

Effects of Polynomial Order and Stabilization Technique on the Large Eddy Simulation of an Axisymmetric Scramjet Combustor using High-Order Flux Reconstruction

T. Sailor Koeplinger¹, Jack R. Edwards²
North Carolina State University, Raleigh, NC, 27606, USA

Seth C. Spiegel³
NASA Glenn Research Center, Cleveland, OH, 44135, USA

The high-order flux reconstruction method is used to analyze the dual-mode operation of an axisymmetric inlet-isolator-combustor flowpath experimentally tested in the ACT-II arc-heated combustion tunnel at the University of Illinois at Urbana-Champaign. A 7-species/3-step reaction mechanism developed at the Air Force Research Laboratory is used in these calculations. Modifying the computation of the backwards rate was found to admit a favorable solution. An isothermal wall of 375 K is assumed. A baseline solution is obtained using a divergence-based filter for solution stabilization on a grid with 20 million solution points and cubic elements (P3, 4th order). This solution compares well with experimental wall pressure measurements, capturing the shock train location correctly while over-predicting the cavity pressure. The cavity pressure discrepancy is caused by the presence of a cavity-wake stabilized instead of a jet-wake stabilized flame. Next, a solution is obtained which uses linear, P1, elements while maintaining the same number of solution points and stabilization mechanism. The P1 solution was found to capture the inlet shocks better but was much more dissipative for unsteady flow features. It was unable to sustain a lit cavity without using ignition aids. Additionally, a stabilization scheme based on a local minimum entropy is implemented for non-reacting flow. This scheme admits a solution with lower dissipation than the divergence filter and requires minimal user-input. The choice of stabilization technique was found to have a much larger effect on the higher-order P3 solution than the 2nd order P1 solution.

I. Nomenclature

a	=	speed of sound
h	=	width of an element
p	=	pressure
ρ	=	mixture density
ρ_i	=	density of species i
P	=	polynomial order of the solution
S	=	Sensor value used in divergence filter
s	=	Numerical entropy function used in entropy limiter
u	=	Vector of conserved variables
$\tilde{u}$	=	Filtered solution vector
$\bar{u}$	=	Mean of solution vector on an element
$\vec{V}$	=	Velocity vector

¹ Research Assistant, Mechanical and Aerospace Engineering, Student Member AIAA

² Professor, Mechanical and Aerospace Engineering, Fellow AIAA

³ Aerospace Engineer, Inlets and Nozzles Branch, Senior Member AIAA

Y	= Species mass fraction
Ω	= Interior domain of an element
δs	= Differential portion of element boundary
$\delta\Omega$	= Boundary of an element

II. Introduction

High-speed air-breathing propulsion (HSABP) engines, such as scramjets and ramjets, are of great interest for applications ranging from improved access to space to military uses. Their operation is difficult to achieve because of the extreme environments which they must withstand and the very short residence times available for the fuel and air to mix and burn. The premier tool for high-fidelity analysis of these flowpaths is large-eddy simulation (LES). These are scale-resolving simulations which aim to capture unsteadiness at scales which have a bulk effect on the flow and model the smaller, more universal, scales of turbulence. These simulations require the use of highly accurate numerical methods to resolve and propagate unsteadiness at the scale close to that of the grid itself. Typically, these methods are designed for structured grids and use large communication and discretization stencils. An alternative family of methods of recent interest are finite-element methods, such as the discontinuous Galerkin (DG), spectral element (SEM), and flux reconstruction (FR) methods. They are desirable for:

- The ability to use mixed-element unstructured grid elements.
- The ability to utilize a much coarser grid than that for an equivalent stencil-based reconstruction method with the same degrees of freedom.
- The ability of these methods to perform “P” scaling, which is to reduce the error of a method by increasing the polynomial order (P) of the solution.
 - For DG the order of accuracy for general nonlinear problems is $P+1$.
 - A “super-accurate” convergence on the order of $2P+1$ is observed theoretically and numerically for specific problems which are linear in nature and with an exact upwinding direction, e.g., acoustic propagation and density-wave advection.

There is a large body of work developing techniques for using these methods to analyze fluid dynamic flows [1-8]. They have been successfully developed and utilized for incompressible and subsonic compressible flows for both steady and unsteady analysis. Many challenges are present when extending these methods to supersonic flows because of shock and other discontinuity-induced instabilities. There has been much work focused on developing better techniques to stabilize these methods when applied to shock-dominated flows, and the development thereof is still an active area of research [9-18].




Figure 1. Diagram of the axisymmetric cavity combustor tested in the ACT-II facility. This figure is taken from Ref. 24. Lengths are given in millimeters.

The ACT-II arc-heated scramjet test facility [19-25] has been used to experiment with different ethylene-fueled HSABP geometries and configurations. It offers pressure tap measurements in addition to a large amount of optical access enabling chemiluminescence and planar laser-induced fluorescence (PLIF) measurements. The ACT-II facility

can reach high-speed flight enthalpy levels but only at relatively low stagnation pressures in the range of 1.0 to 3.0 atm. There is a body of recent and ongoing computational analysis of this facility [26-30]. These studies have investigated the transition from stable ram-mode to unstart, the operational characteristics of the combustor in scram- and ram-mode, the effects of oxygen addition, and the effects of thermal nonequilibrium caused by the arc-heating process.

The case of interest is an axisymmetric cavity combustor in free-jet configuration shown in figure 1. Specifically, this work is an investigation of the stable ram-mode operation of this geometry; including comparing it to previous experimental and computational results. Prior work in Ref. [30] performed an initial analysis of this configuration using the high-order FR method, examining the effect of changing the solution polynomial order on a single grid. Errors in the calculation of species transport properties were found in the code used for that work after it was published; they have been corrected for the present analysis. The inflow conditions for this case are Mach 4.5, a static pressure of 325 Pa, and a static temperature of 503 K, and it is fueled with ethylene to an equivalence ratio (ER) of 1.24. There are 16 circumferential injectors upstream of the cavity, each with a diameter of 0.75mm. Previous numerical studies of this configuration utilized a 6th-order hybrid central-difference/piecewise-parabolic method for discretization of the inviscid fluxes and used computational meshes containing 29 M and 55 M cells [26]. A preceding work applied a high-order FR method to discretize this case using a variety of polynomial orders which ranged from 3.4 M to 40 M solution points [30]. This work uses two FR discretizations. First, a P3 (cubic / 4th order accurate) polynomial on a coarser grid. Second, a P1 (linear / 2nd order accurate) polynomial solution on a finer grid. Both discretizations have approximately 20 million solution points. Additionally, an entropy limiting process is compared against the divergence-based filtering for a non-reacting flow case.

