



Adjoint-Based Aerodynamic Shape Optimisation for Three-Dimensional Hypersonic Configurations

Reece Otto¹, Kyle Damm², and Rowan Gollan³

Abstract

The design of hypersonic airbreathing vehicles involves tightly coupled aerodynamics and propulsion, creating complex design spaces that are most effectively navigated using numerical optimisation techniques. However, achieving high-fidelity, many-parameter optimisation in three dimensions remains difficult while also maintaining efficiency and robustness. To address this challenge, this paper presents an adjoint-based aerodynamic shape optimisation (ASO) framework for three-dimensional hypersonic configurations. The framework integrates high-fidelity CFD analysis, adjoint-based sensitivity evaluation, and a two-level free-form deformation (FFD) parameterisation strategy. The method is demonstrated through the optimisation of a hypersonic lifting-body configuration subject to an internal payload constraint. The configuration was simulated using Reynolds-Averaged Navier–Stokes (RANS) equations and parameterised with 121 design variables. The optimisation required just over four days of runtime on a workstation and achieved a 21.8% improvement in lift-to-drag ratio. Geometric and flow field analyses revealed a reduced frontal area, a slenderised forebody, and waverider-like shock attachment that improved aerodynamic efficiency and weakened side vortices. These results highlight the capability of adjoint-based methods to deliver efficient, high-fidelity optimisation of complex three-dimensional hypersonic vehicles, representing a step toward practical many-variable design studies in this regime.

Keywords: *Aerodynamic Shape Optimisation, Hypersonics, Adjoint Method, CFD*

Nomenclature

Acronyms

ASO – Aerodynamic Shape Optimisation
CFD – Computational Fluid Dynamics
FFD – Free-Form Deformation
MDO – Multidisciplinary Design Optimisation
RANS – Reynolds-Averaged Navier-Stokes
SNOPT – Sparse Nonlinear OPTimizer

Latin

D – Aerodynamic drag
 $\mathbf{D}$ – Design variables
 g – Constraint functions
 J – Objective function

L – Aerodynamic lift

$\mathbf{R}$ – Flow residuals

S – Design surface

$\mathbf{U}$ – Conserved flow variables

$\mathbf{X}$ – Computational grid coordinates

Greek

λ – Adjoint variables

Subscripts

0 – Initial value

new – Updated value

opt – Optimal value

1. Introduction

Smallsats currently dominate the satellite market, with demand expected to grow substantially over the next decade [1]. At present, most are launched as secondary payloads via ride-share services on large orbital-class rockets. Although cost-effective, this model often involves long lead times and limited access to preferred launch windows and orbital slots, reducing flexibility and delaying time-to-market. To address these challenges, increasing attention is being directed toward dedicated launch

¹*School of Mechanical & Mining Engineering, The University of Queensland, r.otto@uq.edu.au*

²*School of Mechanical & Mining Engineering, The University of Queensland, k.damm@uq.edu.au*

³*School of Mechanical & Mining Engineering, The University of Queensland, r.gollan@uq.edu.au*

systems tailored for small payloads [2]. Among these emerging concepts, scramjet-powered vehicles are particularly promising, offering the potential for higher propulsive efficiency, reusability, and improved operational flexibility [3–6].

Scramjets are high-speed airbreathing propulsion systems, widely regarded as the most viable option for sustained atmospheric flight at hypersonic speeds. Unlike rockets, they draw oxygen directly from the atmosphere rather than carrying onboard oxidisers, which significantly improves fuel efficiency [7] and allows a higher payload mass fraction [8]. Moreover, scramjet-powered vehicles operate more like conventional aircraft, offering improved reusability and manoeuvrability [9]. Collectively, these advantages make scramjet-powered vehicles an attractive solution for reducing launch costs and increasing operational flexibility in smallsat missions.

