

An Assessment of the Laminar Hypersonic Double-Cone Experiments in the LENS-XX Tunnel

Jaideep Ray* and Patrick Blonigan†

Sandia National Laboratories, Livermore, CA 94550, USA

Eric T. Phipps‡ and Kathryn Maupin§

Sandia National Laboratories, Albuquerque, NM 87185

In this paper, we investigate two experimental datasets of laminar hypersonic flows, in vibrational and reactive non-equilibrium, over a double-cone geometry, acquired in CUBRC's LENS-XX expansion tunnel. These datasets have yet to be modeled accurately. A previous paper suggested that this could partly be due to mis-specified inlet conditions. The authors of this paper solved a Bayesian inverse problem to infer the inlet conditions of the LENS-XX test-section and found that in one case they lay outside the uncertainty bounds specified in the experimental dataset. However, the inference was performed using approximate surrogate models. In this paper, we revisit the experimental datasets and perform inversions for the tunnel test-section inlet conditions with a Navier-Stokes simulator. The inversion is deterministic and provides uncertainty bounds on the inlet conditions under a Gaussian assumption. We find that deterministic inversion yields inlet conditions that are at variance with what was stated in the experiments. We also assess the usefulness of the uncertainty bounds provided by the Gaussian approximation of the true posterior distribution, and find it marginally useful for the inter-quartile range. We also present an *a posteriori* method to check the validity of the Gaussian assumption for the posterior distribution. This paper contributes to ongoing work on the assessment of datasets from challenging experiments conducted in extreme environments, where the experimental apparatus is pushed to the margins of its design and performance envelopes.

I. Nomenclature

h'_0	= non-dimensional stagnation enthalpy of the flow
PDF	= probability density function
P'_{Pitot}	= non-dimensional Pitot pressure of the flow
$p(x)$	= surface pressure on the double-cone
$q(x)$	= surface heat flux on the double-cone
x	= distance along the axis of the double-cone
U_∞	= inlet velocity being inferred
U_{norm}	= inlet velocity, used for normalization; set to 10^3 m/s
ρ_∞	= inlet density being inferred
ρ_{norm}	= inlet density, used for normalization; set to 10^{-4} kg/m ³
CI	= Credibility Interval
M_∞	= inlet Mach number
Re	= Unit Reynolds number
T_∞, T_{v_∞}	= inlet temperatures, translational and vibrational
y	= model parameters to be inferred, i.e., an inlet flow condition $\{\rho_\infty, U_\infty\}$
IQR	= inter-quartile range, i.e., the span between the 25 th and 75 th percentiles of a random variable

*Technical Staff, Extreme-Scale Data Science and Analytics, MS 9152, Member, AIAA

†Technical Staff, Extreme-Scale Data Science and Analytics, MS 9152, Member, AIAA

‡Technical Staff, Scalable Algorithms, MS 1318

§Technical Staff, Optimization & UQ, MS 1318

II. Introduction

EXPERIMENTAL datasets, meant for model validation, contain measurements of the phenomena that are the subject of the experiment, as well as a detailed quantitative description of the experimental environment and apparatus (so that the experiment may be modeled). The error/uncertainty in measurements that quantify the phenomena, as well as the experimental environment, are also supplied. Model validation efforts usually take the measurements and their (quantified) errors at face value. However, when simulating extreme environments, e.g., high enthalpy flows in shock tunnels with large cross-sections, one is forced to employ novel techniques, instrumentation, and apparatus, whose interactions may not be known very well, or which might be required to work at the edges of their operational envelope. In such scenarios, the quantifications of experimental errors/uncertainties that accompany a (model) validation dataset may be difficult to estimate. This causes complications when the dataset is used to validate a model — if the model predictions and experimental measurements do not agree, it is difficult to judge whether it is due to a shortcoming of the model or due to a mis-specification of the error/uncertainty bounds on the measurements in the experimental dataset. Ref. [1] describes one such case when an experimental dataset of laminar hypersonic flow over a 25°/55° double-cone, gathered in the Calspan-University at Buffalo Research Center's (CUBRC's) Large Energy National Shock Tunnel (LENS-I) wind tunnel, was used to validate a Navier-Stokes model of the flow. The model predictions failed to match experiments. The authors of Ref. [1] simulated the flow in the LENS-I shock-tunnel (the design was made available to them) and found that the inflow into the test section was in vibrational non-equilibrium (the total enthalpy of the flow was 3.71 MJ/kg). This was not known to the experimentalists, and was not reflected in the experimental dataset. When the correct boundary conditions, i.e., vibrational non-equilibrium, were specified, model predictions had far fewer problems matching experimental measurements [1].

A similar problem was encountered in Ref. [2] when using double-cone data, collected in CUBRC's LENS-XX expansion tunnel, to validate SPARC (Sandia Parallel Aerodynamics and Reentry Code), a Navier-Stokes model for non-equilibrium hypersonic flows. The experimental datasets contained pressure and heat flux measurements over the same 25°/55° double-cone, but the flows in question had far higher total enthalpies (5.44 MJ/kg and 21.77 MJ/kg). The uncertainties in the inflow conditions were provided as part of the experimental datasets. The authors created an ensemble of inflow conditions consistent with the specified uncertainties and simulated flows over the double-cone, with the aim of bracketing the measurements. They found that the distribution of model predictions did not agree with the distribution of measurements — they sometimes overlapped partially but were frequently disjoint. This implied that SPARC could not simulate some of the physics in the dataset. An obvious check would be to simulate the flow inside LENS-XX and compare the predicted inlet conditions (alternatively, the boundary conditions imposed by SPARC) with the specification that accompanied the validation dataset. However, the design details of the LENS-XX expansion tunnel are not publicly known and, consequently, the simulation cannot be performed. In addition, it is known that the flow inside LENS-XX is far more complex than in LENS-I; see Refs [3, 4] for a description. Ref. [5] simulated the flow inside LENS-X (the predecessor to LENS-XX) and found sections with steady and unsteady flows.

In order to address the problem of unknown inlet (inflow) conditions, Ref. [6] employed an inversion approach. They assumed that the inflow was axisymmetric and then, conditional on pressure and heat flux measurements on the fore-cone, as well as the total pressure and enthalpy of the flow, inferred the velocity, density and temperature at the inflow of the test section. The inversion was Bayesian and used a Markov chain Monte Carlo (MCMC) method to estimate the inflow conditions as a 3-dimensional probability density function (PDF). The PDF captured the uncertainty in the estimates due to the limited nature of the measurements, the measurement errors, and the missing physics in the SPARC model. For Case 1 (the 5.44 MJ/kg flow), they found that the measured and modeled PDFs were disjoint, but the maximum *a posteriori* (MAP) estimate of the inflow conditions (loosely, the most probable estimate) fell within the experimental uncertainty bounds. For the second case (the 21.77 MJ/kg flow, called “Case 4”), the MAP estimate was clearly outside the experimental uncertainty bounds. However, these inferences were drawn using polynomial surrogates for SPARC's predictions, a necessity due to the tremendous computational cost of MCMC. The polynomial surrogates are themselves approximations of the SPARC model, and the estimated inflow conditions are affected by the accuracy of the surrogates.

In this paper, we revisit the inference of the inflow conditions for Case 1 and Case 4 using the Navier-Stokes model in SPARC (and *not* any surrogates). The inference is cast as a deterministic inverse problem and solved using an optimization method that “tunes” the inflow conditions to maximize the agreement between the experimental data and model predictions. The flow separates at a point on the fore-cone of the double-cone configuration, and, just as in Ref. [6], we only use measurements that lie upstream of the separation point. A quantification of the uncertainty in the estimated inflow conditions can be obtained by approximating the PDF (called the “posterior” PDF of the inflow conditions) as a multivariate Gaussian, whose variance-covariance matrix is provided by the inverse of the Hessian of the

objective function, evaluated at the *optimal* point (inflow estimate which maximizes the match between measurements and predictions). In this paper, we also assess the usefulness of the Gaussian approximation by comparing it to the MCMC solution.

This paper has two main contributions. The first is the assessment of the inferences drawn in Ref. [6] regarding the disagreement of inferred inflow conditions with that specified in the experimental dataset. Note that unlike Ref. [6], the native Navier-Stokes model in SPARC (and *not* its surrogate) is used in this study. Our second contribution is an assessment of the Gaussian approximation of the posterior distribution, available with all deterministic inverse solutions, for uncertainty quantification in the highly nonlinear problems encountered in hypersonic flows. An additional novelty of the paper is the first use of adjoint-based sensitivity analysis to solve an inverse problem in the reconstruction of a hypersonic flow; to date, adjoint-based inverse solutions have been used extensively in design optimization and mesh adaptation.

The paper is laid out as follows. In Sec. III, we review the background literature on double-cone experiments and modeling efforts, as well as adjoint-based sensitivity and inverse solutions. In Sec. IV we perform a grid-refinement study and establish certain other numerical details. In Sec. V, we formulate the inverse problem, and in Sec. VI, our results. We first perform an inversion using a dataset measured in LENS-I which has been successfully modeled, and check whether our inversion yields inflow conditions that agree with those specified in the experimental dataset. Thereafter, we proceed to Case 1 and Case 4. We conclude in Sec. VII.

III. Literature Review

A. The Double-Cone Problem

Double-cone experiments have been conducted in the LENS-I [7, 8] and LENS-XX [9] facilities. The hypersonic flow sets up a complex shock pattern (see numerical Schlieren in Fig. 1 (left)), a separation zone, and complex shock-shock and shock-boundary layer interaction downstream of the separation zone. The flow is laminar throughout. The fore-cone is a 25° half-angle cone, which abuts a truncated 55° aft-cone. The fore-cone sets up an attached shock, whereas the aft-cone creates a detached shock. The two shocks interact to create a triple point and a transmitted shock that impinges on the aft-cone. The adverse pressure gradient set up by the transmitted shock feeds upstream through the (subsonic) boundary layer and causes the flow over the fore-cone to separate at $x \approx 0.05\text{m}$. The separated flow reattaches on the aft-cone. The separation zone, in turn, causes a separation shock that interacts with the transmitted shock and modifies it. Due to the existence of triple points, the region near the intersection of the two cones contains contact surfaces(vortex sheets), as well as a separation zone. These are clearly labeled in Fig. 1(left).

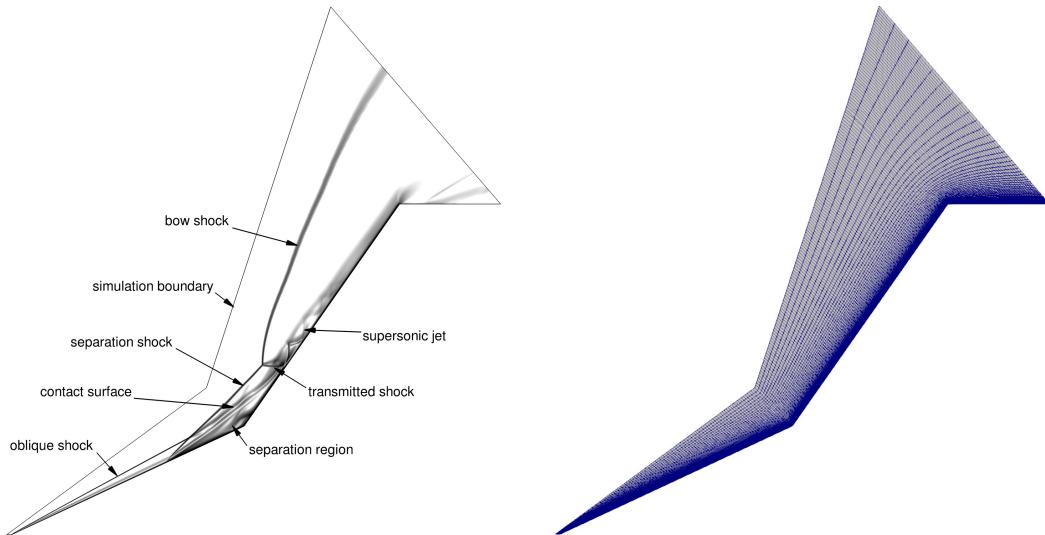


Fig. 1 Left: A numerical Schlieren of the flow over a double-cone, taken from Ref. [6]. Right: A stretched grid on which Computational Fluids Dynamics simulations were performed.

Double-cone experiments: Refs. [7, 8] describe double-cone experiments in the LENS-I shock tunnel. When the experiments were modeled using the Navier-Stokes equations, it was found that the model over-predicted heat-flux but did well with predictions of the pressure on the fore-cone, as well as the separation zone. Note that only one of the experiments, called “Run 35” (see specifications in Table 1), was modeled under the assumptions of an axisymmetric flow at the inlet of the shock tunnel that was also in vibrational and reactive equilibrium. However, it was later discovered that the inflow into the test-section of the shock tunnel was in strong vibrational non-equilibrium and was, in addition, frozen [1]. This discovery was made when the authors of Ref. [1] simulated the flow inside the entire tunnel and examined the flow state at the test-section inlet. This necessarily requires knowledge of the geometry of the nozzle, which was made available to the authors. When the inlet flow conditions (henceforth, inflow conditions) were adjusted to reflect vibrational non-equilibrium, Computational Fluid Dynamics (CFD) simulations using the Navier-Stokes model for Run 35 matched experimental data quite well [1]. The LENS-I experiments were conducted in a N_2 environment at very low flow density and total enthalpy, and a perfect-gas model, with no reactions, was quite appropriate.

The double-cone experiments were also conducted in the LENS-XX *expansion* tunnel [9], where temperatures do not exceed about 3000K [3]. Consequently, the thermochemical non-equilibrium issues that affected the LENS-I inflow conditions do not arise. Six double-cone experiments at different flow conditions (total enthalpies and pressures) were conducted. Two of these experiments, called Case 1 and Case 4, will be investigated in this paper and are summarized in Table 1. Attempts to model these six LENS-XX double-cone experiments have met with varying levels of success. We describe these next.