III. Numerical Approach

A. Governing Equations

The present work solves the Favre-averaged compressible Navier-Stokes equations extended to chemically reactive mixtures of thermally perfect gases by the addition of species continuity equations and reaction source terms. Species enthalpies are obtained using McBride/Chemkin-type polynomial curve fits. Subgrid viscous, diffusion, and heat conduction terms are modeled implicitly through the numerical dissipation associated with de-aliasing and solution stabilization techniques. The mixture molecular viscosity is calculated using Wilke's law [31]. The Prandtl and Schmidt numbers are both assumed to be 0.72. The chemical source terms are calculated explicitly using a finite-rate kinetics model explained in Section IV.A. Laminar chemistry is used with no turbulence-chemistry-interaction model.

B. The Flux Reconstruction (FR) Method

The code used presently is a modified version of the Glenn Flux Reconstruction (GFR) code [32-34] being developed at NASA Glenn Research Center. It is being used with the appropriate correction functions to recover a nodal DG discretization. The solution point locations are chosen to be Gauss quadrature for the increased accuracy of that choice over other options. The interfacial fluxes are evaluated at a separate set of flux points located on the element boundaries. Interfacial inviscid fluxes are determined using Edwards' Low Diffusion Flux Splitting Scheme (LDFSS) [35], while the viscous, diffusive, and conductive fluxes are handled via the Bassi-Rebay 2 (BR2) method [36]. Time advancement is done via explicit Runge-Kutta (RK) methods: the present investigation uses both a 3-stage/3rd-order RK3-SSP method [39] and an 11-stage/3rd-order RK method [40] depending on which is found to perform most efficiently for a given run case.

C. Stabilization Scheme I – A Divergence Based Filter

As with all high-order methods, it is necessary to use a shock-capturing scheme to stabilize the solution in the presence of shocks and other discontinuities. The present approach consists of two components: an indicator which identifies regions with shocks, and a dissipative filtering mechanism for stabilizing the computation. The method presented here is a combination of the de-aliasing filter present in GFR [33] along with concepts from the shock-capturing scheme implemented in NCSU's REACTMB solver [26].

First, for indicating troubled regions, a divergence-based indicator, $S_{divergence}$ (Eq. 1), is taken from the form of the Ducros switch used in REACTMB. The motivation behind this is that shocks will have a uniquely strong signal in the velocity divergence, and the magnitude of the divergence can provide a scale for how much dissipation to add for shocks of different strengths. Additionally, this indicator will not add excessive dissipation to contact discontinuities, such as those between species in a multi-component flow. The divergence portion is calculated as:

$$S_{divergence} = \left(\frac{2k|\nabla \cdot \vec{V}|}{a} \cdot \frac{h}{P+1} \right)^4 \quad (1)$$

where k is a tunable scalar. The present work uses $k = 0.75$ if the divergence is positive, and $k = 0.25$ if the divergence is negative.

This divergence indicator is then multiplied with a jump indicator [10], S_{jump} (Eqs. 2-4), which is a measure of how discontinuous the pressure field is along the boundary of an element based on the jumps in pressure across cell faces. The parameters used in this study are $\Psi_{min} = -3.0$, $\Psi_{max} = -1.0$.

$$M_{jump} = \log_{10} \left(\max \left(\frac{1}{||\delta\Omega||} \int_{\delta\Omega} \left| \frac{(p_1 - p_2)}{0.5(p_1 + p_2)} \right| ds \right), 2^{-23} \right) \quad (2)$$

$$S_{jump} = \begin{cases} 0 & M_{jump} < \Psi_{min} \\ \frac{1 + \sin \xi}{2} & \text{else} \\ 1 & M_{jump} > \Psi_{max} \end{cases}, \quad \xi = \frac{\pi (M_{jump} - 0.5(\Psi_{min} + \Psi_{max}))}{\Psi_{max} - \Psi_{min}} \quad (3,4)$$

$$S = S_{divergence} \cdot S_{jump} \quad \text{if } S < 0.02, \text{ set } S = 0.0 \quad (5)$$

The filtering mechanism (Eqs. (6-8)) used in GFR [32] is inspired by Hesthaven [1]. It provides a way of adding dissipation by means of damping high-order modes in the frequency space of the solution on an element, while leaving the cell average (P0 projection) alone. The specific modal damping coefficients, $\sigma(n)$, are determined according to an exponential formula, since aliasing errors and related instabilities tend to manifest in the higher frequency modes of a solution.

$$\sigma(n) = \exp[-\alpha\eta^c], \quad \text{for } n = 1, \dots, P+1 \quad \text{where } \eta = \frac{n-1}{P} \quad (6)$$

Here, c and α are parameters which control the filter's distribution and amount the highest mode is damped, respectively. The computations in this work use $c = 4$ and $\alpha = 16.0$.

A diagonal matrix, Σ , consisting of the modal damping coefficients is multiplied by the Vandermonde matrix, V , to create a filtering matrix, F , which applies the modal damping coefficients to the nodal solution when multiplied by the nodal solution vector, u .

$$F = V^{-1}\Sigma V, \quad \text{where } \Sigma = \delta_{i,j}\sigma(i) \quad \text{for } i, j = 1, \dots, P+1 \quad (7)$$

One modification present in the current work is a means of dynamically scaling the filter dynamically in-situ. To do so, we use the shock sensor (S in Eq. 5) to interpolate between the full-strength filter and the identity operator. The application of the filter on the solution is given as:

$$\tilde{u} = (S)Fu + (1 - S)Iu \quad (8)$$

This is mathematically equivalent to interpolating between the default values of the modal damping coefficients and unity. Finally, the indicator will only be used to filter elements if the value of S is above a threshold, in this case 0.02. This was found to greatly decrease the dissipation of the scheme by ignoring the lower intensity signal present turbulent regions of the flow.