Despite this promise, achieving routine and reliable hypersonic flight presents significant engineering challenges. A central difficulty lies in integrating the propulsion system with the airframe, driven by the large engine size required to sustain hypersonic speeds [10]. This necessity leads to tightly coupled vehicle configurations in which components must often perform multiple roles under competing requirements. The resulting subsystem interactions can strongly influence overall performance, creating a complex design space with non-intuitive trade-offs. Such design spaces are most effectively navigated using numerical optimisation techniques guided by models that capture the dominant performance drivers across disciplines. This integrated approach—known as Multidisciplinary Design Optimisation (MDO)—represents the state of the art in hypersonic vehicle design [11].

Although effective, the use of MDO is often constrained by significant computational overhead. Accurate performance evaluation typically requires coupling multiple high-fidelity physics solvers, and adequate vehicle characterisation can lead to large, high-dimensional design spaces. To achieve practical turnaround times, designers are forced to compromise on modelling fidelity and geometric complexity. Common approaches include using panel methods for aerodynamic analysis, approximating propulsion with one-dimensional models, and restricting the number of design variables to around ten [12–14]. While such simplifications improve computational efficiency, they also limit potential performance gains. Lower-fidelity models may overlook critical physical phenomena, and overly constrained design spaces reduce the opportunity to discover high-performing configurations.

In contrast, high-fidelity, many-parameter MDO has been computationally feasible in lower-speed aircraft design for decades [15–17]. This progress is largely attributed to the use of the adjoint method, a powerful technique for efficiently computing design sensitivities [18]. Traditional gradient estimation methods—such as finite differences—treat the system as a black box, requiring one function evaluation per design variable. This becomes prohibitively expensive in aerodynamic design, where each function evaluation may involve solving a full computational fluid dynamics (CFD) problem. The adjoint method, by contrast, enables the computation of the entire gradient vector at a cost comparable to a single function evaluation, regardless of the number of design variables. Owing to this scalability, the adjoint method has become a popular choice for many-parameter design optimisation.

Despite its efficiency benefits, the use of the adjoint method in hypersonic MDO has so far been relatively limited. Recent work has instead focused on hypersonic aerodynamic shape optimisation (ASO) as an intermediate step, employing adjoint-based approaches to compute design sensitivities [19–23]. These studies have considered up to 34 design variables, yet research in lower-speed regimes indicates that hundreds may be necessary to fully exploit the potential of ASO, even for relatively simple three-dimensional geometries such as wings [24]. Furthermore, most hypersonic ASO efforts to date have been restricted to two-dimensional cases. While some extensions to three dimensions exist, they have revealed robustness challenges—particularly in handling surface and mesh deformation during the optimisation process [21, 23].

In this work, an adjoint-based ASO framework is developed for three-dimensional hypersonic configurations. The framework enables efficient, high-fidelity optimisation of practical geometries with many design variables, representing an important advance toward overcoming current limitations in hypersonic MDO. Its capability is demonstrated through the optimisation of a hypersonic lifting-body configuration

subject to an internal volume constraint, highlighting both the potential performance gains and the robustness of the approach.

2. Aerodynamic Shape Optimisation Methodology

An overview of the aerodynamic shape optimisation routine implemented in this work is shown in Figure 1. The process begins with a baseline grid X_0 and an initial set of design variables D_0 , which are used to generate parametric design surfaces S_0 . A steady-state flow solution is then computed, from which the objective function J (e.g. drag or aerodynamic efficiency) and constraint functions g (e.g. heat flux or internal volume) are evaluated. When required by the optimiser, the adjoint-based sensitivity solver computes the derivatives of the objective and constraint functions with respect to the design variables: dJ/dD and dg/dD . If these sensitivities satisfy the convergence criteria, the optimal design D_{opt} is returned and the routine terminates. Otherwise, the optimiser generates an updated set of design variables D_{new} , which are used to deform the design surfaces and surrounding volume mesh, yielding S_{new} and X_{new} . This updated grid is passed back to the flow solver, and the loop continues until the convergence criteria are satisfied. The following sections provide further detail on the components of this optimisation methodology.




