AIAA SciTech Forum
12-16 January 2026, Orlando, FL
AIAA SCITECH 2026 Forum

10.2514/6.2026-2510

An Integrated CFD, Guidance, and Control Approach with Heating Considerations

Darryl J. Williams*, Diganta Bhattacharjee †, Maziar S. Hemati ‡, and Graham V. Candler §

This work presents an integrated framework for simulating controlled trajectories of hypersonic glide vehicles (HGVs) that couples high-fidelity computational fluid dynamics (CFD) with conjugate heat transfer (CHT) and integrated guidance and control (IG&C). By using high-fidelity CFD as feedback within closed-loop controls, this framework enables the generation of optimized trajectories under coupled aerothermodynamic effects. A loosely coupled one-dimensional CHT solver captures transient wall heating along the vehicle surface, while aerodynamic forces and moments are obtained from CFD and augmented by surrogate models for gradient and trim computations. The study demonstrates that integrating CFD in-the-loop with controls reveals discrepancies relative to traditional database-based methods, affecting controller stability, trim, and trajectory feasibility. Dynamic pressure tracking studies show that transient heating can meaningfully alter closed-loop behavior. The time-varying model predictive control, based on a linearized model of the HGV dynamics, requires a sampling time of 10 ms. To reduce the number of CFD evaluations needed at this sampling frequency, appropriate approximations of aerodynamic forces, moments, and heating are used during trimmed flight. The results highlight the importance of consistent physics modeling across all components of the IG&C framework and provide a foundation for future efforts: incorporate aerothermal constraints into guidance, design closed-loop controllers robust to model mismatch/uncertainties, and quantify unsteady effects in controlled hypersonic flight simulations. This framework represents a step toward leveraging high-fidelity, coupled physics-based solvers for multidisciplinary optimization and design of hypersonic systems.

I. Introduction

With advances in high-performance computing, physics-based modeling, and numerical methods, integrated multiphysics approaches to hypersonic design and optimization have become increasingly feasible [1, 2]. However, these methods often depend on empirical or sparsely sampled surrogate models to approximate the underlying physics and vehicle characteristics.

A major constraint on the operability of hypersonic systems arises from the conversion of freestream kinetic energy into viscous heating at the vehicle surface. At hypersonic speeds, strong bow shocks form, driving chemical reactions and the excitation of internal energy modes in the air. In lower-density environments, such as during atmospheric reentry, chemical nonequilibrium effects become significant. These characteristics are unique to hypersonic flows and must be captured to accurately model aerodynamic and aerothermal behavior. The inclusion of high-temperature effects and nonequilibrium chemistry strongly influences both aerodynamic performance and thermal loading, making their representation critical for reliable vehicle design and control.

For the exemplar case of a hypersonic glide vehicle (HGV) considered in this study, the vehicle operates across a wide range of freestream conditions over several minutes of flight. As a result, heating along the vehicle body and within the thermal protection system (TPS) is highly trajectory dependent. Accurate prediction of heating and its impact on the vehicle requires an integrated, multiphysics-informed approach that accounts for the relevant flow-field physics, vehicle controllability along the trajectory, and the TPS material response.

Many existing approaches to optimal trajectory generation for boost-glide vehicles overlook the coupled nature of the physics and the vehicle's controllability [3–5]. Typically, trajectories are generated in an open-loop fashion, where surrogate models for vehicle dynamics and thermal response are treated as exact. Such formulations rely entirely on the

*Graduate Research Assistant, Department of Aerospace and Engineering Mechanics
†Postdoctoral Associate, Department of Aerospace and Engineering Mechanics
‡Associate Professor, Department of Aerospace and Engineering Mechanics, AIAA Associate Fellow.
§McKnight Presidential Endowed Professor, Department of Aerospace and Engineering Mechanics, AIAA Fellow.

surrogate's fidelity in capturing vehicle dynamics and aerothermal behavior, leading to errors that inevitably accumulate over time.

This can result in suboptimal or infeasible trajectories for several reasons:

1) The surrogate model may not sufficiently interpolate across the full operational envelope, introducing integration error.
2) The surrogate may omit relevant physics, introducing uncertainty.
3) The problem formulation may neglect control surfaces and vehicle controllability.

Reasons (1) and (2) motivate the introduction of feedback to correct deviations when model predictions diverge from reality, while reason (3) has direct implications for trajectory feasibility. A common issue in the literature is the conflation of state variables with control variables. For example, Rataczak et al. [6] treat bank angle as a control input for trajectory optimization, while Liu et al. [7] use the rates of angle of attack and bank angle as control inputs. Although this relaxes the assumption of instantaneous angle changes made in Rataczak et al. [6], it still yields unrealistic inputs, because the true control authority of a boost-glide vehicle is governed by control-surface deflections, not directly by aerodynamic angles.

This distinction is critical because physical phenomena such as flow separation can occur during hypersonic flight, introducing nonlinearities in both control authority and vehicle stability. These effects are compounded by an inherent trade-off: larger control inputs reduce aerodynamic efficiency [8, 9]. Without accounting for these control-surface effects, the generated trajectories are suboptimal and may not be feasible. Thus, including control-surface dynamics and aerothermodynamic effects is essential for assessing trajectory feasibility and stability.

This work builds upon the prior study of Williams et al. [10], which introduced an integrated guidance and control (IG&C) framework for a generic hypersonic glide vehicle. That study demonstrated the use of model predictive control (MPC) to execute constrained flight trajectories, such as maintaining constant dynamic pressure while preserving vehicle trim. An aerodynamic database derived from computational fluid dynamics (CFD) was employed; however, aerothermodynamic constraints were not included in the control formulation. Additionally, a radiative equilibrium assumption was used to estimate surface wall temperatures, neglecting the material response of the TPS. Finally, while control-surface effects were incorporated into the aerodynamic model, feedback from the true dynamical system was not provided.

The present work extends that framework by incorporating considerations for aerothermodynamic constraints within the IG&C formulation. Additionally, CFD is used in the loop with the IG&C framework to provide real-time, high-fidelity feedback for the optimal control problem. This coupling enables multiphysics integration directly within the control architecture. For the present demonstration, a one-dimensional, loosely coupled conjugate heat transfer (CHT) model is used to account for transient wall-heating effects along the flight trajectories. In a loosely coupled approach, the flow field is first converged to steady state, and the resulting convective heat flux is used to simulate the thermal response accurately. This has been shown to be a reasonable assumption given the relative time scales of the fluid and solid dynamics. Additional augmentation is introduced based on the non-dimensional heat transfer coefficient, as described in later sections. While this study focuses primarily on coupling between CFD and the thermal response, the framework can be extended to include additional physics such as non-rigid body dynamics or material ablation in future work.

The remainder of this paper is organized as follows: Section II outlines the IG&C formulation and the CFD–CHT coupling strategy; Section III presents simulation results and trajectory comparisons; and Section IV summarizes the findings and discusses future work.

II. Methodology

A. Geometry and Meshing

The mesh generation tools and US3D modules used in this work are outlined in the following section. The geometry selected is a generic HGV with control surfaces. The geometry was created using Crosslink's Python APIs, and the fluid grids were generated in Link3D.

1. CrossLink

The CrossLink software was initially developed at Los Alamos National Laboratory to support the NNSA Stockpile Stewardship Program. It is currently being developed and extended at the University of Minnesota's Computational

Hypersonic Research Laboratory for rapid geometry and mesh generation for hypersonic applications [11]. It features Python APIs for geometry creation and supports mesh refinement to handle complex geometries with significant variation, including cases relevant to this manuscript that involve control surface deflections.

2. Link3D

Grid generation was automated using LINK3D, a parallel topology based mesh generation software developed by GoHypersonic, Inc. [12]. In a topology based grid generation approach, the multi block topology is first sketched to fill the volume, and the exposed block boundaries are assigned to geometry groups. The LINK3D grid smoothing engine then refines the grid while ensuring that the boundaries remain aligned with the assigned geometry. Once a baseline or reference topology is created, it can be applied to any geometry within the parametric family for which it was designed. This topology based method allows for the automatic generation of high quality hexahedral grids that accommodate changes to the mesh surface. The geometry and its mesh are shown in Figure 1.