D. Stabilization Scheme II – An Adaptive Entropy Limiter

An alternative stabilization mechanism is explored which utilizes a local minimum entropy principle on the compressible Euler equations to limit the solution. The present scheme is based on recent works which also utilize entropy-based principles [13-17]. First, we pick a numerical entropy: $s = p\rho^{-\gamma}$. This choice of entropy function was made because it is computationally cheaper than the physical entropy which includes the evaluation of a logarithm. Also, in combination with how it is used to define the damping parameter θ (Eq. 11), the entropy constraint is equivalent to Eq. 44 in Ref [15]. The minimum entropy on an element is determined by finding the minimum value of entropy from sampling both the Gauss -Legendre solution points and the face quadrature points for each cell of

interest. Then, the local minimum entropy limit is taken as the smallest minimum entropy value from a cell and its face neighbors.

The task of enforcing this minimum consists of ensuring that no solution point or flux point of a cell is in violation of the local minimum entropy after update the solution during a Runge-Kutta stage. Additionally, the Zhang-Shu positivity-preserving limiter [12] is applied to the solution prior to the entropy limiting steps. For each RK stage, the order of operations follows a similar operator splitting as in Refs. [16-17]:

- 1) Find the minimum entropy on each cell and its face neighbors.
- 2) Calculate the residuals associated with the inviscid fluxes.
- 3) Separately, compute and store the residual associated with the viscous flux and chemical source terms.
- 4) Update the solution with the inviscid fluxes.
- 5) Apply the positivity and entropy limiters.
- 6) Update the solution with the viscous fluxes and source terms.
- 7) Apply the positivity limiter.

There are a few ways to enforce this minimum. The most common technique was developed by Zhang and Shu [13] and consists of scaling the solution linearly towards its average. Recently, Trojak and Dzanic [17] presented a method which leverages exponential modal filters to accomplish the limiting by means of nonlinearly scaling the parameter α in Eq. 6. Our approach combines aspects from both methods. We use an exponential modal filter as an intermediate point in the path of scaling the solution towards the cell average. First, the solution is interpolated between the original solution and the filtered solution (Eq. 9, top branch). If that is insufficient to limit the cell, then it is interpolated, or scaled, from the filtered solution to the cell average (Eq. 9 bottom branch).

$$\hat{u}(\beta) = \begin{cases} \tilde{u} + \beta(u - \tilde{u}) & \text{if } \beta \leq 1 \\ (\beta - 1)\tilde{u} - (\beta - 2)\bar{u} & \text{if } \beta > 1 \end{cases} \quad (9)$$

Where u is the unlimited input solution on an element, $\bar{u}$ is the cell-average solution, $\tilde{u}$ is the modal-filtered solution, and $\hat{u}$ is the limited output solution. $\beta \in [0, 2]$ is a modified scaling parameter which determines where the limited solution falls in the interpolation between the original value, the filtered solution, and the cell average. It is derived from the standard scaling parameter θ which is the amount the solution must be scaled relative to its cell average to ensure all points of interest are bounded. The formula for computing β is:

$$\begin{aligned} \beta_{guess_{low}} &= \frac{(1 - \theta)(\bar{s} - s)}{(\tilde{s} - s)} \\ \beta_{guess_{high}} &= \frac{2(\tilde{s} - \bar{s}) - \theta(s - \bar{s})}{s - \tilde{s}} \end{aligned} \quad \beta = \begin{cases} \beta_{guess_{low}} & \text{if } \beta_{guess_{low}} \leq 1 \\ \beta_{guess_{high}} & \text{else} \end{cases} \quad (10)$$

where $\bar{s} = s(\bar{u})$ and $\tilde{s} = s(\tilde{u})$. The standard linear scaling parameter is then calculated in a way which makes the present formulation algebraically identical to the limiting parameter from Ref [15]. The θ parameter is found for any points violating the minimum entropy on a cell, and the smallest value from that set is used to limit the element.

$$\theta = \min_{\Omega} \frac{(s_i - s_{min})}{\left((s_i - s_{min}) - \left(\frac{\bar{\rho}}{\rho_i} \right)^{\gamma} (\bar{s} - s_{min}) \right)} \quad (11)$$

The advantage of this technique is that we retain the ability to limit without an iterative method, while also being able to bias the dissipation to favor the higher order modes in which the aliasing errors manifest more strongly.

E. Other De-Aliasing / Stabilization

The divergence filter alone is not sufficient to guarantee the boundedness and positivity of thermodynamic variables such as density, pressure, and temperature. As a way of maintaining a robust solution, we have a separate filtering operation which is triggered by boundedness violations after the divergence filter is applied. It checks all solution and flux points for violations and tags them as a troubled cell if one is found. Troubled cells are then by a filtering operation which is repeatedly applied to the solution until it is bounded, up to a maximum of 5 times. For the present study, the filtering mechanism used is the exponential modal filter with the same parameters as the divergence filter (Section III.c).

The bounds used in this work are found in table 1. The lower bounds are set just below the lowest value that is observed or expected in the solution. There are two upper bounds which have been imposed on the solution: Mach number and temperature. These are very seldomly activated but have been found to be necessary for the robustness of the method. The Mach-number boundedness limit helps to prevent issues in the region of strong expansion outside of the injectors; the high pressures and densities in this region, relative to the lower bounds of those quantities, made the standard boundedness filter less effective. The temperature upper bound helps to limit unphysical reaction rates caused by the interaction of aliasing errors with the chemical kinetics in the turbulent mixing regions. It is worth noting that the upper bounds do not necessarily prevent the given variables from exceeding the limiting value; simply that if the cell average solution breaks those thresholds, it will be reduced to 1st order accuracy.

Table 1. Bounds used to trigger the boundedness filtering.