Modeling of double-cone experiments: Ref. [10] provides a good description of modeling studies of the LENS-XX double-cone experiments, including visualizations of the flowfields. The authors showed that the accurate pressure predictions (i.e., they match LENS-XX measurements) on the fore-cone could be achieved with both perfect-gas and non-equilibrium flow models, but the heat-flux predictions were too low, even when accounting for measurement errors (5% for surface pressure and 7% for surface heat-flux on the double-cone). The separation zone was predicted to be smaller, primarily because of late separation. Studies have explored the correct thermochemical modeling approach to adopt, often with conflicting results [11, 12]. Not surprisingly, the high-enthalpy Case 4 shows the largest differences in $p(x)$ and $q(x)$ predictions, depending on the choice of the thermochemical (equilibrium/frozen/non-equilibrium) modeling of the flow, while Case 1 shows the least. Here $p(x)$ and $q(x)$ are the pressure and heat-flux on the double-cone surface, and x is the axial distance from the double-cone nosetip. In all cases, it is easy to match the pressure on the fore-cone well, while under-predicting the heat-flux. Aft of the reattachment point, no model does particularly well, with errors increasing with the total enthalpy of the flow. In Ref. [13] the authors examined whether a sophisticated model for the coupling between vibrational and reactive (i.e., dissociation) non-equilibrium [14] could bridge the gap between measurements and predictions, but found that it worked no better than Park’s simple model [15]. This lack of sensitivity of the results to more sophisticated coupling between vibrational and chemical non-equilibrium has been corroborated by other studies for the same LENS-XX double-cone experiments [11, 16, 17]. In addition, Ref. [13] corroborated the findings of Ref. [10] inasmuch that pressure on the fore-cone could be predicted quite well, but the heat-flux was under-predicted. The main consequence of changing the thermochemical model has been to change the size of the separation zone [11, 12, 16, 17]. The authors in Ref. [11] used thermochemical models whose parameter values were computed from more fundamental simulations using forced harmonic oscillators and quasi-classical trajectory (QCT) calculations, and obtained a very good comparison with the separation zone size for Case 4; however, they did not show the performance of their model on the other experiments conducted in LENS-XX. However, researchers have also studied whether catalytic effects could explain the disagreement between measurements and predictions [18]. They found that even in the highest-enthalpy case (Case 4, see Table 1), catalytic effects were minor in the separation zone, though they increased the heat-flux predictions aft of the reattachment point.

The six LENS-XX experiments have also been the subject of a “blind” modeling comparison, described in Ref. [19]. Here, five different CFD researchers attempted to model the experiments, with very different results. They disagreed among themselves despite the fact that the constituent models were similar, and they did not agree with the measurements either. The simulations were mesh-converged, so the different outcomes of the modeling effort were attributed to small differences in their numerical methods. All models under-predicted the size of the separation zone, as well as the heat-flux on the double-cone surface. One possible cause of the differences between the five CFD modeling efforts could be the role played by state (or slope) limiters used by the researchers, which could lead to excessive numerical dissipation. This numerical issue was explored in Ref. [18], which found that the effect of numerical viscosity was too small to bridge the gap between model predictions and experimental measurements. All of these studies ignored the effect of uncertainties in the inflow conditions, though Ref. [10] did check whether uncertainties/errors in measurement could be used to explain the mismatch between model predictions and experiments.

Table 1 Freestream conditions and their uncertainties (“Error”). Run 35 was conducted in the LENS-I tunnel and is described in Refs. [1, 24]. Cases 1 and 4 were conducted in LENS-XX and are described in Ref. [10].

Test	U_∞ (m/s)	ρ_∞ (g/m ³)	T_∞ (K)	T_{v_∞} (K)	Re (m ⁻¹)	M_∞	h_0 (MJ/kg)	P_{Pitot} (kPa)	Gas
Error	($\pm 3\%$)	($\pm 7\%$)	($\pm 3\%$)			($\pm 3\%$)	($\pm 5\%$)	($\pm 5\%$)	
Run 35	2545	0.5848	98.27	2562	14.3×10^4	12.59	3.71	3.55	N ₂
Case 1	3246	0.499	175	175	0.14×10^6	12.2	5.44	5.1	Air
Case 4	6479	0.984	652	652	0.20×10^6	12.82	21.77	39.5	Air

Uncertainty quantification and double-cone experiments: The difficulties with the LENS-XX experiments match those encountered with LENS-I, and it is natural to wonder whether a mis-specified inflow condition could again be the cause. With this in mind, the authors of Ref. [2] performed a forward propagation of input uncertainties, as described in Sec. II. The aim was to test whether the distribution of pressure and heat-flux predictions from the ensemble bracket the experimental measurements. They could not, but a $\pm 15\%$ perturbation of the inflow condition at the test-section managed to “cover” the experimental data. This raised the possibility that modifying the inflow conditions, outside of the limits placed by the stated experimental uncertainty bounds, could allow CFD simulations to reproduce experimental measurements.

The most straightforward way to test this possibility would be to simulate the flow in the LENS-XX experiments, exactly as done for LENS-I in Ref. [1], and examine the conditions at the test-section inlet. However, the geometry of the LENS-XX nozzle is not public. Instead, in Ref. [6], the authors developed an inference framework, targeted the three experiments listed in Table 1, and inferred the inflow conditions (density, velocity, temperature and vibrational temperature) based on them. The authors found that the MAP estimates for density and velocity were well within the stated experimental error bounds for Run 35, but not for Case 1 and Case 4. They also found that the measurements used in the inference (surface pressure on the double-cone $p(x)$, surface heat-flux $q(x)$, the total enthalpy of the flow h_0 and the Pitot pressure P_{Pitot}) carried very little information on the inflow temperature of the flow which, consequently, could not be estimated. They also investigated whether shortcomings of the thermochemical models could explain the model versus experiment disagreements and found weak evidence for it. However, when the “correct” inflow conditions, as inferred from the fore-cone measurements, were used in CFD predictions, the disagreement between predictions and measurements remained downstream of the separation point, implying that while Bayesian inference found some of the causes of modeling errors, it was far from a comprehensive solution.

Previous uncertainty quantification studies of double-cone experiments performed by (some of) the authors of the current paper [2, 6], used SPARC, a Navier-Stokes simulator that is described in detail in Ref. [6]. It is a second-order accurate finite-volume compressible flow simulator. It has been formulated for reacting gases in thermochemical non-equilibrium. It can accommodate multiple gaseous species, per a chemical mechanism, and evolves them separately. In this paper, we use a 5-species mechanism for air, with 17 reactions [20]. SPARC accommodates structured and unstructured meshes, though we only use the former for double-cone simulations. A MUSCL scheme [21] and a minmod limiter is used for stability for shock-laden flows. SPARC solves the unsteady form of the governing equations, and we simulate over a long period of time, at first-order accuracy in time, to reach steady-state (as required in this paper). Grid resolution [22] and solution verification [23] studies have been performed on the models in SPARC.

B. Adjoint-Based Inverse Solvers

Adjoint-based sensitivity analysis is used to compute gradients for partial differential equation-constrained optimization problems (strictly, they can also be used for problems governed by algebraic systems), including design and inverse problems. Adjoint sensitivity analysis has been used for solving inverse problems for a wide range of applications, including fluid dynamics of turbo-machinery components dynamics [25]. For hypersonic CFD, adjoint-based sensitivity analysis has been [26], seismic wave propagation [27], and ice sheet used for design optimization, adaptation, and uncertainty propagation, but not to solve inverse problems to the author’s knowledge.

NASA has two CFD codes which support adjoint-based design optimization and mesh adaptation for inviscid and viscous hypersonic flows. CART3D, a Cartesian grid inviscid flow solver, leverages the adjoint for optimization and mesh adaptation on a wide range of geometries and flight conditions [28]. CART3D has been applied to numerous inviscid hypersonic flows including those around reentry capsules [29] and asteroids [30]. Another NASA solver,

FUN3D, has an adjoint capability that has been used for shape optimization and mesh adaptation on a wide range of viscous hypersonic flows, including non-equilibrium flows [31–33]. It is a discrete adjoint capability which leverages hand-differentiated Jacobians for verification and validation purposes [32].

The adjoint capability of Stanford University’s SU2 open-source CFD solver has been demonstrated on a range of inviscid and viscous hypersonic flows [34]. Initially, SU2 leveraged a continuous adjoint for mesh adaptation and shape design of reentry vehicles and scramjet inlets [35–37], but more recent work leverages a discrete adjoint capability [38, 39].

Outside of SU2 and NASA, recent work by Damm et al. demonstrated a discrete adjoint solver for scramjet inlet design [40]. Finally, Lockwood et al. developed and demonstrated a discrete adjoint to solve optimization problems as part of uncertainty propagation studies on various flow field properties in non-equilibrium hypersonic flows around bluff bodies [41–43].

IV. Sensitivity Analysis

In this section, we describe the effect of grid resolution on the pressure $p(x)$ and heat-flux $q(x)$ on the surface of the double-cone, as well as on their sensitivities S_α to perturbations in inlet conditions, $\alpha \in \{\rho_\infty, U_\infty, T_\infty, T_{v_\infty}\}$. We define the sensitivity S_α as

$$S_\alpha = \frac{\alpha}{Z} \frac{\partial Z}{\partial \alpha} \quad (1)$$

where $Z \in \{p(x), q(x)\}$ are two of four Quantities of Interest (QoI) in this paper. Thus, this is Local Sensitivity Analysis, with the gradients in Eq. 1 computed for the nominal test-section inlet conditions tabulated in Table 1. The method for computing these gradients is described in Appendix A. The other two QoIs are the total enthalpy h_0 and Pitot pressure P_{Pitot} of the incoming flow, both of which are independent of the mesh resolution. The flow is assumed to be axisymmetric, and we consider three different grids *viz.*, 128×256 (called the “Coarse” mesh), 256×512 (“Medium” mesh) and 512×1024 (“Fine” mesh). These are the same meshes which were used in the grid resolution studies in Ref. [22].

Fig. 2 (top row) plots the surface pressure $p(x)$ and its sensitivities on the left and the surface heat-flux (and its sensitivities) on the right. The computations are performed on the Fine mesh using the inlet conditions corresponding to Case 1. A 5-species 17-reaction chemical mechanism is used, and the model can accommodate flows in vibrational non-equilibrium. The point where the flow separates, at $x = 0.065\text{m}$ is denoted with a vertical line. The QoI is plotted with a solid line, and is represented by the y-axis; the sensitivities are plotted with dashed lines and are represented by the 2nd y-axis. We see that $p(x)$ is constant in the laminar, attached flow region, as are its sensitivities. In contrast, $q(x) \propto x^{-1/2}$, though its sensitivities are also constant upstream of the separation point. Downstream of the separation *i.e.*, of the vertical line, we see significant variations in the sensitivities of $p(x)$ and $q(x)$ primarily due to reattachment of the separated flow and its interactions with complex shock reflections. In the bottom row of the same figure, we plot the same QoIs and their sensitivities, but computed on the Coarse mesh. We see large differences in the sensitivities downstream of the separation point, but not upstream of it. This is not surprising since small changes in mesh resolution can change the location where flow or shocks impinge on a surface. In addition, we see that the sensitivity with respect to temperature, *i.e.*, S_{T_∞} , is very small compared to the sensitivities with respect to ρ_∞ and U_∞ . This is not very surprising either, given that most of the energy in the incoming hypersonic flow is kinetic, not thermal.

The figure leads us to two conclusions:

- 1) Since the sensitivity of $p(x)$ and $q(x)$ to T_∞ is so small, it will be infeasible to infer it (*i.e.*, T_∞) from surface measurements on the double-cone (an observation also borne out in Ref. [6]). Consequently, we will not include it as a quantity to be inferred, making our inverse problem a two-dimensional one, *i.e.*, we estimate (ρ_∞, U_∞) .
- 2) Given the immense changes in sensitivities downstream of the separation zone due to changes in grid resolution, our inverse problem will use only those model predictions that lie upstream of the separation region, *i.e.*, $x \leq 0.065\text{m}$ for Case 1. In addition, since $p(x)$ and $q(x)$ show peaked values downstream of the separation points, any comparisons with measurements will be susceptible to changes in grid resolution.

The computations in this paper are performed on the “Fine” mesh. The mesh is a stretched one (Fig. 1 (right)), with grid points clustered near the double-cone surface.