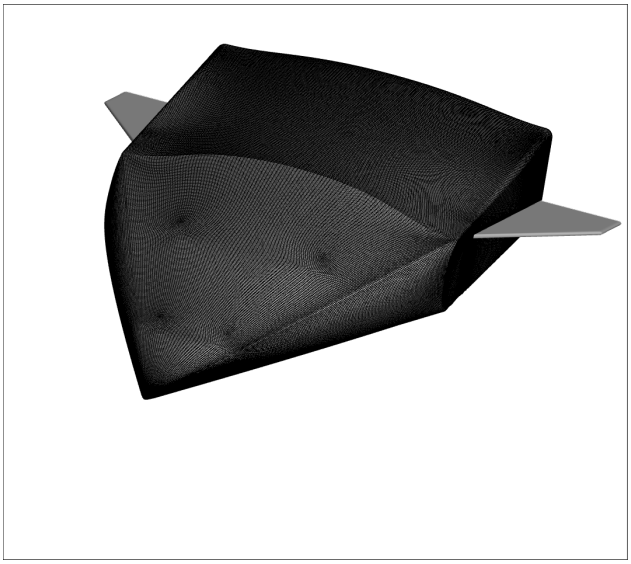

(a) +15° control surface deflection of the full geometry. The surface mesh of the control surfaces are omitted.

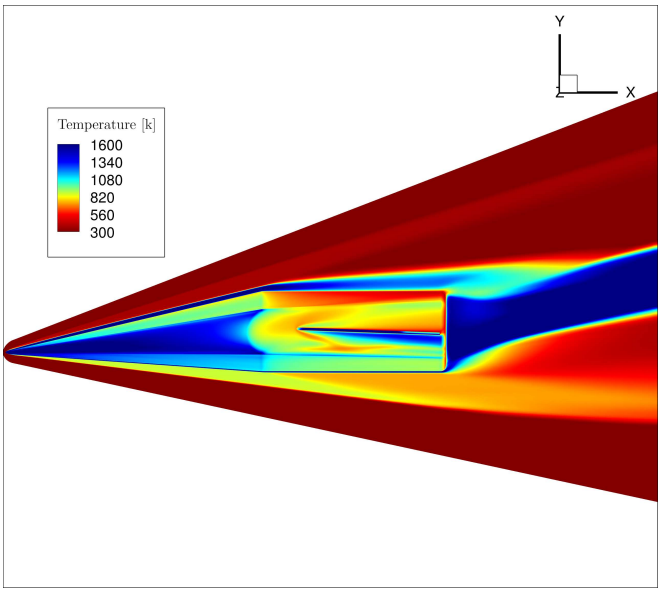

(b) Surface Wall Temperature

Fig. 1 Mesh and geometry of the generic BGV used in this study. The geometry was created using CrossLink's Python API, and the grid was generated using Link3D.

3. US3D Multi-mesh Plugin

In scenarios where geometries, flow conditions, and boundary conditions evolve dynamically, efficiently generating computational grids is critical. Static grids are often inadequate for problems that require high accuracy in regions with rapid geometric changes or across parameter spaces that vary widely. In addition, creating static grids demands significant effort and domain expertise, which can redirect resources from data generation toward grid construction. The Multi-Mesh tool addresses these challenges by enabling adaptive grid interpolation and management across parameters such as Mach number, angle of attack, and rotational motion.

The Multi-Mesh plugin is a framework designed to streamline simulations involving geometries with well defined parametric descriptions. It supports multi parameter mesh transformations, allowing more than one parameter to define the geometry. New grids are generated efficiently using Non Uniform Rational B Splines (NURBS) for interpolation. With a Fortran based backend and Python interfaces, the tool provides a robust and scalable platform for applications such as aerodynamics and fluid structure interactions. Key use cases include trajectory simulations and, as highlighted in this manuscript, the rapid creation of aerodynamic databases.

The Multi-Mesh plugin integrates seamlessly with the US3D flow solver. It employs NURBS interpolation for dynamic geometric transformations and features a modular design for compatibility with Python workflows or compiled

Fortran plugins.

NURBS based interpolation is used to compute control point meshes from the input anchor meshes. NURBS interpolation uses weights defined by the *knot vector*, which specifies the non uniform spacing of the meshes in the parametric space, as well as the degree of interpolation. For each parameter, interpolation is performed by constructing a linear system of basis function weights to compute control points orthogonal to that parameter. This process is repeated for each parameter, with subsequent parameters interpolating through the control point result of the previous parameter. The tensor product of the interpolations across all parameters yields the final result and is a symmetric operation, meaning the result is unique and independent of the parameter ordering.

B. Flow Solver

To evaluate the aerodynamic forces and moments outlined in Equation (8), we generate CFD flow fields using US3D. US3D [13] is an unstructured finite volume code that approximately solves the weak form of the compressible Navier–Stokes equations. In this work, for simplicity, we assume that the fluid is a perfect gas and that the flow is laminar. This choice reduces the computational cost and is justified because the integrated forces are largely insensitive to the chemical assumptions of the flow field and the perfect gas assumption leads to an overestimation of the vehicle surface temperature.

Given that we are concerned with the longitudinal dynamics of the generic vehicle, the fluid domain assumes symmetry along the vehicle centerline. For each run, we drive the solution to a steady state using a hybrid DPLR–FMPR first order implicit method [14]. Its construction and numerics allow rapid convergence of flow fields by leveraging the strong directionality of the physics in the high aspect ratio elements near walls. These elements are required to resolve the thin boundary layers along the vehicle surface. The boundary conditions consist of supersonic inflow and outflow, with a symmetry boundary condition applied along the vehicle centerline in the x–y plane.

The primary objective here is to demonstrate control of generic BGVs with control surfaces during the glide phase. As a result, data were generated for plausible flight conditions.

We used the following parameter ranges to aid in constructing the surrogate model used in the IG&C framework:

- $\delta \in \{-15, -14, \ldots, 14, 15\}^\circ$
- $\alpha \in \{-5, -4, \ldots, 19, 20\}^\circ$
- $V_\infty \in \{2.0, 2.5, \ldots, 4.5, 5.0\}$ km/s
- $h \in \{20, 30, \ldots, 60, 70\}$ km

Here, δ is the fin deflection, α is the angle of attack, V_∞ is the free stream velocity, and ρ_∞ is the free stream density. The densities correspond to the altitudes defined by h, which are taken from the standard atmospheric model [15].

The Multi-Mesh framework described in Section II.A.3 was used to interpolate between anchor meshes for each fin deflection. Integration with US3D for aerodynamic database generation is outlined in Algorithm 1.

C. Conjugate Heat Transfer

A one-dimensional heat conduction solver is loosely coupled with US3D to simulate transient heating at the vehicle surface. After each steady-state flow-field evaluation, the temperature distribution in the TPS is advanced using a finite-volume approach. The CFD-computed surface heat flux is used to apply the boundary condition through Fourier's law, and time integration proceeds using a conservative sub-time step. The surface boundary condition accounts for both convective and radiative fluxes through the Stefan–Boltzmann law, while the back-face of the TPS is modeled as adiabatic. The heat flux supplied by CFD is non-dimensionalized using freestream quantities to compute the heat transfer coefficient C_H. During each sub-time step, a new wall temperature is obtained using the updated surface heat flux, informed by C_H. This process is repeated until the sub-time steps span the full integration interval, which in this work is equal to the controller update time Δt_c. The updated wall-temperature distribution is then mapped back onto the fluid-grid boundary for the next CFD evaluation. The TPS material is POCO graphite, with properties taken from Taylor and Groot [16] (sample 2) and listed in Table 1. The one-dimensional transient heat conduction equation with

Table 1 Thermal Properties for POCO graphite.

