 (2007) 483–490.
- [25] K. E. Niemeyer, C.-J. Sung, M. P. Raju, *Combust. Flame* 157 (2010) 1760–1770.
- [26] K. E. Niemeyer, C.-J. Sung, *Combust. Flame* 161 (2014) 2752–2764.
- [27] D. L. Chapman, L. K. Underhill, *J. Chem. Soc., Trans.* 103 (1913) 496–508.
- [28] S. W. Benson, *J. Chem. Phys.* 20 (1952) 1605–1612.
- [29] J. D. Ramshaw, *Phys. Fluids* 23 (1980) 675.
- [30] U. Maas, S. B. Pope, *Combust. Flame* 88 (1992) 239–264.
- [31] S. Pope, *Combust. Theory Model.* 1 (1997) 41–63.
- [32] S. R. Tonse, N. W. Moriarty, N. J. Brown, M. Frenklach, *Israel J. Chem.* 39 (1999) 97–106.
- [33] F. Christo, A. Masri, E. Nebot, S. Pope, *Proc. Combust. Inst.* 26 (1996) 43–48.
- [34] L. Liang, J. G. Stevens, J. T. Farrell, *Proc. Combust. Inst.* 32 (2009) 527–534.
- [35] L. Liang, J. G. Stevens, S. Raman, J. T. Farrell, *Combust. Flame* 156 (2009) 1493–1502.
- [36] K. He, I. P. Androulakis, M. G. Ierapetritou, *Chem. Eng. Sci.* 65 (2010) 1173–1184.
- [37] H. Yang, Z. Ren, T. Lu, G. M. Goldin, *Combust. Theory Model.* 17 (2012) 167–183.
- [38] L. Tosatto, B. A. V. Bennett, M. D. Smooke, *Combust. Flame* 160 (2013) 1572–1582.
- [39] X. Gou, Z. Chen, W. Sun, Y. Ju, *Combust. Flame* 160 (2013) 225–231.
- [40] W. Sun, Z. Chen, X. Gou, Y. Ju, *Combust. Flame* 157 (2010) 1298–1307.
- [41] F. Contino, H. Jeanmart, T. Lucchini, G. D’Errico, *Proc. Combust. Inst.* 33 (2011) 3057 – 3064.
- [42] F. Contino, T. Lucchini, G. D’Errico, C. Duynslaegher, V. Dias, H. Jeanmart, Simulations of advanced combustion modes using detailed chemistry combined with tabulation and mechanism reduction techniques, Technical Report, SAE Technical Paper 2012-01-0145, 2012.
- [43] Z. Ren, Y. Liu, T. Lu, L. Lu, O. O. Oluwole, G. M. Goldin, *Combust. Flame* 161 (2014) 127–137.
- [44] Z. Ren, C. Xu, T. Lu, M. A. Singer, *J. Comput. Phys.* 263 (2014) 19–36.