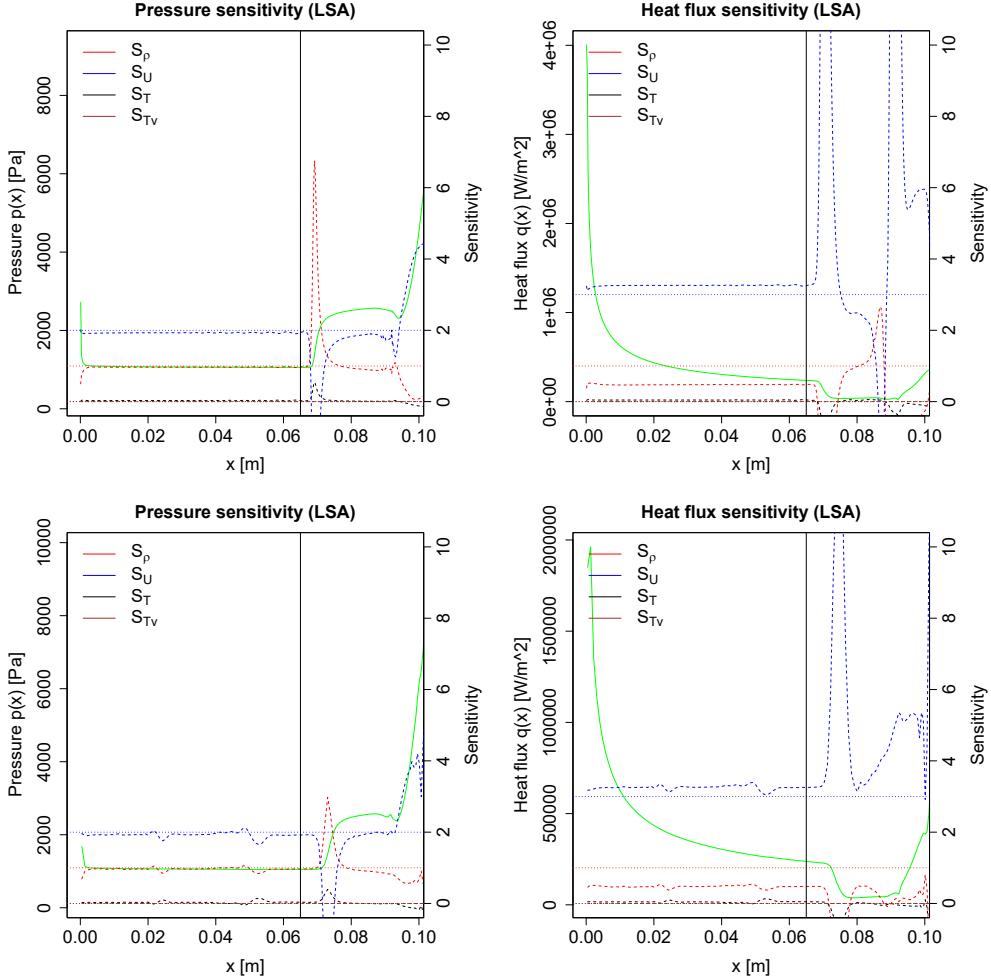


Fig. 2 Top: Predicted pressure $p(x)$ (left) and heat-flux $q(x)$ (right) for Case 1, plotted with a green line. Predictions were obtained on the Fine mesh. The point of flow separation is denoted with a vertical line. On the second y-axis we plot the sensitivity S_α of the relevant quantity of interest w.r.t. $(\rho_\infty, U_\infty, T_\infty, T_{v_\infty})$. **Bottom:** The same plots but computed on the Coarse mesh.

V. Formulation of the Inverse Problem

In this section, we pose the inverse problem for the flow density (ρ_∞) and velocity (U_∞) at the test-section inlet, conditional on pressure ($p(x)$) and heat-flux ($q(x)$) measurements, as well as the total enthalpy (h_0) and Pitot pressure (P_{Pitot}) of the incoming flow. Fig. 2 plots the pressure $p(x)$ and heat-flux $q(x)$ on the double-cone surface for Case 1. The point of flow separation is also plotted with a vertical line. We employ only those measurements that are upstream of the separation point, since the flow physics there is well understood. The probes (sensors) that collect measurements in this part of the flow are referred to as the “calibration” probes. Fig. 2 also shows that while the pressures at the calibration probes are constant, the heat-flux $q(x)$ shows a variation in space. In addition, $p(x)$ and $q(x)$ assume very different values, and need to be scaled down to the same order-of-magnitude to be usable in an inverse problem.

For pressure, the scaling s_p is defined as the reciprocal of the mean of all measured pressures \tilde{p}_i , $i = 1 \dots N_p$, where N_p is the number of pressure calibration probes. Henceforth a tilde ("~") denotes a measured quantity. Since the heat-flux $q(x)$ varies as $x^{-1/2}$ (see Chapter 9 in Ref. [44]), we compute a distance-weighted $Q(\mathbf{x}_i; u, y) = q(\mathbf{x}_i; u, y) \times \sqrt{x_i}$ which is approximately constant over the attached laminar flow over the fore-cone. The scaling s_Q is defined as the reciprocal of the mean of \tilde{Q}_i , $i = 1 \dots N_q$. N_q is the number of calibration probes for heat-flux measurements. The scalings s_h and s_P are the reciprocals of \tilde{h}_0 and \tilde{P}_{Pitot} . N_q and N_p vary for the three experiments in Table 1. N_p is 5 for Run 35, 3 for Case 1, and 4 for Case 4, while N_q is 9 for Run 35, 15 for Case 1, and 13 for Case 4.

Posing the optimization problem: Since the problems of interest are steady-state flows, we consider the following deterministic optimization problem,

$$\min_y g(u, y) = 0 \quad \text{s.t.} \quad f(u, y) = 0, \quad (2)$$

where $u \in \mathbb{R}^n$ is the discrete flow solution consisting of the unknown degrees-of-freedom across the mesh, $y \in \mathbb{R}^m$ is a set of model parameters (for us $y = \{\rho_\infty, U_\infty\}$), $f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ is the discretized residual, and $g : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^+$ is a scalar-valued objective function. While SPARC implements several discretization approaches, this work focuses exclusively on the cell-centered finite volume approach using block-structured meshes, in which case u contains the value of the degrees-of-freedom at each cell center (density of each chemical species, flow velocity, temperature, and molecular vibration temperature), and f is the finite volume residual. The objective function g quantifies the misfit between simulation quantities-of-interest and experimentally measured values:

$$g(u, y) = \sum_{i=1}^{N_q} [s_Q(Q(\mathbf{x}_i; u(y)) - \tilde{Q}_i)]^2 + \sum_{i=1}^{N_p} [s_p(p(\hat{\mathbf{x}}_i; u(y)) - \tilde{p}_i)]^2 + [s_h(h_0(y) - \tilde{h}_0)]^2 + [s_P(P_{Pitot}(y) - \tilde{P}_{Pitot})]^2. \quad (3)$$

Here $Q(\mathbf{x}_i; u(y))$ and $p(\hat{\mathbf{x}}_i; u(y))$ are the (distance-weighted) heat flux and pressure computed at a set of probe locations $\{\mathbf{x}_i\}$ and $\{\hat{\mathbf{x}}_i\}$, respectively (the probe locations are not, in general, the same), $h_0(y)$ is the total enthalpy of the flow, $P_{Pitot}(y)$ is the Pitot pressure of the flow, \tilde{Q}_i , \tilde{p}_i , \tilde{h}_0 , and \tilde{P}_{Pitot} are the corresponding measured values, and s_Q , s_p , s_h and s_P are scaling values chosen for each problem to equilibrate the contribution of each term to the overall objective function.

Equation (2) falls into the category of PDE-constrained optimization, for which numerous methods have been developed in the literature. In this work, we focused on the so-called reduced space approach that transforms Eq. (2) into an unconstrained problem

$$\min_y h(y), \quad h(y) = g(u(y), y) \quad \text{s.t.} \quad f(u(y), y) = 0 \quad (4)$$

where evaluation of $h(y)$ first involves solving the nonlinear equations $f(u(y)) = 0$ for u , then evaluating the original objective function $g(u(y))$. Under suitable differentiability assumptions for f and g , the corresponding reduced gradient can be computed as

$$\nabla h(y) = - \left(\frac{\partial f}{\partial y} \right)^T w + \left(\frac{\partial g}{\partial y} \right)^T \quad \text{s.t.} \quad \left(\frac{\partial f}{\partial u} \right)^T w = \left(\frac{\partial g}{\partial u} \right)^T \quad (5)$$

which involves solving a single linear system for the adjoint variable $w \in \mathbb{R}^n$. A variety of optimization approaches are available for solving Eq. (4). In this work, we use the truncated conjugate gradient method with trust-region globalization [45, Ch. 4] and BFGS Hessian approximation [45, Ch. 6] as provided by the ROL package [46].

Given a set of inflow conditions y , each step of the optimization method requires computing a new steady-state solution u satisfying $f(u, y) = 0$. Due to numerical difficulties in computing these steady-state solutions, SPARC employs implicit time integration methods, starting from an initial uniform flow, to compute them by solving $M\dot{u} + f(u, y) = 0$, where $M \in \mathbb{R}^{n \times n}$ is the mass matrix, until steady-state is reached with $\dot{u} = 0$. These equations are discretized using the first-order Backward Euler (BDF1) time discretization, with the resulting nonlinear equations solved at each time step using Newton's method based on an approximate Jacobian matrix that does not include second-order stencil terms nor viscous cross terms. Typically, these systems are only approximately solved by only applying a small, fixed number of Newton iterations, and thus is not time-accurate.

Computing sensitivities/gradients: Once a steady-state solution is computed, the objective function gradient is then given by Eq. (5). We employ Automatic Differentiation (AD) provided by the Sacado package [47–49] to analytically compute $\partial f / \partial y$, $\partial g / \partial y$, $(\partial f / \partial u)v$ and $(\partial g / \partial u)v$ for any vector $v \in \mathbb{R}^n$ without hand-coding. Using this capability, we then employ graph coloring [50–52] provided by the Zoltan package [53] to construct the true Jacobian $\partial f / \partial u$, which is then explicitly transposed to construct the adjoint operator $(\partial f / \partial u)^T$. Since the heat-flux and pressure computations involve a similar stencil as the finite volume residual, the same coloring procedure is used to construct $(\partial g / \partial u)^T$. While numerous solver strategies are available for solving Eq. (5), in this work we apply GMRES [54] using a block tri-diagonal/line implicit solver [55] applied to the transpose of SPARC's approximate Jacobian as a preconditioner. Due to ill-conditioning of the true adjoint matrix, we furthermore apply multiple steps of iterative refinement to obtain a solution with small residual. Details of the adjoint calculation are in Appendix B.

Estimation of the Gaussian posterior: Under the assumptions that the heat-flux, pressure, total enthalpy, and Pitot pressure at the calibrated values differ from the experimental measurements by additive, Gaussian noise, and their dependence on the calibrated parameters is approximately linear in the vicinity of the optimal parameter values, the posterior of the calibrated parameters is approximately Gaussian, and solution of Eq. (2) is equivalent to a maximum likelihood estimation procedure for the mean of the posterior. It is straightforward to show the covariance of the posterior is then given by the inverse Hessian of the negative log-likelihood, i.e., the objective function Eq. (3). Even if the posterior distribution is not a Gaussian, it can be approximated as one, provided the uncertainty is small (see Appendix A of Ref. [56] for a derivation, including the use of finite-difference Hessians to approximate non-Gaussian distributions as Gaussians). At the completion of the calibration process, we estimate the Hessian through a first order finite difference approximation, differencing the gradient Eq. (5) with a relative finite difference step size of 10^{-7} , and then compute its inverse through Gaussian elimination. The relative finite difference step size was determined empirically, with successively smaller step sizes until the Hessian became insensitive to it.

A discussion of the posterior distribution is incomplete without a prior. Our inversion scheme is unconstrained, as we suspect that $\{\rho_\infty, U_\infty\}$ may lie outside the bounds of the measurement errors specified by the experimental dataset, but have no principled way of deciding how large the deviation may be. In a Bayesian interpretation, one may state we specify a non-informative prior e.g., a uniform distribution with infinite bounds or a Jefferey's prior on the mean of the Gaussian posterior distribution.

VI. Results

A. Verification of the Inversion Method using Run 35

We first show the utility of our inversion technique, described in Sec. V, to estimate the inlet conditions of the test-section, using pressure and heat-flux measurements on the double-cone. Run 35 is used to test the utility of the inversion technique because (1) the flow has been successfully modeled [1], i.e., the stated experimental conditions can reproduce experimental measurements and (2) the experiment was conducted in an N_2 environment, which, along with its low total enthalpy (see Table 1), ensures that complications due to dissociated and reactive flow do not arise. Note that the flow conditions stated in Table 1 are at variance with the original paper where the experiment was first described [7]. The reason is as follows. The inflow conditions for Run 35 (also called Case D in Ref. [7]) were computed from shock tunnel measurements under the assumption that the flow was in equilibrium i.e., $T_\infty = T_{v_\infty}$. When the authors of Ref. [1] simulated the flow inside the nozzle of the LENS-I shock tunnel (where Run 35 was conducted), they found that the flow entering the test section was in vibrational non-equilibrium i.e., $T_\infty \neq T_{v_\infty}$, and frozen, i.e., the vibrational energy was partitioned out and did not play much of a role in the pressure and heat-flux measurements on the double-cone. The frozen flow inflow conditions are in Ref. [1] and are used in this paper, and the flow model assumes an ideal gas model. In addition, the surface temperature of the double-cone was maintained at 296.11 K, instead of 300 K for Cases 1 and 4.

This problem has also been solved using surrogate models and MCMC in Ref. [6], and we compare the exact posterior distribution from the MCMC solution versus the one obtained from deterministic inversion via Gaussian assumptions. In Fig. 3(top left), we plot the joint $(\rho_\infty/\rho_{norm}, U_\infty/U_{norm})$ PDF from the MCMC as gray symbols, with the Gaussian PDF from the deterministic inversion overlaid as contours. The MAP estimate (from the MCMC solution) and the optimal one (from the deterministic inversion) are also plotted. In Fig. 3 (bottom left) we plot the PDFs from the MCMC and Gaussian posteriors for the normalized inflow density. The Gaussian posterior PDF is too wide to fit in the figure and appears as a horizontal line. In Fig. 3 (top right) we plot the PDFs for the normalized inflow velocity and see the large discrepancy between the MCMC and Gaussian PDFs. The nominal inflow conditions (from Ref. [1]), the MAP values from MCMC and the optimal value from the deterministic inversion are also plotted and agree well.

In Table 2, we summarize the posterior PDFs and compare with the experimental specification. We see that mean and median values for the normalized inflow density and velocity, as computed using MCMC and our inversion capability, are very close to each other and certainly well within the experimental error bounds. The inter-quartile range (IQR) for U_∞/U_{norm} as computed using MCMC and our method are comparable, but the difference becomes very evident at the tail of the distribution, i.e., 90% CI, where the thin-tailed Gaussian distribution incurs a large error; the true (i.e., MCMC) posterior is far more compact. In the case of ρ_∞/ρ_{norm} , which is more difficult to infer, the IQR and 90% CI computed using our new method are too large to be useful.