Property	Correlation	A	B	C	D
Thermal diffusivity $\alpha(T)$	$\alpha(T) = Ae^{-BT^C} + D$	4.461×10^{-4}	7.321×10^{-2}	5.675×10^{-1}	9.785×10^{-6}
Thermal conductivity $k(T)$	$k(T) = Ae^{-BT^C} + D$	8.566×10^{1}	4.172×10^{-5}	1.468	3.686×10^{1}

Algorithm 1 CFD Aero-database Generation

Initialization: Given the lower and upper limits of (α, h, V, δ), define a four dimensional hypercube.
Set: `Mass` = 1000 (mass of vehicle), `nevals` = 400 (number of CFD iterations), and `nit` = 0 (iteration counter).
while Aero-database is not complete **do**
  Update Inflow:
  Input: (α, h, V)
  Output: ρ, updated inflow conditions
  Call `atmospheric_model(`h`)` to compute ρ_∞ and T_∞
  Call `reset_BC(`$V_\infty, \rho_\infty, \alpha, T_\infty$`)` to update boundary conditions
  Interpolate Mesh:
  Input: δ
  Output: Interpolated mesh configuration
  Call `multi_mesh_interpolate(`δ`)`
  Compute CG and I_y:
  Input: Wall centroid locations
  Output: Center of gravity (CG) and moment of inertia I_y
  Call `compute_CG_and_I()` (with fin deflection)
  Converge Aerodynamic Forces and Moments:
  Initialize `nit` $\leftarrow 1$
  while `nit` $\leq$ `nevals` **do**
    `cfl ramping`
    Perform CFD evaluation using `cfl`
    `nit` $\leftarrow$ `nit` + 1
  end while
end while

temperature-dependent thermal diffusivity $\alpha(T)$ is

$$\frac{\partial T}{\partial t} = \frac{\partial}{\partial x}\left(\alpha(T)\,\frac{\partial T}{\partial x}\right), \tag{1}$$

where $\alpha(T) = k(T)/(\rho c_p)$ and $k(T)$ is the temperature-dependent thermal conductivity. The boundary conditions are

$$-k\,\frac{\partial T}{\partial x}\bigg|_{x=0} = 0, \qquad \text{(adiabatic back-face)} \tag{2a}$$

$$-k\,\frac{\partial T}{\partial x}\bigg|_{x=L} = q_w - \sigma\epsilon T_w^4, \qquad \text{(convective–radiative surface)} \tag{2b}$$

where $\sigma = 5.67 \times 10^{-8}$ W/(m^{2}K^4) is the Stefan–Boltzmann constant, $\epsilon = 0.89$ is the surface emissivity, and $T_w = T_N$ is the wall temperature. Using a second-order central difference for the surface gradient with ghost cell T_{N+1} gives

$$\frac{\partial T}{\partial x}\bigg|_{x=L} \approx \frac{T_{N+1} - T_{N-1}}{2\Delta x}, \tag{3a}$$

Substituting into Eq. (2b),

$$-k(T_w)\,\frac{T_{N+1} - T_{N-1}}{2\Delta x} = q_w - \sigma\epsilon T_w^4, \tag{3b}$$

Solving for the ghost temperature,

$$T_{N+1} = T_{N-1} + \frac{2\Delta x}{k(T_w)}\left(q_w - \sigma\epsilon T_w^4\right), \tag{3c}$$

The convective heat flux q_w is defined using the heat transfer coefficient as

$$q_w = C_H \, \rho_\infty V_\infty C_{p\infty} (T_\mathrm{r} - T_w), \tag{4a}$$

$$C_H = \frac{q_w}{\rho_\infty V_\infty C_{p\infty} (T_\mathrm{r} - T_w)}, \tag{4b}$$

$$C'_H = C_H \, \rho_\infty V_\infty C_{p\infty}, \tag{4c}$$

$$q_w = C'_H (T_\mathrm{r} - T_w), \tag{4d}$$

where ρ_∞, V_∞, and $C_{p\infty}$ are the freestream density, velocity, and specific heat. The recovery temperature is

$$T_\mathrm{r} = T_\infty + \frac{1}{2}\sqrt{Pr}\,\frac{V_\infty^2}{C_{p\infty}}, \tag{5}$$

with $Pr = 0.72$ for air. The domain $0 \le x \le L$ is discretized into N uniformly spaced nodes with $\Delta x = L/N$. Node indices are $i = 1, \ldots, N$, where $i = 1$ is the back-face and $i = N$ is the surface. The finite-volume form of Eq. (1) is

$$\frac{\partial T_i}{\partial t} = \frac{1}{\Delta x}\left(F_{i+1/2} - F_{i-1/2}\right), \tag{6a}$$

where the flux across the right interface $i + 1/2$ is

$$F_{i+1/2} = -\alpha_{i+1/2}\,\frac{T_{i+1} - T_i}{\Delta x}, \tag{6b}$$

and the flux across the left interface $i - 1/2$ is defined analogously. The interface diffusivity is computed using arithmetic averaging:

$$\alpha_{i+1/2} = \frac{1}{2}\left(\alpha(T_i) + \alpha(T_{i+1})\right), \tag{6c}$$

which also applies for the $i - 1/2$ interface using nodes $(i - 1, i)$. Substituting the flux definitions gives

$$\frac{\partial T_i}{\partial t} = \frac{1}{\Delta x^2}\Big[\alpha_{i+1/2}\left(T_{i+1} - T_i\right) - \alpha_{i-1/2}\left(T_i - T_{i-1}\right)\Big], \tag{6d}$$

The discrete system in Eq. (6d) is advanced in time using the forward Euler:

$$T_i^{n+1} = T_i^n + \Delta t_s \left.\frac{\partial T_i}{\partial t}\right|^n, \tag{7a}$$

where n denotes the time index and the sub-step Δt_s is taken such that

$$\Delta t_c = \sum_{k=1}^{n} \Delta t_s. \tag{7b}$$

A conservative, stability-limited sub-time step is chosen as

$$\Delta t_s \le \frac{1}{2}\,\frac{\Delta x^2}{\alpha_\mathrm{max}}, \tag{7c}$$

where α_max is the maximum thermal diffusivity in the domain and the factor of $1/2$ provides additional margin for numerical stability.

D. Integrated Guidance and Control

The MPC framework used in this work is derived from [10], with a change to the control variable. Instead of using control surface deflection as the input, the control variable is taken to be the deflection *rate*. A second modification is the use of the full aerodynamic database to compute the trim characteristics of the vehicle.

1. Equations of Motion

The longitudinal dynamics of the hypersonic vehicle are simulated using

$$\begin{aligned}
\dot{V} &= -\frac{D}{m} - g\sin\gamma, \\
\dot{\alpha} &= q - \frac{L}{mV} + \frac{g}{V}\cos\gamma, \\
\dot{\gamma} &= \frac{L}{mV} - \frac{g}{V}\cos\gamma, \\
\dot{q} &= \frac{\tau_y}{I_y},
\end{aligned} \tag{8}$$

where V is the speed, α is the angle of attack, γ is the flight path angle, q is the pitch rate, and I_y is the pitch-axis moment of inertia. The aerodynamic quantities L, D, and τ_y denote lift, drag, and pitching moment, respectively.

2. Guidance Computation

The guidance scheme is adapted from Kelley et al. [17], who derived an approximately optimal control law for constant dynamic-pressure glide at a specified lift-to-drag ratio, including the maximum. Several modifications were introduced such that the guidance accounts for control surface effects, pitch dynamics, and includes corrective terms ensuring the dynamic-pressure error converges to zero [18]. The modified guidance law is reproduced here

$$\gamma = \frac{1}{\left(1 - \frac{V^2}{2g_0\rho(h)}\frac{d\rho}{dh}\right)}\left(-\frac{1}{\left(\frac{L}{D}\right)} + \frac{e_\gamma \Delta\bar{q}}{\rho(h) g_0 V}\right) \tag{9}$$

Here, e_γ is a limiting term ensuring $\sin\gamma \approx \gamma$ holds (i.e., γ is a small-angle) while shaping the dynamic-pressure convergence; $\Delta\bar{q}$ is the dynamic-pressure error; and $\frac{d\rho}{dh}$ is the derivative of atmospheric destiny computed using the standard atmosphere model. Note that g_0 here denotes a constant acceleration due to gravity but we have applied this guidance law in conjunction with a gravity model that captures variations due to the altitude, flight speed, and flight path angle (see Eq. (10)) for our simulations. The algorithm implementing this guidance law is summarized in Algorithm 2.