	Temperature	Density	Pressure	Mach
Lower Bound	250.0 K	0.0023 kg/m ³	325.0 Pa	N/A
Upper Bound	4000.0 K	N/A	N/A	5.5

There are also areas of constant filtering, termed sponge filters, defined to provide extra dissipation to particularly troublesome regions of the flow. The included regions are the injectors, a small region outside of the injectors, and the last 20 mm of the geometry before the outflow plane. The region outside of the injectors is defined by extending 2 mm before and after the injector orifice and 1 mm off the wall. The sponge filter takes on the form of the de-aliasing filter in Eqs. 6-8 (with $S = 1$), and its parameters are: $c = 2, \alpha = 36.0$. The sponge filters are disabled when the entropy limiting filter is used, because the solution remains robust without them. Finally, as a mechanism for stabilizing the flow during startup, a weak global filter is used on every cell in the fluid domain with the parameters $c = 6$ and $\alpha = 0.1$.

IV. Cavity Combustor Case Setup

A. Chemical Mechanism

The chemical model used in the present study is a three-step, seven-species finite-rate mechanism (see table 2), tuned to match ignition delay in shock-tube experiments at the Air Force Research Laboratory [39]. This is significantly different than the previous computations of this geometry that used a 32-species skeletal mechanism. The reason for the change in chemical model is to balance the desire for representative chemistry with the computational cost of performing these simulations. The chemical source terms are being integrated explicitly, which is possible because of the restrictive Courant-Friedrichs-Lewy (CFL) conditions associated with explicit DG methods.

Table 2 Reactions and coefficients for the AFRL mechanism in Ref. [39]. Form: $AT^B e^{-\theta/T}$

Reaction	A	B	θ
$C_2H_4 + O_2 \leftrightarrow 2CO + 2H_2$	1.04E+13	0	21,511
$2CO + O_2 \leftrightarrow 2CO_2$	7.21E+07	1.92	16,478
$2H_2 + O_2 \leftrightarrow 2H_2O$	9.47E+16	-0.9	1.00

The majority of the results presented in this work were obtained using a modified backwards rate calculation; the standard method of calculating backward rates was unable to achieve a stable flowpath which is detailed in Section V.a. It was found that omitting the conversion of the equilibrium constant from the partial pressure-based constant, K_p , to the concentration-based K_c , resulted in a solution which is both stable and compares favorably with experimental wall pressures. The effect of this change is twofold. First, it biases the initial breakdown of ethylene into CO and H₂ heavily in favor of the products. This change can be favorable for this case because previous studies using the FFCM-30 mechanism [27] have shown that the presence of atomic oxygen from the arc heating process causes the fuel to break down and ignite much quicker than it would in a clean air flow. Second, the modification causes an overestimation of the backwards rate for the oxidation of CO and H₂. This reduces the stiffness of the mechanism and lowers the heat release. Given the highly simplified form of the global chemical mechanism, and the fact that it was fitted to data collected at significantly higher operating pressures than the present case, we suggest that this modification is acceptable to the extent that it enables the method to produce a stable solution to investigate the spatial discretization.

The equilibrium rate K_c is computed as:

$$K_c = \exp(\Delta g) \left(\frac{p_{atm}}{RT} \right)^{\Delta n} \quad (12)$$

Where Δg is the change in Gibbs free energy from the reactants to the products and Δn is the difference in the reaction order for the backwards rate minus the forwards. The modification which we use is to set $\Delta n = 0$, and thus ignore the effects of reactions producing greater or fewer moles of gas than they consume. Finally, the equilibrium rate is then used to compute the modified backwards rate, k_b , from the forwards rate, k_f , as:

$$k_b = \frac{k_f}{K_c} \quad (13)$$

While this is non-physical in the context of an equilibrium rate for an elementary reaction, we tentatively accept it as a modeling decision given the nature of the global mechanism and the improvement in the solution behavior seen for this case.

B. Coarse Grid

The primary grid used in this study is an unstructured voxel grid generated in Fidelity Pointwise that is shown in figure 2. It consists of a near-wall T-Rex layer, a freestream of cubic voxels, and an unstructured tetrahedral/pyramidal region connecting the two. The voxel grid format is chosen because it facilitates the notion of a freestream length scale which is most comparable to work done on multi-block structured meshes. The freestream element width is 1.8mm, and the height of the first cell from the wall is 0.08mm. In total there are approximately 600,000 elements in the mesh. The mesh is elevated to 4th order using the high-order mesh capabilities of Pointwise to accurately represent the curved surfaces of the combustor geometry. The average distance between solution points is used as an effective length scale for a mesh given a particular P order. It can be calculated as $h / (P+1)$ for a mesh with a grid spacing of h . For a P3, 4th order, solution; the effective freestream length scale is 0.45mm and a discretization containing 20.4 million total solution points.




Figure 2. Centerline views (top, bottom left) and an axial cut (bottom right) of the finite-element grid used in this study.

C. Fine Grid

As seen in figure 3, a second higher-resolution grid was generated for this study to enable running a P1 simulation with the same length scale and number of solution points as a P3 solution on the coarser mesh. Following that target, the freestream element width is 0.9 mm and the spacing of the first cell off the wall is 0.04 mm. There are




Figure 3. Slices of the fine grid which is used to run P1 solutions with a similar number of total solution points and freestream length scale to a P3 solution on the coarse mesh.

approximately 4 million elements in this grid resulting in 22 million solution points at polynomial order P1. Since this mesh is only being run with linear elements which cannot take advantage of surface curvature, the mesh was not elevated to high order. The computational expense of running this mesh at P1 has parity with that of the coarser mesh at P3 in terms of solution time per cpu hour. Each iteration takes approximately 2.5 times longer to calculate but the explicit timesteps are also larger by about the same factor.

D. Run Case and Boundary Conditions

The flow conditions used for this work match the stable ram-mode operation conditions studied in Refs. [26-30] with an ER of 1.24. The solution was initialized to the main inflow conditions everywhere except for the injectors and cavity region, which were both initialized to their stagnation conditions. The main and injector flow conditions are given in table 3 and correspond to case A from Ref. [26]. The modification of the inflow composition from that work is neglected, since the present chemistry does not include atomic oxygen or nitric oxide. The wall is treated as isothermal with a temperature of 375 K. This wall temperature was determined by examining the previous results which used a conjugate heat transfer model. The outflow is a supersonic outflow.