We conclude that the deterministic inversion can infer inflow conditions correctly, i.e., within the experimental error bounds, and within the IQR and 90% CI computed from MCMC's posterior distribution, if the model is capable of computing the flow. The IQR computed from the Gaussian approximation of the posterior is comparable for inflow

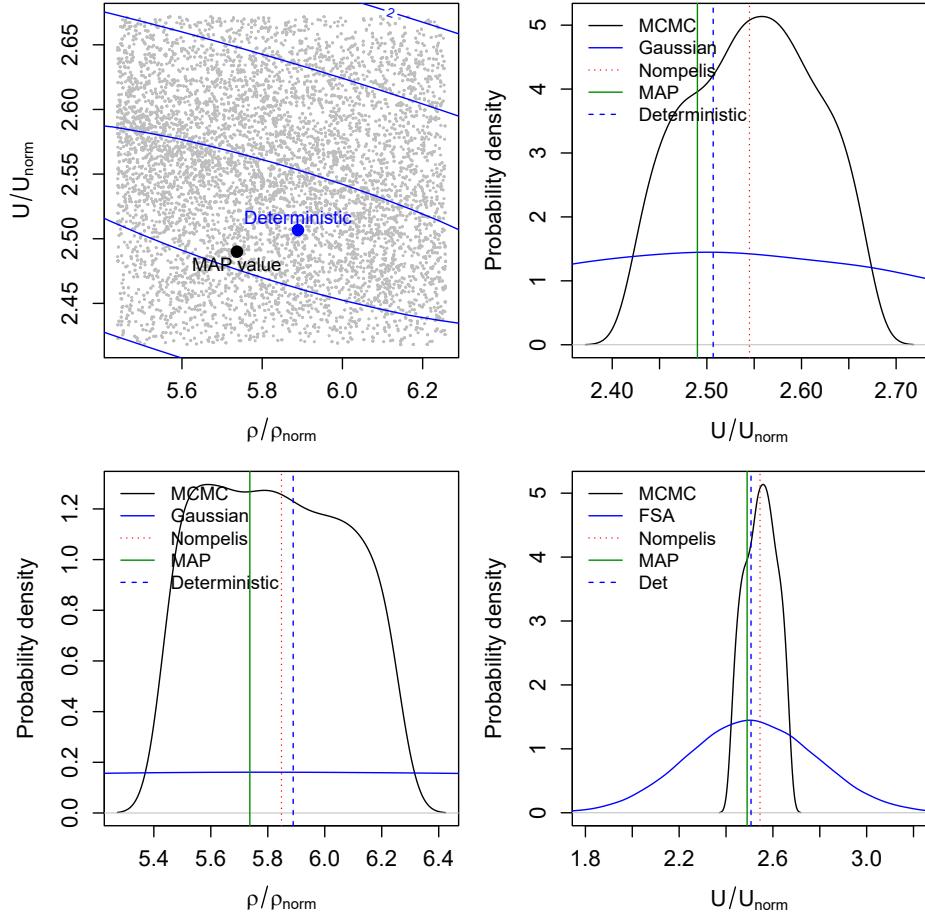


Fig. 3 Gaussian posterior PDFs from the deterministic inversion, compared to the ones from MCMC, for Run 35. Top left: Joint $(\rho_\infty/\rho_{norm}, U_\infty/U_{norm})$ PDF from MCMC plotted with gray symbols overlaid with contours from the Gaussian PDF. The MAP estimate and the Optimal from the deterministic inversion are also plotted. Bottom left: Marginal posterior PDFs for ρ_∞/ρ_{norm} , from MCMC and the Gaussian approximation from the deterministic inversion. The Gaussian approximation is too wide and appears as a straight line. The nominal/experimental specification (“Nompelis”; dotted vertical line), the MAP estimate (“MAP”; solid vertical line) and the optimal estimate (“Opt”; vertical dashed line) from deterministic inversion are also plotted, and are close. Top right: The PDFs from MCMC and deterministic inversion, for the normalized inflow velocity (zoomed in). Bottom right: The same as the figure above, but zoomed out to show the extent of the Gaussian posterior distribution.

quantities that can be inferred with confidence (U_∞/U_{norm} in our case), but the 90% CI is overly wide. For inflow quantities that are difficult to estimate from the double-cone measurements (ρ_∞/ρ_{norm} in our case), only the mean estimate is useful, with both the IQR and 90% CI being overly wide.

B. Case 1

Having tested our inversion capability on low-enthalpy Run 35, which had been previously modeled numerically [1], we proceed to Case 1. Case 1 is of moderate enthalpy (see Table 1), and we model it with a five-species, 17-reaction mechanism and assume vibrational non-equilibrium, i.e., the governing equations have two separate temperatures, one each for translational and vibrational energy. The flow at the inlet is assumed to be in vibrational equilibrium ($T_\infty = T_{v_\infty}$). The experiment was conducted in the LENS-XX expansion tunnel. An inversion for the inflow conditions has been performed for this experimental dataset, using Bayesian inference, surrogate models, and MCMC [6]. The study found that the inferred inflow density (MAP value) was within the stated experimental error bounds but the inferred inflow

Table 2 Summarization of posterior PDFs for ρ_∞/ρ_{norm} and U_∞/U_{norm} for Run 35. The experimental specification is also provided. The uncertainty in the experimental measurements is 7% for inflow density and 3% for inflow velocity. $\rho_{norm} = 1.0 \times 10^{-4} \text{ kg/m}^3$ and $U_{norm} = 1.0 \times 10^3 \text{ m/s}$. IQR stands for Inter-Quartile Range and CI for Credibility Interval.

Summary	ρ_∞/ρ_{norm}			U_∞/U_{norm}		
	MCMC	Deterministic	Experiment	MCMC	Deterministic	Experiment
Mean	5.83	5.89	5.848 (5.43, 6.26)	2.548	2.506	2.545 (2.47, 2.62)
Median	5.82	5.89		2.55	2.506	
IQR	(5.62, 6.03)	(4.24, 7.54)		(2.5, 2.6)	(2.32, 2.69)	
90% CI	(5.47, 6.21)	(1.88, 9.9)		(2.44, 2.65)	(2.06, 2.95)	

velocity coincided with the 90% CI.

In Fig. 4, we plot the joint $(\rho_\infty/\rho_{norm}, U_\infty/U_{norm})$ PDF from the MCMC solution in Ref. [6] (gray points), and compare it with the bivariate Gaussian distribution obtained using our deterministic inversion method. We see much the same behavior as in Sec. VI.A. In the joint distribution plot, we see that the MAP estimate from the MCMC solution and the optimal one from our new method are quite close. The Gaussian distribution for U_∞/U_{norm} is too wide compared to the PDF obtained from MCMC. In the case of ρ_∞/ρ_{norm} , the Gaussian distribution is too wide to be useful, whereas the PDF from the MCMC solution provides some bounds on the estimated inflow density.

It is in the summary of the posterior distributions, in Table 3, that the difference between the MCMC solution and our current inversion becomes clear. The optimal estimate for U_∞/U_{norm} is well outside the experimental bounds, as are the mean and median computed from the MCMC solution. The optimal ρ_∞/ρ_{norm} computed using our deterministic inversion is also outside the experimental error bounds, though the mean and median from the MCMC solution falls inside them. The IQR and 90% CI from the Gaussian approximation, for ρ_∞/ρ_{norm} , are too wide compared to the MCMC solution, and are unphysical (too low a density); this is similar to Sec. VI.A. For U_∞/U_{norm} , the IQRs are somewhat comparable but the 90% CI from the Gaussian posterior is too wide.

In Fig. 5 we plot the pressure and heat-flux predictions using the nominal and estimated inflow conditions. The measured values are also plotted. It is clear, and unsurprising, that the estimated inflow conditions accurately predict the pressure and heat-flux at the calibration probes. However, when comparing with all the probes, the nominal inflow conditions provide a better agreement. This is best captured by the root mean square error (RMSE) of the predictions (versus the measured values) provided in Table 4. The predictions at the “calibration” probes show unequivocal improvement when computed with the estimated inflow conditions, while the opposite is true when all the probes are considered (most of the probes lie on the aft-cone where complex shock interactions, flow separation and re-attachment take place). Since the flow over the fore-cone is simple (attached and laminar), this points towards serious model-form error. This is corroborated by Fig. 5 (left) which shows the large discrepancy between predicted and measured pressure, regardless of the inflow conditions used for computing the predictions.

To conclude, both the MCMC and our deterministic inversion suggest that the experimental specification of the inflow velocity may be erroneous, i.e., the flow at the inlet is faster than stated and lies outside the experimentally specified range. The inference of the inflow density is rather unclear, as MCMC and the deterministic inversion provide opposite suggestions, with our current estimate lying well outside the experimental measurement error bounds. There is significant model-form error and the predictions on the aft-cone show large errors, regardless of the inflow conditions used to compute them.

C. Case 4

Finally, we address the high-enthalpy Case 4 (see Table 1 for the flow conditions). The flow is modeled with a 5-species, 17-reaction mechanism and is assumed to be in vibrational and reactive non-equilibrium. The inflow is assumed to be in equilibrium i.e., $T_\infty = T_{v_\infty}$. The flow has extensive regions of vibrational and chemical non-equilibrium in the vicinity of the double-cone [10]. This experimental dataset has also been studied using Bayesian inference in Ref. [6], which found that the most probable value of both the inflow density and velocity (i.e., the MAP estimate) lay outside the experimental error bounds supplied with the dataset. However, since the Bayesian analysis was performed

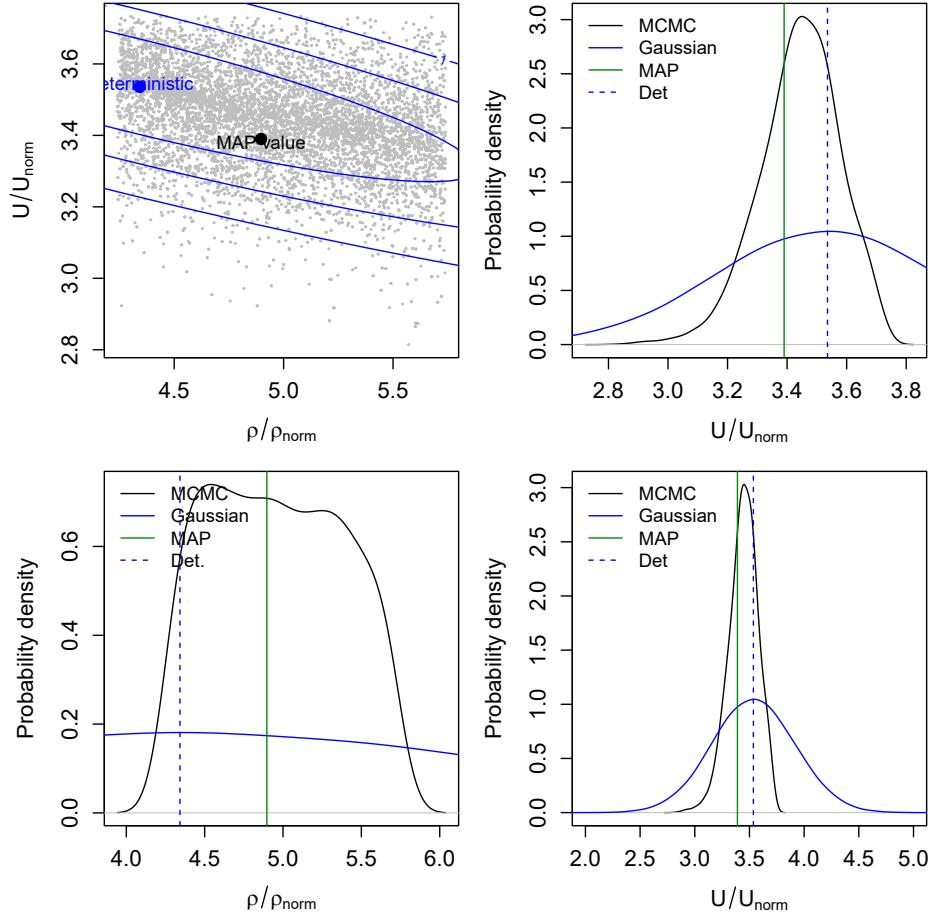


Fig. 4 Gaussian posterior PDFs from the deterministic inversion, compared to the ones from MCMC, for Case 1. Top left: Joint $(\rho_\infty/\rho_{\text{norm}}, U_\infty/U_{\text{norm}})$ PDF from MCMC plotted with gray symbols overlaid with contours from the Gaussian PDF. The MAP estimate and the Optimal from the deterministic inversion are also plotted. Bottom left: Marginal posterior PDFs for $\rho_\infty/\rho_{\text{norm}}$, from MCMC and the Gaussian approximation from the deterministic inversion. The Gaussian approximation is too wide and appears as an approximately horizontal line. The MAP estimate (“MAP”; solid vertical line) and the optimal estimate (“Opt”; vertical dashed line) from deterministic inversion are also plotted, and are close. Top right: The PDFs from MCMC and deterministic inversion, for the normalized inflow velocity. Bottom right: The same as the figure above, but zoomed out to show the extent of the Gaussian posterior distribution.

using surrogate models, we repeat it here using exact, if deterministic inversion, with an approximate Gaussian posterior distribution

In Fig. 6 we plot the joint PDF of $(\rho_\infty/\rho_{\text{norm}}, U_\infty/U_{\text{norm}})$ from the MCMC as gray sample points. The bivariate Gaussian distribution from the deterministic inversion is plotted with contours. We see that the optimal point from the deterministic inversion is close to the MAP value (Fig. 6 (top left)). The marginal distribution from MCMC for U_∞/U_{norm} bears some resemblance to the Gaussian approximation from our deterministic inversion (Fig. 6 (top right)), and the mean is very close to the MAP value. In Fig. 6 (bottom left) we plot the marginal density of $\rho_\infty/\rho_{\text{norm}}$ from MCMC, as well as the Gaussian approximation. The MAP value is close to the mean of the Gaussian, but the Gaussian is too wide. Comparing with Case 1 (Sec. VI.B), we see that the Gaussian approximation resembles the MCMC solution much better in Case 4.