Algorithm 2 Guidance Synthesis

Input: Current flight conditions, weight W, initial γ
while not converged **do**
  Find (α, δ) pairs satisfying:
  - $\tau_y \approx 0$ (zero pitching moment)
  - $L \approx W\cos\gamma$ (lift balances weight component)
  Of all possible (α, δ) pairs, choose pair corresponding to maximum (L/D)
  Assume $\gamma \ll 1$, so that $\gamma \approx \sin\gamma$
  Compute γ_{i+1} via Eq. (9)
  if $|\gamma_{i+1} - \gamma_n| <$ tol **then**
    return commanded control surface deflection δ_c
  else
    Update $\gamma_n \leftarrow \gamma_{i+1}$
  end if
end while

E. Adaptive Updating with CFD

The controller requires updated vehicle states on very small time scales ($\tau \sim O(10^{-2})$ s). Using CFD to evaluate aerodynamic forces at every time step over a full trajectory, which can last several minutes, would require an impractically large number of CFD evaluations. To mitigate this computational burden, it is observed that the angle of attack α and control surface deflection δ_c dominate the variation in aerodynamic forces. This motivates the adaptive update scheme summarized in Algorithm 3.

Algorithm 3 Adaptive Aerodynamic Coefficient Update Scheme

Input: Current state $x_n = [V_n, \alpha_n, \gamma_n, q_n]$, time step Δt, previous coefficients C_n
Run CFD to steady state and compute aerodynamic coefficients $C_n = \{C_L, C_D, C_M\}$
Record $(\alpha_n, \delta_n, C_n)$
Compute control command δ_c using LQR based on current state x_n
Integrate equations of motion over Δt to obtain x_{n+1}
if $|\alpha_{n+1} - \alpha_n| > 0.01\,\alpha_n$ **or** $|\delta_{n+1} - \delta_n| > 0.01\,\delta_n$ **then**
    Run CFD at x_{n+1} to compute new coefficients C_{n+1}
else
    Estimate C_{n+1} by scaling C_n using the change in dynamic pressure
end if
Return: Updated coefficients C_{n+1}

In practice, this scheme substantially reduces the number of CFD evaluations while preserving accuracy in regions where the flight conditions evolve rapidly. The current threshold of ±1% was chosen as a pragmatic upper limit on allowable changes in flight configuration before the aerodynamic force estimates become too uncertain. An analogous procedure is used for updating the heat transfer coefficient C_H when capturing the time-accurate heating along the vehicle.

F. CFD Feedback for Closed-Loop Controls

This section outlines the integration and coupling framework and describes how the pre-computed database is used online once CFD provides feedback for the closed-loop control of the generic HGV.

1. Use of Aerodynamic Database: Controller and Guidance Synthesis

Although CFD is used to compute HGV states and provide feedback to the controller for tracking reference states, the pre-computed database is used for two reasons.

1) Analytical Derivatives. The derivatives of aerodynamic forces and moments with respect to angle of attack and fin deflection δ_c are needed in the linear time varying controller formulation, where the nonlinear dynamics of Eq. (8) are linearized about the current vehicle state. Because the derivatives are not analytically available, they must be computed numerically. For simplicity, polynomials are used as surrogate models to interpolate pre-computed aerodynamic coefficients, enabling rapid evaluation of analytical derivatives. Computing these online using a finite-difference approach would be prohibitively expensive due to the poor scaling of computing this numerically.

2) Trim Computations. For similar computational savings, required trim tuples (α, δ), as used in the guidance algorithm (Algorithm 2), are approximated offline. Figure 2 shows the dependence of trim on Mach number and freestream density. The database was generated without aerothermal effects from the TPS material, so force mismatches are expected due to boundary layer variations caused by wall temperature assumptions. These trim curves are approximated by polynomials and used to model vehicle trim behavior.

2. Online Usage: CFD in the Loop

Figure 3 illustrates how data is interpolated from the polynomial curves online. The procedure is as follows:

1) Initialize the state for a given altitude and velocity. Guidance takes current vehicle states and computes reference states that aim to track a dynamic pressure. A nearest-neighbor approach selects the closest density and velocity corresponding to a trim curve (see Figure 2).
2) Reference states from guidance are used along with a linearized predictive model. The linearization is about the reference vehicle state, but sensitivities are derived from nearest-neighbor polynomial representations of trim configurations. This is solved as an optimal control problem for MPC to generate a command output $\dot{\delta}$.
3) CFD is converged to steady state, and real aerodynamic loads are combined with the commanded deflection rate to integrate the equations of motion and obtain new vehicle states X_{t+1}. The process repeats until trajectory termination.

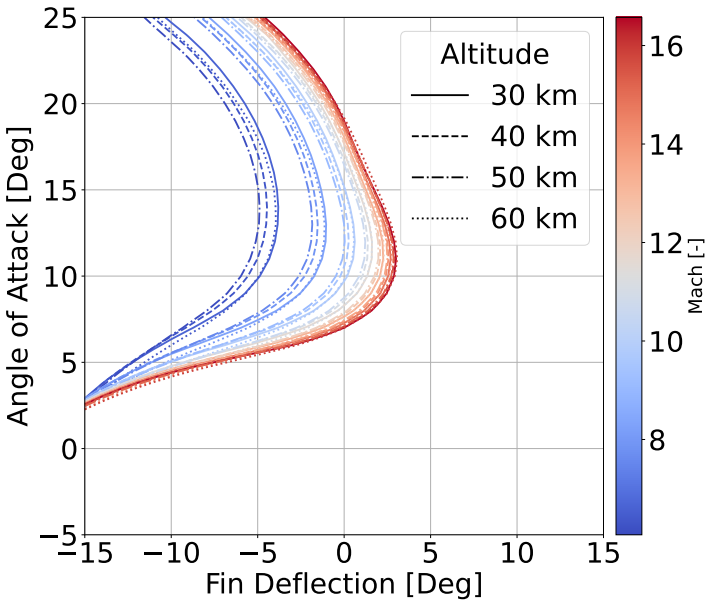

Fig. 2 Trim configuration dependence on density and Mach number for demonstration HGV. Each contour line denotes zero-pitch moment. A subset of densities is shown using different line styles.

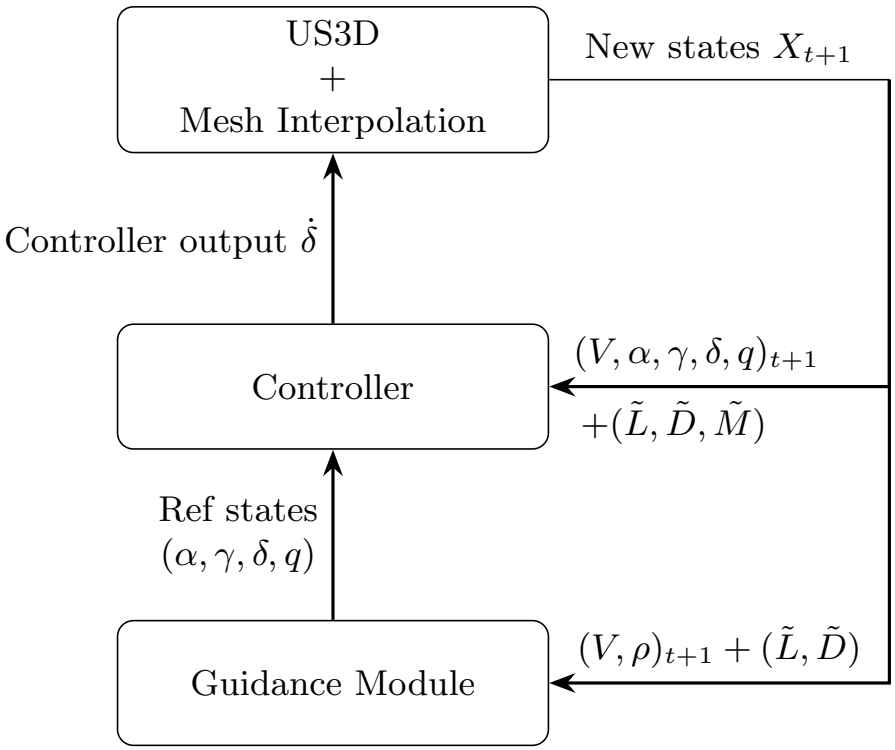

Fig. 3 Schematic overview of CFD in the loop with guidance and controls. Aerodynamic variables denoted with a tilde $\tilde{C}$, are derived from interpolation of the offline-computed database.