- [45] T. Lu, C. K. Law, *Combust. Flame* 146 (2006) 472–483.
- [46] K. E. Niemeyer, C.-J. Sung, *Combust. Flame* 158 (2011) 1439–1443.
- [47] D. Goodwin, N. Malaya, H. Moffat, R. Speth, Cantera: an object-oriented software toolkit for chemical kinetics, thermodynamics, and transport processes. Version 2.1a1, <https://code.google.com/p/cantera/>, 2014.
- [48] M. Sjöberg, J. E. Dec, W. Hwang, Thermodynamic and chemical effects of EGR and its constituents on HCCI autoignition, Technical Report, SAE Technical Paper 2007-01-0207, 2007.
- [49] H. J. Curran, P. Gaffuri, W. J. Pitz, C. K. Westbrook, *Combust. Flame* 114 (1998) 149–177.
- [50] H. J. Curran, P. Gaffuri, W. J. Pitz, C. K. Westbrook, *Combust. Flame* 129 (2002) 253–280.
- [51] H. Curran, P. Gaffuri, W. Pitz, C. Westbrook, n-heptane, detailed mechanism, version 2, https://www-pls.llnl.gov/?url=science_and_technology-chemistry-combustion-n_heptane_version_2, 2004.
- [52] N. Peters, G. Paczko, R. Seiser, K. Seshadri, *Combust. Flame* 128 (2002) 38–59.
- [53] G. P. Smith, D. M. Golden, M. Frenklach, N. W. Moriarty, B. Eiteneer, M. Goldenberg, C. T. Bowman, R. K. Hanson, S. Song, W. C. Gardiner Jr, GRI-Mech 3.0, 1999.
- [54] S. M. Sarathy, S. Park, B. W. Weber, W. Wang, P. S. Veloo, A. C. Davis, C. Togbe, C. K. Westbrook, O. Park, G. Dayma, Z. Luo, M. A. Oehlschlaeger, F. N. Egolfopoulos, T. Lu, W. J. Pitz, C.-J. Sung, P. Dagaut, *Combustion and Flame* 160 (2013) 2712 – 2728.
- [55] Y. Yang, J. Dec, N. Dronniou, B. Simmons, *SAE Int. J. Fuels Lubr.* 3 (2010) 725–741.
- [56] T. Tsujimura, W. J. Pitz, F. Gillespie, H. J. Curran, B. W. Weber, Y. Zhang, C.-J. Sung, *Energy Fuels* 26 (2012) 4871–4886.

	ϕ	RPM	T_0 (K)	p_0 (kPa)
NH1	0.5	1500	330	101.325
NH2	1.0	1500	330	101.325
NH3	1.2	1500	330	101.325

Table 1: Initial conditions for single cell n-heptane engine simulations at HCCI conditions adapted from Liang et al. [34]. Simulations begin at 150° before top dead center (BTDC) and end at 150° after top dead center (ATDC) to approximate the engine cycle between intake valve closing and exhaust valve opening.

Species	Number of neighbors	Number of reactions	Target species importance
CO	63	96	local
CO ₂	25	31	local
HO ₂	337	763	global
n-C ₇ H ₁₆	36	146	local
O ₂	352	624	global
OH	519	1051	global

Table 2: The number of neighbors, reactions, and target species behavior related to some commonly used DRGEP target species.

Target set	Maximum error (%)	Average error (%)
{n-C ₇ H ₁₆ , CO, HO ₂ }	22.5	2.06
{n-C ₇ H ₁₆ , CO, HO ₂ , OH}	12.5	1.52
{RII = 3}	10.6	1.69
{RII = 4}	5.34	1.31
{RII = 5}	5.45	1.16

Table 3: Error statistics for constant-volume autoignition studies of n-heptane at 5 atm and 20 atm, 640 K to 1200 K, and $\phi = \{0.5, 1, 2\}$. The DRGEP threshold was set to $\varepsilon_{\text{DRGEP}} = 10^{-4}$ in all cases.

Target set	Maximum error (%)	Average error (%)
{i-C ₅ H ₁₁ OH, CO, HO ₂ }	12.3	4.43
{i-C ₅ H ₁₁ OH, CO, HO ₂ , OH}	9.51	3.40
{RII = 3}	6.86	2.83
{RII = 4}	6.58	2.33
{RII = 5}	5.31	2.00

Table 4: Error statistics for constant-volume autoignition studies of isopentanol at 5 atm and 20 atm, 700 K to 1100 K, and $\phi = \{0.38, 0.6, 1.0\}$. The DRGEP threshold was set to $\varepsilon_{\text{DRGEP}} = 5 \times 10^{-3}$ in all cases.

	ϕ	RPM	T_0 (K)	p_0 (kPa)	EGR %
IP1	0.38	1200	405	101.325	0
IP2	0.38	1200	364	140.000	0
IP3	0.38	1200	333	200.000	37

Table 5: Initial conditions for single cell isopentanol engine simulations at HCCI conditions adapted from Yang et al. [55]. Simulations begin at 158° BTDC and end at 122° ATDC.