In Table 5, we summarize the posterior distributions from MCMC and deterministic inversion and compare with experimental results. *In case of U_∞/U_{norm} , the mean from MCMC and deterministic inversion, as well as the MAP value from MCMC (which is very close to the mean of the Gaussian posterior) lie outside the experimental error bounds.*

Table 3 Summarization of posterior PDFs for ρ_∞/ρ_{norm} and U_∞/U_{norm} for Case 1. The experimental specification is also provided. The uncertainty in the experimental measurements is 7% for inflow density and 3% for inflow velocity. $\rho_{norm} = 1.0 \times 10^{-4} \text{ kg/m}^3$ and $U_{norm} = 1.0 \times 10^3 \text{ m/s}$. IQR stands for Inter-Quartile Range and CI for Credibility Interval. Note that the MAP value for inflow velocity in the MCMC solution is different from the mean, and is actually quite close to the mean of the Gaussian approximation.

Summary	ρ_∞/ρ_{norm}			U_∞/U_{norm}		
	MCMC	Deterministic	Experiment	MCMC	Deterministic	Experiment
Mean	4.96	4.33	4.999 (4.64, 5.34)	3.44	3.54	3.246 (3.15, 3.34)
Median	4.95	4.33		3.45	3.54	
IQR	(4.6, 5.31)	(2.86, 5.79)		(3.36, 3.53)	(3.28, 3.79)	
90% CI	(4.33, 5.64)	(0.73, 7.96)		(3.21, 3.65)	(2.92, 4.16)	

This also holds true for ρ_∞/ρ_{norm} , indicating that the inferred inflow conditions do not agree with the stated ones in the experimental dataset. In this respect, this outcome tallies with the one from the previous work [6]. Comparing the Gaussian marginals with their MCMC counterparts, we see that the IQR computed from the Gaussian approximation of U_∞/U_{norm} is somewhat similar to that from MCMC, but the 90% CI is too wide. For ρ_∞/ρ_{norm} , the Gaussian is simply too wide to be useful. Again, we see that variables that can be inferred with moderate uncertainty (U_∞/U_{norm} in our case) will admit a useful Gaussian approximation.

D. Discussion

The results above show that our deterministic inversion method infers inflow conditions (the mean of the Gaussian posterior distribution) that matches the inferences drawn with MCMC in Ref. [6]. For Run 35, which was conducted in the LENS-I shock tunnel and has been modeled successfully, the inferred inflow conditions lie within the experimental error bounds of the inflow specification that accompanied the dataset. For the LENS-XX experiments, which have yet to be modeled with good accuracy, our inversion method leads us to believe that the inflow conditions may have been mis-specified in both experimental datasets. The Bayesian analysis in Ref. [6] reaches much the same conclusion for Case 4, but the opposite one for Case 1. However, this conclusion is dependent on the accuracy of the surrogate models of the Navier-Stokes flow simulator, which could be the reason for this disagreement. In contrast, our method uses the Navier-Stokes simulator natively. The studies above also show that the Gaussian approximation of the posterior distribution is not very accurate. When a quantity can be inferred with some confidence, e.g., U_∞/U_{norm} , the approximate Gaussian posterior can provide a useful IQR (in comparison to MCMC), but the tails of the distribution e.g., 90% CI, are always too large to be useful. The usefulness of a Gaussian approximation to the posterior distribution when the uncertainty in a quantity is small is well-known (see Ref. [56], Appendix A) and its applicability can be monitored, without access to the true posterior distribution from MCMC, using the methods in Ref. [57]. Next, we apply these methods to check the appropriateness of the Gaussian posterior assumption for our problem.

Checking the Gaussian assumption: The check for the appropriateness of a Gaussian assumption for the posterior distribution is based on cross-validation arguments. The inflow estimates discussed above were drawn from $N_p + N_q + 2$ observations, corresponding to N_p calibration probes for pressure, N_q for the heat-flux, the total enthalpy h_0 and Pitot pressure P_{Pitot} . Consider estimating inflow conditions (ρ^*, U^*) using a subset of the observations (in our case, 75% of

Table 4 Root-mean-square error (RMSE) for predictions of pressure and heat flux, using the nominal and estimated inflow conditions. The optimized inflow conditions did not improve predictions at all probes.

Variable	Calibration probes		All probes	
	Nominal	Inferred	Nominal	Inferred
Pressure (N/m ²)	2.6×10^1	1.46×10^1	1.8×10^3	1.9×10^3
Heat flux (W/m ²)	7.6×10^4	1.25×10^4	8.68×10^4	1.5×10^5

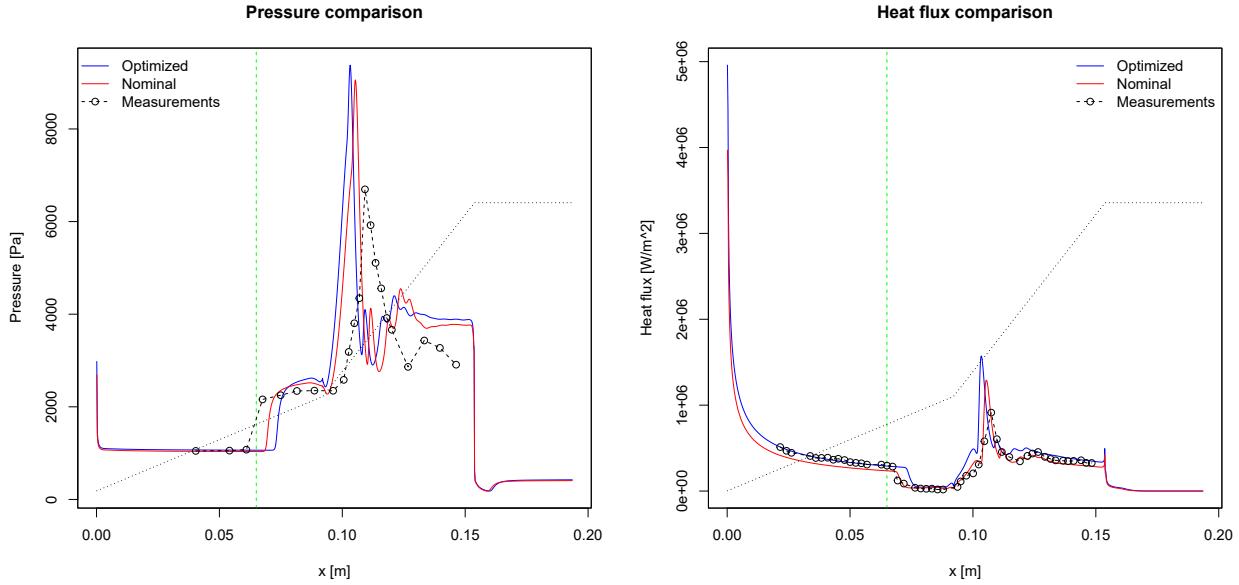


Fig. 5 Predictions of pressure (left) and heat flux (right) using the optimized (blue) and nominal (red) values of (ρ_∞, U_∞) for Case 1. The measured values at the probes are also plotted. Calibration does not improve the predictive skill of the model for the full flow, though the predictions are far better for the region used to estimate the inflow conditions.

the observations), selected randomly. This can be repeated K times, using a different, randomly selected subset of the observations each time, to yield K different set of inflow estimates, $(\rho^*, U^*)_k, k = 1 \dots K$. The scatter in $(\rho^*, U^*)_k$ reflects the uncertainty in the estimates due to the limited nature of the observations from which they were drawn. In addition, the scatter will be larger (wider) than the true uncertainty distribution which could be computed using Bayesian inference (as we did in Ref. [6]), as the Bayesian computation would use the full set of $N_p + N_q + 2$ observations. If the Gaussian posterior is far wider than the scatter in $(\rho^*, U^*)_k$, it would be evidence that the Gaussian assumption for the posterior distribution is inappropriate.

Note that for our problem, while we have N_p and N_q calibration probes, they do not carry independent information. The pressure and heat-flux measurements can be reduced to two constants (modulo measurement error) under a self-similar transformation (see Ref. [6]). Thus, we have 4 independent sources of observations, corresponding to pressure, heat-flux, h_0 and P_{Pitot} . We perform a 4-way cross-validation (i.e., $K = 4$) by selecting 3 measurement modalities at a time and computing $(\rho^*, U^*)_k, k = 1 \dots 4$. In Fig. 7, we plot the inflow estimates $(\rho^*, U^*)_k$ as symbols.

Table 5 Summarization of posterior PDFs for $\rho_\infty/\rho_{\text{norm}}$ and U_∞/U_{norm} for Case 4. The experimental specification is also provided. The uncertainty in the experimental measurements is 7% for inflow density and 3% for inflow velocity. $\rho_{\text{norm}} = 1.0 \times 10^{-4} \text{ kg/m}^3$ and $U_{\text{norm}} = 1.0 \times 10^3 \text{ m/s}$. IQR stands for Inter-Quartile Range and CI for Credibility Interval. Note that the MAP value for inflow velocity in the MCMC solution is different from the mean, and is actually quite close to the mean of the Gaussian approximation.

Summary	$\rho_\infty/\rho_{\text{norm}}$			U_∞/U_{norm}		
	MCMC	Deterministic	Experiment	MCMC	Deterministic	Experiment
Mean	9.186	8.619	9.84 (9.15, 10.52)	6.8	6.95	6.479 (6.23, 6.67)
Median	9.169	8.619		6.834	6.95	
IQR	(8.57, 9.81)	(6.29, 10.94)		(6.67, 6.96)	(6.5, 7.36)	
90% CI	(8.0, 10.4)	(2.97, 14.3)		(6.38, 7.09)	(5.93, 7.96)	

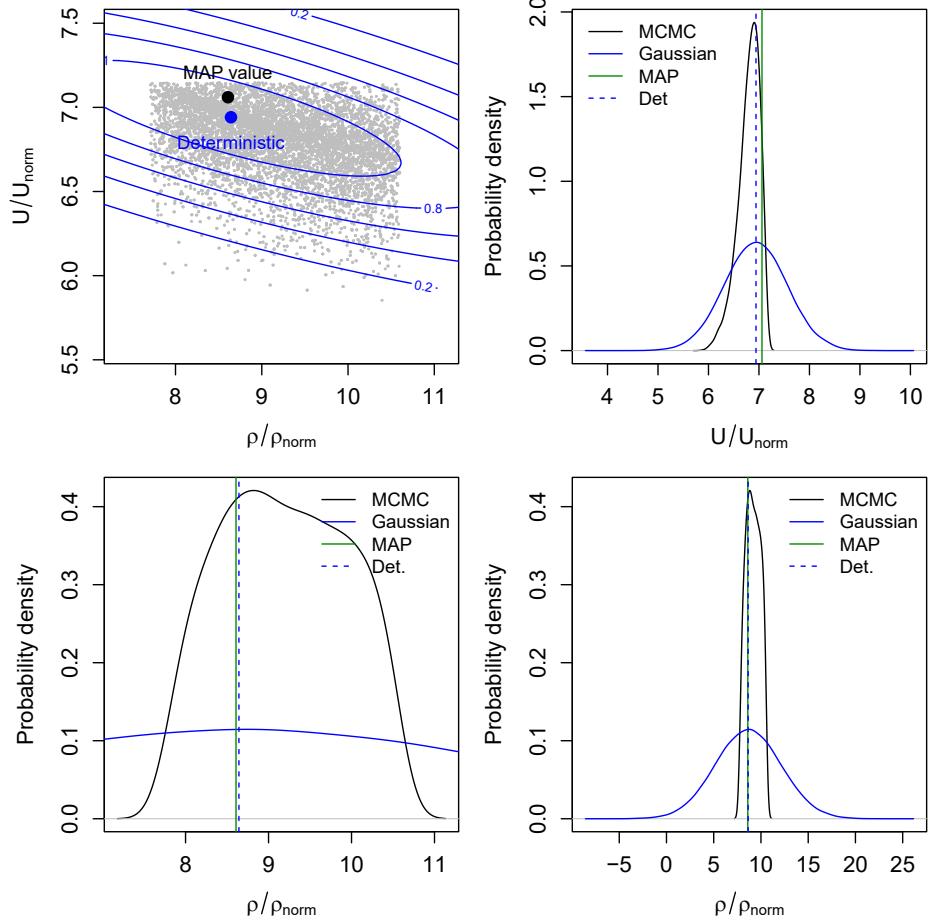


Fig. 6 Gaussian posterior PDFs from the deterministic inversion, compared to the ones from MCMC, for Case 4. Top left: Joint $(\rho_\infty/\rho_{norm}, U_\infty/U_{norm})$ PDF from MCMC plotted with gray symbols overlaid with contours from the Gaussian PDF. The MAP estimate and the Optimal from the deterministic inversion are also plotted. Bottom left: Marginal posterior PDFs for ρ_∞/ρ_{norm} , from MCMC and the Gaussian approximation from the deterministic inversion. The Gaussian approximation is too wide and appears as an approximately horizontal line. The MAP estimate (“MAP”; solid vertical line) and the optimal estimate (“Opt”; vertical dashed line) from deterministic inversion are also plotted, and are close. Top right: The PDFs from MCMC and deterministic inversion, for the normalized inflow velocity. Bottom right: The same as the figure to the left, but zoomed out to show the extent of the Gaussian posterior distribution.