III. Results and Discussion

This section outlines results and discussion for simulations of controlled trajectories using CFD in the loop with closed-loop controls. The numerical experiments provide both quantitative and qualitative assessments of the integrated guidance and control framework when CFD is used for feedback. Results are organized into three parts: (1) influence of controller update time, (2) demonstration of guided trajectories in CFD, and (3) influence of modeling mismatches with aerothermal considerations. All experiments use the same initial conditions of 45 km altitude and 4 km/s velocity.

A. Required Sampling Time

Throughout these numerical experiments, it was found that using CFD instead of interpolating the aerodynamic database led to pitch instabilities. Over the predictive horizon of the controller, force mismatches between interpolation and CFD caused the previously stable sampling time of $\Delta t_c = 100$ ms to become insufficient for the MPC linear model of the longitudinal dynamics. These mismatches required a more frequent controller update, as shown in Figure 4.

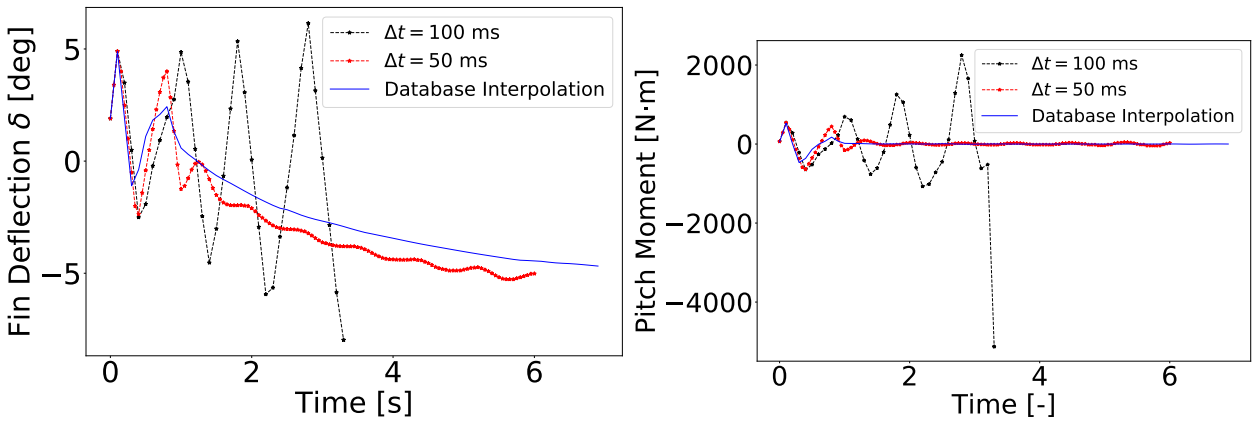

(a) Fin deflection vs. time **(b) Pitch moment vs. time**

Fig. 4 Comparison between controller update frequencies for dynamic pressure tracking. Database interpolation is updated every 100 ms for comparison.

A sampling time of $\Delta t_c = 50$ ms provided acceptable tracking during the initial portion of the transient trajectory. However, performance degraded later in the flight, and pitch instabilities emerged farther along the trajectory, as shown in Figure 5. Between approximately 55 and 60 seconds, abrupt changes in the reference states occur due to the nearest-neighbor method used for selecting trim curves in the guidance algorithm. The combination of database-derived trim estimates and the linearized HGV model ultimately required a reduced sampling time of $\Delta t_c = 10$ ms to maintain stability.

Reducing the sampling time from 100 ms to 10 ms imposes negligible computational overhead on the controller itself and only a marginal increase in CFD cost. Although more frequent updates increase the number of CFD evaluations, the adaptive CFD heuristic described in Sub-Section II.E offsets most of the additional expense. Improvements to the interpolation scheme may further relax the required update rate.

B. Dynamic Pressure Tracking Performance

A sweep of dynamic pressure tracking was performed. The tuned tracking parameters e_γ are summarized in Table 2. Here, e_γ is a tunable parameter that governs the aggressiveness of dynamic pressure tracking. Higher values correspond to more aggressive tracking, whereas lower values allow more relaxed tracking.

Table 2 Tunable dynamic pressure tracking parameter e_γ as a function of deviation from nominal dynamic pressure.

Dynamic Pressure Tracking [%]	-20	-10	0	10	20
e_γ	0.06	0.08	0.09	0.08	0.10

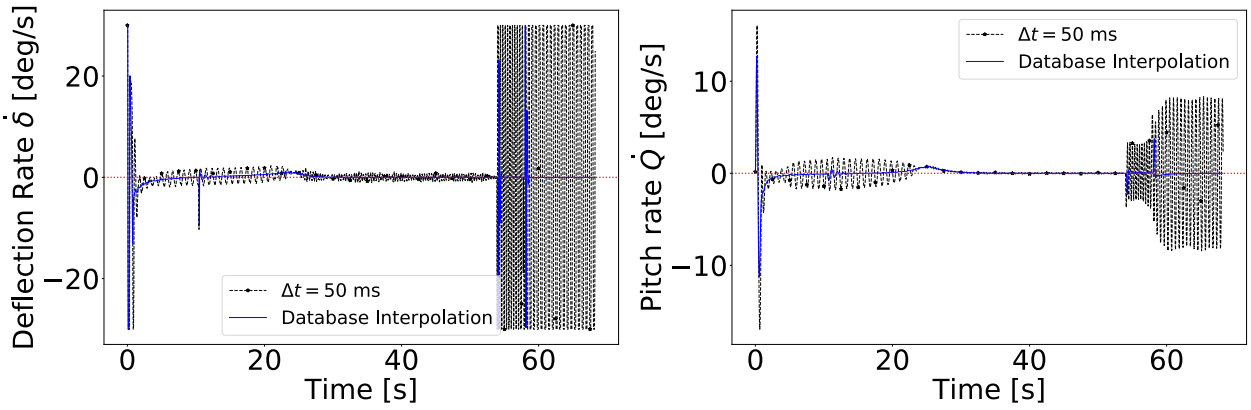

(a) Fin deflection rate vs. time

(b) Pitch rate vs. time

Fig. 5 Longer-time comparison of controller update frequencies for dynamic pressure tracking.

The resulting CFD-flown trajectories are shown in Figure 6. As in Figure 5, periodic spikes remain, caused by the discontinuous handover between trim points. A notable trend is that increased dynamic pressure cases during transient conditions tend to produce more stable or less oscillatory trajectories. Although the initial altitude and velocity are identical, the initial flight-path angles differ. For the cases requiring a reduction in dynamic pressure, the flight-path angle decreases, which requires the controller to command less lift. This low-lift configuration reduces maneuverability and results in more abrupt corrective motion of the control surface. Outside the transient, the angle of attack increases over time to maintain constant dynamic pressure.

However, the results in this subsection are obtained under the same radiative equilibrium assumption used to create the aerodynamic database. The next subsection relaxes this assumption for further comparison.

C. Aerothermal Considerations: Influence of Modeling Mismatches

The prior section focused on tracking performance using CFD in closed-loop control. Here, we examine how physical modeling assumptions influence the resulting trajectory. The analysis focuses on the +10% dynamic pressure case and compares several physical models against the assumptions used to generate the aerodynamic database.

A comparison of these cases is shown in Figure 7. All cases remain stable under a 10 ms controller update. However, as noted earlier, interpolation of trim curves introduces discontinuities that manifest as spikes in the vehicle dynamics. The conjugate heat transfer cases show larger deviations than the database or radiative equilibrium cases. This indicates that changes in wall temperature affect vehicle stability.