Fig 1. Schematic of the aerodynamic shape optimisation routine.

2.1. Geometric Parameterisation and Deformation

The optimisation process begins with the generation of a parametric geometry, which the optimiser adjusts to search for better-performing designs. In CFD-based optimisation, the computational mesh must also deform consistently with the evolving surface geometry to preserve continuity. In this work, an integrated strategy is employed using the two-level free-form deformation (FFD) method of Gagnon and Zingg [25].

Using this method, the structured volume mesh is represented as a network of B-spline volumes, one for each block in the domain, with the design surfaces forming their boundaries as B-spline surfaces. Direct manipulation of the surface control points can introduce issues such as creases or overlap, and it is also difficult to assign intuitive design variables prior to parameterisation [26]. Instead, the surface

control points are embedded within a higher-level B-spline volume, or FFD hull. By repositioning the hull control points, these multi-patch surfaces can be deformed consistently while retaining smoothness, and reducing the number of required design variables.

Since the design surfaces and volume mesh are both represented by B-spline volumes, displacing the surface control points naturally deforms the mesh in their vicinity. However, this local deformation alone is insufficient, as large surface changes can severely degrade mesh quality or even cause self-intersections. To maintain robustness, the surrounding volume mesh control points are also adjusted in response to surface displacements. In this work, the Inverse Distance Weighting (IDW) method [27] is used to propagate the surface deformation smoothly throughout the computational domain.

An example of the implemented geometry framework is shown in Figures 2 and 3, where a cylindrical geometry is morphed into a re-entry capsule. Figure 2a presents the original cylinder, with a quarter section embedded in an FFD hull whose control points are highlighted in red. By repositioning these points and reflecting the deformation, the capsule surface shown in Figure 2b is obtained. The corresponding deformation of the volume mesh is illustrated in Figure 3. Figures 3a and 3b demonstrate how the IDW method propagates the surface control point displacements into the volume control net, while Figures 3c and 3d show that the resulting computational mesh remains smooth with minimal loss of quality.




Fig 2. Two-level free-form deformation of a cylindrical geometry into a re-entry capsule. The geometry is parameterised as a quarter-section, and the red points denote the FFD control net.

2.2. Flow Solver

The aerodynamic objective and constraint functions are evaluated within the design loop using flow solutions obtained from *Eilmer* [28]. *Eilmer* is an open-source, multi-physics hypersonic flow solver capable of modelling both transient and steady problems, including turbulence and high-temperature gas effects. The governing equations are discretised using the finite-volume method and solved on multi-block, structured or unstructured, body-fitted grids.

In the present work, the Reynolds-Averaged Navier–Stokes (RANS) equations are solved under the assumption of an ideal-air gas model. Convergence to steady state is accelerated using a Jacobian-Free Newton–Krylov (JFNK) method [29]. To ensure robustness and efficiency within the optimisation loop, an adaptive CFL strategy is employed, physicality checks are performed on the flow state, and each solution is warm-started from the previous one.




Fig 3. Deformation of the computational mesh when a cylindrical geometry is morphed into a re-entry capsule.

2.3. Shape Sensitivity Calculation

Shape sensitivities, defined as the derivatives of the objective and constraint functions with respect to the design variables, are required by the gradient-based optimisation algorithm to navigate the design space. In CFD-based optimisation, these derivatives must capture the dependence of the function of interest f on the flow solution variables U and the grid coordinates X , both of which vary with the design variables D . This relationship can be expressed as:

$$f = f(X(D), U(D)) \quad (1)$$

Using the adjoint method, the total derivative of f with respect to D is expressed as:

$$\frac{df}{dD} = \left. \frac{\partial f}{\partial D} \right|_U + \lambda^T \left. \frac{\partial R}{\partial D} \right|_U \quad (2)$$

where the flow is assumed to be at steady state, or equivalently, the flow residuals are zero, $R = \partial U / \partial t = 0$. The partial derivatives $\partial f / \partial D|_U$ and $\partial R / \partial D|_U$ represent the sensitivity of shape perturbations on the function and residuals while holding flow state U fixed. The adjoint variables λ are determined by solving the adjoint equations:

$$\left[\frac{\partial R}{\partial U} \right]^T \lambda = - \left[\frac{\partial f}{\partial U} \right]^T \quad (3)$$