Table 3. Flow conditions specified at the inflow boundaries.

	Type	Mach	Pressure	Temperature	Y_{O_2}	Y_{N_2}	$Y_{C_2H_4}$
Main Inflow Conditions	Supersonic	4.5	351.9 Pa Static	502.9 K Static	0.233	0.767	0.0
Injector Conditions	Characteristic	0.01	54120 Pa Total	300.0 K Static	0.0	0.0	1.0

V. Results

A. Issues Using the 3-Step Mechanism

The first set of results presented here use the un-modified AFRL 3-step mechanism. This mechanism admits a very intermittent response from the combustor. The intermittency is characterized by periods of stable combustion interrupted by the upstream propagation of the flame causing strong flashback events. Contours displaying the progression of one such event are shown in figure 4. These events send a large pressure wave upstream from the cavity and lead to fuel burning around and upstream of the injectors. Subsequently, as the flame from the flashback dies down, it is nearly extinguished, relying on hot gases in the cavity recirculation region to relight the combustor. These events occur repeatedly until the flame is unable to reignite, and the combustor blows out entirely. Subsequent results use the modified equilibrium rate computation as detailed in Section IV.a.




Figure 4 Snapshots of a flashback event seen when using the un-modified chemical mechanism. The period shown is from the pseudo-steady operation, through the propagation of flames upstream of the cavity and progressing to an extinguished cavity.

B. Baseline results

As the baseline run condition a P3 (4th order) solution is run on the coarse mesh (1.8mm spacing) using the divergence-based filter. This setup is chosen based off the polynomial order study conducted on this mesh in Ref. [30] which demonstrated that the P3 solution was a good tradeoff between computational cost and scale-resolving ability. The procedure used to start the solution was to: initialize a non-reacting flow through the combustor, then enable the chemical source terms with an ignition aid, and finally turn off the ignition aid once a thermal throat has formed, and the combustor is operating firmly in ram-mode. The ignition aid is a limiter in the finite-rate calculation which scales the species concentrations to mimic a pressure of no less than 10 kPa and evaluates the reaction rates at a temperature of no less than 2,000 Kelvin. The solution is then run for a total of 7 ms with the last 4 ms used to create a time-average of the solution.




Figure 5. Comparison of the inlet and isolator flow features between the 4th order GFR solution and a 6th order REACTMB finite-volume solution [26].

Co-plots of representative instantaneous solutions from the present work and a previous study using NCSU's REACTMB high-order finite-volume solver [26] are shown in figure 5. Beginning at the inlet on the left side of the graphics, these solutions show improved capturing of the inlet shocks compared to Ref [30] because of the improved filtering coefficients. Downstream, the stabilization scheme admits decent resolution of the shock-boundary layer interaction (SBLI) at the start of the isolator and does not inhibit the transition to turbulence as the flow encounters the adverse pressure gradient emanating from the combustor. These results show a broad qualitative agreement with the REACTMB solution. The shocks themselves are better resolved in the finite-volume scheme; this is attributed to both the challenge of minimally filtering a finite-element solution so that it may represent a step function without excessive numerical oscillations and to the slightly coarser discretization in the present study: an effective length scale of $h / (P + 1) = 0.45$ mm, compared to finite-volume cells of width 0.3 mm.




Figure 6. Midplane contour plots of an instantaneous snapshot of the baseline solution (P3 solution on the coarse mesh using the divergence filter).

Moving further downstream, plots from a representative instantaneous snapshot are shown in figure 6. The flowpath is operating in a ram-mode with the flow going subsonic just after the injectors. The combustor is in a cavity-wake stabilized mode of operation, with most of the heat release occurring either in the cavity shear layer, or in the core flow in the latter half of the cavity. This agrees with the clean-air solutions obtained in Ref. [24] Case A. The turbulent structures are visualized in the transverse, or out-of-plane, velocity contour at the bottom of figure 6. The temperature contours in combination with the out-of-plane velocity contours allow us to ascertain that the scale of the structures being resolved is very close to the resolution of the solution points within an element. This implies that the stabilization methods are successful at admitting a low dissipation, or equivalently, high-bandwidth, solution in the powerhead region.

In the time-average plots presented in figure 7, we can clearly see the cavity stabilized flame structure which is similar to case A presented in Ref. [26]. Additionally, the location of the thermal throat between the injectors and the cavity, denoted by the white sonic line, extends slightly farther downstream than in that previous analysis. We attribute this to the change in the chemical mechanism compared to previous studies. The previous work included atomic oxygen in the inflow to account for dissociation from the arc-heating process. This significantly reduced the ignition delay time and anchors the heat release to the injectors creating a jet-wake stabilized flame. Thus, there is an earlier pressure rise at the injectors and results in a throat which is slightly farther upstream. This change in ignition delay

time also affects the mode of operation of the combustor and could be a reason why the present analysis results in a cavity-wake stabilized instead of a jet-wake stabilized flame.




Figure 7. Time-averaged views of the baseline solution. The sonic line is marked in white on the bottom most contour plot.

Figure 8 shows a comparison between the experimental wall pressure measurement and the time-average P3 solution. There is a good agreement in the shock-train location, but a deviation in the cavity peak pressure which results in a slight over-prediction of the outflow pressure. There is also a deficit in the pressure upstream of the injectors at an axial coordinate of approximately 250mm. These discrepancies can at least partially be attributed to the difference between the cavity-wake stabilized flame presently observed and the jet-wake stabilized flame which is believed to occur physically.




Figure 8. Comparison between the experimental wall pressure readings and the time-averaged solution obtained with GFR.

C. Fine Grid P1 Results

The linear, P1, solution on the finer grid is initialized in the same manner as the baseline solution; non-reacting flow initialization, source terms enabled with an ignition aid, ignition aid disabled once ram-mode is established. This simulation was run to compare the effects of changing the polynomial order of a solution independent of the effective length scale of the solution. For the implementation in GFR the P1 solution on the finer mesh also has the same computational cost for a given advancement of solution time as the baseline P3 analysis.