Tracking error with respect to the RANS baseline is also notable. Although skin-friction increases due to the turbulent boundary layer, the controller whose linear model assumes laminar flow with radiative equilibrium still maintains trim. Despite this, the CFD solution fails during trim-curve switching. The failure is preceded by a sudden, large downward pitch moment. This appears to result from an overcorrection triggered by the discontinuity during updates, which is amplified by force and trim mismatches between radiative equilibrium and CHT. Further investigation is required.

It is also noted that the CHT errors increase during rapid control-surface motion. Because no interpolation is performed between successive deformed grids, wall temperatures are not mapped between meshes in a way that preserves thermal gradients. This assumption is not valid even with the shortened sampling time of 10-100 ms. Numerical error may accumulate due to strong temperature gradients in high aspect ratio boundary-layer grid cells. Future work will quantify and mitigate this error.

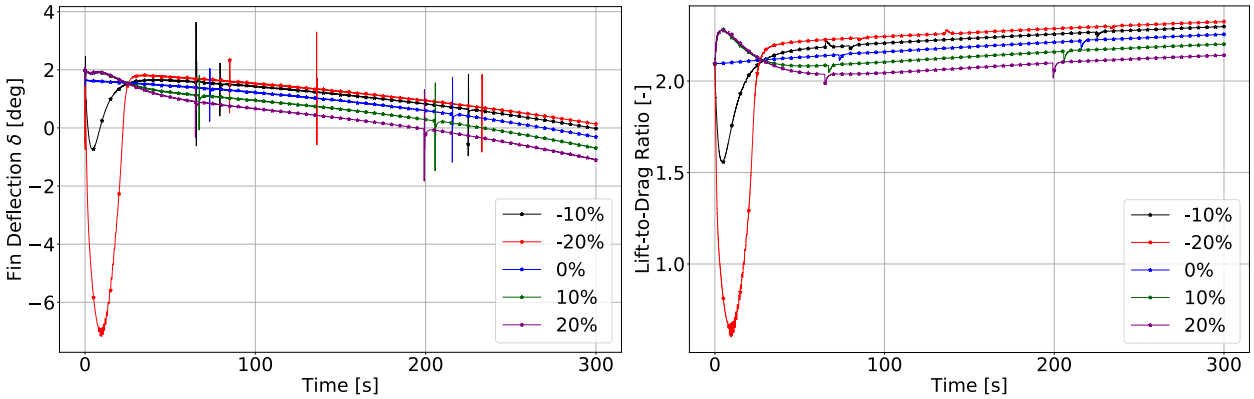

(a) Fin deflection vs. time

(b) Lift-to-drag ratio vs. time

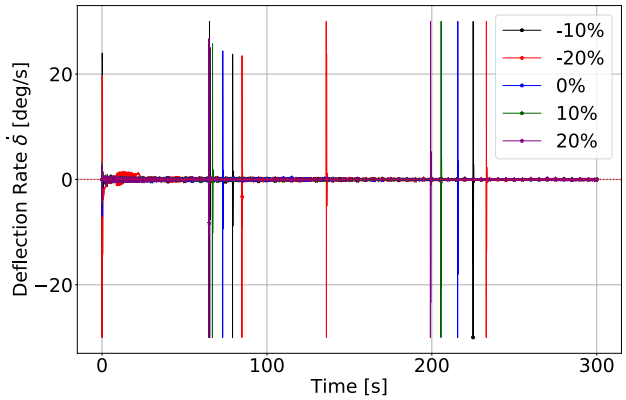

(c) Fin deflection rate vs. time

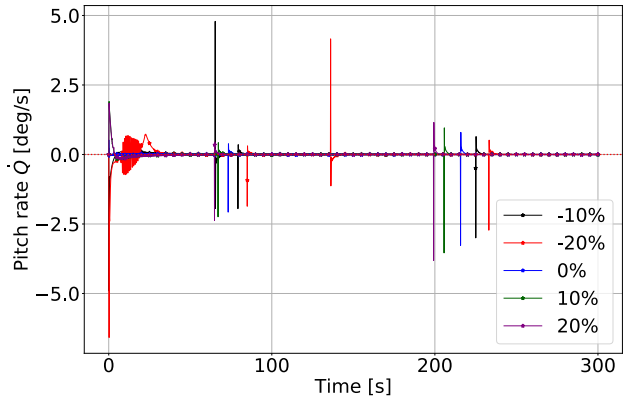

(d) Pitch rate vs. time

Fig. 6 Tracking performance with varying dynamic pressure. Percentages denote deviations from initial conditions $h = 45$ km and $V_\infty = 4$ km/s.

IV. Conclusion and Future Work

This work demonstrates an integrated framework that couples numerical optimization, high-fidelity CFD, speed-up gains from appropriate approximations, and surrogate modeling to simulate controlled trajectories of generic HGVs. The primary advantage of using high-fidelity CFD within closed-loop control for hypersonic glide vehicles is the ability to provide realistic feedback that reduces uncertainty in vehicle stability, trim, and trajectory feasibility. In addition, this approach increases confidence in control-surface efficacy under physical phenomena such as flow separation and high-temperature aerothermodynamic effects. The results highlight the challenges and opportunities of using CFD as the feedback source within an MPC-based integrated guidance and control architecture. The trajectory simulations show that using CFD in the loop introduces force and moment discrepancies relative to database interpolation, which directly influence controller stability and update-rate requirements. A controller update time of 10 ms is required to maintain stable tracking, in contrast to the 100 ms interval sufficient when using an aerodynamic database. These results indicate that the MPC horizon and linearized model become increasingly sensitive to modeling mismatches when higher-fidelity aerodynamics drive the closed-loop behavior. Dynamic pressure tracking studies reveal that trim-curve discontinuities introduce periodic spikes in the vehicle dynamics that persist across update frequencies. Under transient conditions, increased dynamic pressure tends to produce more stable trajectories, whereas low-lift configurations reduce control authority and increase sensitivity to oscillatory artifacts. When aerothermal assumptions are varied, conjugate heat transfer cases deviate more strongly from database-derived models using radiative equilibrium. These deviations suggest that wall-temperature evolution and heating models can meaningfully alter closed-loop behavior even when the vehicle remains nominally trimmed. Therefore, while lower-fidelity aerodynamic methods can reproduce general trends in lift-to-drag ratios compared to high-fidelity CFD, the differences in absolute values are large enough to influence guidance and control performance.This underscores the importance of consistent physics modeling across all