In practice, evaluating the shape sensitivities using the adjoint method involves three main steps:

1. Compute the partial derivative terms in Equations 2 and 3.
2. Solve the adjoint equations for λ .
3. Assemble the total derivative in Equation 2 from the computed partial derivatives and adjoint solution.

In this work, all partial derivatives are computed using complex-step differentiation [30], which provides near machine-precision accuracy for sufficiently small step sizes. Constructing the adjoint operator—defined as the transpose of the flow Jacobian, $[\partial R / \partial U]^T$ —is computationally intensive, as the Jacobian contains $(NM) \times (NM)$ elements, where N is the number of conserved variables and M the number of cells. To reduce cost, residual derivatives are computed only over each cell's stencil of influence, and operator assembly is parallelised using the same domain decomposition as the flow solver [22]. The resulting sparse linear system is then solved using a preconditioned, restarted Generalised Minimal RESidual (GMRES) algorithm [31]. The preconditioner is formed as a first-order approximation of the transposed flow Jacobian, using an Incomplete Lower-Upper factorisation with zero fill (ILU(0)). In this work, we found that solver robustness and efficiency were improved by assembling the preconditioner with a different flux calculator than that used for the full adjoint operator. Further gains are achieved by warm-starting each adjoint solve from the previous iteration within the design loop, analogous to the flow solver.

2.4. Optimisation Algorithm

Several optimisation frameworks exist to couple flow and sensitivity solvers with optimisation algorithms. In this work, the open-source package pyOptSparse [32] is used, which provides access to a variety of third-party algorithms. The chosen algorithm is the Sparse Nonlinear OPTimizer (SNOPT) [33], a gradient-based method built on Sequential Quadratic Programming (SQP). SQP solves the nonlinear optimisation problem through a sequence of quadratic programming (QP) subproblems, with the objective approximated quadratically and the constraints linearised. Each QP solution defines a search direction, along which a line search is performed. The Hessian is approximated using a quasi-Newton scheme, requiring only first-order gradient information. SNOPT terminates when the optimality and feasibility tolerances are satisfied, corresponding to sufficiently small Lagrangian gradients and constraint violations within prescribed limits.

3. Aerodynamic Optimisation of a Hypersonic Lifting Body

To demonstrate the capabilities of the developed optimisation framework, a case study is carried out on a representative hypersonic lifting-body configuration. The baseline geometry, shown in Figure 6, shares the major dimensions of the configuration studied by Zhang et al. [34]. An internal payload, shown in yellow, is included to ensure the design remains physically realistic and to impose a practical constraint on the optimisation, preventing degenerate geometries that may be aerodynamically efficient but infeasible in practice.

The study considers a single design point corresponding to Mach 8 atmospheric flight at an altitude of 40km, consistent with the conditions of Zhang et al. [34] and summarised in Table 1. Flow solutions are obtained by solving the Reynolds-Averaged Navier–Stokes (RANS) equations with the one-equation Spalart–Allmaras (S–A) turbulence model [35]. The computational mesh, shown in Figure 5, consists of 46 blocks and approximately 613,000 cells, generated using GridPro v8.0 [36].




Fig 4. Baseline lifting body and internal payload geometry.

Table 1. Free-stream flow conditions for lifting body optimisation.