Figure 9. Progression of the P1 solution on the fine grid from the cutoff of the ignition aid to combustor blowoff.

When attempting to ignite this flowpath, we find that the cavity is unable to sustain a flame. The ignition aid was run for 2.2 ms before it was disabled. After that, the flame falls back from the injectors to the cavity and eventually blows out entirely. Figure 9 shows the progression of the solution from the end of the ignition process through the flame extinguishing. While there is no steady-state solution to compare to, there are some key differences between the P1 and P3 solutions. A comparison of instantaneous snapshots is presented in figure. 10. The snapshot for the P1 solution was taken 0.1 ms after the ignition aid was cutoff since that is when the flame is in the most comparable configuration to the P3 solution.

The P1 solution shows a distinct increase in the scale of the turbulent structures, particularly in the region downstream of the injectors and in the cavity shear layer. The lack of smaller scale turbulent structures suppresses the mixing of the fuel and oxidizer and is likely to contribute to less complete and/or intense combustion. Additionally, there is colder flow in the recirculating corner of the cavity. This can be an effect of both the reduced combustion intensity and of a higher level of heat flux into the isothermal walls because of the lower order discretization. Both the reduced mixing efficiency and the cold cavity are plausible reasons why the P1 solution is unable to hold a flame.




Figure 10. Comparison of the P1 (top) and P3 (bottom) solutions. Both solutions are set up to have approximately the same number of solution points and the same distance between solution points such as to isolate the effects of polynomial order. Both simulations have approximately the same computational cost, with the P1 solution able to take larger timesteps, but with each timestep being more expensive to compute.

Conversely, the resolution of the shocks and the SBLI in the inlet and front part of the isolator is significantly improved compared to the higher order solution. The solution resolves those flow features more finely and shows less numerical artificing where the filter is being applied. Additionally, the boundary layer on the compression surface of the inlet is slightly smaller for the P1 solution, which implies that there is less filtering needed to capture the leading edge of the shocks at the inflow plane. This behavior is the reverse of what we observe with the turbulent flow structures. The dissipation from the shock capturing scheme is applied at the element-wide level. Since the P1 solution uses finer elements to keep the solution point spacing constant, the filter applies less dissipation overall to the region. The P3 solution is attempting to use a cubic polynomial to approximate the step function of a shockwave, and using a filter to ensure monotonicity and positivity of the solution will cause it to get smeared across the larger P3 elements. Without having a method to resolve the shock at a sub-cell level, applying a cell-based shock-capturing scheme to the larger elements used for high-order solutions will almost certainly contribute to a greater degree of dissipation to the solution as a whole.

D. Entropy Filtering Method

The entropy-based limiting method is presently investigated for the nonreacting flow configuration. Both the divergence filter and the entropy limiter are run on the coarse grid with a P3 solution and the fine grid with a P1 solution. All solutions were run to a statistically steady state.

Figure 11 compares the stabilization methods for the coarse P3 grid. Beginning with the inlet shock structures, any numerical method which uses a dissipative mechanism to capture shocks will create an enlarged boundary layer at an attached shock because the dissipation used to capture shocks will also affect the momentum / temperature flux from the wall. This layer is visible here on the compression surface of the inlet. Comparing the width of that layer between the two filtering methods in figure 11, the entropy limiting procedure is a less dissipative representation of the attached shock at the inflow plane. Moving slightly downstream, both the expansion corner at the start of the isolator and subsequent SBLI are resolved with less numerical dissipation. Thus, the entropy limiter results in a smaller isolator boundary layer. The start of the SBLI separation shock is also clearer. Conversely, the divergence filter admits smoother representations of the weaker shocks, most visibly for the SBLI reattachment shock.




Figure 11. Comparison between the divergence filter and the entropy limiter for the coarse mesh with P3 solutions.

The next point for comparison is the cavity mixing region, which is where we see the development of turbulent structures for the non-reacting case. The divergence sensor is activated by the expansion of the fuel from the injectors, so it applies excess dissipation which delays the onset of unsteady features in the flow. This can be observed in figure 11 by comparing the out of plane velocity contours from the two methods. The entropy limiter admits both a quicker turbulence transition and more intense unsteadiness where the shear layer impinges with the cavity closeout ramp. These observations imply that the solver is operating with a lower degree of numerical dissipation using the entropy limiter than the divergence filter.




Figure 12. Comparison between the stabilization techniques for a P1 solution on the fine mesh.

The same mixing test was done for the finer mesh with a P1 solution, shown in figure 12. For this case, there are much fewer differences between the action of the two techniques, with the entropy limiting showing some very slight advantages in resolving unsteadiness in the cavity shear layer. This suggests that the 2nd order P1 solution is less sensitive to the choice of stabilization technique than the higher order solutions.

VI. Conclusion

The present work has shown the FR-DG method equipped with a divergence-based filter and boundedness limiter is able to conduct high-order scale-resolving analysis of a dual-mode scramjet combustor which compares well with experimental results. The largest discrepancies between the present results and the experiment can be explained by the simplified modeling of chemical source terms leading to a differing flame structure, and consequentially, the heat release and pressure distribution in the cavity. A secondary fine grid was made to isolate the effects of the polynomial order of accuracy on the solutions obtained. This grid was designed to run at the same cost and length scale as the original coarse grid but using a P1 (2nd order) solution instead of a P3 (4th order) reconstruction. This analysis revealed an important tradeoff between low-order and high-order solutions with equivalent solution point length scales when using cell-based shock-capturing schemes for flows with shocks: lower polynomial orders can capture shocks more cleanly because of finer mesh resolutions and are less sensitive to the filtering method; however, they strongly dissipate small-scale turbulent flow structures needed for accurate large eddy simulations. In this case, the additional dissipation is enough to prevent the flame stabilization; resulting in a combustor blowout after the ignition aid is disabled. Finally, an alternative filtering method is explored which utilizes a local minimum entropy principal to limit the solution. The entropy-based method shows promise as a way of admitting a solution with lower dissipation and requiring no problem-dependent user inputs. Future work will focus on applying the entropy limiting to the chemically reacting flow and determining the sensitivity of the solution to degrees of freedom afforded by the method (e.g., choice of entropy functional, interpolated filter).