These are computed for Run 35. The ρ^* estimates are plotted in the left column and the U^* in the right column. The marginalized posterior distributions for ρ_∞ and U_∞ are plotted as box-and-whisker plots, for the MCMC solution from Ref. [6] on top and the Gaussian posterior below. We see that ρ^* and U^* both show scatters that are wider than the IQR for MCMC posterior distribution (top row of Fig. 7). In contrast, the bottom row, which contains the distribution obtained under a Gaussian assumption, shows that ρ_k^* and U_k^* are clustered around the median, and are considerably narrower than the IQR. While four $(\rho^*, U^*)_k$ samples cannot adequately represent a distribution, it is clear that the Gaussian posteriors are far wider than the scatter in $(\rho^*, U^*)_k$, (qualitatively) suggesting that the use of Gaussian posteriors is inappropriate for this problem. While this was very evident from Fig. 3, Fig. 4 and Fig. 6, where the Gaussian posterior was compared with the one obtained using MCMC, the cross-validation test described above does not require one to compute an exact posterior. It also explains why only the Gaussian IQR for U_∞/U_{norm} was (marginally) useful, while the rest of the summary statistics were not. However, given the ease with which the Gaussian posterior can be computed and then checked for the appropriateness of the Gaussian assumption (via cross-validation), it is a convenient tool for quantifying the uncertainty in any estimates obtained using our deterministic inversion method.

Inversions on a coarse mesh: Our inversion approach can be used to check the consistency of an experimental dataset, e.g., to check whether certain measurements are consistent with specifications that accompany a dataset, e.g., the boundary conditions. This is especially powerful if the check can be performed with a subset of the measurements where the accuracy of the flow model is not in doubt — in our case, the attached laminar flow on the fore-cone. Bayesian inversion can be quite involved, as it requires the construction of surrogate models, as well as deciding on a stopping criterion for the MCMC procedure. In contrast, the inversion methodology described here can be encoded into a flow simulator and used widely; the basic theory is identical to optimizing an objective function. However, the inversion procedure can be quite computationally expensive and it is open to question whether the same quantities could be inferred on a coarser mesh, thus reducing computational cost.

In order to test this hypothesis, we repeated the inversion study described in Sec. VI.A on the Coarse and Medium meshes (see Sec. IV), and compare the marginalized posterior distributions for ρ_∞/ρ_{norm} and U_∞/U_{norm} in Fig. 8. We see that the differences between the distributions are minuscule. This is not entirely unexpected. The inferred inflow conditions are affected by the measurement errors in the surface pressure $p(x)$ and heat-flux $q(x)$, along with any errors in total enthalpy of the flow and Pitot pressure (which are mesh-independent). If these errors dominate the discretization errors introduced by an overly coarse mesh, inferences drawn quickly on a coarse mesh may provide an acceptable, if preliminary, diagnosis of the quality / consistency of an experimental dataset.

Constrained versus unconstrained inference The comparison between the Gaussian posterior and the one from the MCMC solution, as shown in Figs. 3, 4 and 6, is not entirely fair. The Bayesian inference employed a uniform distribution $\mathcal{U}(a, b)$ to serve as an informative prior for ρ_∞ and U_∞ , which constrained the parameter-space that was explored. This constraint arose from practical considerations as training data spanning a finite domain in the parameter-space had to be generated before a surrogate model could be created. In contrast, our deterministic inversion was unconstrained and did not require any knowledge of a bounded domain where the optimal value of (ρ_∞, U_∞) lay. In Bayesian parlance, therefore, $b = -a = \infty$ and the deterministic inversion starts with less information. If any prior information, in the form of bounds on ρ_∞ and U_∞ are known *a priori*, the Gaussian posterior can be truncated accordingly. It would certainly render the posterior more specific.

VII. Conclusion

In this paper, we have developed a deterministic inversion method to check the consistency of an experimental dataset. As we begin to validate models against challenging test cases, with experimental data acquired in extreme environments, the possibility of measurement errors creeping into the model validation dataset becomes uncomfortably high. This is because the experimental apparatus is often operated at the edges of its operational envelope in these experiments, and their behavior under those conditions is not always known. In many cases, it may be possible to partition an experimental dataset into two and infer one from the other. If the inferred quantities lie within the experimental error bounds of the corresponding measurements, the dataset can be deemed consistent. This does require that the model be an accurate one.

This approach was used in Ref. [6] to analyse the LENSI (Run 35) and LENSI-XX double-cone datasets (Case 1 and Case 4) studied in this paper. They found that Run 35 was consistent, but Cases 1 and 4 were not. However, that study relied on (approximate) surrogate models of the Navier-Stokes flow model that we use. Our studies in Sec. VI, conducted using the deterministic inversion method that uses a Navier-Stokes model natively, largely bears out the older results. If the measurements on the double-cone are correct, then the inlet conditions need to be different, provided the inlet flow is axisymmetric. The measurements used in this study (as well as in Ref. [6]) are collected from the attached laminar-flow region of the fore-cone, where the flow physics are relatively well-known.

The deterministic inference method is formulated as an optimization problem whose objective function is given by the (scaled) sum-of-squares difference between the experimentally measured quantities and those predicted by the CFD simulator. The optimization method is solved using a gradient-based approach, where the gradient of the objective function is computed using adjoint sensitivities facilitated by differentiation of the CFD solver using automatic differentiation. Since the inference method is deterministic, it only directly provides an estimate of the inferred parameters. However, a posterior distribution can be estimated under Gaussian assumptions with the inferred parameters as the mean of the distribution and the covariance given by the inverse of the Hessian of the objective function.

The mean of the posterior distribution obtained by our new method tallied with the MAP value, as well as the mean and median computed from the posterior density obtained from the Bayesian inverse solution in Ref. [6]. That held true for all three test cases in this paper. The Gaussian approximation to the posterior density was only marginally useful — in case of a quantity that could be inferred with some confidence (e.g., U_∞/U_{norm}), the inter-quartile range provided by the Gaussian approximation was comparable to the one from the Bayesian solution computed with MCMC (in Ref. [6]).

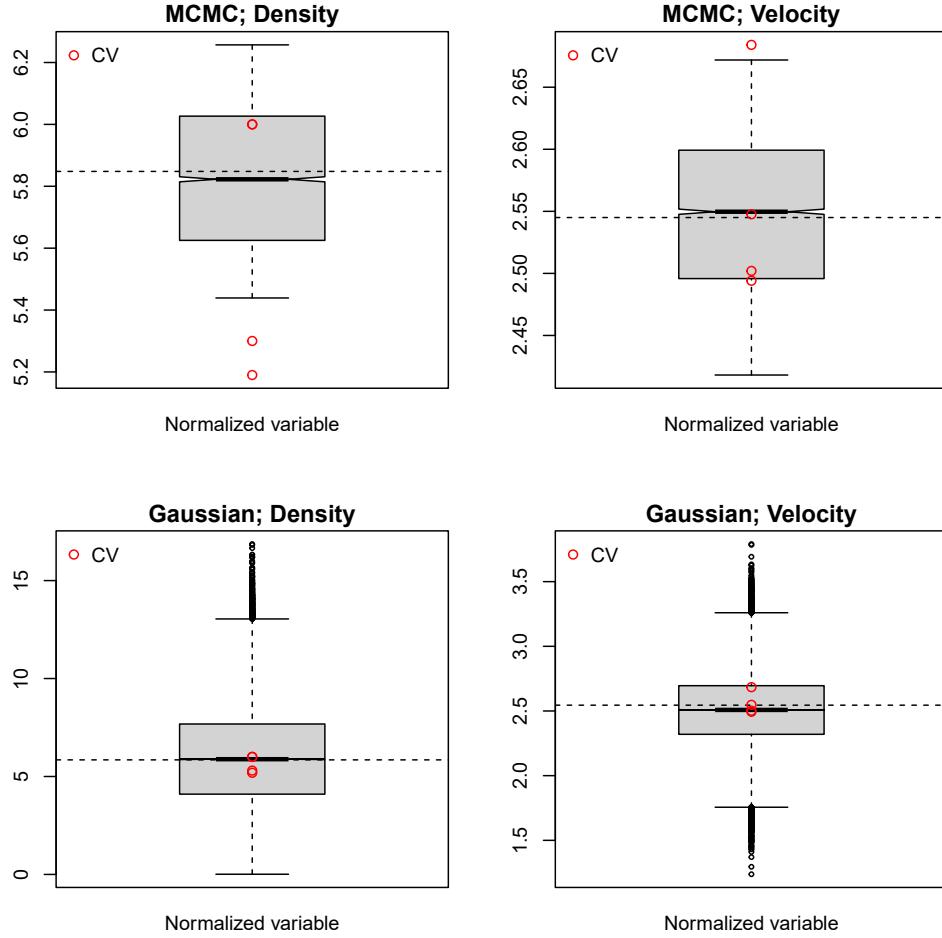


Fig. 7 The “cross-validation” (CV) estimates $(\rho^*, U^*)_k, k = 1 \dots 4$ for Run 35 plotted as symbols, with ρ_∞ on the left and U_∞ on the right. The posterior distribution from a Bayesian inference of (ρ_∞, U_∞) (from Ref. [6]) are plotted as box-and-whisker plots on the top row. The bottom row has the posterior distributions obtained with a Gaussian assumption. It is clear that the Gaussian posterior is overly wide (compare the vertical axes of the Bayesian and deterministic estimates).

The tail of the Gaussian distribution was too wide. For a quantity that was more difficult to estimate e.g., ρ_∞/ρ_{norm} , the Gaussian approximation to the posterior was not helpful.

We find evidence that the inlet conditions for Case 1 could be different from what was stated in the experimental dataset. This conclusion is conditional on the model (Park’s, [15]) used for vibrational non-equilibrium, but studies have shown that using a more sophisticated model does not improve predictions materially [11, 13, 16]. Case 4 experiences vibrational and chemical non-equilibrium, and detailed studies have shown that existing thermochemistry models do not reproduce measurements well [10–12]. The inlet conditions, as inferred by us, are definitely outside the range of stated experimental error.

This raises the question of what could be the cause of the discrepancy between model predictions and measurements. All CFD studies to date have assumed that the inlet flow is axisymmetric and steady, but the experimental dataset available to us does not allow us to test an alternative hypothesis. More research is definitely warranted, as is more experimental data. Repeated runs, that would allow us to assess the run-to-run variability and measurements on different sides of the double-cone, that would allow us to test the assumption of axisymmetric flow, would most likely top a modeler’s wish-list. One interesting study [58] looked at the effect of non-equilibrium flow *at the inlet* and found that it could increase the size of the separation bubble, without materially changing the heat-flux and pressure profiles on

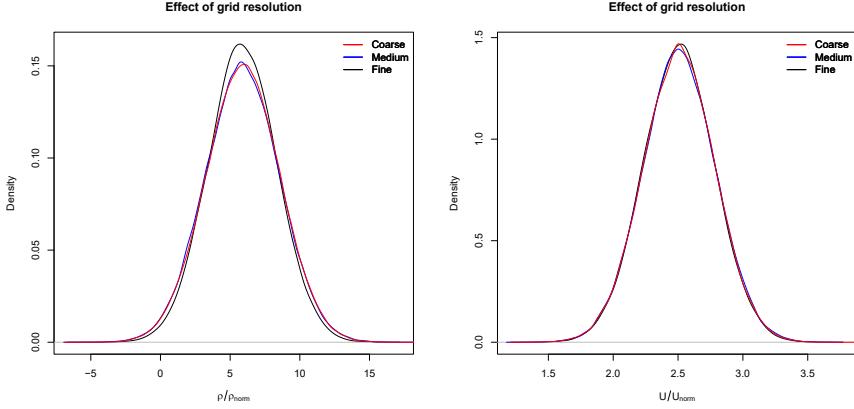


Fig. 8 Marginalized posterior distributions for $\rho_\infty/\rho_{\text{norm}}$ (left) and U_∞/U_{norm} (right), for Run 35 (Sec. VI.A), computed on a Coarse (128×256), Medium (256×512) and Fine (512×1024) mesh. There is barely any difference in the posterior distributions.

the fore-cone. However, the study was performed for an O_2 environment in the LENS-I *shock* tunnel; in contrast, the LENS-XX *expansion* tunnel is designed to suppress non-equilibrium processes upstream of the test-section inlet.

A. Forward Sensitivities

Referring to (2), the double-cone problems of interest are steady-state flows given by solving $f(u, y) = 0$. However, due to numerical difficulties in computing these steady-state solutions, SPARC employs implicit time integration methods, starting from an initial uniform flow, to compute them by solving numerically the following coupled system of ordinary differential equations:

$$M\dot{u} + f(u, y) = 0, \quad (6)$$

where $M \in \mathbb{R}^{n \times n}$ is the mass matrix. For computing steady-state solutions, SPARC employs the Backward Euler/BDF1 time discretization, resulting in the following nonlinear equations that must be solved at each time step k :

$$\tilde{f}_k(u_{k+1}, y) \equiv \frac{1}{\Delta t_k} (Mu_{k+1} - Mu_k) + f(u_{k+1}, y) = 0, \quad k = 0, 1, \dots \quad (7)$$

The above time step equations are solved using Newton's method

$$\left(\frac{1}{\Delta t_k} M + J(u_{k+1}^l, y) \right) \Delta u_{k+1}^l = -\tilde{f}_k(u_{k+1}^l, y), \quad u_{k+1}^{l+1} = u_{k+1}^l + \Delta u_{k+1}^l, \quad l = 0, 1, \dots, \quad (8)$$

where $u_{k+1}^0 = u_k$ and J is an approximation to the Jacobian matrix $\partial f / \partial u$ that does not include second-order stencil terms nor viscous cross terms. For steady-state solutions, (7) is only approximately solved by applying a small, fixed number of Newton iterations (often one) in (8), and thus is not time-accurate. We thus refer to this approach as pseudo-transient time stepping as it is essentially a form of pseudo-transient continuation [59].