Property	Symbol	Value	Units
Velocity	V_∞	2.54	km/s
Pressure	p_∞	278	Pa
Temperature	T_∞	251	K
Angle of attack	α	8.0	degrees

3.1. Objective and Constraint Formulation

As mentioned, the objective of this optimisation study is to maximise the lift-to-drag ratio while constraining the internal volume to accommodate a payload. The aerodynamic forces are obtained by integrating the pressure and viscous contributions over the design surface:

$$\mathbf{F}_{\text{net}} = \iint_S -p\hat{\mathbf{n}}dA + \iint_S \boldsymbol{\tau}_w dA \quad (4)$$

where S is the design surface, p is the pressure, $\hat{\mathbf{n}}$ is the outward surface normal, dA is the differential surface element, and $\boldsymbol{\tau}_w$ is the viscous shear stress vector. The lift and drag forces are then obtained by projecting $\mathbf{F}_{\text{net}}$ onto the directions normal and parallel to the free-stream velocity vector, respectively.




Fig 5. Computational grid for baseline lifting body configuration.

To accommodate the internal payload, a geometric constraint is imposed between the frustum and the sampled points on the vehicle surface. For each surface point, an effective radius is defined relative to the frustum axis and compared against the corresponding frustum radius at the same axial location. If the frustum radius is larger, the frustum would protrude through the surface at that point, constituting a constraint violation. For a given surface point i , this constraint is evaluated as

$$g_i(\mathbf{D}) = \begin{cases} r_{\text{frust}} - r_{\text{surf}}, & r_{\text{frust}} > r_{\text{surf}} \\ 0, & \text{otherwise} \end{cases} \quad (5)$$

where r_{surf} is the effective radius of the sampled surface point and r_{frust} is the radius of the frustum at the same axial location.

Directly enforcing a constraint at every surface point is inefficient, as it would require computing gradients for a large number of individual functions. Moreover, summing these constraints can lead to poor scaling characteristics and hinder optimisation convergence. To address this, the distance constraints are aggregated using the discrete Kreisselmeier–Steinhauser (KS) function [37–39], which blends all individual constraints into a single, smooth, and well-bounded function. This aggregated constraint is then enforced to remain less than or equal to zero throughout the optimisation.

3.2. Shape Parameterisation

The geometry was parameterised using the two-level FFD method described in Section 2.1. The chosen FFD hull, shown in Figure 6, is a tri-cubic B-spline volume with $10 \times 6 \times 2$ control points. Directly

parameterising each coordinate of every control point would yield 360 design variables, but this approach risks control-point crossover and the creation of impractically thin features. To improve robustness and maintain manageability, the parameterisation was structured by grouping control points, which still provided sufficient flexibility to capture meaningful shape changes.

Control of the overall length was constrained, while additional degrees of freedom were assigned to modify the width and axial profile through translation and scaling of x -control planes. The thickness profile was adjusted through vertical scaling and translation of paired control points in the y -direction. Further design variables allowed axial and vertical translation of the internal payload, as well as its pitching angle. In total, this parameterisation resulted in 121 design variables.




Fig 6. Parameterisation of the lifting-body geometry using a two-level FFD scheme. Control points directly manipulated by the optimiser are shown in red.

3.3. Optimisation Results

The results of the optimisation study are presented in this section. The optimisation was performed using SNOPT with optimality and feasibility tolerances of 1×10^{-5} and 1×10^{-4} , respectively. The objective and constraint histories are shown in Figure 7. Convergence was achieved in 42 design iterations, requiring just over four days of runtime on a workstation¹. The histories show a clear plateau in the objective function, while the constraint violation approaches zero by the final iterations, indicating a well-converged solution.

From Table 8, the optimisation study produced a 21.8% improvement in L/D . This gain was driven primarily by a reduction in drag, which outweighed a nearly 20% decrease in lift. In some design scenarios, such a reduction in lift may be undesirable; however, this can be addressed by reformulating the optimisation problem to constrain lift while minimising drag.