Acknowledgments

We acknowledge the computational resources provided by the Department of War's High Performance Computing Modernization Program. Additionally, resources supporting this work were provided by the NASA High-End Computing (HEC) Program through the NASA Advanced Supercomputing (NAS) Division at Ames Research Center.

We acknowledge the use of the GFR software, provided by the Inlets and Nozzles Branch at NASA Glenn Research Center.

References

- [1] Hesthaven, J. S. and Warburton, T., *Nodal Discontinuous Galerkin Methods: Algorithms, Analysis, and Applications*, Vol. 54 of *Texts in Applied Mathematics*, Springer New York, 2008.
- [2] Bassi, F. and Rebay, S., "A High Order Discontinuous Galerkin Method for Compressible Turbulent Flows," *Discontinuous Galerkin Methods: Theory, Computation, and Application*, edited by B. Cockburn, G. E. Karniadakis, and C.-W. Shu, *Lecture Notes in Computational Science and Engineering*, Springer, 2000, pp. 77–88.
- [3] Cockburn, B., Shu, C.-W., "The local discontinuous Galerkin method for time-dependent convection-diffusion systems," *SIAM J. Numer. Anal.* 35 (6) (1998) 2440–2463
- [4] Cockburn, B., "Discontinuous Galerkin Methods for Computational Fluid Dynamics," *Encyclopedia of Computational Mechanics*, edited by E. Stein, R. Borst, and T. J. R. Hughes, Wiley, 2004. <https://doi.org/10.1002/0470091355.ecm053>
- [5] Huynh, H., "A Flux Reconstruction Approach to High-Order Schemes Including Discontinuous Galerkin Methods," *AIAA Paper* 2007-4079, Jun 2007.
- [6] Huynh, H., "A Reconstruction Approach to High-Order Schemes Including Discontinuous Galerkin for Diffusion," *AIAA Paper* 2009-403, Jan 2009.
- [7] Huynh, H., "High-Order Methods Including Discontinuous Galerkin by Reconstructions on Triangular Meshes," *AIAA Paper* 2011-44, Jan 2011.
- [8] Hoskin, D.S., Van Heyningen, R.L., Nguyen, N.C., Vila-Pérez, J., Harris, W.L., Peraire, J., *Discontinuous Galerkin Methods for Hypersonic Flows. Progress in Aerospace Sciences*, vol. 146, Apr. 2024
- [9] D'Aferio, F. M., Mihaescu, M., and Hanifi, A., "Very-Low-Dissipation, Physics-Based, and Parameter-Free Shock Capturing for Discontinuous Galerkin Methods," *Physics of Fluids*, Vol. 36, No. 12, 2024, p. 126115. <https://doi.org/10.1063/5.0241108>
- [10] Barter, G. E. (Garrett E.), "Shock Capturing with PDE-Based Artificial Viscosity for an Adaptive, Higher-Order Discontinuous Galerkin Finite Element Method," Thesis. Massachusetts Institute of Technology, 2008.
- [11] Moro, D., Nguyen, N. C., and Peraire, J., "Dilation-based Shock Capturing for High-order Methods," *International Journal for Numerical Methods in Fluids*, Vol. 82, No. 7, 2016, pp. 398–416. <https://doi.org/10.1002/fld.4223>
- [12] Zhang, X., and Shu, C.-W., "On Positivity-Preserving High Order Discontinuous Galerkin Schemes for Compressible Euler Equations on Rectangular Meshes," *Journal of Computational Physics*, Vol. 229, No. 23, 2010, pp. 8918–8934. <https://doi.org/10.1016/j.jcp.2010.08.016>