SPARC has numerous linear solvers available for solving (8), such as fixed-point solvers, preconditioned Krylov solvers, and sparse-direct solvers. For the problems of interest here, SPARC employs a small number of fixed-point iterations using its block tri-diagonal/line implicit solver [55]. Of relevance to the adjoint approach described later, SPARC can also apply GMRES to (8) using exact Jacobian-vector products provided by the tangent forward AD mode implemented with Sacado, preconditioned by a small number of fixed-point iterations of the block tri-diagonal solver. Finally, to accelerate convergence to steady-state, the time step sizes Δt are gradually increased over the course of the time integration. Ultimately, the time integration stops when the norm of the nonlinear residual $\|f(u_k, y)\|_2$ becomes sufficiently small, or a maximum number of time steps is exceeded.

SPARC supports computing sensitivities of the flow solution u with respect to model parameters such as inflow boundary conditions, which was used to generate the sensitivities of pressure and heat-flux as shown in Fig. 2. This is

done by differentiating (6) with respect to y ,

$$M\dot{Z} + \frac{\partial f}{\partial u}(u, y)Z + \frac{\partial f}{\partial y} = 0 \quad (9)$$

where $Z \equiv \partial u / \partial y \in \mathbb{R}^{n \times m}$. After BDF1 time discretization, this becomes

$$F_k(Z_{k+1}, u_{k+1}, y) \equiv \frac{1}{\Delta t_k}(MZ_{k+1} - MZ_k) + \frac{\partial f}{\partial u}(u_{k+1}, y)Z_{k+1} + \frac{\partial f}{\partial y}(u_{k+1}, y) = 0. \quad (10)$$

This is implemented in SPARC by directly modifying its time integration implementation to also compute $\{Z_k\}$, where the relevant partial derivatives are computed by Sacado. Note that by using the tangent forward mode of AD, $(\partial f / \partial u)Z + \partial f / \partial y$ is computed directly without explicitly forming $\partial f / \partial u$, with a cost proportional to m . Furthermore, SPARC supports several approaches for solving (7) and (10). The method used here solves (7) and (10) together at each time step (equivalent to the simultaneous corrector method [60] in the literature). As with the forward solution, a small number of Newton iterations are used each time step with inexact linear solves using the SPARC native (approximate) Jacobian. The corresponding Newton system is

$$\left(\frac{1}{\Delta t_k} M + J(u_{k+1}^l, y) \right) \Delta Z_{k+1}^l = -F_k(Z_{k+1}^l, u_{k+1}^l, y), \quad Z_{k+1}^{l+1} = Z_{k+1}^l + \Delta Z_{k+1}^l, \quad l = 0, 1, \dots, \quad (11)$$

which is solved along with (8). Note that Newton matrix is the same in (8) and (11), and thus Δu_{k+1}^l and ΔZ_{k+1}^l are computed together by a single linear system solve with multiple right-hand-sides. Furthermore, we have found the same run-schedule for determining time-step sizes of the flow solution typically works well for the forward sensitivity system when using the combined approach, making the approach fairly robust. Finally, this approach provides the forward solution and its sensitivity with respect to m parameters in about $m + 1$ times the cost of the forward solution alone.

B. Adjoint Sensitivities

While the forward sensitivity approach described above amounts to a direct differentiation of the forward simulation time integrator, a different approach is required for the adjoint sensitivity. As shown in Eq. (5), the adjoint, w is governed by the following linear equation,

$$\left(\frac{\partial f}{\partial u}(u, y) \right)^T w = \left(\frac{\partial g}{\partial u}(u, y) \right)^T, \quad (12)$$

where u is the solution to $f(u, y) = 0$ in (2).

Due to severe ill-conditioning, directly solving (12) for w using SPARC's available linear solvers is unable to compute w with a sufficient level of accuracy. Instead, we employ a quasi-Newton approach which requires solving a sequence of linear systems of the form

$$J^T(u_\infty, y) \Delta w_{k+1}^l = - \left(\frac{\partial f}{\partial u}(u, y) \right)^T w_{k+1}^l + \left(\frac{\partial g}{\partial u}(u, y) \right)^T, \quad (13)$$

where J^T is the approximate adjoint operator computed from the approximate Jacobian used in SPARC's forward simulation. This approach, coupled with a line search-based Newton method from the NOX [61] nonlinear solver package and SPARC's block tri-diagonal solver, was found to be effective for the Run 35 and Case 1 versions of the double-cone problem. By further specifying use of the true adjoint as the Newton matrix, a true Newton's method is obtained:

$$\left(\frac{\partial f}{\partial u}(u, y) \right)^T \Delta w_{k+1}^l = - \left(\frac{\partial f}{\partial u}(u, y) \right)^T w_{k+1}^l + \left(\frac{\partial g}{\partial u}(u, y) \right)^T. \quad (14)$$

We found this approach to be effective for all double-cone problems, including Case 4, using at each Newton iteration GMRES preconditioned by the block tri-diagonal solver applied to the approximate adjoint J^T . In both cases, multiple Newton iterations were required to obtain a solution to the adjoint system with a small residual due to the poor conditioning of the true adjoint matrix. We note that Newton's method applied to a linear system is equivalent to iterative refinement, and when GMRES is used as the linear solver, it is equivalent to restarted-GMRES.

The various partial derivatives appearing in (14) are computed using automatic differentiation with Sacado which, as described above, is designed to differentiate the parallel CFD residual assembly in SPARC using the tangent forward

mode of AD. However, because of the transpose, the tangent forward mode of AD cannot be used to implicitly compute the adjoint residual $(\partial f / \partial u)^T w - (\partial g / \partial u)^T$. While reverse mode AD can be used to compute this implicitly without constructing $(\partial f / \partial u)^T$, and Sacado provides reverse mode AD tools, those tools are not able to differentiate the parallel computations implemented in SPARC’s residual assembly. Thus, we explicitly form $\partial f / \partial u$ and $\partial g / \partial u$ using the forward AD tools in Sacado along with graph coloring, and then explicitly transpose them.

Graph coloring is a well-known method for computing sparse Jacobian matrices using forward mode AD techniques designed for computing dense Jacobian-vector products [50]. In brief, it constructs a matrix V with a minimal number of columns so that $\partial f / \partial u$ can be easily extracted from $W = (\partial f / \partial u)V$. The idea is to compute W using tangent forward mode AD by seeding the derivative directions for u with columns of V . There are many approaches that have been developed for constructing V , but a common approach is to compute a distance-2 coloring of the bipartite graph whose vertices are the rows and columns of the matrix, with an edge between a given row and column if the corresponding entry of the matrix is nonzero. The distance-2 coloring then partitions the columns into groups of structurally orthogonal columns (i.e., columns with no overlap in nonzeros). Each group corresponds to one column of V and is given by the columns of the identity matrix corresponding to the columns in the group. Once W is computed, it is straightforward to compute $\partial f / \partial u$ given the list of columns associated with each color/group. In this work, the coloring method is implemented by the Zoltan/Zoltan2 package [53] using its parallel distance-2 coloring algorithms [51, 52]. Since the heat-flux and pressure terms of the objective function (3) use the same finite volume stencil as the residual evaluation, the same coloring information can be applied for computing $\partial g / \partial u$ (note the total enthalpy and Pitot pressure terms do not depend on u and hence do not contribute to $\partial g / \partial u$).

Once w is computed, the objective gradient is computed via

$$\nabla h(y) = - \left(\frac{\partial f}{\partial y}(u, y) \right)^T w + \left(\frac{\partial g}{\partial y}(u, y) \right)^T. \quad (15)$$

Since the dimension of y is small (only 2 in this case), we compute the needed partial derivatives, $\partial f / \partial y$ and $\partial g / \partial y$, using traditional forward mode AD techniques using Sacado and then perform the needed vector linear algebra operations. Unlike the forward sensitivity method described earlier, we found the cost of computing w using the above techniques, as well as assembling the gradient ∇h to be insignificant compared to the forward solution (typically around 1% of the total runtime).

Acknowledgments

This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy’s National Nuclear Security Administration under contract DE-NA0003525.

This article has been authored by an employee of National Technology & Engineering Solutions of Sandia, LLC under Contract No. DE-NA0003525 with the U.S. Department of Energy (DOE). The employee owns all right, title and interest in and to the article and is solely responsible for its contents. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this article or allow others to do so, for United States Government purposes. The DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<https://www.energy.gov/downloads/doe-public-access-plan>).

References

- [1] Nompelis, I., Candler, G. V., and Holden, M. S., “Effect of vibrational nonequilibrium on hypersonic double-cone experiments,” *AIAA Journal*, Vol. 41, No. 11, 2003, pp. 2162–2169. <https://doi.org/10.2514/2.6834>.
- [2] Kieweg, S., Ray, J., Weirs, V. G., Carnes, B., Dinzl, D., Freno, B., Howard, M., Rider, W. J., and Smith, T., “Validation Assessment of Hypersonic Double-Cone Flow Simulations using Uncertainty Quantification, Sensitivity Analysis, and Validation Metrics,” *AIAA 2019 Aerosciences Conference, AIAA SciTech 2019 (AIAA-2019-2278)*, 2019. <https://doi.org/10.2514/6.2019-2278>.
- [3] MacLean, M., Dufrene, A., Wadhams, T., and Holden, M., “Numerical and experimental characterization of high enthalpy flow in an expansion tunnel facility,” *48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition, Aerospace Sciences Meetings, (AIAA-2010-1562)*, 2010. <https://doi.org/10.2514/6.2010-1562>.

- [4] Dufrene, A., MacLean, M., Parker, R., Wadhams, T., and Holden, M., "Characterization of the new LENS expansion tunnel facility," *48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition (AIAA-2010-1564)*, 2010. <https://doi.org/10.2514/6.2010-1564>.
- [5] Nompelis, I., Candler, G., Holden, M., and Wadhams, T., "Numerical simulation of high-enthalpy experiments in the LENS X expansion tube facility," *42nd AIAA Aerospace Sciences Meeting and Exhibit, Aerospace Sciences Meetings, (AIAA-2004-1000)*, 2004. <https://doi.org/10.2514/6.2004-1000>.
- [6] Ray, J., Kieweg, S., Dinzl, D., Carnes, B., Weirs, V., Freno, B., Howard, M., Smith, T., Nompelis, I., and Candler, G. V., "Estimation of Inflow Uncertainties in Laminar Hypersonic Double-Cone Experiments," *AIAA Journal*, Vol. 58, No. 10, 2020, pp. 4461–4474. <https://doi.org/10.2514/1.J051609>.
- [7] Holden, M., "Experimental studies of laminar separated flows induced by shock wave/boundary layer and shock/shock interaction in hypersonic flows for CFD validation," *38th Aerospace Sciences Meeting and Exhibit, Aerospace Sciences Meetings, (AIAA-2000-0930)*, 2000. <https://doi.org/10.2514/6.2000-930>.
- [8] Harvey, J., Holden, M., and Wadhams, T., "Code validation study laminar shock/boundary layer and shock/shock interactions in hypersonic flow Part B: Comparison with Navier–Stokes and DSMC solutions," *39th Aerospace Sciences Meeting and Exhibit, Reno, NV, U.S.A (AIAA-2001-1031)*, 2001. <https://doi.org/10.2514/6.2001-1031>.
- [9] MacLean, M., Holden, M. S., and Dufrene, A., "Measurements of real gas effects on regions of laminar shock wave/boundary layer interaction in hypervelocity flows," Tech. rep., CUBRC Inc., 2014. https://www.cubrc.org/_iassets/docs/laminar-xx-paper.pdf.
- [10] Youssefi, M. R., and Knight, D., "Assessment of CFD capability for hypersonic shock wave laminar boundary layer interactions," *Aerospace*, Vol. 4, No. 25, 2017. <https://doi.org/10.3390/aerospace4020025>.
- [11] Holloway, M. E., Chaudhry, R. S., and Boyd, I. D., "Assessment of Hypersonic Double-Cone Experiments for Validation of Thermochemistry Models," *Journal of Spacecraft and Rockets*, Vol. 59, No. 2, 2022, pp. 389–400. <https://doi.org/10.2514/1.A35052>.
- [12] Kianvashrad, N., and Knight, D. D., "Nonequilibrium Effects on Prediction of Aerothermodynamic Loading for a Double Cone," *AIAA Journal*, Vol. 57, No. 7, 2019, pp. 2946–2963. <https://doi.org/10.2514/1.J057883>.
- [13] Hao, J., Wang, J., and Lee, C., "Numerical simulation of high-enthalpy double-cone flows," *American Institute of Aeronautics and Astronautics Journal*, Vol. 55, No. 6, 2017, pp. 2471–2475. <https://doi.org/10.2514/1.J055746>.
- [14] Marrone, P. V., and Treanor, C. E., "Chemical relaxation with preferential dissociation from excited vibrational levels," *Physics of Fluids*, Vol. 6, No. 9, 1963, pp. 1215–1221. <https://doi.org/10.1063/1.1706888>.
- [15] Park, C., "On convergence of computation of chemically reacting flows," *23rd AIAA Aerospace Sciences Meeting (AIAA-1985-0247)*, 1985. <https://doi.org/10.2514/6.1985-247>.
- [16] Gao, J., Hao, J., Wang, J., and Lee, C., "Effect of Thermochemical Nonequilibrium Modeling on High-Enthalpy Double-Cone Flow," *Journal of Spacecraft and Rockets*, Vol. 58, No. 4, 2021, pp. 1243–1247. <https://doi.org/10.2514/1.A35002>.
- [17] Zuo, F.-Y., and Hu, S.-L., "Thermochemical non-equilibrium effects on aerothermodynamic prediction of laminar double-cone flow," *Acta Astronautica*, Vol. 182, 2021, pp. 179–188. <https://doi.org/10.1016/j.actaastro.2021.01.058>, URL <https://www.sciencedirect.com/science/article/pii/S0094576521000692>.
- [18] Nompelis, I., and Candler, G. V., "US3D predictions of double-cone and hollow cylinder and cylinder-flare at high enthalpy," *44th AIAA Fluid Dynamics Conference (AIAA-2014-3366)*, 2014. <https://doi.org/10.2514/6.2014-3366>.
- [19] Candler, G. V., "Next-generation CFD for hypersonic and aerothermal flows (Invited)," *22nd AIAA Computational Fluid Dynamics Conference, AIAA AVIATION Forum, (AIAA 2015-3048)*, 2015. <https://doi.org/10.2514/6.2015-3048>.
- [20] Park, C., "Review of chemical-kinetic problems of future NASA missions I: Earth entries," *Journal of Thermophysics and Heat Transfer*, Vol. 7, No. 3, 1990, pp. 385–398. <https://doi.org/10.2514/3.431>.
- [21] van Leer, B., "Towards the ultimate conservative difference scheme. V. A second-order sequel to Godunov's method," *Journal of Computational Physics*, Vol. 32, No. 1, 1979, pp. 101–136. [https://doi.org/10.1016/0021-9991\(79\)90145-1](https://doi.org/10.1016/0021-9991(79)90145-1), URL <https://www.sciencedirect.com/science/article/pii/0021999179901451>.
- [22] Carnes, B., Weirs, V. G., and Smith, T., "Code verification and numerical error estimation with application to model validation of laminar, hypersonic flow over a double cone," *AIAA 2019 Aerosciences Conference, AIAA SciTech 2019 (AIAA-2019-2175)*, 2019. <https://doi.org/10.2514/6.2019-2175>.