The comparison of surface geometry and pressure distribution is shown in Figure 8. The optimised configuration exhibits a reduced frontal area and a tighter fit of the outer mold line (OML) around the frustum. In addition, the frustum angle is rotated to align more closely with the oncoming flow, lowering the projected frontal area and concentrating most of the enclosed volume in the upper portion

¹Workstation specifications: Dual Intel Xeon Silver 4216 CPUs @ 2.10 GHz (32 cores total), 196 GB RAM.

of the vehicle, where surface pressures are lower. A more slender nose and forebody profile reduce leading-edge pressures, while the thinner forebody produces a more rapid expansion over the upper surface—both contributing to additional drag reduction.

Further aerodynamic insight is provided in Figure 9, which compares the Mach number fields of the baseline and optimised configurations. The optimised design features a more concave lower surface and a leading-edge profile that lies closer to the shock. Similar to a hypersonic waverider [40], this configuration more effectively separates high-pressure air beneath the vehicle from the low-pressure region above, thereby improving overall aerodynamic efficiency. It also weakens the vortex observed in the baseline near the leading-edge base, as less high-pressure air spills over the upper surface.




Fig 7. Objective and constraint function history for lifting body optimisation study.

Table 2. Comparison of aerodynamic performance between baseline and optimised lifting bodies.

Configuration	Drag D (N)	Lift L (N)	Lift-on-drag L/D
Baseline	939	3190	3.40
Optimised	626	2589	4.14
Difference (%)	-33.3	-18.8	+21.8






Fig 8. Comparison of baseline and optimised surface geometry and pressure distribution.



(a) Baseline configuration.



(b) Optimised configuration.

Fig 9. Flow fields of baseline and optimised lifting bodies.

4. Conclusions

In this work, an adjoint-based aerodynamic shape optimisation framework was developed for three-dimensional hypersonic configurations. The framework integrates high-fidelity CFD analysis, adjoint-based sensitivity evaluation, and a two-level free-form deformation parameterisation strategy, enabling efficient optimisation of complex geometries with many design variables.

The framework was applied to optimise the aerodynamic performance of a hypersonic lifting body under an internal volume constraint. The configuration was simulated using high-fidelity RANS and parameterised with 121 design variables. The optimiser converged in 42 design iterations, requiring just over four days of runtime on a workstation, with this efficiency made possible by the adjoint method.

The optimisation achieved a 21.8% improvement in the lift-on-drag ratio, driven primarily by drag reduction. The optimised geometry exhibited a reduced frontal area, a more slender forebody, and an internal payload aligned more closely with the oncoming flow. Flow field comparisons showed the optimised configuration achieved high aerodynamic efficiency similar to a waverider, with the shock lying close to the leading edge to separate the high-pressure lower surface from the low-pressure upper surface. This reduced spillage of high-pressure air onto the upper surface, weakening the side vortices and further reducing drag.

Overall, this study demonstrates that adjoint-based optimisation can be applied efficiently and effectively to improve the aerodynamic performance of realistic three-dimensional hypersonic configurations. The developed framework represents a step toward enabling high-fidelity, many-variable optimisation in the hypersonic regime, helping to advance future design practices.