- [13] Zhang, X., and Shu, C.-W., “A Minimum Entropy Principle of High Order Schemes for Gas Dynamics Equations,” *Numerische Mathematik*, Vol. 121, No. 3, 2012, pp. 545–563. <https://doi.org/10.1007/s00211-011-0443-7>
- [14] Wu, K., Zhang, X., and Shu, C.-W., “High Order Numerical Methods Preserving Invariant Domain for Hyperbolic and Related Systems,” <https://doi.org/10.48550/ARXIV.2512.09116>
- [15] Lv, Y., and Ihme, M., “Entropy-Bounded Discontinuous Galerkin Scheme for Euler Equations,” *Journal of Computational Physics*, Vol. 295, 2015, pp. 715–739. <https://doi.org/10.1016/j.jcp.2015.04.026>
- [16] Ching, E. J., Johnson, R. F., Burrows, S., Higgs, J. P., and Kercher, A. D., “Positivity-Preserving and Entropy-Bounded Discontinuous Galerkin Method for the Chemically Reacting, Compressible Navier-Stokes Equations,” 2023. <https://doi.org/10.2514/6.2023-0664>
- [17] Trojak, W., and Dzanic, T., “Positivity-Preserving Discontinuous Spectral Element Methods for Compressible Multi-Species Flows,” *Computers & Fluids*, Vol. 280, 2024, p. 106343. <https://doi.org/10.1016/j.compfluid.2024.106343>
- [18] Huynh, H. T., “Shock Capturing via Limiting for High-Order Methods Including Discontinuous Galerkin and Flux Reconstruction,” *Computers & Fluids*, Vol. 300, 2025, p. 106743. <https://doi.org/10.1016/j.compfluid.2025.106743>
- [19] Smith, S., Gruber, M., Steiner, R., Collatz, M., and Mathur, T., “Development and Calibration of an Axisymmetric Direct-Connect Supersonic Combustion Flowpath” AIAA Paper 2009-5035, Aug. 2009. <https://doi.org/10.2514/6.2009-5035>
- [20] Gruber, M. R., Smith, S., and Mathur, T., “Experimental Characterization of Hydrocarbon-Fueled, Axisymmetric, Scramjet-Combustor Flowpaths,” AIAA Paper 2011-2311, June 2012. <https://doi.org/10.2514/6.2011-2311>
- [21] Landsberg, W., Vanyai, T., McIntyre, T., and Veeraragavan, A., “Experimental Scramjet Combustion Modes of Hydrocarbon Mixtures at Mach 8 Flight Conditions,” *AIAA Journal*, Vol. 58, No. 12, 2020, pp. 5117–5122. <https://doi.org/10.2514/1.J059856>
- [22] Baccarella, D., Liu, Q., Lee, T., Hammack, S., and Do, H., “The Supersonic Combustion Facility ACT-2,” AIAA Paper 2017-0103, Jan. 2017. <https://doi.org/10.2514/6.2017-0103>
- [23] Baccarella, D., Liu, Q., McGann, B., Lee, T., and Do, H., “Experimental Study of Heat Release Induced Choking in a Supersonic Circular Combustor,” AIAA Paper 2018-1139, Jan. 2018. <https://doi.org/10.2514/6.2018-1139>
- [24] Liu, Q., Baccarella, D., McGann, B., and Lee, T., “Dual-Mode Operation and Transition in Axisymmetric Scramjets,” *AIAA Journal*, Vol. 57, No. 11, 2019, pp. 4764–4777. <https://doi.org/10.2514/1.J058391>
- [25] Liu, Q., Baccarella, D., McGann, B., and Lee, T., “Cavity-Enhanced Combustion Stability in an Axisymmetric Scramjet Model,” *AIAA Journal*, Vol. 57, No. 9, 2019, pp. 3898–3909. <https://doi.org/10.2514/1.J058204>
- [26] Hash, C., Drummond, P., Edwards, J., Kato, N., Liu, Q., and Lee, T., “Numerical Simulation of Stable and Unstable Ram-Mode Operation of an Axisymmetric Ethylene-Fueled Inlet-Isolator-Combustor Configuration,” *Combustion and Flame*, Vol. 242, No. 8, 2022, Paper 112157. <https://doi.org/10.1016/j.combustflame.2022.112157>
- [27] Hash, C., Drummond, P., Edwards, J., Kato, N., Liu, Q., and Lee, T., “Numerical Simulation of Unstart Processes in an Axisymmetric Scramjet Combustor,” AIAA Paper 2022-1849, Jan. 2022. <https://doi.org/10.2514/6.2022-1849>
- [28] Richardson, S., Edwards, J. R., Hash, C. A., Drummond, P. M., Lee, T., and Kato, N., “Numerical Simulation of Scram-Mode Operation of an Axisymmetric Combustor in an Arc Tunnel,” *Journal of Propulsion and Power*, Vol. 41, No. 4, 2025, pp. 429–439. <https://doi.org/10.2514/1.B39754>
- [29] Richardson, S., Edwards, J.R., Gessman, I., Lim, J., Paganini, A., and Lee, T., “Numerical Simulation of Oxygen-Enhanced Combustion in an Arc-Heated Scramjet Facility,” *AIAA Journal*, vol. 0, no. 0, Jan. 2026. <https://doi.org/10.2514/1.J066099>
- [30] Koeplinger, T.S., Hash, C., Edwards, J.R., Spiegel, S., “High-Order Flux Reconstruction for the Implicit Large Eddy Simulation of an Axisymmetric Scramjet Combustor”, AIAA Paper 2026-1754, Jan. 2026. <https://doi.org/10.2514/6.2026-1754>
- [31] Wilke, C. R., “A Viscosity Equation for Gas Mixtures,” *Journal of Chemical Physics*, Vol. 18, No. 4, 1950, pp. 517–519. <https://doi.org/10.1063/1.1747673>
- [32] Spiegel, S. C., DeBonis, J. R., and Huynh, H. T., “Overview of the NASA Glenn Flux Reconstruction Based High-Order Unstructured Grid Code,” San Diego, California, USA, 2016. <https://doi.org/10.2514/6.2016-1061>
- [33] Spiegel, S. C., Borghi, M. R., and Yoder, D. A., “Large Eddy Simulations of a Single-Injector Cooling Flow Using the High-Order Flux Reconstruction Method,” San Diego, CA & Virtual, 2022. <https://doi.org/10.2514/6.2022-1813>
- [34] Spiegel, S. C., Yoder, D. A., DeBonis, J. R., Huynh, H. T., Heinlein, G., Borghi, M. R., and Georgiadis, N. J., “New Capabilities and Improvements to the High-Order Glenn Flux Reconstruction Code,” Orlando, FL, 2025. <https://doi.org/10.2514/6.2025-0061>
- [35] Edwards, J. R., “A Low-Diffusion Flux-Splitting Scheme for Navier-Stokes Calculations,” *Computers & Fluids*, Vol. 26, No. 6, 1997, pp. 635–659. [https://doi.org/10.1016/S0045-7930\(97\)00014-5](https://doi.org/10.1016/S0045-7930(97)00014-5)
- [36] Bassi, F. and Rebay, S., “A High Order Discontinuous Galerkin Method for Compressible Turbulent Flows,” *Discontinuous Galerkin Methods: Theory, Computation, and Application*, edited by B. Cockburn, G. E. Karniadakis, and C.-W. Shu, Lecture Notes in Computational Science and Engineering, Springer, 2000, pp. 77–88
- [37] Gottlieb, S. and Shu, C.-W., “Total Variation Diminishing Runge-Kutta Schemes,” *Mathematics of Computation*, Vol. 67, No. 221, 1998, pp. 73–85. <https://doi.org/10.1090/S0025-5718-98-00913-2>
- [38] Al Jahdali, R., Dalcin, L., Boukharfane, R., Nolasco, I. R., Keyes, D. E., and Parsani, M., “Optimized Explicit Runge–Kutta Schemes for High-Order Collocated Discontinuous Galerkin Methods for Compressible Fluid Dynamics,” *Computers & Mathematics with Applications*, Vol. 118, 2022, pp. 1–17. <https://doi.org/10.1016/j.camwa.2022.05.006>
- [39] Hassan, E. A., Hageman, M. D., Knadler, M. S., and Lucas, S., “Error In Shock Tube Ignition Delay Time Predictions Due to Bifurcation and Boundary Layer Effects,” Orlando, FL, 2025. <https://doi.org/10.2514/6.2025-2140>