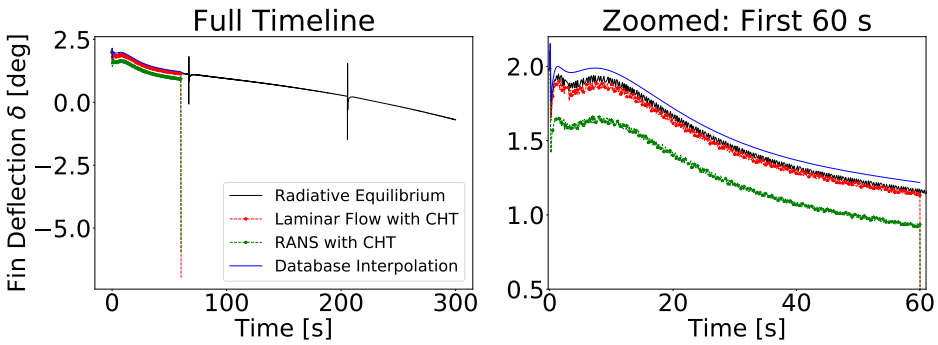

(a) Fin deflection vs. time

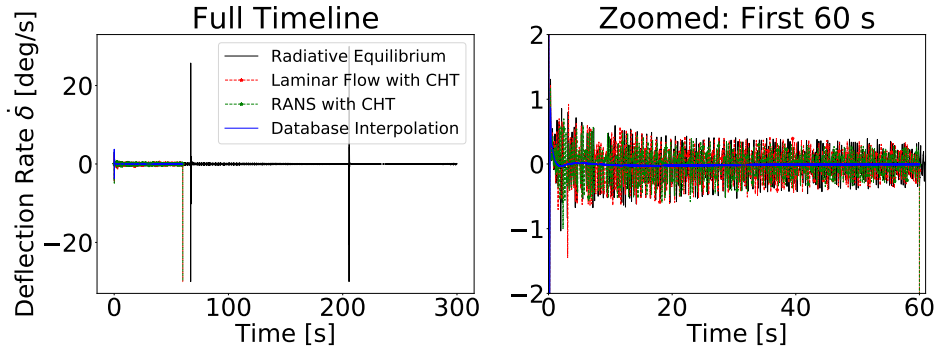

(b) Fin deflection rate vs. time

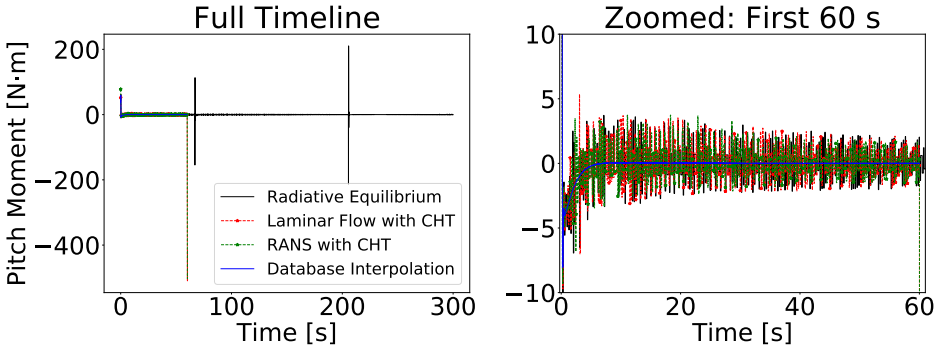

(c) Pitch moment vs. time

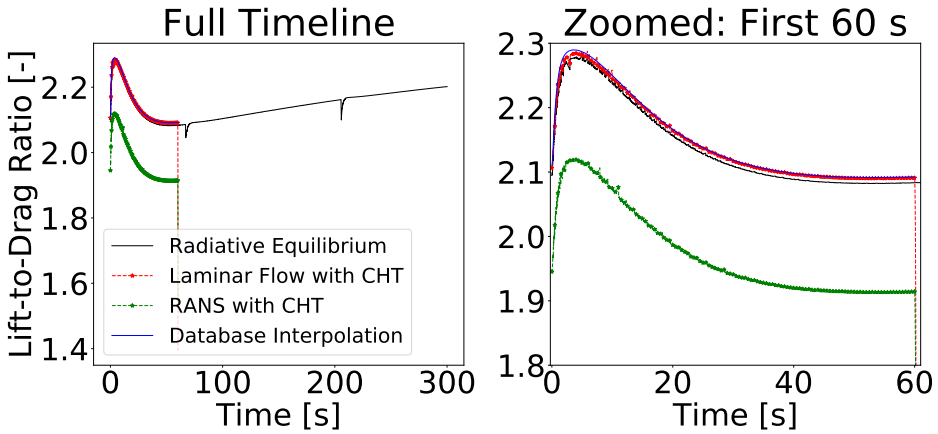

(d) Lift-to-drag ratio vs. time

Fig. 7 Comparison of tracking performance for a +10% increase in dynamic pressure from initial conditions $h = 45$ km and $V_\infty = 4$ km/s. Cases include CFD with conjugate heat transfer (with and without turbulence), radiative equilibrium, and nearest-neighbor interpolation of the aerodynamic database.

components of the IG&C pipeline.

Future efforts will focus on improving the consistency and robustness of the integrated IG&C framework with aerothermal considerations. Key directions include:

Improved physical modeling and uncertainty quantification. We will investigate uncertainties associated with aerothermal modeling assumptions, including CHT, radiative equilibrium, and chemical nonequilibrium effects. Understanding how these uncertainties propagate through the MPC prediction model is essential for robust closed-loop operation.

Incorporation of aerothermal constraints into IG&C. Heating constraints will be added to the IG&C framework with the goal of generating multi-objective trajectories balancing dynamic pressure tracking and minimal heat load. This will enable comparisons between trajectories generated using engineering estimates of convective heat flux and those using high-fidelity CFD and CHT. The primary goal is to quantify the sensitivity of optimized trajectories to aerothermal modeling fidelity. General thoughts on the implementation is outlined in the appendix in section .A.

Consistent thermal-state mapping between fluid and solid domains. Improved methods for transferring wall temperatures and heat fluxes between successive deformed meshes will be developed. Correct mapping is needed to prevent unphysical changes in thermal gradients during rapid control-surface motion.

Database-free or reduced-database dependence strategies. To reduce dependence on large aerodynamic databases, we will explore trim-curve approximation using sparse CFD sampling and local flow-field interpolation. This may reduce discontinuities and improve controller smoothness, while also alleviating the need to generate large numbers of CFD evaluations offline.

Low-density and rarefied-flow modeling. Future simulations will extend into regimes where real-gas effects, near-non-continuum behavior, and possible slip conditions near the nose and control-surface leading edges become important.

Investigation of unsteady aerodynamics within 10–100 ms timescales. High-frequency unsteady effects may contribute to control-surface-induced oscillations, especially when tight update intervals are required by the MPC. Quantifying these unsteady contributions will help refine both controller design and CFD update heuristics. This also supports improved quantification of uncertainty between modeling assumptions and the acceptable fidelity levels for simulations of controlled trajectories.

Collectively, these efforts aim to create a predictive, aerothermodynamic informed IG&C framework that minimizes reliance on databases, maintains closed-loop stability under high-fidelity physics, and more accurately captures the interdisciplinary coupling between aerodynamics, heating, and guidance and control for hypersonic flight.

Appendix

A. Proposed Modified Guidance with Thermal Constraint

We seek to impose a lower bound on the reference flight-path angle to satisfy a thermal constraint. Two approaches are formulated and we are working on implementing these in our framework. In both cases the goal was to track dynamic pressure while reducing surface heating along the trajectory.

We begin by defining the necessary heat-flux quantities. The derivation is organized in the subequations below.

Effective heat flux. The effective surface heat flux is

$$q_{\text{eff}} = C'_H(T_\infty - T_w) \; - \; \sigma\epsilon T_w^4. \tag{10a}$$

Root finding. Define

$$F(\gamma) = \frac{\partial q_{\text{eff}}}{\partial t}, \tag{10b}$$

and obtain the optimal flight–path angle from

$$F(\gamma^*) = 0. \tag{10c}$$

Time derivative of the heat flux. Applying the product rule gives

$$F(\gamma) = C_H\, C_p \left[\frac{\partial \rho}{\partial t}\, V\, (T_r - T_w) + \rho\, \frac{\partial V}{\partial t}\, (T_r - T_w) + \rho\, V \left(\frac{\partial T_r}{\partial t} - \frac{\partial T_w}{\partial t} \right) \right] - \; 4\sigma\epsilon T_w^3 \frac{\partial T_w}{\partial t}. \tag{10d}$$

Dynamic dependence on γ. The dependent variables satisfy

$$\begin{aligned} \frac{\partial h}{\partial t} &= V \sin\gamma, \\ \frac{\partial V}{\partial t} &= -\frac{D}{m} - g_{\text{eff}} \sin\gamma, \\ \frac{\partial T_r}{\partial t} &= \sqrt{0.72}\, \frac{V}{C_p}\, \frac{\partial V}{\partial t}, \\ \frac{\partial \rho}{\partial t} &= \frac{\partial \rho}{\partial h}\, \frac{\partial h}{\partial t}, \\ \frac{\partial T_w}{\partial t} &= \frac{dT_w}{dt}, \end{aligned}$$

with

$$g_{\text{eff}} = g_0 - \frac{V^2 \cos^2\gamma}{R + h}.$$

Derivatives with respect to γ. Differentiation yields

$$\frac{\partial}{\partial \gamma} \left(\frac{\partial h}{\partial t} \right) = V \cos\gamma, \qquad \frac{\partial g_{\text{eff}}}{\partial \gamma} = \frac{2V^2 \sin\gamma \cos\gamma}{R + h},$$

$$\frac{\partial}{\partial \gamma} \left(\frac{\partial V}{\partial t} \right) = -\frac{\partial g_{\text{eff}}}{\partial \gamma} \sin\gamma - g_{\text{eff}} \cos\gamma.$$