References

1. Novaspace. *Prospects for the Small Satellite Market, 10th edition* Market intelligence report. Dec. 2024. <https://nova.space/hub/product/prospects-for-the-small-satellite-market/> (2025).
2. Peeters, W., Damp, L. & Williams, P. Launching Smallsats: The Example of Southern Launch. *New Space* **8**, 201–212 (2020).
3. Smart, M. K. & Tetlow, M. R. Orbital Delivery of Small Payloads Using Hypersonic Airbreathing Propulsion. *Journal of Spacecraft and Rockets* **46**, 117–125 (Jan. 2009).
4. Preller, D. & Smart, M. K. Reusable Launch of Small Satellites Using Scramjets. *Journal of Spacecraft and Rockets* **54**, 1317–1329 (2017).
5. Sam, L., Idithsaj, P. T., Nair, P. P., Suryan, A. & Narayanan, V. Prospects for Scramjet Engines in Reusable Launch Applications: A Review. *International Journal of Hydrogen Energy* **48**, 36094–36111. ISSN: 0360-3199 (2023).
6. McTaggart, L. & Vanyai, T. *Commercial Potential Evaluation of Scramjet Powered Vehicles for Access-to-Space* in *AIAA SCITECH 2023 Forum* AIAA 2023-2685 (2023).
7. Kerrebrock, J. L. *Aircraft Engines and Gas Turbines* 2nd ed, 402. ISBN: 0262111624 (MIT Press, Cambridge, MA, 1992).
8. Smart, M. K. Scramjets. *The Aeronautical Journal* **111**, 605–619 (2007).
9. Dissel, A., Kothari, A., Raghavan, V. & Lewis, M. *Comparison of HTHL and VTHL Airbreathing and Rocket Systems for Access to Space* in *40th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit* AIAA 2004-3988 (2004).
10. Bowcutt, K. G. *Physics Drivers of Hypersonic Vehicle Design* in *22nd AIAA International Space Planes and Hypersonics Systems and Technologies Conference* AIAA 2018-5373 (2018).
11. Bowcutt, K. G. Multidisciplinary Optimization of Airbreathing Hypersonic Vehicles. *Journal of Propulsion and Power* **17**, 1184–1190 (2001).
12. Jazra, T. *Optimisation of Hypersonic Vehicles for Airbreathing Propulsion* PhD thesis (School of Mechanical and Mining Engineering, The University of Queensland, 2010).
13. Preller, D. *Multidisciplinary Design and Optimisation of a Pitch Trimmed Hypersonic Air-breathing Accelerating Vehicle* PhD thesis (School of Mechanical and Mining Engineering, The University of Queensland, 2018).