- [23] Freno, B. A., Carnes, B. R., and Weirs, V. G., “Code-verification techniques for hypersonic reacting flows in thermochemical nonequilibrium,” *Journal of Computational Physics*, Vol. 425, 2021, p. 109752. <https://doi.org/10.1016/j.jcp.2020.109752>.
- [24] Harvey, J., Holden, M., and Candler, G., “Validation of DSMC/Navier-Stokes Computations for Laminar Shock Wave/Boundary Layer Interactions Part 3,” *36th AIAA Thermophysics Conference*, 2003. <https://doi.org/10.2514/6.2003-3643>.
- [25] Petra, N., Martin, J., Stadler, G., and Ghattas, O., “A Computational Framework for Infinite-Dimensional Bayesian Inverse Problems, Part II: Stochastic Newton MCMC with Application to Ice Sheet Flow Inverse Problems,” *SIAM Journal on Scientific Computing*, Vol. 36, No. 4, 2014, pp. A1525–A1555. <https://doi.org/10.1137/130934805>, URL <https://doi.org/10.1137/130934805>.
- [26] Iollo, A., Ferlauto, M., and Zannetti, L., “An Aerodynamic Optimization Method Based on the Inverse Problem Adjoint Equations,” *Journal of Computational Physics*, Vol. 173, No. 1, 2001, pp. 87–115. [https://doi.org/https://doi.org/10.1006/jcph.2001.6845](https://doi.org/10.1006/jcph.2001.6845), URL <https://www.sciencedirect.com/science/article/pii/S0021999101968457>.
- [27] Bui-Thanh, T., Ghattas, O., Martin, J., and Stadler, G., “A Computational Framework for Infinite-Dimensional Bayesian Inverse Problems Part I: The Linearized Case, with Application to Global Seismic Inversion,” *SIAM Journal on Scientific Computing*, Vol. 35, No. 6, 2013, pp. A2494–A2523. <https://doi.org/10.1137/12089586X>, URL <https://doi.org/10.1137/12089586X>.
- [28] Nemeć, M., Aftosmis, M., Murman, S., and Pulliam, T., “Adjoint formulation for an embedded-boundary Cartesian method,” *43rd AIAA Aerospace Sciences Meeting and Exhibit*, 2005, p. 877. <https://doi.org/10.2514/6.2005-877>.
- [29] Nemeć, M., and Aftosmis, M., “Aerodynamic shape optimization using a Cartesian adjoint method and CAD geometry,” *24th AIAA Applied Aerodynamics Conference*, 2006, p. 3456. <https://doi.org/10.2514/6.2006-3456>.
- [30] Nemeć, M., Aftosmis, M. J., and Brown, P. G., “Numerical prediction of meteoric infrasound signatures,” *Planetary and Space Science*, Vol. 140, 2017, pp. 11–20. <https://doi.org/10.1016/j.pss.2017.03.003>.
- [31] Nielsen, E. J., Lu, J., Park, M. A., and Darmofal, D. L., “An implicit, exact dual adjoint solution method for turbulent flows on unstructured grids,” *Computers & Fluids*, Vol. 33, No. 9, 2004, pp. 1131–1155. <https://doi.org/10.1016/j.compfluid.2003.09.005>.
- [32] Nielsen, E. J., and Kleb, W. L., “Efficient construction of discrete adjoint operators on unstructured grids using complex variables,” *AIAA journal*, Vol. 44, No. 4, 2006, pp. 827–836. <https://doi.org/10.2514/6.2005-324>.
- [33] Thompson, K. B., “Aerothermodynamic Design Sensitivities for a Reacting Gas Flow Solver on an Unstructured Mesh Using a Discrete Adjoint Formulation,” Ph.D. thesis, North Carolina State University, 2017.
- [34] Economou, T. D., Palacios, F., Copeland, S. R., Lukaczyk, T. W., and Alonso, J. J., “SU2: An Open-Source Suite for Multiphysics Simulation and Design,” *AIAA Journal*, Vol. 54, No. 3, 2016, pp. 828–846. <https://doi.org/10.2514/1.J053813>, URL <https://doi.org/10.2514/1.J053813>.
- [35] Copeland, S., Lonkar, A., Palacios, F., and Alonso, J., “Adjoint-Based Goal-Oriented Mesh Adaptation for Nonequilibrium Hypersonic Flows,” *51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, 2013. <https://doi.org/10.2514/6.2013-552>, URL <https://arc.aiaa.org/doi/abs/10.2514/6.2013-552>.
- [36] Copeland, S. R., “A Continuous Adjoint Formulation for Hypersonic Flows in Thermochemical Nonequilibrium,” Ph.D. thesis, Stanford University, 2015.
- [37] Kline, H., and Alonso, J., “Adjoint of generalized outflow-based functionals applied to hypersonic inlet design,” *AIAA Journal*, Vol. 55, No. 11, 2017, pp. 3903–3915. <https://doi.org/10.2514/1.J055863>.
- [38] Albring, T. A., Sagebaum, M., and Gauger, N. R., “Efficient aerodynamic design using the discrete adjoint method in SU2,” *17th AIAA/ISSMO multidisciplinary analysis and optimization conference*, 2016, p. 3518. <https://doi.org/10.2514/6.2016-3518>.
- [39] Munguía, B. C., Maier, W. T., Needels, J. T., and Alonso, J. J., “Goal-Oriented Mesh Adaptation for Flows in Thermochemical Nonequilibrium,” *23rd AIAA International Space Planes and Hypersonic Systems and Technologies Conference*, 2020. <https://doi.org/10.2514/6.2020-2400>, URL <https://doi.org/10.2514/6.2020-2400>.
- [40] Damm, K. A., Gollan, R. J., Jacobs, P. A., Smart, M. K., Lee, S., Kim, E., and Kim, C., “Discrete adjoint optimization of a hypersonic inlet,” *AIAA Journal*, Vol. 58, No. 6, 2020, pp. 2621–2634. <https://doi.org/10.2514/1.J058913>.
- [41] Lockwood, B., and Mavriplis, D., “Parameter Sensitivity Analysis for Hypersonic Viscous Flow Using a Discrete Adjoint Approach,” *48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*, 2010. <https://doi.org/10.2514/6.2010-447>, URL <https://arc.aiaa.org/doi/abs/10.2514/6.2010-447>.

- [42] Lockwood, B., Rumpfkeil, M., Yamazaki, W., and Mavriplis, D., “Uncertainty quantification in viscous hypersonic flows using gradient information and surrogate modeling,” *49th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, 2010, p. 885. <https://doi.org/10.2514/6.2011-885>.
- [43] Lockwood, B., Anitescu, M., and Mavriplis, D., “Mixed aleatory/epistemic uncertainty quantification for hypersonic flows via gradient-based optimization and surrogate models,” *50th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*, 2012, p. 1254. <https://doi.org/10.2514/6.2012-1254>.
- [44] Schlichting, H., and Gersten, K., *Boundary Layer Theory*, 8th ed., Springer, 2000.
- [45] Nocedal, J., and Wright, S., *Numerical Optimization*, Springer, 2006.
- [46] Kouri, D., Ridzal, D., and von Winckel, G., *The ROL Project Website*, 2020 (accessed May 22, 2020). URL <https://trilinos.github.io/rol.html>.
- [47] Phipps, E. T., and Gay, D. M., “Sacado Automatic Differentiation Package,” <http://trilinos.org/packages/sacado/>, 2018.
- [48] Phipps, E. T., Bartlett, R. A., Gay, D. M., and Hoekstra, R. J., “Large-Scale Transient Sensitivity Analysis of a Radiation-Damaged Bipolar Junction Transistor via Automatic Differentiation,” *Advances in Automatic Differentiation*, edited by C. H. Bischof, H. M. Bücker, P. D. Hovland, U. Naumann, and J. Utke, Springer, 2008, pp. 351–362. https://doi.org/10.1007/978-3-540-68942-3_31.
- [49] Phipps, E. T., and Pawlowski, R. P., “Efficient Expression Templates for Operator Overloading-based Automatic Differentiation,” *Recent Advances in Algorithmic Differentiation*, edited by S. Forth, P. Hovland, E. Phipps, J. Utke, and A. Walther, Springer, 2012.
- [50] Gebremedhin, A. H., Manne, F., and Pothen, A., “What Color Is Your Jacobian? Graph Coloring for Computing Derivatives,” *SIAM Review*, Vol. 47, No. 4, 2005, pp. 629–705. <https://doi.org/10.1137/S0036144504444711>.
- [51] Bozdag, D., Gebremedhin, A. H., Manne, F., Boman, E. G., and Catalyurek, U. V., “A framework for scalable greedy coloring on distributed-memory parallel computers,” *Journal of Parallel and Distributed Computing*, Vol. 68, No. 4, 2008, pp. 515–535. <https://doi.org/10.1016/j.jpdc.2007.08.002>.
- [52] Bozdag, D., Catalyurek, U. V., Gebremedhin, A. H., Manne, F., Boman, E. G., and Ozguner, F., “Distributed-Memory Parallel Algorithms for Distance-2 Coloring and Related Problems in Derivative Computation,” *SIAM Journal on Scientific Computing*, Vol. 32, No. 4, 2010, pp. 2418–2446. <https://doi.org/10.1137/080732158>.
- [53] Boman, E. G., Çatalyürek, Ü. V., Chevalier, C., and Devine, K. D., “The Zoltan and Isorropia Parallel Toolkits for Combinatorial Scientific Computing: Partitioning, Ordering and Coloring,” *Scientific Programming*, Vol. 20, 2012, p. 713587. <https://doi.org/10.3233/SPR-2012-0342>.
- [54] Saad, Y., and Schultz, M. H., “GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems,” *SIAM Journal on Scientific and Statistical Computing*, Vol. 7, No. 3, 1986, pp. 856–869. <https://doi.org/10.1137/0907058>.
- [55] Wright, M. J., Candler, G. V., and Bose, D., “Data-Parallel Line Relaxation Method for the Navier-Stokes Equations,” *AIAA Journal*, Vol. 36, No. 9, 1998, pp. 1603–1609. <https://doi.org/10.2514/2.586>.
- [56] Yuen, K.-V., *Bayesian Methods for Structural Dynamics and Civil Engineering*, John Wiley and Sons (Asia) Pte Ltd, 2010.
- [57] Dovi, V. G., Paladino, O., and Reverberi, A. P., “Some remarks on the use of the inverse Hessian matrix of the likelihood function in the estimation of statistical properties of parameters,” *Applied Mathematics Letters*, Vol. 4, No. 1, 1991, pp. 87–90. [https://doi.org/10.1016/0893-9659\(91\)90129-J](https://doi.org/10.1016/0893-9659(91)90129-J).
- [58] Ninni, D., Bonelli, F., Colonna, G., and Pascazio, G., “On the influence of non equilibrium in the free stream conditions of high enthalpy oxygen flows around a double-cone,” *Acta Astronautica*, Vol. 201, 2022, pp. 247–258. <https://doi.org/10.1016/j.actaastro.2022.09.017>, URL <https://www.sciencedirect.com/science/article/pii/S0094576522004829>.
- [59] Kelley, C. T., and Keyes, D. E., “Convergence Analysis of Pseudo-Transient Continuation,” *SIAM Journal on Numerical Analysis*, Vol. 35, No. 2, 1998, pp. 508–523. <https://doi.org/10.1137/S0036142996304796>.
- [60] Maly, T., and Petzold, L. R., “Numerical methods and software for sensitivity analysis of differential-algebraic systems,” *Applied Numerical Mathematics*, Vol. 20, No. 1-2, 1996, pp. 57–79.
- [61] Pawlowski, R., “NOX Object-Oriented Nonlinear Solutions Package,” https://trilinos.github.io/nox_and_loca.html, 2022.