Thus,

$$\begin{aligned} \frac{\partial}{\partial \gamma} \left(\frac{\partial T_r}{\partial t} \right) &= \sqrt{0.72}\, \frac{V}{C_p}\, \frac{\partial}{\partial \gamma} \left(\frac{\partial V}{\partial t} \right), \\ \frac{\partial}{\partial \gamma} \left(\frac{\partial \rho}{\partial t} \right) &= \frac{\partial \rho}{\partial h}\, \frac{\partial}{\partial \gamma} \left(\frac{\partial h}{\partial t} \right). \end{aligned}$$

Analytic derivative $F'(\gamma)$. Applying the product rule to (10d),

$$F'(\gamma) = C_H\, C_p \left[\left(\frac{\partial}{\partial \gamma} \frac{\partial \rho}{\partial t} \right) V (T_r - T_w) + \rho \left(\frac{\partial}{\partial \gamma} \frac{\partial V}{\partial t} \right) (T_r - T_w) + \rho V \left(\frac{\partial}{\partial \gamma} \frac{\partial T_r}{\partial t} \right) \right]. \tag{10e}$$

Newton iteration. The optimal angle obeys

$$\gamma_{k+1} = \gamma_k - \frac{F(\gamma_k)}{F'(\gamma_k)}, \tag{10f}$$

with convergence when $|F(\gamma_k)| < \text{tol}$.

Optimal solution. Thus,

$$\gamma_{k+1} = \gamma_k - \frac{F(\gamma_k)}{F'(\gamma_k)} : \quad k \to \infty \tag{10g}$$

The solution γ^* to Eq. (10g) provides a lower bound on the flight–path angle before descent would cause the heat flux to increase. In practice, this constraint is rarely active along the trajectory because the wall–temperature rate of change, dT_w/dt, is sufficiently large in magnitude leaving this lower bound mostly inactive throughout the trajectory. Despite this, we introduce two modifications to the guidance scheme to account for a multi–objective formulation in which the vehicle tracks dynamic pressure while simultaneously minimizing thermal loading.

1) Use γ^* when heating constraints become active; otherwise use the nominal guidance law in Eq. (9).
2) Replace the constant dynamic-pressure tracking gain e_γ with a functional form $e_\gamma = e_\gamma(T_{\text{max wall}})$ or $e_\gamma = e_\gamma(q_{\text{max wall}})$ so that the tracking aggressiveness respects the thermal lower bound.

Acknowledgments

The author would like to thank Dr.Travis Drayna, Dr.Matt Bartkowicz, and Dr.Joseph Brock for help in compiling and meshing tools need for the study.

This work was supported by the Air Force Office of Scientific Research under grant number FA9550-1-22-0004. The views and conclusions contained herein are those of the authors and should be not interpreted as representing the official policies or endorsements, either expressed or implied, of the AFOSR or the U.S. Government.

References

[1] Herrmann, E. G., Cox, A., and Mavris, D., "Multidisciplinary Design and Analysis of a Hypersonic Glide Vehicle With Trimmed Aerodynamics," *AIAA SCITECH 2025 Forum*, 2025. doi:10.2514/6.2025-1338, URL https://arc.aiaa.org/doi/abs/10.2514/6.2025-1338.

[2] Rataczak, J. A., Chaudhry, R. S., McMahon, J. W., and Boyd, I. D., "Investigation of Surface-Catalycity Effects on Hypersonic Glide Vehicle Trajectory Optimization," *Journal of Spacecraft and Rockets*, Vol. 61, No. 4, 2024, pp. 1019–1031. doi:10.2514/1.A35764, URL https://doi.org/10.2514/1.A35764.

[3] Mines, J. M., and Patterson, A. P., "Hypersonic Trajectory Design and Analysis Under Uncertainty Using Linear Covariance," *AIAA SCITECH 2025 Forum*, 2025. doi:10.2514/6.2025-2299, URL https://arc.aiaa.org/doi/abs/10.2514/6.2025-2299.

[4] Shi, Y., and Wang, Z., "A Deep Learning-Based Approach to Real-Time Trajectory Optimization for Hypersonic Vehicles," *AIAA Scitech 2020 Forum*, 2020. doi:10.2514/6.2020-0023, URL https://arc.aiaa.org/doi/abs/10.2514/6.2020-0023.

[5] Wang, Z., and Grant, M. J., "Autonomous Entry Guidance for Hypersonic Vehicles by Convex Optimization," *Journal of Spacecraft and Rockets*, Vol. 55, No. 4, 2018, pp. 993–1006. doi:10.2514/1.A34102, URL https://doi.org/10.2514/1.A34102.

[6] Rataczak, J. A., Chaudhry, R. S., McMahon, J. W., and Boyd, I. D., "Investigation of Surface Catalycity Effects on Hypersonic Glide Vehicle Trajectory Optimization," *Journal of Spacecraft and Rockets*, Vol. 61, No. 4, 2024, p. 1019 1031.

[7] Liu, Z., Cai, Y., Liu, Y., Deng, Y., and Jiang, H., "Trajectory Optimization of Hypersonic Glide Vehicle Based on Pseudospectral Method Assisted by BP Neural Network," *2023 China Automation Congress (CAC)*, 2023, pp. 7652–7657. doi:10.1109/CAC59555.2023.10450472.

[8] Bowcutt, K. G., "Multidisciplinary Optimization of Airbreathing Hypersonic Vehicles," *Journal of Propulsion and Power*, Vol. 17, No. 6, 2001, pp. 1184–1190. doi:10.2514/2.5893, URL https://doi.org/10.2514/2.5893.

[9] Starkey, R. P., and Lewis, M. J., "Critical Design Issues for Airbreathing Hypersonic Waverider Missiles," *Journal of Spacecraft and Rockets*, Vol. 38, No. 4, 2001, pp. 510–519. doi:10.2514/2.3734, URL https://doi.org/10.2514/2.3734.

[10] Williams, D., Bhattacharjee, D., Drayna, T. W., Bartkowicz, M., Hemati, M., and Candler, G. V., "Integrated Guidance and Control of Generic Hypersonic Glide Vehicles Using Computational Fluid Dynamics," *AIAA SCITECH 2025 Forum*, 2025. doi:10.2514/6.20250264, URL https://arc.aiaa.org/doi/abs/10.2514/6.20250264.

[11] Drayna, T. W., Stam, H. R., , Long, L. J., Lang, L. M., and Pate, J. P., *Crosslink V.0.11.x User Manual*, 2022. doi:10.2172/1888178, URL https://www.osti.gov/biblio/1888178.

[12] Bartkowcz, M., Haag, C., Nitzkorski, Z., Gidzak, V., and Brock, J., "LINK3D - Grid generation software overview and application," 2023.

[13] Candler, G. V., Johnson, H. B., Nompelis, I., Gidzak, V. M., Subbareddy, P. K., and Barnhardt, M., "Development of the US3D code for advanced compressible and reacting flow simulations," *53rd AIAA Aerospace Sciences Meeting*, 2015. doi:10.2514/6.2015-1893.

[14] Nompelis, I., Drayna, T., and Candler, G., *Development of a Hybrid Unstructured Implicit Solver for the Simulation of Reacting Flows Over Complex Geometries*, 2012. doi:10.2514/6.2004-2227, URL https://arc.aiaa.org/doi/abs/10.2514/6.2004-2227.

[15] *U.S. Standard Atmosphere, 1976*, 1976. URL https://ntrs.nasa.gov/api/citations/19770009539/downloads/19770009539.pdf.

[16] Taylor, R. E., and Groot, H., *Thermophysical properties of Poco graphite*, Defense Technical Information Center, 1978.

[17] Kelley, H. J., Cliff, E. M., and Lutze, F. H., "Boost glide range optimal guidance," *Optimal Control Applications and Methods*, Vol. 3, No. 3, 1982, p. 293 298.

[18] Bhattacharjee, D., Jaeger, S., and Hemati, M., *Integrated Guidance and Control of Quasi-Equilibrium Hypersonic Gliding Using Model Predictive Control*, 2025. doi:10.2514/6.2025-3550, URL https://arc.aiaa.org/doi/abs/10.2514/6.2025-3550.