14. Mackle, K. *Co-design of Hypersonic Vehicle Shape and Trajectory* PhD thesis (School of Mechanical and Mining Engineering, The University of Queensland, 2024).
15. Martins, J. R. R. A. *A Coupled-Adjoint Method for High-Fidelity Aero-Structural Optimization* PhD thesis (Stanford University, 2002).
16. Kennedy, G. J. *Aerostructural Analysis and Design Optimization of Composite Aircraft* PhD thesis (Graduate Department of Applied Science and Engineering, University of Toronto, 2012). ISBN: 978-0-494-79469-2.
17. Kenway, G. K. W. *A Scalable, Parallel Approach for Multi-Point, High-Fidelity Aerostructural Optimization of Aircraft Configurations* PhD thesis (Graduate Department of Aerospace Engineering, University of Toronto, 2013).
18. Martins, J. R. R. A. Aerodynamic design optimization: Challenges and perspectives. *Computers & Fluids* **239**, 105391. ISSN: 0045-7930 (2022).
19. Eyi, S. & Yumuşak, M. Aerothermodynamic Shape Optimization of Hypersonic Blunt Bodies. *Engineering Optimization* **47**, 909–926 (2015).
20. Kline, H. L. & Alonso, J. J. Adjoint of Generalized Outflow-Based Functionals Applied to Hypersonic Inlet Design. *AIAA Journal* **55**, 3903–3915 (Nov. 2017).
21. Kline, H. L. *The Continuous Adjoint Method for Multi-fidelity Hypersonic Inlet Design* PhD thesis (Department of Aeronautics and Astronautics, Stanford University, 2017).
22. Damm, K. A., Gollan, R. J., Jacobs, P. A., Smart, M. K., Lee, S., Kim, E. & Kim, C. Discrete Adjoint Optimization of a Hypersonic Inlet. *AIAA Journal* **58**, 2621–2634 (2020).
23. Novotny, N. L., Rumpfkeil, M. P., Camberos, J. A., Nielsen, E. J. & Diskin, B. Aerodynamic Exergy-Based Analysis and Optimization of the Generic Hypersonic Vehicle Using FUN3D. *Journal of Aircraft* **61**, 972–987 (2024).
24. Lyu, Z., Kenway, G. K. W. & Martins, J. R. R. A. Aerodynamic Shape Optimization Investigations of the Common Research Model Wing Benchmark. *AIAA Journal* **53**, 968–985 (2015).
25. Gagnon, H. & Zingg, D. W. Two-Level Free-Form and Axial Deformation for Exploratory Aerodynamic Shape Optimization. *AIAA Journal* **53**, 2015–2026 (2015).
26. Lee, C., Koo, D. & Zingg, D. W. Comparison of B-Spline Surface and Free-Form Deformation Geometry Control for Aerodynamic Optimization. *AIAA Journal* **55**, 228–240 (2017).
27. Witteveen, J. & Bijl, H. *Explicit Mesh Deformation Using Inverse Distance Weighting Interpolation in 19th AIAA Computational Fluid Dynamics* AIAA 2009-3996 (June 2009).
28. Gibbons, N. N., Damm, K. A., Jacobs, P. A. & Gollan, R. J. Eilmer: An Open-source Multi-physics Hypersonic Flow Solver. *Computer Physics Communications* **282**, 108551. ISSN: 0010-4655 (2023).
29. Damm, K. A., Gibbons, N. N., Jacobs, P. A. & Gollan, R. J. *Application of a Jacobian-Free Newton-Krylov Method to the Simulation of Hypersonic Flows in AIAA Scitech 2023 Forum* AIAA 2023-2295 (Jan. 2023).
30. Martins, J. R. R. A., Sturdza, P. & Alonso, J. J. The complex-step derivative approximation. *ACM Trans. Math. Softw.* **29**, 245–262. ISSN: 0098-3500 (Sept. 2003).
31. Saad, Y. *Iterative Methods for Sparse Linear Systems* 92–95 (SIAM, Philadelphia, Pennsylvania, 2003).
32. Wu, N., Kenway, G., Mader, C. A., Jasa, J. & Martins, J. R. R. A. pyOptSparse: A Python Framework for Large-scale Constrained Nonlinear Optimization of Sparse Systems. *Journal of Open Source Software* **5**, 2564 (2020).
33. Gill, P. E., Murray, W. & Saunders, M. A. SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization. *SIAM Review* **47**, 99–131 (2005).
34. Zhang, B., Feng, Z., Xu, B. & Yang, T. Efficient Aerodynamic Shape Optimization of the Hypersonic Lifting Body Based on Free Form Deformation Technique. *IEEE Access* **7**, 147991–148003 (Oct. 2019).
35. Spalart, P. & Allmaras, S. *A one-equation turbulence model for aerodynamic flows in 30th Aerospace Sciences Meeting and Exhibit* AIAA 1992-439 (Jan. 1992).
36. Eiseman, P. R., Ebenezer, S. J., Sudharshanam, V. & Anbumani, V. *GridPro: Automatic Multiblock Grid Generation System* 2021. <http://www.gridpro.com>.

37. Kreisselmeier, G. & Steinhauser, R. *Systematic Control Design by Optimizing a Vector Performance Index* in *Computer Aided Design of Control Systems* (ed Cuenod, M. A.) (Pergamon, 1980), 113–117. ISBN: 978-0-08-024488-4.
38. Martins, J. R. A. & Poon, N. M. K. *On Structural Optimization Using Constraint Aggregation* in *Proceedings of the 6th World Congress on Structural and Multidisciplinary Optimization* (Rio de Janeiro, Brazil, May 2005).
39. Kennedy, G. J. & Hicken, J. E. Improved constraint-aggregation methods. *Computer Methods in Applied Mechanics and Engineering* **289**, 332–354. ISSN: 0045-7825 (2015).
40. Anderson, J. *Hypersonic and High-Temperature Gas Dynamics, Second Edition* ISBN: 1563477807 (American Institute of Aeronautics and Astronautics, Nov. 2016).