(a) Constant volume ignition delays at 20 atm

(b) Temperature trace of constant-volume autoignition simulation at 900 K, 20 atm, $\phi = 0.5$

Figure 1: In (a) the constant-volume ignition delays using two static target species sets for n-heptane at 20 atm, with $\epsilon_{\text{DRGEP}} = 10^{-4}$ and a mass fraction cutoff of 10^{-30} are displayed. The static target species set without OH has a maximum percent error of 22 % while the set including OH has a maximum error of just 12.5 %. In (b) the temperature trace of the two target sets at 900 K, 20 atm, $\phi = 0.5$ is compared to that of the full mechanism; by adding OH to the target species set, the percent error in ignition delay is decreased from 17 % to just 1.8 %.

(a) Ignition delay error in crank angle degrees

(b) Normalized simulation wall time

Figure 2: Comparison of (a) ignition delay error and (b) simulation wall time normalized by the longest wall time in each case for single-cell n-heptane engine simulations at the HCCI conditions listed in Table 1. The two static target species sets were used with $\varepsilon_{\text{DRGEP}} = 10^{-4}$ and a mass fraction cutoff of 10^{-30} .

Figure 3: Temperature trace of two static target species sets for the constant-pressure reactor ignition of a stoichiometric n-heptane/air mixture by a pilot stream of the H radical. The mass flow rate of the pilot stream is 1:100 with that of the n-heptane/air mixture.

(a) An example diagram of a species and a neighbor in the DRGEP graph structure. The dark and light grey arrows correspond to the lines in (b), i.e. to the rows and columns of the adjacency matrix respectively. Species *B* corresponds to a representative neighbor of *A*, and is marked on the adjacency matrix.

(b) The adjacency matrix

Figure 4: An example of (a) the DRGEP graph and (b) adjacency matrix structure used in this work created using the GRI 3.0 mechanism [53].

(a) 700 K ignition case

(b) 1000 K ignition case

Figure 5: The Column Sum of commonly used target species for the sample n-heptane constant-volume ignition cases.

Figure 6: A sample DRGEP reduction with $\varepsilon_{\text{DRGEP}} = 10^{-5}$ for two choices of target species. The large black circles indicate (a) $\text{n-C}_7\text{H}_{16}$ and (b) OH. Arrow width and color indicate the magnitude of the direct interaction coefficient between species.

(a) 700 K ignition case

(b) 1000 K ignition case

Figure 7: The Row Average of commonly used target species for the sample n-heptane constant-volume ignition cases.

(a) 700 K ignition case

(b) 1000 K ignition case

Figure 8: The RII of commonly used target species for the sample n-heptane constant-volume ignition cases.

(a) First stage ignition

(b) Second stage ignition

(c) Post ignition

Figure 9: Target species selection for a three RII target species criterion with mass fraction cutoff of 10^{-8} for the 700 K ignition case exhibiting two-stage ignition response.

(a) Ignition

(b) Post ignition

Figure 10: Target species selection for a three RII target species criterion with mass fraction cutoff of 10^{-8} for the 1000 K ignition case exhibiting single-stage ignition response.

Figure 11: Temperature trace of two static target species sets compared to that of a three RII target species criteria for the constant-pressure reactor ignition of a stoichiometric n-heptane/air mixture by a pilot stream of the H radical. The mass flow rate of the pilot stream is 1:100 with that of the n-heptane/air mixture.

Figure 12: Comparison of the ignition delay error (in crank angle degrees) and simulation wall time normalized by the longest time in each case for single-cell HCCI simulations of n-heptane at conditions listed in Table 1.

Figure 13: Comparison of the ignition delay error (in crank angle degrees) and simulation wall time normalized by the longest time in each case for single-cell HCCI simulations of isopentanol at conditions in Table 5.

Figure 14: A simple example of reaction pathways that may lead to RII biasing. Species B_1 – B_4 participate in reactions with A , and hence are its direct neighbors.